

# Political Issue Extraction Model: A Novel Hierarchical Topic Model That Uses Tweets By Political And Non-Political Authors

Aditya Joshi<sup>1,2,3</sup> Pushpak Bhattacharyya<sup>1</sup> Mark Carman<sup>2</sup>

<sup>1</sup>IIT Bombay, India

<sup>2</sup>Monash University, Australia

<sup>3</sup>IITB-Monash Research Academy, India

{adityaj, pb}@cse.iitb.ac.in, mark.carman@monash.edu

## Abstract

People often use social media to discuss opinions, including political ones. We refer to relevant topics in these discussions as political issues, and the alternate stands towards these topics as political positions. We present a Political Issue Extraction (PIE) model that is capable of discovering political issues and positions from an unlabeled dataset of tweets. A strength of this model is that it uses twitter timelines of political and non-political authors, and affiliation information of only political authors. The model estimates word-specific distributions (that denote political issues and positions) and hierarchical author/group-specific distributions (that show how these issues divide people). Our experiments using a dataset of 2.4 million tweets from the US show that this model effectively captures the desired properties (with respect to words and groups) of political discussions. We also evaluate the two components of the model by experimenting with: (a) Use to alternate strategies to classify words, and (b) Value addition due to incorporation of group membership information. Estimated distributions are then used to predict political affiliation with 68% accuracy.

## 1 Introduction

Political discussions in social media contain contentious topics (called ‘**political issues**’), and alternate stands with respect to these issues (called ‘**positions**’). We present a topic model that discovers political issues and positions in tweets. Our model is called **Political Issue Extraction (PIE)**

**model.** The input is the twitter timelines of authors (*i.e.*, the user who created the tweet), and political affiliation information for a subset of authors (*i.e.*, political authors. Antonym: Non-political authors). Political and non-political authors contribute to formation of topics, whereas only political authors contribute to position that a group is likely to take. Since our dataset consists of tweets from the US, political affiliation can be one of the three groups: ‘*Democrats*’, ‘*Republicans*’ or ‘*Unknown*’.

For every tweet, we estimate two latent variables: issue and position. To discover topics related to issues and positions, we classify words in a tweet in three categories: issue words, position words and emoticons. Instead of document-specific distributions as in LDA, we include a hierarchy of author-specific and group-specific position distributions in our model. This hierarchy estimates three distributions for each topic: global position, position of a given political group and position of a specific author.

We evaluate our model by (a) validating our topics against standard topic lists, (b) considering different strategies of splitting words into the three categories, and (c) validating how the model benefits from the group information. Finally, we use our model to predict political affiliation of authors.

Models based on LDA by Jo and Oh (2011), Mei et al. (2007); Lin and He (2009) extract sentiment-coherent topics. Past work related to political opinion has been reported by Gayo-Avello, Metaxas, and Mustafaraj (2011); Conover et al. (2011); O’Connor et al. (2010); Wong et al. (2013); Yano, Cohen, and Smith (2009); Lin, Xing, and Hauptmann (2008);

Wang, Mohanty, and McCallum (2005), and more recently by Benton et al. (2016). Two of them close to our work are by Grimmer (2010); Fang et al. (2012). Our model improves upon them in three ways:

1. In PIE model, position words depend on both issue and position latent variables (as opposed to only the latter in prior work),
2. In PIE model, a novel hierarchical author/group-wise distribution is considered instead of document-wise distribution.
3. To the best of our knowledge, PIE model is the first that operates at the author level by using complete author timelines of both political and non-political authors, and affiliation information of a subset of authors (only political authors).

The rest of the paper is organized as follows. We present the structure and estimation procedure of PIE model in Section 3 and discuss our experiment setup in Section 4. The evaluation is in Section 5. Finally, we conclude and point to future work in Section 6.

## 2 Related Work

Our model is based on Latent Dirichlet Allocation (LDA) given by Blei, Ng, and Jordan (2003). Models based on LDA by Jo and Oh (2011), Mei et al. (2007), Lin and He (2009), and Zhao et al. (2010) present approaches to extract sentiment-coherent topics in datasets.

