

A Latent Variable Model for Viewpoint Discovery from Threaded Forum Posts

Minghui Qiu

School of Information Systems
Singapore Management University
Singapore

minghui.qiu.2010@smu.edu.sg

Jing Jiang

School of Information Systems
Singapore Management University
Singapore

jingjiang@smu.edu.sg

Abstract

Threaded discussion forums provide an important social media platform. Its rich user generated content has served as an important source of public feedback. To automatically discover the viewpoints or stances on hot issues from forum threads is an important and useful task. In this paper, we propose a novel latent variable model for viewpoint discovery from threaded forum posts. Our model is a principled generative latent variable model which captures three important factors: viewpoint specific topic preference, user identity and user interactions. Evaluation results show that our model clearly outperforms a number of baseline models in terms of both clustering posts based on viewpoints and clustering users with different viewpoints.

1 Introduction

Threaded discussion forums provide an important social media platform that allows netizens to express their opinions, to ask for advice, and to form online communities. In particular, responses to major sociopolitical events and issues can often be found in discussion forums, which serve as an important source of public feedback. In such discussion threads, we often observe heated debates over a controversial issue, with different sides defending their viewpoints with different arguments. For example, after the presidential debate between Barack Obama and Mitt Romney, there were heated discussions in online forums like CreateDebate¹ where some people expressed their support for Obama while some

others have their opposition to him. For a user who is not closely following an event or issue, instead of going through all the existing posts in a long thread, she may want to quickly get an overview of the major viewpoints and arguments given by the different sides. For policy makers who want to obtain public feedback on social issues from social media, it is also desirable to automatically summarize the viewpoints on an issue from relevant threads. In this paper, we study the problem of modeling and discovering different viewpoints in forum threads.

Recently there has been some work on finding contrastive viewpoints from text. The model proposed by Paul et al. (2010) assumes viewpoints and topics are orthogonal dimensions. Another model proposed by Fang et al. (2012) assumes that documents are already grouped by viewpoints and it focus on identifying contrastive viewpoint words under the same topic. However, these existing studies are not based on interdependent documents like threaded forum posts. As a result, at least two important characteristics of threaded forum data are not considered in these models. (1) **User identity**: The user or publisher of each forum post is known, and a user may publish several posts in the same thread. Since the same user's opinion on an issue usually remains unchanged, posts published by the same user are likely to contain the same viewpoint. (2) **User interactions**. A thread is like a conversation, where users not only directly comment on the issue under discussion but also comment on each other's posts. Users having different viewpoints may express their disagreement or even attack each other while users having the same viewpoint often support each other.

¹<http://www.createdebate.com/>

The interaction expressions in forum posts may help us infer the relation between two users and subsequently infer the viewpoints of the corresponding posts.

In this paper, we propose a novel latent variable model for viewpoint discovery from threaded forum posts. Our model is based on the following observations: First, posts with different viewpoints tend to focus on different topics. To illustrate this point, we first apply the Latent Dirichlet Allocation (LDA) model (Blei et al., 2003) on a thread about “will you vote Obama” and obtain a set of topics. This thread comes from a data set that has each user’s viewpoint annotated. Using the ground truth viewpoint labels, we group all posts published by users with viewpoint 1 (or viewpoint 2) and compute the topic proportions. The two topic distributions are shown in Figure 1. We can see that indeed the two viewpoints each have some dominating topics. Our second observations is that the same user tends to hold the same viewpoint. In our model, we use a user-level viewpoint distribution to capture this observation, and the experiments show that it works better than assuming a global viewpoint distribution. Third, users with the same viewpoint are likely to have positive interactions while users with different viewpoints tend to have negative interactions. Using a sentiment lexicon, we can first predict the polarity of interaction expressions. We then propose a novel way to incorporate this information into the latent variable model. In summary, we capture the three observations above in a principled generative latent variable model. We present the details of our model in Section 3.




Figure 1: Topic distributions of two viewpoints for the thread “will you vote Obama?” The dotted line is the average topic probability.

We use two tasks to evaluate our model. In the first task, we evaluate how well posts with different viewpoints are separated. In the second task, we

evaluate how well our model is able to group users with different viewpoints. For both tasks, we compare our model with an existing model as well as a few degenerate versions of our model. The results show that our model can clearly outperform the baselines in terms of three evaluation metrics. The experiments are presented in Section 5.

The contributions of our work are threefold: (1) We identify the importance of using user interactions to help infer viewpoints in forum posts. (2) We propose a principled latent variable model to jointly model topics, viewpoints and user interactions. (3) We empirically verify the validity of the three assumptions in our model using real data sets.

