Ideology Detection for Twitter Users with Heterogeneous Types of Links

Yupeng Gu   Ting Chen   Yizhou Sun
University of California, Los Angeles, CA
{ypgu,tingchen,yzsun}@cs.ucla.edu

Bingyu Wang
Northeastern University, Boston, MA
rainicy@ccs.neu.edu

ABSTRACT

The problem of ideology detection is to study the latent (political) placement for people, which is traditionally studied on politicians according to their voting behaviors. Recently, more and more studies begin to address the ideology detection problem for ordinary users based on their online behaviors that can be captured by social media, e.g., Twitter. As far as we are concerned, however, the vast majority of the existing methods on ideology detection on social media have oversimplified the problem as a binary classification problem (i.e., liberal vs. conservative). Moreover, though social links can play a critical role in deciding one's ideology, most of the existing work ignores the heterogeneous types of links in social media. In this paper we propose to detect *numerical* ideology positions for Twitter users, according to their *follow*, *mention*, and *retweet* links to a selected set of politicians. A unified probabilistic model is proposed that can (1) explain the reasons why links are built among people in terms of their ideology, (2) integrate heterogeneous types of links together in determining people's ideology, and (3) automatically learn the quality of each type of links in deciding one's ideology. Experiments have demonstrated the advantages of our model in terms of both ranking and political leaning classification accuracy. It is shown that (1) using multiple types of links is better than using any single type of links alone to determine one's ideology, and (2) our model is even more superior than baselines when dealing with people that are sparsely linked in one type of links. We also show that the detected ideology for Twitter users aligns with our intuition quite well.

Keywords

Ideology detection, heterogeneous network, social network analysis

1. INTRODUCTION

Ideology detection, i.e., ideal point estimation, dates back to early 1980s, where political scientists first cast light on

ways of revealing politicians' political affiliation using their roll call voting data [22]. Ideal points can be regarded as embeddings of political figures into a latent Euclidean space, and the positions are good representations of their political ideologies. Decades later, the estimation of ideal points has attracted interest from the computer science domain. Researchers begin to use more complicated, statistical models to conduct the estimation on roll call data [7, 8, 12].

Recently, more and more studies pay attention to ideology detection for users on social media, which captures rich information for ordinary citizens in addition to political figures. For example, based on text information, topic models and sentiment analysis approaches are utilized to classify users based on their opinions on different topics [21, 10], and sometimes text features are directly fed into standard classification models such as support vector machines and boosted decision trees to classify users' political affiliations [6, 25, 20]. Based on users' public profile information, such as marriage status, age and income, classifiers are trained to get users' ideology category [1]. Based on link information, various network-based methods are proposed to infer user's ideology [9, 27, 26, 29, 2].

However, there are two major limitations of the existing literature. First, most of these approaches oversimplify the ideology detection problem as a binary classification problem, i.e., classify people as liberal or conservative, ignoring the fact that people's ideology lies in a very broad spectrum. For example, from our estimations in Fig. 1 we can see that Anderson Cooper (@andersoncooper) is almost in the middle and only a little bit to the left, while Rachel Maddow (@maddow) is much more liberal. Second, despite the successful utilization of link information in determining one's ideology, most of the works ignore the heterogeneous link types in social media, which leads to significant information loss. As shown in Fig. 1, by looking only at the *follow* links, there is no way to differentiate the ideology between Rachel Maddow (@maddow) and Megyn Kelly (@megynkelly); however from their *retweet* links, it is quite clear the two have very different ideology. Take Glenn Greenwald (@ggreenwald) as another example: he has equal number of links to Barack Obama (@BarackObama) and Rand Paul (@RandPaul) in our network. Only by knowing *retweet* is a more important link type than *follow*, we are able to predict Greenwald's ideology as liberal. It requires a model that can integrate various link types (i.e., relations) and assign them with different weights to provide a better overall ideology estimation. Admittedly, we may pre-define weights on different link types so as to obtain an integrated weighted graph,

but it either requires strong domain knowledge on defining the link weights, or takes huge amount of time to do multiple rounds of training and select the set of weights with the best held-out performance, which makes it infeasible for real world networks.
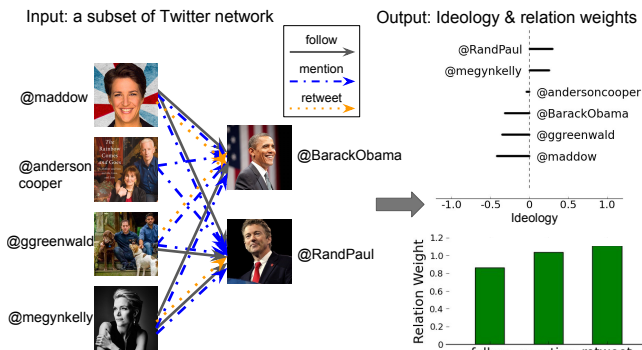
Figure 1: An illustration of ideology detection problem in Twitter network.

In this paper, we propose a unified probabilistic model, *ML-IPM* (Multi-Link Ideal Point Estimation Model), to detect *numerical* ideology positions for a subset of Twitter users, according to their *follow*, *mention*, and *retweet* links. Our approach has several advantages. First, it utilizes social choice theory to explain the reasons why links are built among people in terms of their ideology. Second, it is able to combine heterogeneous types of links in determining people's ideology, with different weights for each link type. Third, the strength of each link type can be automatically learned according to the observed network. Experiments demonstrate the advantages of our model in terms of both ranking and political leaning prediction accuracy. It is shown that (1) using multiple types of links is better than using any single type of links alone to determine one's ideology, and (2) our model is even better than baselines when dealing with people that are sparsely linked in one type of links. We also show that the detected ideology for Twitter users aligns with our intuition quite well. The main contributions of our work are summarized as below:

- We propose to detect *numerical* ideology for Twitter users using link information only. In particular, heterogeneous types of links including *follow*, *mention*, and *retweet* are considered.

- We propose a unified probabilistic model, *ML-IPM*, to solve the problem, which can (1) explain why links are built according to social choice theory and (2) consider different types of links by modeling their strength in determining one's ideology.