The past work in analytics related to political opinion can be broadly classified in three categories. The first category predicts the outcome of an election. Gayo-Avello, Metaxas, and Mustafaraj (2011) predict the election outcome for a pair of Democrat and Republican candidates, while Metaxas, Mustafaraj, and Gayo-Avello (2011) aggregate sentiment in tweets in order to map it to votes. Conover et al. (2011) use a network graph of authors and their known political orientations, in order to predict the political orientation. The second category of work in the political domain deals with correlation of sentiment with real-world events. O'Connor et al. (2010) correlate sentiment expressed in text to time series of real-world events. Wong et al. (2013) derive

the consistency between tweeting and retweeting behaviour and real-world sentiment about an event. Gerrish and Blei (2011) present the ideal point topic model that correlates votes by legislators to bills being passed along with sentiment in the texts of these bills.

Our PIE model falls in the third category: extraction of political topics from a dataset. In this respect, the work closest to ours is by Grimmer (2010) and Fang et al. (2012). Grimmer (2010) present a hierarchical topic model to understand how senators explain their work in their press releases. Assuming a single topic per press release, the topic of a document is derived from an author-specific distribution over topics. Fang et al. (2012) divide words in a political statement as topic words and opinion words, based on POS tags, and assume different distributions for the two. Like these two works, we assume a single topic per tweet, and divide words into categories based on POS tags. Our model improves upon these two works in the following key ways:

- A richer latent variable structure. In our model, the opinion words depend on BOTH topic and opinion latent variables (as opposed to only the latter). This structure allows our model to generate topics corresponding to political positions, which is not achieved in the past work.
- The author-topic distribution and hierarchy of author-sentiment distributions (as noted in Section 1). In case of our models, authors are arranged into groups and this group-wise distribution is tightly linked to the structure of the model.

## 3 PIE Model

In this section, we introduce the PIE model. We first discuss the rationale behind the design. We then describe the structure of the model. Following that, we present details of the input/output and the method for estimation of distributions.

### 3.1 Design Rationale

The primary goal of the model is to discover topics related to political issues and two positions per issue. To be able to discover issues and corresponding positions, PIE model considers two latent variables

$p$  and  $i$ . Topic  $i$  represents the identifier for a political issue while the pair  $p$ - $i$  represents the identifier for a position.

The first component of the model is derived from the **nature of data**. A tweet contains two kinds of content words: “objective” words that describe an entity and “subjective” words that express opinion towards an entity. Emoticons can be thought of as a third category of sentiment words. We represent these three kinds of words as three observed variables: topic words  $t$ , opinion words  $o$  and emoticons  $e$ . A topic word  $t$  is derived from the topic  $i$ , an opinion word is derived from the topic  $i$  and sentiment  $p$  pair while an emoticon  $e$  is derived from sentiment  $p$  alone.

The second component of the model is derived from the **nature of the problem**. In a political setting, people are organized in political groups while the opinion of a group towards a topic is a result of individual users in the group. To model this nature, we use a hierarchical structure that relates global sentiment, group-wise sentiment and author-specific sentiment. Our model contains a hierarchy of distributions  $\psi_i$ ,  $\psi_{ig}$  and  $\psi_{iu}$  indicating global, group-wise and author-specific sentiment respectively, towards political issue indicated by  $z$  for a user  $u$  who belongs to group  $g$ .

### 3.2 Structure

The model is shown in Figure 1. The **input** is a dataset of  $T$  tweets, each annotated with one among  $U$  authors(/users of twitter). Also, each author has exactly one out of  $G$  political affiliations. We represent a tweet as a collection of three types of words: (a) *Issue words* that describe a political issue, (b) *Position words* that express a position towards the issue, and (c) *Emoticons*. We do not consider hashtags as a special case because of Cohen and Ruths (2013) who show that hashtags are not strong indicators of political orientation. Two latent variables are defined for each tweet: issue  $i$  and position  $p$ . **The topics corresponding to issue  $i$  represent political issues, while topics corresponding to pairs of issue  $i$  and position  $p$  represent political positions.** The assumption that a tweet has exactly one issue and position is reasonable due to limited length of tweets. The arrows from  $i$  and  $p$  that lead to  $t$ ,  $o$  and  $e$  realize the role of the three categories of content

terms as follows:

1. The topic words  $t$  describe a political issue and hence, are based only on the topic  $i$ .
2. The opinion words  $o$  express position towards a political issue and hence, are based on both the issue  $i$  and position  $p$ .
3. The sentiment of an emoticon  $e$  does not depend on the issue being talked about and hence, is based only on position  $p$ .