2 Related Work

There are a few different lines of work that are related to our work. For discovering different viewpoints from general text, Paul et al. (2010) used the model proposed by Paul and Girju (2010) to jointly model topics and viewpoints. They assume these two concepts are orthogonal and they do not consider user identity. In comparison, our model has the notion of topics and viewpoints, but we explicitly model the dependency of topics on viewpoints, i.e. we assume each viewpoint has a topic distribution. We also consider author identities as an important factor of our model. Fang et al. (2012) proposed a model that also combines topics and viewpoints. But they assume that documents are already grouped by viewpoints, which is not the case for forum posts. Therefore, their model cannot be directly applied to forum posts.

There has also been some work on finding viewpoints from social media. Somasundaran and Wiebe (2010) studied how to identify stances in online debates. They used a supervised approach for classifying stances in ideological debates. In comparison, our model is an unsupervised method. The same authors proposed an unsupervised method which relies on associations of aspects with topics indicative of stances mined from the Web for the task (Somasundaran and Wiebe, 2009). In contrast, our model is also an unsupervised one but we do not rely on any external knowledge.

Part of our work is related to detecting agreement/disagreement from text. For this task, nor-

mally supervised methods are used (Galley et al., 2004; Abbott et al., 2011), which require sufficient labeled training data. In our work, since we deal with different languages, we use a lexicon-based approach that does not need training data. Recently, Mukherjee and Liu (2012) proposed an unsupervised model to extract different types of expressions including agreement/disagreement expressions. However, our focus is not to detect agreement/disagreement expressions but to model the interplay between agreement/disagreement expressions and viewpoints. The work by Mukherjee and Liu (2012) can potentially be combined with our model.

Another line of related work is subgroup detection, which aims to separate users holding different viewpoints. This problem has recently been studied by Abu-Jbara and Radev (2012), Dasigi et al. (2012), Abu-Jbara et al. (2012) and Hassan et al. (2012), where a clustering based approach is used. Lu et al. (2012) studied both textual content and social interactions to find opposing network from online forums. In our experiments we show that our model can also be used for subgroup detection, but meanwhile we also directly identify viewpoints, which is not the goal of existing work on subgroup finding or opposing network extraction.

3 Model

3.1 Motivation

Before we formally present our latent variable model for viewpoint discovery, let us first look at the assumptions we would like to capture in the model.

Viewpoint-based topic distribution: The first assumption we have is that different viewpoints tend to touch upon different topics. This is because to support a viewpoint, users need to provide evidence and arguments, and for different viewpoints the arguments are likely different. To capture this assumption, in our model, we let each viewpoint have its own distribution of topics. Given the viewpoint of a post, the hidden topic of each word in the post is chosen according to the corresponding topic distribution associated with that viewpoint.

User identify: The second assumption we have is that the same user tends to talk from the same viewpoint, although there are also users who do not

clearly have a viewpoint. In our model, we assume that there is a user-level viewpoint distribution. For each post by a user, its viewpoint is drawn from the corresponding viewpoint distribution.

User interaction: An important difference between threaded forum posts and regular document collections such as news articles is that posts in the same thread form a tree structure via the “reply-to” relations. Many reply posts start with an expression that comments on a previous post or directly addresses another user. These interaction expressions may carry positive or negative sentiment, indicating an agreement or a disagreement. For example, Table 1 shows the interaction expressions from a few sample posts with words such as “correct,” “agree,” and “delusional,” implying the polarity of the interaction expressions. The polarity of these interaction expressions can help us infer whether two posts or two users hold the same viewpoint or not. In our model, we assume that the polarity of each interaction expression can be detected. Details of how we perform this detection are in Section 3.4.

| | Post |
|---|--|
| + | You are correct . Obama got into office w/ everything ... I agree with your post Dan. Obama is so ... |
| - | Most of your post is delusional , especially the part ... Are you freaking nutz? Palin is a BIMBO! |

Table 1: Sample posts with positive (+) and negative(-) interactions.

While the way to capture the first two assumptions discussed above is fairly standard, modeling user interactions is something new. In our model, we assume that the polarity of an interaction expression is generated based on the viewpoint of the current post and the viewpoint of post(s) that the current post replies to. The intuition is that if the viewpoints are the same, we are more likely to see a positive interaction whereas if the viewpoints are different we are more likely to see a negative interaction.

