Polarity Detection for Twitter Users

Group Name: BHE6 Group Number: 5

Junheng Hao
805025633
jhao@cs.ucla.edu

Song Jiang
605222804
songjiang@cs.ucla.com

Tao Meng
705225623
tmeng@cs.ucla.edu

Yewen Wang
905229899
wyw10804@cs.ucla.edu

Zhiping Xiao
604775684
patricia.xiao@cs.ucla.edu

ABSTRACT

In this project, we study the problem of political polarity detection, which aims to predict a person's political leaning according to his/her behaviors on social network. We designed a conditional random field model that can classify a person into two partisans: democratic and republican, conducted experiment on a politicians' twitter dataset released in 2016 and compared with existing baselines, and also collected a new twitter dataset with more complete information. Our experiment results indicate our model is effective.

Keywords

polarity detection, conditinal random filed, social network analysis

1. INTRODUCTION

The polarity detection task in this project is to detect a person's political leanings, i.e., to predict which political partisan a person belongs to. This problem is mainly studied on politicians, traditional ways usually modeling this problem according to politicians' voting behavior, but as the booming of social network, we can capture the politicians' behaviors on social media and modeling the polarity detection problem based on the obtained information.

There are some existing work dealing with this polarity detection problem, traditional methods mainly takes the person's voting behaviors into consideration[7], and some latest models pay attention to people's polarity tendency on social network[5]. Since the second types of model is more powerful, we follow its idea and study the polarity detection problem according to people's online behaviors that captured by social media.

In this project, we propose a conditional random field model to detect the political polarity of twitter users, according to their online relationships on twitter including follow and be followed, and their online behaviors on twitter including mention and be mentioned, and retweets and be-retweeted. In our model, we regard each person as a node and use the follow and be followed relationship to construct a directed graph, then we apply the behaviors as attributes of each node. We verified our model's effectiveness on a twitter dataset of the members of 113th(2013-2015) U.S. congress, and also collects a more completed and advanced twitter dataset that can be applied in political polarity detection task.

Our main contributions are as follows:

1. We designed a conditionally-trained, generative directed graphical model, that can effectively detect a person's political polarity

2. We collected a dataset consisting of 115th U.S. congress (2017-2019) members' behaviors on twitter

The rest parts of this report is developed as follows: Section 2 presents the formulated problem, section 3 introduces related work and lists several baseline methods, section 4 illustrates our designed model, section 5 shows our data preparation, experiment settings and results, section 6 concludes the whole project and listed several future work directions.

2. PROBLEM FORMALIZATION

With timeliness and breadth, social medias plays a vital role in recent political ecology. Politicians share their opinions on social media, which can spread to the public within a short time period affect peoples' voting behaviours [5, 7, 3, 8, 9]. For example, Donald Trump shares his comments on Twitter, which helps people understand him deeper. Ko Wen-je and Han Kok-ju, two new politic stars, present both their politics and lives on Facebook. Their activities on such social medias attract millions of supports, and thus bring them successes in the mayor election. Therefore, detecting the polarity of users on social media is an essential issue to help politicians to understand their situations.

In this project, we investigate polarity detection problem via predicting a Twitter user's party preference, i.e., Republican or Democratic, based on his/her retweet, follow and mention behaviours. Formally, given a Twitter user Z_j with his/her behaviours $B(Z_j)$, our goal is to infer the polarity Z_j. However, the complicated and heterogeneous relations between Twitter users makes it non-trivial to model them in a unified way. To tackle this challenge, we propose a

DOI:

novel graphic probabilistic method to detect people's polarity via their behaviours on social media. For each user, our method first models and captures his/her complicated behaviours and relations on Twitter through a joint probability distribution. Afterwards, our model infers this user's polarity via optimizing such distribution. Although defined on Twitter network, our approach is very general to other social networks.

3. RELATED WORK & BASELINE METHODS

3.1 Related Work

Recently some mining approaches have been using information from social media (or social networks) to analyze users' political leaning. Classification/clustering have been widely applied in ideology detection, which predict the political party a user belongs to or favors (or congregate users into several clusters) based on user features. Such features includes personal information (age, marriage, etc), text (hashtags, punctuations, comments) and connections (actions such as like, retweeting, following, etc). Previous research work has shown effectiveness by adopting weighted user features [6, 1] and matrix factorization [4]. 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 methods cannot be directly applied in our task.