The model estimates two sets of distributions: one for words and another for the user groups. The three categories of words lead to three distributions that are estimated by the model: (a) **Issue word-issue distribution**  $\vec{\eta}_i$ , (b) **Position word-issue-position distribution**  $\vec{\phi}_{ip}$ , and (c) **Emoticon-position distribution**  $\vec{\chi}_e$ . Since an emoticon may not completely belong to either of the political positions, we do not rely on a mapping to a position but consider a distribution of emoticons over positions. This incorporates more intricate forms of opinion expression like sarcasm. In addition to these term-specific distributions, the PIE model estimates author/group-specific distributions: (a) **Author-Issue distributions**:  $\vec{\theta}_g$  is the probability of an issue with respect to a group  $g$ ,  $\vec{\theta}_u$  is the probability of an issue with respect to an author  $u$ , (b) **Author-position distributions**:  $\vec{\psi}_i$  is the probability of a position with respect to an issue  $i$ ,  $\vec{\psi}_{ig}$  is the probability of a position with respect to an issue  $i$  and group  $g$ , while  $\vec{\psi}_{iu}$  is the probability of a position with respect to issue  $i$  and author  $u$ . Variables and distributions in the model are in Table 1.

The generative process of the corpus can be described as follows<sup>1</sup>:

<sup>1</sup>‘Dir’ in the generative process denotes a Dirichlet prior.

1. For each issue  $i$ , select  $\vec{\eta}_i \sim \text{Dir}(\gamma)$ , and  $\vec{\psi}_i \sim \text{Dir}(\beta_1)$
2. For each position  $p$  select  $\vec{\phi}_p \sim \text{Dir}(\delta_1)$ ,  $\vec{\phi}_{ip} \sim \text{Dir}(\delta_2 \vec{\phi}_p)$ , and  $\vec{\chi}_p \sim \text{Dir}(\epsilon)$
3. For each group  $g$  select  $\vec{\theta}_g \sim \text{Dir}(\alpha_1)$ , and  $\vec{\psi}_{ig} \sim \text{Dir}(\beta_2 \vec{\psi}_i)$
4. For each author  $u$  select  $\vec{\theta}_u \sim \text{Dir}(\alpha_2 \vec{\theta}_g)$ , and  $\vec{\psi}_{iu} \sim \text{Dir}(\beta_3 \vec{\psi}_{ig})$
5. For each tweet  $k$  select
  - (a) topic  $i_k \sim \vec{\theta}_{u_k}$  and sentiment  $p_k \sim \vec{\psi}_{i_k, u_k}$
  - (b) all topic words,  $t_{kj} \sim \vec{\eta}_{i_k}$
  - (c) all opinion words,  $o_{kj} \sim \vec{\phi}_{i_k, p_k}$
  - (d) all emoticons,  $e_{kj} \sim \vec{\chi}_{p_k}$