- A scalable learning algorithm is provided to learn both the ideology of users and the strength of each link type. Experimental results on a Twitter subnetwork demonstrate the advantages of our model over the state-of-the-art approaches.

2. PRELIMINARIES AND PROBLEM DEFINITION

In this section, we introduce some preliminaries and define our problem.

2.1 Ideal Point / Ideology

Ideal point (also known as political ideology) is a political term which indicates the political leaning of politicians and provides quantitative measures of congressmen across time [22, 23, 24]. It is a real-valued vector which refers to the position on the political spectrum, a real line or a high dimensional space that symbolizes political dimensions. In U.S. people usually refer "left" as democracy or liberality, and "right" as conservatism. Therefore, in one dimensional case, a person whose ideal point is negative (left of the origin) is treated as liberal. The more negative a person's ideal point is, the more liberal he/she is considered to be. Ideology is believed to be an important measure of political behaviors, and ideology estimation is essential in many political science applications.

2.2 Problem Definition

Before the definition of our problem, we first look at some vivid cases in order to motivate our work. Suppose P_1 and P_2 are two political parties in the following examples.

Case 1. User U is a fan of politics and she is extremely interested in political news. Because politicians in P_1 tweet more frequently and U wants to keep track of the latest news and policies, she decides to follow most of the politicians of party P_1 but a small fraction of politicians in P_2. However, most of her retweets come from people in P_2, as she is more happy about their words and opinions. Given the *follow* and *retweet* links from U to these politicians, can we predict which party does she agree with more? More specifically, can we provide U's relative position among all the users that appeared in the network?

Case 2. User V hates P_1, and he will argue with politicians of P_1 by *mentioning* them on Twitter to show his displeasure. Sometimes he also mentions politicians in P_2 to support their policies. As expected, V follows most politicians in P_2, but few politicians in P_1. In this scenario, can we predict that the position of V lies close to the positions of people in P_2 on the political spectrum, based on the *follow* and *mention* links issued by him?

In this work, we focus on the problem of ideology detection for Twitter users. Formally, given a heterogeneous Twitter subnetwork with R types of links ($R = 3$ in our case), our goal is to (1) learn the multidimensional ideology for all users presented in the network and (2) determine the strength of each type of links in the process of ideology detection, which can best explain the observed network. Although defined on Twitter network, our approach is very general to other social networks.

3. APPROACH

In this section, we introduce our solution to the proposed problem. We start from ideology model under a single link type, then introduce how to extend the model when multiple types of links exist, and finally introduce the learning algorithm.

3.1 Ideology Estimation Model via Single Link Type

As in traditional ideal point models, each user has a position in a K-dimensional space, which represents his/her ideology. For a politics-related network, ideology can help explain the reason for link generation, which is a reflection

of people's online behaviors. Take *follow* link as an example: the proximity of two users' positions in the latent ideology space indicates a high probability that they have many politician friends (followees) in common, and vice versa. Because Twitter users are free to follow anyone they like, we assume that each possible *follow* link from u_i to v_j is the outcome of u_i's choice of following v_j or not. Each choice corresponds to a utility function which measures a user's preference over the two choices. Rational users are expected to make the decision that maximizes their utility. Intuitively, the utility function should consider the distance between the positions of two users in the latent ideology space, since people are more likely to follow others whose political opinions are close to themselves. We discuss why a link is formed from both the link-sender and link-receiver's points of view:

The role of link-sender. From a link-sender's (say user u_i) point of view, one has to make a choice on whether to follow user v_j or not, which depends on the two utility functions defined on the two options. She will choose the option that maximizes her utility function, which is related to her ideology position $\boldsymbol{p}_i \in \mathbb{R}^K$.

The role of link-receiver. From a link-receiver's (say user v_j) point of view, each of them is associated with two placements, corresponding to the two options of whether she will receive a link from a link sender or not. We use $\boldsymbol{\psi}_j \in \mathbb{R}^K$ and $\boldsymbol{\zeta}_j \in \mathbb{R}^K$ to denote the positions of the two placements, which can be observed by other users.

We now define the utility functions for a user u_i's choice of (1) "following" a target user v_j ($U_i(\boldsymbol{\psi}_j)$) and (2) not "following" a target user j ($U_i(\boldsymbol{\zeta}_j)$), by extending the model in [5] to the scenario of social networks:

$$
\begin{aligned}
U_i(\boldsymbol{\psi}_j) = u_i(\boldsymbol{\psi}_j) + \eta_{ij}, \text{ if } u_i \text{ follows } v_j \\
U_i(\boldsymbol{\zeta}_j) = u_i(\boldsymbol{\zeta}_j) + \nu_{ij}, \text{ if } u_i \text{ does not follow } v_j
\end{aligned} \tag{1}
$$

where $u_i(\boldsymbol{\psi}_j) = -||\boldsymbol{p}_i - \boldsymbol{\psi}_j||^2$, $u_i(\boldsymbol{\zeta}_j) = -||\boldsymbol{p}_i - \boldsymbol{\zeta}_j||^2$, and η_{ij}, ν_{ij} are random noises drawn from the standard Gumbel distribution, which is commonly used in discrete choice theory [3]. In other words, the two utility functions are defined as the negative of squared distance function between user u_i's ideology position and the positions of the two possible placements of user v_j, plus a random noise.

According to the utility functions, user u_i will follow v_j if $U_i(\boldsymbol{\psi}_j) > U_i(\boldsymbol{\zeta}_j)$. Using the same manipulation in [5], the probability that user u_i decides to follow user v_j is given by

$$
p(u_i \to v_j) = P(U_i(\boldsymbol{\psi}_j) > U_i(\boldsymbol{\zeta}_j)) = \sigma(\boldsymbol{p}_i \cdot \boldsymbol{q}_j + b_j) \tag{2}
$$

where $\boldsymbol{q}_j = 2(\boldsymbol{\psi}_j - \boldsymbol{\zeta}_j)$, $b_j = ||\boldsymbol{\zeta}_j||^2 - ||\boldsymbol{\psi}_j||^2$ and $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function. $\boldsymbol{q}_j$ can be interpreted as the image of user v_j viewed by others, and b_j can be regarded as a bias term for v_j. Usually we can only estimate $\boldsymbol{q}_j$ and b_j instead of $\boldsymbol{\psi}_j$ and $\boldsymbol{\zeta}_j$, and it is adequate since we are interested in the ideology $\boldsymbol{p}_i$ and the probability of a link can be calculated based on $\boldsymbol{q}_j$ and b_j.