Another branch of research work focuses on estimating continuous political ideology. Researchers have been mining the rich information in the social network structures together with tweet labeling to obtain an accurate ideology score for users on social networks. There are some observations and assumptions on ideology detection across the social networks, such as the fact that one user may be positively affected and related by its friends or followees and retweeted text, which are likely to represent his/her political ideology. Barberaá [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. Gu [5] extends the task on heterogeneous social network with different types of links by using on embedding based methods and adaptive link weight in social networks. These two approaches serve as state-of-the-art methods for ideology detection, which will be explained in methodology.

3.2 Baseline methods

Following [5], we consider the following methods as baselines and current state-of-the-art methods.

AVER The simplest baseline where the ideology of a user is the average score of his/her outgoing neighbors. Each Republican is assigned an ideology score of 1, and each Democrat is assigned a score of -1.

B-IPM [2] Bayesian Spatial Following model (also named Bayesian Ideal Point Estimation in the paper) that considers ideology as a latent variable, whose value can be inferred by examining which politics actors each user is following. The key assumption of this model is that Twitter users prefer to follow politicians whose position on the latent ideological dimension are similar to theirs. The probability that user i follows a political account j is then formulated as a logit model,

$$p(e_{ij} = 1|\alpha_j, \beta_i, \gamma, \theta_i, \phi_j) = \text{logit}^{-1}\left(\alpha_j + \beta_i - \gamma||\theta_i - \phi_j||^2\right) \tag{1}$$

where α_j measures the popularity of user j and β_i measures the political interest of user i. $\theta_i, \phi_j \in \mathbb{R}$ is the ideal point of user i and j, which means that the decision is a function of the squared Euclidean distance in the latent ideological dimension between user i and j (γ is a normalizing constant). The likelihood function $p(y|\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \gamma)$ is maximized in B-IPM. Though Maximum-likelihood estimation methods are usually intractable given the large number of parameters involved, Barberaá use a Hamiltonian Monte Carlo algorithm and employ a hierarchical setup that considers each of the four sets of parameters as draws from four common population distributions: $\alpha_j \sim \mathbf{N}(\mu_\alpha, \nu_\alpha)$, $\beta_i \sim \mathbf{N}(\mu_\beta, \nu_\beta)$, $\theta_i \sim \mathbf{N}(\mu_\theta, \nu_\theta)$ and $\phi_j \sim \mathbf{N}(\mu_\phi, \nu_\phi)$. Briefly, B-IPM adopts both user features (in/out degrees) and latent ideological variable to model the probability for user connections (follow relations).

SL-IPM [5] In the Ideology Estimation Model via Single Link Type Model (denoted as SL-IPM), each user has a position in a K-dimensional space, which represents his/her ideology features. The binary status (presence/absence) of a social link (we assume that only one type of link is involved) from u_i to u_j is modeled as a Bernoulli event as,

$$p(e_{ij} = 1) = \sigma\left(\mathbf{p}_i \cdot \mathbf{q}_j + b_j\right) \tag{2}$$

where $\mathbf{p}_i \in \mathbb{R}^K$ and $\mathbf{q}_j \in \mathbb{R}^K$ are user embeddings, b_j is bias(popularity) term and $\sigma(\cdot)$ is sigmoid function. Thus, the log-likelihood of observing the whole network G is,

$$\log p(G) = \sum_{e_{ij}=1} \log p(e_{ij} = 1) + \sum_{e_{ij} \in S_-} \log(1 - p(e_{ij} = 1)) \tag{3}$$

where S_- is negative sample sets to accelerate the computation due to the large numbers of negative sampling space.

ML-IPM [5] Based on SL-IPM model as embeddings for users, Ideology Estimation Model via Multiple Link Types (denoted as ML-IPM), adopts multiple link types in the networks ($r = 1, 2, \ldots, R$). For example, in Twitter networks, we have `retweet`, `follow` and `mention` relations between users. Equation 2 in ML-IPM settings changes to

$$p(e_{ij}^{(r)} = 1) = \sigma\left(\mathbf{p}_i \cdot \mathbf{q}_j^{(r)} + b_j^{(r)}\right) \tag{4}$$

where $\mathbf{q}_j^{(r)}$ and $b_j^{(r)}$ become relation-specific. Combining all types of link in the networks, the objective function of ML-IPM is,

$$\begin{aligned} \mathcal{J}(\mathbf{P}, \mathbf{Q}, \mathbf{B}, \mathbf{w}) = \sum_{r=1}^{R} w_r \cdot \Bigg(&\sum_{e_{ij}^{(r)}=1} \log p(e_{ij}^{(r)} = 1) \\ &+ \sum_{e_{ij}^{(r)} \in S_-^{(r)}} \log\left(1 - p(e_{ij}^{(r)} = 1)\right) \Bigg) \\ \text{s.t.} \left(\Pi_{r=1}^{R} w_r\right)^{1/R} = 1 & \end{aligned} \tag{5}$$

where $\mathbf{w}$ are weight parameters that can be learned iteratively with other optimizing variables $\mathbf{P}, \mathbf{Q}, \mathbf{B}$.