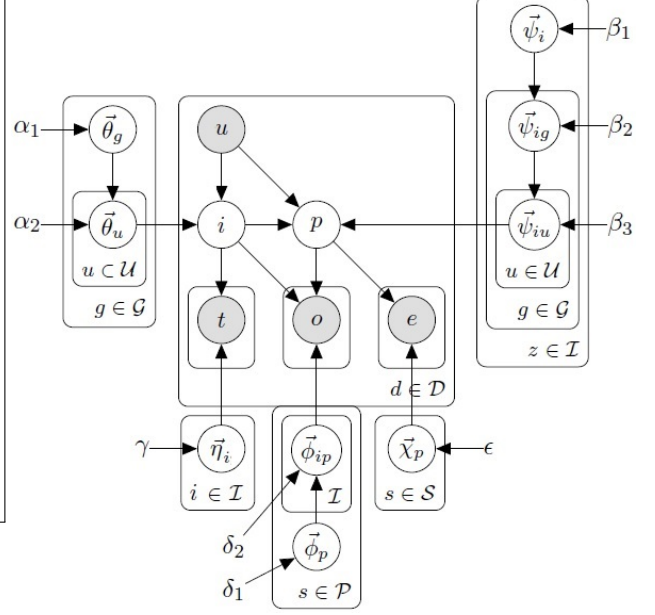


Figure 1: PIE Model: Plate Diagram

### 3.3 Estimation

In the PIE Model, we need to estimate word-topic distributions namely  $\vec{\eta}_i$ ,  $\vec{\chi}_e$  and  $\vec{\phi}_o$  and author-specific distributions  $\vec{\theta}_u$  and  $\vec{\psi}_{i,u}$ . The estimation of the joint probability distribution is computationally intractable. Hence, we use Gibbs sampling by Casella and George (1992) to estimate the underlying distributions. For computational efficiency, we use moment matching for estimation. The sampling algorithm runs for a pre-determined number of iterations. We implement a block sampler based on Heinrich (2005) that samples  $p_k$  and  $i_k$  of the  $k^{\text{th}}$  tweet together and results in faster convergence. The joint probability of  $p_k$  and  $i_k$  is given by:

$$P(p_k, i_k | u_k, \vec{p}_{-k}, \vec{i}_{-k}) \propto \theta_{i_k | u_k} \psi_{p_k | i_k, u_k} \left( \prod_j \eta_{t_{kj} | i_k} \right) \left( \prod_j \phi_{o_{kj} | p_k, i_k} \right) \left( \prod_j \chi_{e_{kj} | p_k} \right)$$

where all parameters  $\theta, \psi, \eta, \phi$  and  $\chi$  are estimated withholding information regarding the previous assignment to the  $k^{\text{th}}$  tweet. The generative story is omitted in the current version, due to lack of space.

The word-specific distributions are estimated as:

$$\hat{\eta}_{t|i} = \frac{N_{t,i}^{(n)} + \gamma \frac{1}{V^{(n)}}}{N_i^{(n)} + \gamma}$$

$$\hat{\theta}_{i|u} = \frac{N_{i,u}^{(i)} + \alpha_2 \theta_{i|g(u)}}{N_u^{(i)} + \alpha_2}, \quad \hat{\theta}_{i|g} = \frac{N_{i,g}^{(i)} + \alpha_1 \frac{1}{V^{(i)}}}{N_g^{(i)} + \alpha_1}$$

$$\hat{\phi}_{o|p,i} = \frac{N_{o,p,i}^{(o)} + \delta_2 \phi_{o|p}}{N_{p,i}^{(o)} + \delta_2}, \quad \hat{\phi}_{o|p} = \frac{N_{o,p}^{(o)} + \delta_1 \frac{1}{V^{(o)}}}{N_p^{(o)} + \delta_1}$$

Random Variables	
$u, g$	Author of a tweet and Group of the author
$i, p$	Issue and position of a tweet
$t, o$	Issue/Position-word in a tweet
$e$	Emoticon in a tweet
Distributions	
$\vec{\theta}_{u/g}$	Dist. over issues for author $u$ / group $g$
$\vec{\psi}_{i,u/g}$	Dist. over positions for issue $i$ and author/group
$\vec{\eta}_i$	Dist. over topic-words for issue $i$
$\vec{\phi}_{i,p}$	Dist. over opinion-words for issue-position pair
$\vec{\chi}_p$	Dist. over emoticons for position $p$
Hyper-parameters	
$\alpha, \beta$	Concentration par. issue/position dist.
$\gamma$	Concentration par. for issue-word dist.
$\delta$	Concentration par. for position-word dist.
$\epsilon$	Concentration par. for emoticon dist.
Counts	
$N_{t,i}^{(t)}$	Frequency of topic-word $t$ in tweets for topic $i$
$N_{i,u}^{(i)}$	Frequency of tweets on topic $i$ by author $u$
$V^{(t)}$	Vocabulary size for topic words

Table 1: Glossary of Variables/Distributions used

$$\hat{\chi}_{e|p} = \frac{N_{e,p}^{(e)} + \epsilon \frac{1}{V^{(e)}}}{N_p^{(e)} + \epsilon}$$

where the count notation can be read as follows:

$N_{e,p}^{(e)}$  denotes the number of times emoticon  $e$  occurs within tweets assigned to position  $p$  across the corpus and  $V^{(e)}$  is the size of the emoticon vocabulary. The equations show that for speed and ease of implementation, we use a simple approximation to the group-wide Dirichlet mean parameter  $\vec{\phi}_{o|p}$  rather than estimating expected table counts within a Chinese Restaurant Process as given by Griffiths and Tenenbaum (2004). (We leave an investigation of more precise parameter estimation to future work.)