After the probability of a directed link between two users is determined, we are able to write down the log-likelihood of observing a network G as

$$
\begin{aligned}
l(G) &= \log \Big(\prod_{(i,j)} \sigma_{ij}^{I_{[i \to j]}} (1 - \sigma_{ij})^{1-I_{[i \to j]}} \Big) \\
&= \sum_{(i,j): i \to j} \log \sigma_{ij} + \sum_{(i,j): i \nrightarrow j} \log(1 - \sigma_{ij})
\end{aligned} \tag{3}
$$

where $I_{[\cdot]}$ is the indicator function, $i \to j$ means there is a link from u_i to v_j, and $i \nrightarrow j$ means there are no links from u_i to v_j. Here σ_{ij} is short for $\sigma(\boldsymbol{p}_i \cdot \boldsymbol{q}_j + b_j)$ in Eq. 2.

Obviously not every pair of people has a link in between. In fact, links in social networks are rather sparse. Therefore we will not go through every possible pair of links in the likelihood above. Instead, we select all the existing links and sample the same number of non-existing links. We denote the set of existing links as S_+ and the sampled set of non-existing links as S_-. In addition, in the case where multiple links may exist between users (e.g., a user can *mention* another user several times), our model can be generalized by simply treating multiple links as multiple occurrences of a single edge. Suppose e_{ij}^+ is the number of links from u_i to v_j and e_{ij}^- is the number of sampled non-existing links from u_i to v_j (typically $e_{ij}^- = 1$), then the log-likelihood above becomes

$$
l(G) = \sum_{(i,j) \in S_+} e_{ij}^+ \cdot \log \sigma_{ij} + \sum_{(i,j) \in S_-} e_{ij}^- \cdot \log(1 - \sigma_{ij}) \tag{4}
$$

3.2 Ideology Estimation Model via Multiple Link Types

We now address the challenge of utilizing multiple types of links for ideology detection. In a heterogeneous network, nodes can be connected via different types of relations. On Twitter, people can *follow*, *mention* or *retweet* others. It naturally forms three different types of links, and different link types certainly have different interpretations.

In order to utilize the knowledge in the heterogeneous network across all kinds of relations, our ultimate model is a weighted combination of models using all types of links. According to our previous assumption, a link-sender u_i will have the same intrinsic ideal points $\boldsymbol{p}_i$ across all link types. However we posit that the roles of link-receivers are different for different relations, as the images of users will change when observed by different types of behaviors. For example, u_i can easily decide to *follow* v_j but hesitates to *retweet* from v_j. Therefore we use the same $\boldsymbol{p}_i$ for user u_i for all types of links, and modify $\boldsymbol{\psi}_j$, $\boldsymbol{\zeta}_j$ to $\boldsymbol{\psi}_j^{(r)}$, $\boldsymbol{\zeta}_j^{(r)}$ which are specific to the relation r. As a result, $\boldsymbol{q}_j$ and b_j will be changed to $\boldsymbol{q}_j^{(r)}$ and $b_j^{(r)}$ as well. In consideration of the heterogeneity in different types of links, we also add a relation weight w_r which represents the relative importance of the links in the corresponding relation type r. Besides, we will use the average log likelihood of each link in each type of relation in order to balance the scale of different link types.

Denoting $\boldsymbol{P} = \{\boldsymbol{p}_i\}, \boldsymbol{Q} = \{\boldsymbol{q}_j^{(r)}\}$ and $\boldsymbol{B} = \{b_j^{(r)}\}$, we define our objective as

$$
\begin{aligned}
&l(\boldsymbol{G}|\boldsymbol{P}, \boldsymbol{Q}, \boldsymbol{B}) \\
&= \sum_{r=1}^{R} w_r \cdot \frac{\sum\limits_{(i,j) \in S_+^{(r)}} e_{r,ij}^+ \log \sigma_{r,ij} + \sum\limits_{(i,j) \in S_-^{(r)}} e_{r,ij}^- \log(1 - \sigma_{r,ij})}{\sum\limits_{(i,j) \in S_+^{(r)}} e_{r,ij}^+ + \sum\limits_{(i,j) \in S_-^{(r)}} e_{r,ij}^-}
\end{aligned}
$$

$$\tag{5}$$

where $e_{r,ij}^+$ is the number of existing links from u_i to v_j via relation r, $e_{r,ij}^-$ is the number of sampled non-existing links from u_i to v_j via relation r and $\sigma_{r,ij} = \sigma(\boldsymbol{p}_i \cdot \boldsymbol{q}_j^{(r)} + b_j^{(r)})$ for short. We will use these notations in the remaining of this section. The constraint we put on w_r is $w_r > 0$, $r = 1, \cdots, R$ and $\prod_{r=1}^{R} w_r = 1$, which forces the geomet-

ric mean[1] of $w_1, \cdots, w_R$ to be 1. In the remaining of the paper we will use the R-dimensional vector $\boldsymbol{w} = (w_1, \cdots, w_R)$ to denote the vector of network weights. For succinctness, we will use N_r to denote $\sum_{(i,j) \in S_+^{(r)}} e_{r,ij}^+ + \sum_{(i,j) \in S_-^{(r)}} e_{r,ij}^-$, i.e., the total number of links in relation r.