4. METHODOLOGY

4.1 Model Formalization

We have three kinds of information in the data set: retweet, follow and mention. Previous work like [5] treat them equivalently and assigned different weights, however, in this project, we believe they should not be treated equivalently.

In fact, every user has connections or relations in real world social network. Twitter is a social media, and it can somehow reflect this network. One can easily figure out that follow is the most fundamental thing to define such a relation in Twitter, and retweet and mention are more like common behaviours. So we will use follow information to construct the graph: if there is a link between user a and user b, we make a link between a and b.

Since we regrad the retweet and mention as behaviours, we should focus on the behaviours with polarity information. In the dataset we do not have the text data for the retweet and mention, so we cannot do text classification to learn if a behaviour has polarity information. Noting that congressmen are a group of people close related to politics, we believe the behaviours related to congressmen are more likely to convey polarity information.

When we only consider the behaviours related to congressmen, we have 8 classes of behavious: a user {retweet, mention, is retweeted by, is mentioned by} a congressman in $\{D, R\}$. We have construct the graph according to the follow information, so here we only care about a congressman is in D or R rather than which congressman the user is related to. We will treat the behaviours in the same class equivalently, and totally we will have 8 different behaviours.

The logic of these things is like: the polarity of a user is determined by the social relations, and the behaviours of a user is affected by the polarity. So what we want is trying to predict the users' polarity according to the observed behaviours, to solve which we should use conditional probability to model that. Thus, we choose conditional random field (CRF) model.

We denote $\mathbf{B} = \{B_1, B_2, ..., B_N\}$ as the set of observed behaviours related to congressmen, and B_i is the behaviours set of user i. Each of the behaviour is one of the 8 kinds of defined behaviour. $\mathbf{z} = \{z_1, z_2, ..., z_N\}$ is the polarity of every user and $z_i \in \{0, 1\}$. $E = \{(i, j) \mid i \text{ follows } j \lor j \text{ follows } i\}$ is the edge set. The conditional probability is given by

$$P(\mathbf{z}|\mathbf{B}, \Theta) \propto \left(\prod_{(i,j)\in E} f(z_i, z_j) \right) \prod_{i=1}^{N} \left(\prod_{b \in B_i} b(z_i) \right),$$

with parametric potential functions

$$\Theta = (\lambda, \mu),$$

$$f(z_i, z_j) = \exp(\lambda \mathbf{1}[z_i = z_j]),$$

$$b(z_i) = \exp\left(\frac{\mu_b z_i}{|B_i|} \right).$$

4.2 Training And Prediction

In training, the objective is set to be maximize the marginal conditional probability with respect to congressmen (labeled users):

$$\Theta^* = \arg\max_{\Theta} P(\mathbf{z_l}|\mathbf{B}, \Theta),$$

where $\mathbf{Z_l}$ is the labels of labeled users. We can training our model by standard CRF marginal inference algorithm, namely, variable elimination algorithm to get the marginal probability. We then apply gradient descent to train the parameters.

In prediction, we do the maximum a posteriori (MAP) inference to get the prediction:

$$\mathbf{z}_u^* = \arg\max_{\mathbf{z}_u} P(\mathbf{z}_u|\mathbf{B}, \mathbf{z}_l, \Theta).$$

The MAP inference problem for general CRF can be solved by loopy belief propagation algorithm.

4.3 Evaluation

Our model can predict the polarity of unlabeled users, but we do not have a method to evaluate the performance of our model. One way to deal with this problem is split the congressmen (labeled users) into training set and development set.

We treat the congressmen in development set exactly same as the unlabeled users: we do not provide their label in prediction, and also we do not use their label to define observed behaviours. We will use the accuracy of development set to evaluate the model.

The baseline models we will use predict the congressmen labels with all the other information except for the label of this current user. This is exactly the situation that we have only one congressman in the development set and all the others in training set. After prediction the baseline models compute the overall accuracy or AUC, which is same as we do a $k-$fold cross validation when k is the number of congressmen. But we don't do this not only because our model is heavy and cannot afford such a large k, but also our model is much more robust that we do not need so much information to predict the the labels. So we will use the $k-$fold cross validation to evaluate our model with some common value of k.