The author/group-specific distributions are estimated in a hierarchical manner as follows:

$$\hat{\theta}_{i|u} = \frac{N_{i,u}^{(i)} + \alpha_2 \theta_{i|g(u)}}{N_u^{(i)} + \alpha_2}, \quad \hat{\theta}_{i|g} = \frac{N_{i,g}^{(i)} + \alpha_1 \frac{1}{V^{(i)}}}{N_g^{(i)} + \alpha_1}$$

$$\hat{\psi}_{p|i,u} = \frac{N_{p,i,u}^{(p)} + \beta_3 \psi_{p|i,g(u)}}{N_{i,u}^{(p)} + \beta_3}, \text{ where :}$$

$$\hat{\psi}_{p|i,g} = \frac{N_{p,i,g}^{(p)} + \beta_2 \psi_{p|i}}{N_{i,g}^{(p)} + \beta_2}, \quad \hat{\psi}_{p|i} = \frac{N_{p,i}^{(p)} + \beta_1 \frac{1}{V^{(p)}}}{N_i^{(p)} + \beta_1}$$

The notation here is the same as for the word-issue distributions, except that the counts are now at the “tweet-level” rather than the “word-level”, i.e.  $N_{i,u}^{(i)}$  indicates number of tweets by author  $u$  assigned the topic  $i$ . Note again the use of simple estimates for the group-wide parameters  $\vec{\theta}_g$  and  $\vec{\psi}_{z,g}$ .

## 4 Experiment Setup

We create a dataset of tweets using Twitter API (<https://dev.twitter.com/>). The authors whose timelines will be downloaded are obtained as follows. We first obtain a list of famous Democrats and Republicans using sources like about.com, The Guardian and Fanpagelist. This results in a list of 32 Republicans and 46 Democrats. We expand this list by adding randomly selected friends of these twitter handles. (The choice of “friends” as opposed to “followers” is intentional.) We then download complete twitter timelines of all authors (Twitter sets the upper limit to 3200 tweets). The resultant dataset consists of 2441058 tweets. Dirichlet hyperparameters and values of  $I=35$  and  $P=2$  are experimentally determined. We set priors on position words using a word list of 6789 words given by McAuley and Leskovec (2013). Function words and 25 most frequent words are removed.

Segregation strategy	Coherence
POS-based	<b>0.468</b>
POS-based+PMI Collocns.	0.436
Subjectivity-based	0.451
POS+Subjectivity-based	0.457

**Table 2:** Average topic coherence per topic for different strategies of word segregation

## 5 Evaluation

To validate the efficacy of our model, our evaluation addresses the following questions:

- What impact do components of the model have, on its ability to discover these issues and positions? (*Section 5.1*)
- What political issues and positions does the model discover? (*Section 5.2*)
- Once we discovered political issues, positions and group-wise distribution, can the model be used to predict political affiliation? (*Section 5.3*)

### 5.1 Impact of Model Components on Performance

We evaluate two key components of PIE model, namely, segregation of words and hierarchy of author-group distributions.

**Segregation of words:** A key component of PIE model is the strategy to decide whether a word is an issue word or position word. We experiment with following alternatives to do this: (a) **POS-based segregation** as done in Fang et al. (2012) using twitter POS tagger Bontcheva et al. (2013). We experimentally determine the optimal split as: nouns as issue words, and adjectives, verbs and adverbs as position words, (b) **POS-based+PMI-based collocation handler** to include n-grams, using Bird (2006), (c) **Subjectivity-based segregation** classifies words present in the subjectivity word list by McAuley and Leskovec (2013) as position words, (d) **POS+Subjectivity-based segregation** where we first categorize nouns as issue words, and then look for other words in the subjectivity word list. In order to select the best strategy of segregation, we compute topic coherence metric  $Cv$  using Palmetto by Röder, Both, and Hinneburg

	Average cosine similarity		$\Delta$
	With	Without	
<b>Within members of the same group</b>			
Demo.-Demo.	<b>0.261</b>	0.253	0.01
Repub.-Repub.	<b>0.014</b>	0.013	0.001
<b>Within members of different groups</b>			
Demo.-Repub.	<b>0.108</b>	0.113	-0.005
Repub.-Demo.	<b>0.040</b>	0.042	-0.002

**Table 3:** Average cosine similarity between author-position distributions, with and without group membership information

(2015) for all topics. This metric uses normalized PMI. Average coherence per topic is shown in Table 2. We observe the highest value of 0.468 in case of the POS-based strategy. The remaining subsections report results for our experiments with this strategy.