Identification. Because the probability in Eq. 2 involves the product of two sets of parameters, there would be an identifiability issue where the objective will be exactly the same if we multiply every $\boldsymbol{p}_i$ by a non-zero constant c and divide every $\boldsymbol{q}_j^{(r)}$ by the same constant. In other words we need to control the scale of model parameters $\boldsymbol{P}$ and $\boldsymbol{Q}$. In our work we add an l_2 regularization term on $\boldsymbol{P}$ and $\boldsymbol{Q}$. We also add the same regularization on $\boldsymbol{B}$ to avoid overfitting. Therefore Eq. (5) becomes

$$
\begin{aligned}
&l(\boldsymbol{G}|\boldsymbol{P},\boldsymbol{Q},\boldsymbol{B}) \\
&= \sum_{r=1}^{R} w_r \cdot \frac{\sum\limits_{(i,j) \in S_+^{(r)}} e_{r,ij}^+ \log \sigma_{r,ij} + \sum\limits_{(i,j) \in S_-^{(r)}} e_{r,ij}^- \log(1 - \sigma_{r,ij})}{N_r} \\
&\quad - \frac{\mu}{2} \left(||\boldsymbol{P}||_F^2 + \sum_{r=1}^{R} ||\boldsymbol{Q}^{(r)}||_F^2 + \sum_{r=1}^{R} ||\boldsymbol{b}^{(r)}||_2^2 \right)
\end{aligned}
$$
$$(6)$$

Here $||\cdot||_F$ denotes the Frobenius norm of a matrix: $||A||_F = \left(\sum_{m,n} A_{mn}^2 \right)^{1/2}$ and $\mu > 0$ is a parameter that controls the effect of regularization terms.

3.3 Optimization

We propose a two-step algorithm to optimize the model parameters. The first step is updating $\boldsymbol{w}$ fixing all other parameters. The second step is updating $\boldsymbol{P}, \boldsymbol{Q}, \boldsymbol{B}$ given $\boldsymbol{w}$. The two steps will run iteratively until convergence.

Update $\boldsymbol{w}$. Fortunately we are able to find a closed form solution for network weights $\{w_r\}_{r=1}^R$. Maximizing the objective function (Eq. 6) w.r.t w_r is equivalent to minimizing

$$
J(\boldsymbol{w}) = \sum_{r=1}^{R} w_r \cdot L_r \quad \text{s.t.} \quad \sum_{r=1}^{R} \log w_r = 0 \tag{7}
$$

where

$$
L_r = -\frac{\sum\limits_{(i,j) \in S_+^{(r)}} e_{r,ij}^+ \log \sigma_{r,ij} + \sum\limits_{(i,j) \in S_-^{(r)}} e_{r,ij}^- \log(1 - \sigma_{r,ij})}{N_r} > 0 \tag{8}
$$

is a constant with respect to w_r, which can be considered as the loss function associated with relation r. We can find the optimal value of w_r by setting the derivative of Lagrangian w.r.t. w_r to zero, which leads to

$$
\lambda = \left(\prod_{r=1}^{R} L_r \right)^{1/R} \tag{9}
$$

and

$$
w_r = \frac{(\prod_{r=1}^{R} L_r)^{1/R}}{L_r}, \quad \forall r = 1, \cdots, R \tag{10}
$$

The value of w_r in each round can be interpreted as the geometric mean of the loss for all link types divided by the loss for link type r. Intuitively, in order to minimize the objective function, we should assign small weights to relations with big loss (L_r), namely smaller average log likelihood. In

---

[1] The geometric mean of several positive variables $w_1, \cdots, w_R$ is given by $(\prod_{r=1}^{R} w_r)^{1/R}$.

other words, we will assign a big weight to a relation that is consistent with the current model (i.e., with a big average log likelihood function), which agrees with our intuition.

Update $\boldsymbol{P}, \boldsymbol{Q}, \boldsymbol{B}$. Given $\boldsymbol{w}$, all other parameters are updated using gradient ascent algorithm.

Time complexity. In each iteration the time complexity of updating $\boldsymbol{w}$ is $O(K \cdot \sum_{r=1}^{R} N_r)$, and updating $\boldsymbol{P}, \boldsymbol{Q}, \boldsymbol{B}$ requires $O(K \cdot \sum_{r=1}^{R} N_r)$operations as well, where K is the dimension of the ideology vector. Therefore the time complexity of our model is linear to the number of edges in the network in particular. Hence, our proposed algorithm is efficient and can be scaled to large networks.

4. EXPERIMENTS

In this section we will show the experimental results based on a subset of Twitter network. We will first introduce how to construct the subnetwork. Then we will compare our method with several baselines, in terms of both ranking accuracy and political leaning classification. Finally, we conduct extensive case studies to show the estimated ideology is meaningful and agrees with our intuition very well.

4.1 Data Preparation

In order to prepare our dataset, we start from politicians and track the users that are connected to them. We first collect the list of all the members of the 113^{th} U.S. congress (2013-2015). We manually search their names on Twitter and match a name to an account if the account is verified and we have enough confidence to show he/she is the corresponding congressman/congresswoman. This step will provide us a list of 487 politicians.

We then use Twitter's streaming API and REST API to collect their followees and followers. In consideration of efficiency, we collect at most 5,000 followers and followees for every congressman. On one hand, in order to select politics-related users, we set a threshold $t = 20$ where we only keep users who follow at least t congressmen or are followed by at least t congressmen. These accounts are likely to be enthusiastic about political issues. On the other hand, we also include around 10,000 random users who follow 3~5 politicians to include more peripheral (less politics-related) Twitter users. Our approach will be evaluated on this Twitter subnetwork with these users as vertices.

There are various reasons apart from ideology why people interact with each other on Twitter, such as they know each other in real life. Therefore, preparing the set of link-receivers is crucial in that the decision to follow them or not should reveal other users' political interests. In order to alleviate the effect of random factors in our model, we restrict the set of link-receivers to be U.S. congress members; in other words, the social links between two ordinary citizens will not be taken into consideration. This strategy allows us to minimize other social factors behind link generation in our model, and it will make our model more accurate.

Here is how we construct the social network for different relations: for *follow* link type, we will add an edge from u_i to v_j to the network if a user u_i follows a politician v_j. For *mention* and *retweet* link types, we collect part[2] of recent tweets for all the users. If u_i posts a tweet which mentions/retweets from a politician v_j, we add an edge from

---

[2] At most 3,200 tweets for each user are available due to API limits.

u_i to v_j to the *mention/retweet* network. If a user does not mention (retweet from) any politician, he/she will be an isolated node and we will simply remove the user from the corresponding network. Note that in contrast to the follow links, a user can mention or retweet another user for multiple times. In this case we simply assume multiple edges exist between them, as discussed in Section 3.1. More details of our dataset can be found in Table 1.