5. EXPERIMENT

5.1 Data Plan

5.1.1 Data Preparation

Twitter has a tremendous amount of data, for which it provides an API that is freely available to all developers (refer to https://developer.twitter.com/).

Considering the deadline's conflict with the rate limit, we might use Yupeng Gu's 2014 dataset put online (http://web.cs.ucla.edu/˜ypgu/) if the latest dataset is not ready in time.

Yupeng's dataset consists of three kinds of binary relations that naturally exist in Twitter, namely: friend, mention, and retweet. However, over the past few years, Twitter has updated a lot. Examining the latest API we found five kinds of binary relations in total: friend, retweet, mention, reply, favorite.

If strictly following Yupeng's data-colletion strategy, the data is collected by:

1. Collecting a list of congressman, together with their Twitter accounts and their party, as labeled data;

2. Crawling at most N accounts those politicians are following on Twitter, and at most N ($N = 5000$ in Yupeng's paper) most recent followers of each congress man;

3. Among all the unlabeled users, keep those who follow at least t congressman or at least followed by t congressman ($t = 20$ in Yupeng's dataset, turns out to be approximately $30,000$ valid users under this category), as core users who are enthusiastic about political issues;

4. In addition, also include M additional users who only follow 3 to 5 congressman ($M = 10,000$ in Yupeng dataset), as peripheral users;

5. Crawl the most recent Tweets of those labeled and unlabeled users, analyze those Tweets to collect the relation of *retweet*, and the relation of *mention*, among the users in the list (labeled, and selected unlabeled ones).

If we are to use the newly crawled data, the last step should also include extracting *favorites* and *reply* information.

Furthermore, although this dataset didn't include text data, we could actually crawl the text data of that period if needed, once we figure out at what time he crawled those data (to make them match).

5.1.2 Data-Crawling Tools

There's a useful tool, Tweepy (https://tweepy.readthedocs.io/en/v3.5.0/api.html), which is used by the crawlers of both Yupeng's and ours.

The only challenging part of data collection is the speed issue. Nowadays Twitter is more and more strict on crawlers. According to some blogs and forum discussions left on the internet years ago, we are pretty sure that back then Twitter allows several crawls per second. But now, it only allows 1 crawl per minute in average (refer to https://developer.twitter.com/en/docs/basics/rate-limiting.html), which is quite annoying.

Getting more tokens would be a good idea, but unfortunately Twitter has some restrictions on it, like, it is hard to parallel and crawl a user's full history using multiple token's, and, one phone number could be tied to no more than 3 accounts, each account could have no more than 3 new tokens a day, etc.

Yupeng used Tweepy for his dataset collection. The Tweepy API is in general the best API of Tweeter-crawling purpose we've found so far. However, Twitter's policies, such as the rate limits and time limits, changes a lot during the past few years, and thus we need to keep alert to all potential hazards in the online references.

In March, 2019, Twitter's tricky rules in regards with crawling Tweet contents include:

- Once the rate-limit is hit, waiting for 15 minutes is necessary, and if keep on trying, any unsuccessful tryout will be counted by Twitter, and thus further slower down the speed;

- If an account has status "protected", then crawling its Tweets will end up in endless waiting time;

- For Tweets (or in other words, user-timeline API visit), rate-limit is quite casual, and is seldom hit if the crawlers take 3 to 5 seconds break;

- For a user's favorite Tweets, strictest rate-limit applies;

- Only the most recent $3,200$ Tweets of a user is available for crawling, and we have no access to earlier Tweets; but for favorite, seems that there's no such limits.

Some of the rules are well-delivered in the official documentation, while some others are either not obvious, or simply being there without any announcement. But all of them are equally important in terms of data-collection.

5.1.3 Dataset

Since Yupeng released his dataset, we could use it for our project. Another advantage of using his dataset is that, it'll make our comparison with his model's performance much easier, since his dataset was used for Twitter Ideology analysis task as well. It also provides many baseline results from the same dataset.

5.1.4 Text data analysis

We also did some preliminary exploration into the text content of the Tweets from congressman in both parties, and concluded some interesting findings, include:

1. Generally speaking, a Democratic congressman posts shorter Tweets than a Republican congressman. The longest Tweet from Republican side is approximately twice the length of the upper-bound of the Democratic side's posts.

2. Democratic congressman uses emoji more than the Republican group, and more often do we see them use multiple emoji icons together in a sentence.

3. Generally speaking, Democratic party is more "active" on Twitter, in terms of, for example, more vivid ways of expression, and that many Democratic congressman would comment on the same Tweet or same article, multiple times, successively; while the Republicans in general have a longer way of expression and less amount of comments on one single issue. (Of course exemptions are everywhere, such as Trump.)