**Efficacy of author-group distributions:** To evaluate the benefit of our hierarchy of distributions, we obtain average cosine similarity between author-position distributions of the two political groups. For every author with known political affiliation, we first obtain the cosine similarity between the  $\psi_{iu}$  of the author and the  $\psi_{iu}$  of other authors belonging to his/her own political group. This is then averaged over all authors. This value indicates how different/similar authors of a affiliation are. Table 3 shows these values for different combinations. The columns indicate two scenarios: when political affiliation information is used (‘with’) and when it is not used (‘without’) during the estimation. The rows indicate the four possible scenarios. The average cosine similarities are not symmetric, by design. The sign on  $\Delta$  shows that incorporation of political affiliation information makes authors of the same group more similar to each other, and authors of different groups less similar to each other, as desired.

## 5.2 Qualitative Evaluation

The issues extracted from our model are represented by the topics formed using topic words. We list some of the topics extracted using PIE model in Table 4. Each cell contains top 5 words of each topic with a manually assigned description in boldface. These topics are the political issues underlying our dataset. The issues discovered are “*health in-*

<b>Insurance</b>	<b>Gun laws</b>	<b>Crime</b>
health insurance people care plan	people gun laws guns control	scene police man fire suspect
<b>Abortion</b>	<b>Security/War</b>	<b>Employment</b>
abortion baby babies freedom women	attack video security police forces	workers job wages jobs people
<b>Immigration</b>	<b>Economy</b>	<b>Climate</b>
workers immigration stories patriot politics	tax jobs debt taxes spending	climate people change warming years
<b>Marriage</b>	<b>Election</b>	<b>Disasters</b>
people freedom marriage rights women	bill vote campaign state election	acres fire weather snow storm

**Table 4:** Top words in Political Issues discovered by PIE

*urance, abortion, security, employment, gun laws, immigration, economy, climate, marriage, election, disasters, crime and government”*. In addition to these, topics beyond political issues are also observed, as expected. These include **sports** (*game, team, season, year, football*), **promotional online content** (*{blog, showcase, article, courtesy, support}* or *{photo, photos, video, entry, album}*), etc. Manual identification of political issues from the set of retrieved topics is necessary because approaches like considering top-k probable topics may not work. For example, social media concepts such as followers or promotional online content occur more frequently than immigration and abortion. We validate that **our political issues appear in at least one out of three online lists of issues** from Gallup.com, About.com and Ontheissues.com in Table 5.

The alternate positions that people take, are shown in Table 6. Each box consists of a political issue written in boldface and top five words in topics corresponding to alternate sentiment. These topics show what we mean by “*alternate*” positions and that they are not merely positive or negative. In case of the political issue “*abortion*”, the con-

	A	Gallup	O		A	G	O		A	G	O
Abortion		✓	✓	Government		✓		Economy	✓	✓	
Climate	✓	✓		Gun Laws	✓		✓	Election			
Crime			✓	Immigration		✓	✓	Employment		✓	✓
Disasters	✓			Insurance		✓	✓	Security/War	✓		✓
Marriage	✓	✓	✓								

**Table 5:** Comparison of Our Political Issues with three Online Lists of Political Issues from About.com (A), Gallup (G) and OnTheIssues (O)


Abortion		Security/War	
Join	Prolife	Killed	Military
Religious	Killed	Syrian	Illegal
Stand	Born	Military	Russian
Support	Unborn	Fast	Targeting
Conservative	Aborted	Furious	Back
Gun laws		Immigration	
Illegal	Dont	Join	Top
Free	Free	Support	Enter
Dont	Stop	Back	Check
Vote	Illegal	Stand	Stop
Stop	Give	Proud	Join
Insurance		Marriage	
Pay	Check	Back	Gay
Federal	Hear	Don	Religious
Signed	Here	Lost	Political
Paid	Call	Liberal	Free
Uninsured	Hope	Great	**