Relation	follow	mention	retweet
Number of links	1,764,956	2,395,813	718,124
Total number of users	46,477	34,775	30,990

Table 1: Statistics for Twitter Dataset

4.2 Performance Evaluation

4.2.1 Baseline Methods

We compare our Multiple Link Types Ideal Point Estimation Model (ML-IPM) with the following baseline methods:

- The simplest baseline where the ideology of a user is the average score of her outgoing neighbors. Each Republican is assigned an ideology score of 1, and each Democrat is assigned a score of -1. We denote this baseline method as AVER.

- The Bayesian Ideal Point Estimation Model (B-IPM) [2] to be introduced in Section 5. Although the author does not mention their generalization to relations other than *follow*, we adopt the model for other types of links for comparison.

- Our Single Link Type Ideal Point Estimation Model (SL-IPM) where only one type of link is present in the social network, as introduced in Section 3.1.

- A special case of our model ML-IPM where the weights for different types of links are fixed (ML-IPM-fixed). In this case the weights for different link types are uniformly distributed, namely $w_1 = w_2 = \cdots = w_R = 1$.

The research problem in [27] is also similar to our task (with single link type), however their approach is not scalable to the size of our dataset (the time complexity is $O(n^2)$ for each iteration, where n is the number of users). Therefore we will not consider their method for comparison.

4.2.2 Evaluation Measures

In our experiments, we will evaluate the ranking and classification accuracy to demonstrate the effectiveness of our ideal point estimation model.

Ranking. In order to evaluate the effect of continuous ideology, we design the ranking evaluation based on 100 manually labeled users, with integer labels from 1 (most liberal) to 5 (most conservative). Here we evaluate the pairwise accuracy between Twitter users, where a pair of users is considered correct if the order of their 1-dimensional ideologies aligns with the order of manual labels (for example, a pair of users (u_a, u_b) where u_a is labeled as "1" and has an ideology of -0.9, and u_b is labeled as "4" and has an ideology of 0.2). We use five different sets of random initialization for the model parameters, and report the mean and standard deviation on a total of 3,857 pairs of users.

Classification. In the classification task, we will classify users as liberal or conservative based on the ideology we have inferred from the dataset. To obtain the ground truth of some users in our dataset, we crawl congress people's party affiliation from the Internet (which is open to public). In addition, we collect the political bias for 100 popular newspaper accounts[3], and we also take advantage of the labeled users in our previous task. These multi-dimensional ideal points are used to train a logistic regression classifier. The classification performance is measured by the Area Under ROC Curve (AUC), and is averaged over 10 different runs by sampling different training data. For all the experiments, we use five different sets of random initialization for the model parameters.

Method	Ranking Accuracy	Classification AUC
AVER (*follow*)	0.427	0.523
AVER (*mention*)	0.446	0.558
AVER (*retweet*)	0.474	0.587
B-IPM (*follow*)	0.443 ± 0.102	0.868 ± 0.021
B-IPM (*mention*)	0.433 ± 0.183	0.558 ± 0.064
B-IPM (*retweet*)	0.501 ± 0.127	0.561 ± 0.066
SL-IPM (*follow*)	0.626 ± 0.011	0.953 ± 0.015
SL-IPM (*mention*)	0.623 ± 0.027	0.951 ± 0.018
SL-IPM (*retweet*)	0.637 ± 0.005	0.958 ± 0.005
ML-IPM-fixed	0.655 ± 0.008	0.930 ± 0.035
ML-IPM	$\mathbf{0.663 \pm 0.007}$	$\mathbf{0.986 \pm 0.013}$

Table 2: Experimental Results

Discussion. We report the ranking accuracy and classification AUC in Table 2, where the relation in the bracket represents the type of link used in the corresponding method. From the above table we can see our advantage over the baseline methods in terms of user ranking and classification accuracy. The comparison on a single link type between our approach (rows 7-9) and existing work (rows 1-3,4-6) illustrates our considerable improvement for both tasks. Nevertheless, information from a single network may not be adequate in determining one's ideology: people may *follow* well-known political figures for reasons other than political proximity, and it is not rare that a user keeps criticizing others by *mentioning* them, or *retweets* ironically. Therefore, integrating heterogeneous types of links becomes necessary, and the results are always better if we automatically update the weight for each link type (rows 10-11). In addition, our method is not sensitive to initialization as the standard deviation is small (row 11). In sum, our learned ideology has a higher quality when we combine heterogeneous types of links together, with learned weights of all link types.

4.2.3 Cold-start Problem

Cold start problem is known as the issue in a system where the inference for users without sufficient information will suffer. Our unified approach is also able to tackle the cold-start problem by integrating knowledge from other types of links: users that are sparsely linked by one type of link may be active in other networks. Thus we are able to have a significantly better understanding of the users by utilizing all types of links.

---

[3]Source: http://www.mondotimes.com/newspapers/usa/usatop100.html

To prepare the candidate users, we select all users whose number of outgoing links is less than or equal to a given threshold in the training dataset, and we use classification AUC (as in previous subsection) and link prediction as evaluation measures. In the link prediction task, we observe some social links in the network and predict a set of unobserved links of the same type. In the experiments, we also use five different initializations and 90% of the links (both existing and non-existing) are chosen uniformly at random as training. For the baseline method AVER, we simply define the probability of a link using a modified version of Eq. 2 without bias. In Fig. 2 we compare our unified ML-IPM model (in solid red line) with (1) our model where weights are fixed (in regular dashed green line); (2) SL-IPM (in dotted blue line); (3) B-IPM (in dash-dotted black line); and (4) AVER (in bold dashed magenta line). We can see our unified approach is the best among all baseline methods, especially when the users are extremely sparsely linked.

4.2.4 Parameter Study

Here we study how ideology dimension K affects our model. The effect of K is shown in Table 3. Our method is not sensitive to the choice of K, and the best K should be around 5. For above approaches and experiments, we use $K = 5$ as the dimension of ideology, if not mentioned otherwise.

Dimension K	1	3	5	7
Classification AUC	0.979	0.981	**0.986**	0.977

Table 3: Effect of ideology dimension

4.3 Case Studies

We will analyze the relation weights $\boldsymbol{w} = \{w_r\}$ and visualize Twitter users' ideology in this subsection.