4. The Democratic party and the Republican party are using totally different vocabularies [1]. Each side has approximately $190,000$ different words, with only approximately $75,000$ words appeared on both sides.

After doing a simple word2vec and alignment, we have some further conclusion on the two parties' features:

1. Most words' embedding aligned well, which means that in fact their expressions aren't such different.

2. The top word that couldn't be aligned well is *usphs*, which is the abbreviation of United States Public Health Service. This, we suppose, reflects the major point that the two parties couldn't agree with each other: medical issues.

[1] Here we refer to vocabulary in total, including emoji, etc.

3. However, even if they disagree with each other on *usphs*, the closest words in their embedding, measured by cosine similarity, are almost the same.

5.2 Experiment setting

As mentioned above, we do not have the groundtruth of users U, thus we use the evaluation method mentioned in Section 4.3. In our experiment, we select $k = 5$.

5.3 Results

We report the result of experiments in table 1 . The relation in the bracket represents link type in each corresponding method. From the result table we can observe that our proposed model outperforms baseline methods in the binary classification task. We could also see that the comparison on a single graph of our proposed method and existing methods. The results illustrate the advantage of our proposed model.

Table 1: Experimental Results

Method	AUC
AVER(follow)	52.3
AVER(mention)	55.8
AVER(retweet)	58.7
B-IPM(follow)	86.8
B-IPM(mention)	55.8
B-IPM(retweet)	56.1
Proposed Model	**91.1**

Table 2: Task distribution

Task	Member
1.Data Preprocessing	Patricia Xiao
2. Model Design	Tao Meng, Yewen Wang
3. Model Implementation	Song Jiang
4. Experiments	Junheng Hao, Song Jiang
5. Presentation	Tao Meng, Yewen Wang
6. Writing Report	All

5.4 Case Study

Although in this experiment we do not have the labels of users, we still could analyze their polarity via there actions on Twitter. Table 3 shows four examples we select from all the users. Based on their actions on Twitter (i.e., mention, follow and retweet), we can infer each user's polarity manually and compare the tendency with result of our proposed model. Note that in this case study, all the four users do not hold retweet by Democratic congressman, so we remove this. As we can observe in table 3, user_12,user_29283,user_813286 has some interactions with Democratic congressman while few interactions with Republican congressman. On the other hand, user_610793 are more concern about Republican than Democratic. Our proposed model could capture both this two kinds of users and make corresponding predictions.

6. CONCLUSION

In summary, we designed a conditionally-trained, generative directed graphical model, that can effectively detect a person's political polarity, and collected a dataset consisting of 115th U.S. congress (2017-2019) members' behaviors on twitter. One of the limitation of our work is we simply consider the behaviors information. Furthermore, we could include text information in the future. From the preliminary analysis of the newly-crawled text data, we observed different features in posts from the two parties, as is mentioned before. This piece of information is expected to help with more accurate classification.

7. REFERENCES

[1] F. Al Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. In *Sixth International AAAI Conference on Weblogs and Social Media*, 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. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of twitter users. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 192–199. IEEE, 2011.

[4] S. Gottipati, M. Qiu, L. Yang, F. Zhu, and J. Jiang. Predicting userâĂŹs political party using ideological stances. In *International Conference on Social Informatics*, pages 177–191. Springer, 2013.

[5] Y. Gu, T. Chen, Y. Sun, and B. Wang. Ideology detection for twitter users with heterogeneous types of links. *arXiv preprint arXiv:1612.08207*, 2016.

[6] F. Pla and L.-F. Hurtado. Political tendency identification in twitter using sentiment analysis techniques. In *Proceedings of COLING 2014, the 25th international conference on computational linguistics: Technical Papers*, pages 183–192, 2014.

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

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

[9] F. M. F. Wong, C. W. Tan, S. Sen, and M. Chiang. Quantifying political leaning from tweets, retweets, and retweeters. *IEEE transactions on knowledge and data engineering*, 28(8):2158–2172, 2016.

8. APPENDIX TEAMWORK

The task distribution of our teamwork is in Table 2

Table 3: Actions and predictions of four selected users

User_id	Mention_D	MentionBy_D	Mention_R	MentionBy_D	Retweet_D	Retweet_R	RetweetBy_R	Prediction
12	6	0	0	0	13	0	0	Democratic
29283	27	0	5	0	4	2	0	Democratic
610793	0	0	11	0	2	29	11	Republican
813286	3	78	0	68	1	1	0	Democratic