**Table 6:** Top words in Political Positions Discovered by PIE; \*\* is a popular twitter handle

trasting positions correspond to topics given by *join*, *religious*, *stand*, *support*, *conservative* and *prolife*, *killed*, *born*, *unborn*, *aborted*. It can be seen that the first position gives a religious view whereas the second position presents an emotional appeal with respect to abortion. Similarly, consider the box for “immigration”. One position corresponds to “support” and “stand” by immigrants while the opposing position corresponds to “check” or “stop” immigration. In case of “insurance”, authors are divided into ones who talk about “paying” and the ones who see “hope” in revised insurance policies. Finally, look at the box corresponding to “gun laws”. Both positions contain topics with negative words but differ in a way that one position talks about a “vote” while the opposite position mentions “give”.

Figure 2 shows the absolute difference between the  $P(p|i, g)$  for a topic-party pair. The observations

are intuitive since the issues with **least difference are employment and disasters** while the most contentious are *abortion, election and immigration*.



**Figure 2:** Difference between Political Positions

Approach	Accuracy (%)
Baseline: <i>Gottipati et al. (2013)</i>	60
Log likelihood-based	68

**Table 7:** Comparison of PIE model with past approaches for prediction of Political Affiliation

### 5.3 Application: Prediction of Political Affiliation

Obtaining a set of non-political authors with reliable political affiliation is challenging. We first select authors who were labeled as ‘Unknown’. This means that PIE model did not know about their political affiliation during training. Among these authors, we select ones who have mentioned their political affiliation in their profile description. This results in 25 test authors (out of which 6 are Democrats). We consider two approaches to predict political affiliation:

1. **Baseline:** This baseline is similar to *Gottipati et al. (2013)* except that the vectors in our case

are based on PIE model. We calculate cosine similarities between estimated author-position distribution  $\psi_{iu}$  for test authors and each of the group-position distributions  $\psi_{ig}$  for the two political groups. The predicted affiliation is the group with greater similarity value.

2. **Log likelihood-based:** We use the distribution for words and groups, that have been computed during training. For each test author, we again run our estimation on their tweets twice, once for each group. The goal is to learn  $\psi_{iu}$  and  $\theta_u$  and compute two log-likelihood values, once for each group, and predict the more likely affiliation.

Table 7 compares our approach with a past approach. The log likelihood approach results in the best accuracy of 68%.

## 6 Conclusion & Future Work

In this paper, we presented a Political Issue Extraction (PIE) model to discover political issues and two positions per issue, using a dataset of tweets by US politicians and civilians. Our PIE model represented a tweet as three sets of words: topic words, opinion words and emoticons. To model author-specific distributions, we considered a hierarchical set of distributions.

To evaluate PIE model, we compared multiple strategies to classify words into three categories and showed that POS-based classification gives highest topic coherence. Our model was able to identify: a) topics corresponding to political issues, b) alternate positions that the two parties may take, and c) the issues that are likely to be the most “contentious”. We estimated twelve political issues (such as security, disasters, immigration, etc.) and positions within each. Using cosine similarity within groups, we showed that our model placed members of the same group closer to each other than the ones from the other group, when group information was provided. Our PIE model discovers that abortion, immigration and marriage are among the most contentious political issues. Finally, we also presented findings of a pilot study to predict political affiliation of authors using PIE model, and achieved an accuracy of 68%.

As future work, we wish to be able to automatically identify which of the topics extracted from

our model are political issues. Although the current model can be used for extraction more than two positions in principle, we would like to see if any additional challenges come up in that case. This model can be mapped to identification of controversial topics, brand loyalty, etc.