It is interesting to study the relative weights $\boldsymbol{w}$ for different types of links. Intuitively, a larger w_r implies the relation r is more important in our task. The automatically learned weights are shown in Table 4. We can observe the weight for *retweet* is the highest among all types of links, while the weight for *follow* is the lowest. This is reasonable as following someone on social network is common and does not stand for a very close relationship. For the other two relations, a retweet is likely to show the endorsement of the author of the original tweet; on the other hand, people are free to mention anybody and this behavior might indicate either pleasant or hostility. Therefore *mention* behaviors may not indicate a strong homophily as *retweet* does. In sum, a larger weight represents the significance of the corresponding source of information, and for the specific task of detecting Twitter users' ideology based on their social links, *retweet* is the most crucial relation among all kinds of relations.

Relation r	follow	mention	retweet
Weight w_r	0.866	1.035	1.117

Table 4: Weights of different link types

To further demonstrate the meaning of our learned weights, we randomly add 1,000,000 links between pairs of Twitter user and politician to our dataset. We denote this type of link as "random" and apply our model to the heterogeneous network with four relations (*follow*, *mention*, *retweet* and

random). As a result, the weights of four types of links become 1.02, 1.23, 1.33, 0.59, respectively. The *random* relation has the least importance score, and the trend of other three relations agrees with our previous result. This result shows our model is robust against noise and the learned weights reflect the importance of relations quite well.
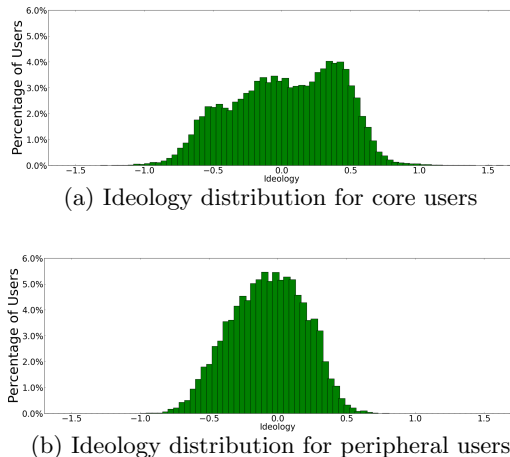

(a) Ideology distribution for core users

(b) Ideology distribution for peripheral users

Figure 4: Ideology Distribution

In the remaining of this section we will visualize the latent ideology position of Twitter users. For visualization purposes we use $K = 1$ in all figures below. Fig. 4 shows the distribution of ideology for all Twitter users in our dataset. We can observe three peaks around $x = -0.5$, $x = 0$ and $x = 0.5$ in Fig. 4(a), which represent the three clusters of liberal, neutral and conservative users respectively. Most of the users with less political interests (Fig. 4(b)) will stay in the middle and form only one cluster. The statistics are reasonable because the distribution of peripheral users is more similar to the public, who are distributed more uniformly along the political spectrum. It is noteworthy that Twitter users are a biased sample of the public: they may be younger and more active on the Internet. Therefore skewness in our detected ideology may exist.

Twitter users may also indicate their locations on their personal profiles. We collect all users in our dataset who claim themselves to live in one of the 50 states in the U.S. (or Washington, D.C.), and calculate the average ideology for each area. Then we are able to map the average score to a color between red and blue. As a result, 9,362 users are identified and 29 states are labeled as red (conservative), as shown in Fig. 5. Note that our colored map may differ from political polls or election results, given the population of Twitter users and the sample variance in the dataset.

In Fig. 6, we also visualize the ideology for some popular media, magazine, newspaper accounts and a few political figures including president Barack Obama, Hillary Clinton and Donald Trump, the two presidential candidates in 2016[4] and so on. The length of each line (x-axis) represents the value of our detected ideology. The text on the y-axis denotes their Twitter screen name, and the color represents the political leaning provided by Mondo Times[5]: red rep-

---

[4]The ideology for Hillary Clinton and Donald Trump is estimated based on their most recent tweets and friends.
[5]http://www.mondotimes.com/, a worldwide news media
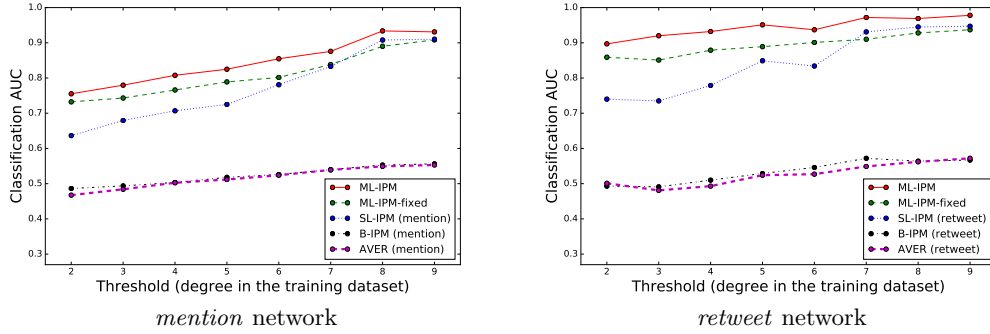
Figure 2: Classification AUC for sparsely linked users in different types of networks.
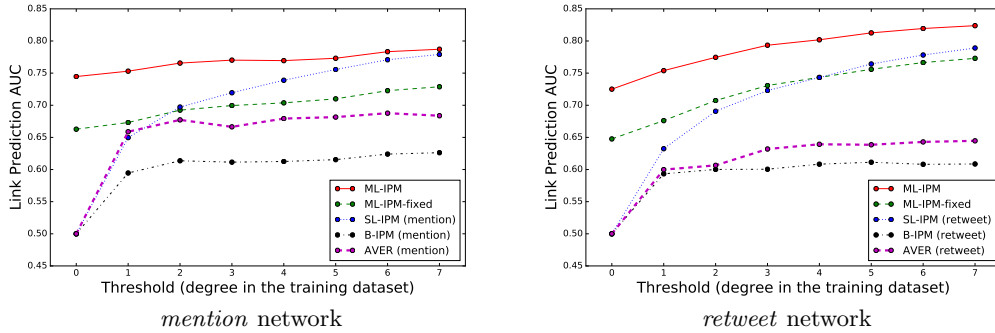