## References

- About.com. Political Issues, url = <http://uspolitics.about.com/od/electionissues/>.
- Benton, A.; Paul, M. J.; Hancock, B.; and Dredze, M. 2016. Collective supervision of topic models for predicting surveys with social media. In *Proceedings of AAAI Conference*.
- Bird, S. 2006. Nltk: the natural language toolkit. In *Proceedings of the COLING-ACL on Interactive presentation sessions*, 69–72. Association for Computational Linguistics.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *the Journal of machine Learning research* 3:993–1022.
- Bontcheva, K.; Derczynski, L.; Funk, A.; Greenwood, M. A.; Maynard, D.; and Aswani, N. 2013. TwitIE: An open-source information extraction pipeline for microblog text. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Association for Computational Linguistics.
- Casella, G., and George, E. I. 1992. Explaining the gibbs sampler. *The American Statistician* 46(3):167–174.
- Cohen, R., and Ruths, D. 2013. Classifying political orientation on twitter: It’s not easy! In *ICWSM*.
- Conover, M. D.; Gonçalves, B.; Ratkiewicz, J.; Flammini, A.; and Menczer, F. 2011. Predicting the political alignment of twitter users. In *Privacy, security, risk and trust, 2011 IEEE third international conference on social computing*, 192–199. IEEE.
- Fang, Y.; Si, L.; Somasundaram, N.; and Yu, Z. 2012. Mining contrastive opinions on political texts using cross-perspective topic model. In *Proceedings of the fifth ACM international conference on Web search and data mining*, 63–72. ACM.



- Gallup.com. Jobs, Government, and Economy remain top US problems, url = <http://www.gallup.com/poll/169289/jobs-government-economy-remain-top-problems.aspx>.
- Gayo-Avello, D.; Metaxas, P. T.; and Mustafaraj, E. 2011. Limits of electoral predictions using twitter. In *ICWSM*.
- Gerrish, S., and Blei, D. M. 2011. Predicting legislative roll calls from text. In *Proceedings of the 28th international conference on machine learning (icml-11)*, 489–496.
- Gottipati, S.; Qiu, M.; Yang, L.; Zhu, F.; and Jiang, J. 2013. Predicting users political party using ideological stances. In *Social Informatics*. Springer. 177–191.
- Griffiths, D., and Tenenbaum, M. 2004. Hierarchical topic models and the nested chinese restaurant process. *Advances in neural information processing systems* 16:17.
- Grimmer, J. 2010. A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis* 18(1):1–35.
- Heinrich, G. 2005. Parameter estimation for text analysis. Technical report, Technical report.
- Jo, Y., and Oh, A. H. 2011. Aspect and sentiment unification model for online review analysis. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 815–824. ACM.
- Lin, C., and He, Y. 2009. Joint sentiment/topic model for sentiment analysis. In *Proceedings of the 18th ACM conference on Information and knowledge management*, 375–384. ACM.
- Lin, W.-H.; Xing, E.; and Hauptmann, A. 2008. A joint topic and perspective model for ideological discourse. In *Machine Learning and Knowledge Discovery in Databases*. Springer. 17–32.
- McAuley, J. J., and Leskovec, J. 2013. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*, 897–908. International World Wide Web Conferences Steering Committee.
- Mei, Q.; Ling, X.; Wondra, M.; Su, H.; and Zhai, C. 2007. Topic sentiment mixture: modeling facets and opinions in weblogs. In *Proceedings of the 16th international conference on World Wide Web*, 171–180. ACM.
- Metaxas, P. T.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (not) to predict elections. In *Privacy, security, risk and trust (PASSAT), 2011 IEEE third international conference on and 2011 IEEE third international conference on social computing (SocialCom)*, 165–171. IEEE.
- O’Connor, B.; Balasubramanyan, R.; Routledge, B. R.; and Smith, N. A. 2010. From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM* 11:122–129.
- Ontheissues.com. Candidates on the Issues, url = <http://www.ontheissues.org/default.htm>.
- Röder, M.; Both, A.; and Hinneburg, A. 2015. Exploring the space of topic coherence measures. In *Proceedings of the eight International Conference on Web Search and Data Mining, Shanghai, February 2-6*.
- Wang, X.; Mohanty, N.; and McCallum, A. 2005. Group and topic discovery from relations and text. In *Proceedings of the 3rd international workshop on Link discovery*, 28–35. ACM.
- Wong, F. M. F.; Tan, C. W.; Sen, S.; and Chiang, M. 2013. Quantifying political leaning from tweets and retweets. In *ICWSM*.
- Yano, T.; Cohen, W. W.; and Smith, N. A. 2009. Predicting response to political blog posts with topic models. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 477–485. Association for Computational Linguistics.
- Zhao, W. X.; Jiang, J.; Yan, H.; and Li, X. 2010. Jointly modeling aspects and opinions with a maxent-lda hybrid. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, 56–65. Association for Computational Linguistics.