Figure 3: Link prediction AUC for sparsely linked users in different types of networks.
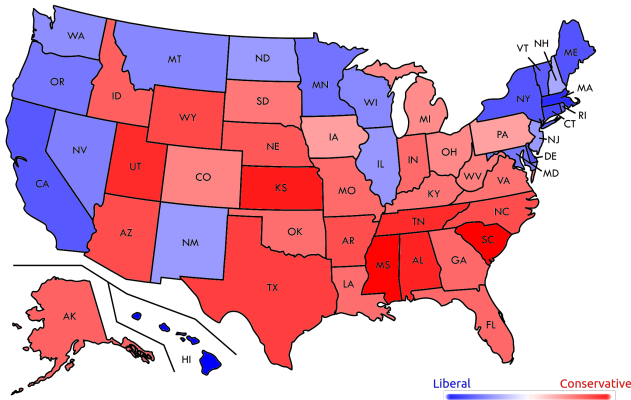

Figure 5: Average ideology for Twitter users in each state. Darker red means more conservative, while darker blue means more liberal.

resents a conservative account; blue stands for a liberal account; and black represents an account with no bias. From this vivid example we can see the position of those accounts on a unidimensional spectrum, and we are able to compare them with politicians. Again, there may be slight difference between the detected ideology and the actual political leaning, as the ideology is detected according to the behavior of these Twitter accounts instead of the editorial content.

directory.

It is also worth analyzing the obvious exceptions in the results. Opposite to the label provided by Mondo Times, Wall Street Journal (WSJ) is considered as a media source leaning to the left in our study. When we investigate the social links of WSJ in our dataset, we find its links to politicians from two parties are balanced. Besides, our estimation refers to only the behaviors of their Twitter account (i.e. whom they follow/mention/retweet) instead of the editorial contents of the journal itself. This mainly explains the reason why our algorithm does not label it as a conservative account. Interestingly, WSJ is shown to be one of the most liberal online media outlets in some other studies [11, 27, 17].

5. RELATED WORK

5.1 Ideology Detection in Roll Call Voting Data

Ideal point models attempt to estimate the position of each lawmaker in the latent political space. Legislative voting is one of the sources for quantitative estimation of lawmakers' ideal points. Poole and Rosenthal [22] were among the first few researchers in political science domain to provide a thorough and rigorous approach for ideology estimation. Afterwards, political scientists have proposed numerous methods to infer political ideology from roll call voting records [23, 13, 24, 16, 15, 5]. Researchers study the public voting record of lawmakers and model the probability of each vote, which is usually described as the interaction of the lawmaker's ideal point and the position of the bill. Along this line of research, computer science researchers extend the ideal point model to a variety of aspects, including applying
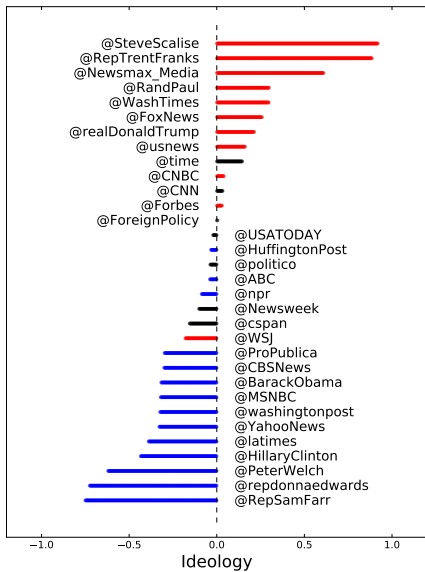
Figure 6: One dimensional ideology for selected users. Red: conservative; Blue: liberal; Black: no bias. Color labels are provided by Mondo Times.

natural language processing and topic modeling techniques on bills [7, 8, 12, 14, 19].

5.2 Ideology Detection in Social Networks

Apart from voting data, recently some approaches have been using information from social networks to analyze user's political leaning. These approaches can be further grouped into two categories according to the task they are dealing with: classification/clustering and continuous ideology estimation. Classification and clustering tasks predict which political party a user belongs to (or congregate users into several clusters), and can be handled from various aspects. When users interact with each other via messages, especially when they have discussions or arguments, methods from natural language processing can be applied to analyze the topics and people's sentiment. Pla and Hurtado [21] analyze users' tweets on Twitter, extract features such as hashtags and punctuations, and determine the polarity of a tweet, and a user's political leaning is simply a weighted average of her tweets. Boutet et al. [4] analyze the characteristics of three main parties in UK election, and predict which party a user supports by the amount of Twitter messages referring to political accounts. Based on debate records on various political topics, Gottipati et al. [10] adopt a matrix factorization approach to discover users' attitudes towards different political issues, and use k-means on the user feature vector to congregate users into several clusters. Some researchers extract text features such as hashtags and latent semantic analysis of a conversion and feed them into a standard classifier in order to classify users [25, 6, 20]. Users' public profile information (such as marriage status, age et al.) can also be utilized as inputs of classification models [1]. However, apart from classification tasks, it is also important to make inference about the continuous ideal points, which can be applied to rank ordering and demonstrate the relative difference between people in the same political party. Since a continuous

measure of ideology is often desired for real world applications and our goal in the paper is to detect real-valued ideal points for users in social networks, those classification and clustering methods cannot be directly applied in our task.

With respect to the work on estimating continuous political ideology, researchers have been mining the rich information in network structures to obtain an accurate ideology score for users on social networks. Typically, inference of a user's ideal point is made by exploring her neighbors and her relationship with labeled users (such as politicians whose party affiliation is clearly known to the public). Therefore, a simple yet intuitive approach would be calculating the ratio of the number of Democrats and Republicans that a user follows on the social network [9]. Wong et al. [27, 26] assume liberal people tend to tweet more about liberal events and the same for conservative users. Given the political leaning of social events by labeling associated tweets, they are able to infer the ideology of users who have engaged. Barberá [2] proposes a probabilistic model to describe the likelihood of the social network, in which the probability to observe a link between two users is defined as a function of their ideal points. Although most of the existing methods have utilized the homophily phenomenon [18, 28] in network generation and made appropriate assumptions of the network structure, they neglect the fact that most social networks contain multiple types of interactions between people. Those methods fail to define how close two nodes are when different link types are present. Although experts can pre-define weights for different link types, it is almost impossible to exhaustively list all possible combinations of link weights given the scale of real world data. Moreover, experiments show that the performance of existing approaches will suffer when a user is sparsely linked to others. Our approach, on the other hand, is able to automatically learn the weights for different types of links according to the network structure and determine how important each link type is in an ideology detection task. The cold start problem is also overcome by transferring information learned from one type of link to other link types.

6. CONCLUSION AND FUTURE WORK

In this paper we present a novel approach for ideology detection on Twitter using heterogeneous types of links. Instead of predicting binary party affiliations to users, our work focuses on detecting the continuous ideal points for Twitter users, which is more comprehensive. In addition, we improve over traditional ideology estimation models by integrating information from heterogeneous link types in social networks. Specifically, our model is able to automatically update the importance scores of various relations on Twitter, and these weights are incorporated in the unified framework to achieve an outstanding performance. Moreover, our algorithm has an even better performance for sparsely linked users in separate networks. Finally, we evaluate our model on a subnetwork of Twitter, and the results show our advantage over the baseline methods. Extensive case studies also demonstrate our model's alignment with human intuitions.

One limitation of this work is that we only take links into consideration when deciding one's ideology. In the future, we plan to integrate other content information, such as text, into the current framework, to better understand people's ideology.

7. REFERENCES

[1] F. Al Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. *ICWSM*, 270, 2012.

[2] P. Barberá. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Analysis*, 23(1):76–91, 2015.

[3] M. E. Ben-Akiva and S. R. Lerman. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press, 1985.

[4] A. Boutet, H. Kim, E. Yoneki, et al. What's in your tweets? i know who you supported in the uk 2010 general election. In *ICWSM*, 2012.

[5] J. Clinton, S. Jackman, and D. Rivers. The statistical analysis of roll call data. *American Political Science Review*, 98(02):355–370, 2004.

[6] M. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of twitter users. In *Privacy, Security, Risk and Trust and 2011 IEEE Third Int. Conf. on Social Computing*, pages 192–199. IEEE, 2011.

[7] S. Gerrish and D. M. Blei. Predicting legislative roll calls from text. In *Proc. of the 28th Int. Conf. on Machine Learning (ICML'11)*, pages 489–496, 2011.

[8] S. Gerrish and D. M. Blei. How they vote: Issue-adjusted models of legislative behavior. In *Advances in Neural Information Processing Systems (NIPS'12)*, pages 2762–2770, 2012.

[9] J. Golbeck and D. Hansen. Computing political preference among twitter followers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1105–1108. ACM, 2011.

[10] S. Gottipati, M. Qiu, L. Yang, F. Zhu, and J. Jiang. Predicting user's political party using ideological stances. In *Social Informatics*, pages 177–191. Springer, 2013.

[11] T. Groseclose and J. Milyo. A measure of media bias. *The Quarterly Journal of Economics*, pages 1191–1237, 2005.

[12] Y. Gu, Y. Sun, N. Jiang, B. Wang, and T. Chen. Topic-factorized ideal point estimation model for legislative voting network. In *Proceedings of the 20th Int. Conf. on Knowledge discovery and data mining (KDD'14)*, pages 183–192, 2014.

[13] J. J. Heckman and J. M. Snyder Jr. Linear probability models of the demand for attributes with an empirical application to estimating the preferences of legislators. Technical report, National Bureau of Economic Research, 1996.

[14] M. Iyyer, P. Enns, J. Boyd-Graber, and P. Resnik. Political ideology detection using recursive neural networks. In *Association for Computational Linguistics*, 2014.

[15] S. Jackman. Multidimensional analysis of roll call data via bayesian simulation: identification, estimation, inference, and model checking. *Political Analysis*, 9(3):227–241, 2001.

[16] J. Londregan. Estimating legislators' preferred points. *Political Analysis*, 8(1):35–56, 1999.

[17] J. R. Lott Jr and K. A. Hassett. Is newspaper coverage of economic events politically biased? *Public Choice*, 160(1-2):65–108, 2014.

[18] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pages 415–444, 2001.

[19] V.-A. Nguyen, J. Boyd-Graber, P. Resnik, and K. Miler. Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress. In *Proc. of ACL*, 2015.

[20] M. Pennacchiotti and A.-M. Popescu. A machine learning approach to twitter user classification. *ICWSM*, 11:281–288, 2011.

[21] F. Pla and L. Hurtado. Political tendency identification in twitter using sentiment analysis techniques. In *Proceedings of COLING*, pages 183–192, 2014.

[22] K. T. Poole and H. Rosenthal. A spatial model for legislative roll call analysis. *American Journal of Political Science*, pages 357–384, 1985.

[23] K. T. Poole and H. Rosenthal. Patterns of congressional voting. *American Journal of Political Science*, 35(1):pp. 228–278, 1991.

[24] K. T. Poole and H. Rosenthal. *Congress: A Political-Economic History of Roll Call Voting*. Oxford University Press, 1997.

[25] D. Rao, D. Yarowsky, A. Shreevats, and M. Gupta. Classifying latent user attributes in twitter. In *Proc. of the 2nd Int. workshop on Search and mining user-generated contents*, pages 37–44. ACM, 2010.

[26] F. Wong, C. W. Tan, S. Sen, and M. Chiang. Media, pundits and the us presidential election: Quantifying political leanings from tweets. In *Proc. of the Int. Conf. on Weblogs and Social Media*, 2013.

[27] F. M. F. Wong, C. W. Tan, S. Sen, and M. Chiang. Quantifying political leaning from tweets and retweets. In *ICWSM*, 2013.

[28] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *Proc. of the 20th Int. Conf. on World wide web*, pages 705–714. ACM, 2011.

[29] D. X. Zhou, P. Resnick, and Q. Mei. Classifying the political leaning of news articles and users from user votes. In *ICWSM*, 2011.