3.2 Model description

We use the following notation to represent our data. We consider a set of forum posts published by U different users on the same event or issue, where user u ($1 \leq u \leq U$) has published N_u posts. Let $w_{u,n,l}$ ($1 \leq l \leq L_{u,n}$) denote the l -th word in the n -th post by user u , where $L_{u,n}$ is the number of words

in the n -th post by user u . $w_{u,n,l}$ is represented by an index between 1 and V where V is the vocabulary size. Furthermore, we assume that some of the posts have user interaction expressions, where the polarity of the expression is known. Without loss of generality, let $s_{u,n} \in \{0, 1\}$ denote the polarity of the interaction expression of the n -th post by user u . In addition, for each post that has an interaction expression, we assume we also know the previous post(s) it replies to. (In the case when the current post replies to a user, we assume all that user's existing posts are being replied to.) We refer to these posts as the *parent posts* of the current post.

We assume that there are T topics where each topic is essentially a word distribution, denoted as ψ^t . We also assume that there are Y different viewpoints expressed in the collection of posts. For most controversial issues, Y can be set to 2. Each viewpoint y has a topic distribution θ^y over the T topics. While these T topics are meant to capture the topical differences between viewpoints, since these viewpoints are all about the same issue, there are also some words commonly used by different viewpoints. We therefore introduce a background topic ψ^B to capture these words. Finally, each user u has a distribution over the Y viewpoints, denoted as φ^u .

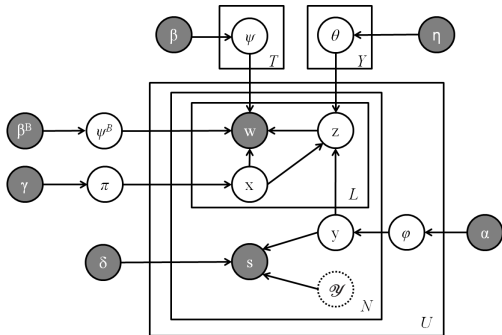


Figure 2: Plate notation of the Joint Viewpoint-Topic Model with User Interaction (JVTM-UI). The dotted circle for \mathcal{Y} means the variables represented by \mathcal{Y} are not new variables but a subset of the y variables.

Figure 2 shows the plate notation of the complete model. We assume the following generation process in our model. When user u generates her n -th post, she first samples a viewpoint from φ^u . Let this viewpoint be represented by a hidden variable $y_{u,n}$. For the l -th word in this post, she first samples an in-

dicator variable $x_{u,n,l}$ from a Bernoulli distribution parameterized by π . If $x_{u,n,l} = 0$, then she draws $w_{u,n,l}$ from ψ^B . Otherwise, she first samples a topic, denoted as $z_{u,n,l}$, according to $\theta^{y_{u,n}}$, and then draws $w_{u,n,l}$ from $\psi^{z_{u,n,l}}$.

Furthermore, if this post is a reply to a previous post or another user, she may first comment on the parent post(s). The polarity of the interaction expression in the post is dependent on the viewpoint $y_{u,n}$ and the viewpoints of the previous post(s). Let us use $\mathcal{Y}_{u,n}$ to denote the set of y variables associated with the parent posts of the current post. The user draws $s_{u,n}$ according to following distribution:

$$p(s_{u,n} = 1 | y_{u,n}, \mathcal{Y}_{u,n}, \delta) = \frac{\sum_{y' \in \mathcal{Y}_{u,n}} \mathbb{I}(y_{u,n} == y') + \delta}{|\mathcal{Y}_{u,n}| + 2\delta},$$

$$p(s_{u,n} = 0 | y_{u,n}, \mathcal{Y}_{u,n}, \delta) = 1 - p(s_{u,n} = 1 | y_{u,n}, \mathcal{Y}_{u,n}, \delta), \quad (1)$$

where $\mathbb{I}(\cdot)$ is 1 if the statement inside is true and 0 otherwise, and $\delta > 0$ is a smoothing parameter.

Finally, we assume that ψ^B , ψ^t , φ^u , θ^y and π all have some uniform Dirichlet priors.

3.3 Inference

We use collapsed Gibbs sampling to estimate the model parameters. In the initialization stage of Gibbs sampling, for a reply post to a recipient, we initialize its corresponding reply polarity s according to all the labeled polarity of interaction words. Specifically, if the majority of labeled interaction words are positive, we set $s = 1$, otherwise we set $s = 0$.

Let \mathbf{Y} denote the set of all y variables, and $\mathbf{Y}_{-(u,n)}$ denote \mathbf{Y} excluding $y_{u,n}$. Similar notation is used for the other variables. We sample $y_{u,n}$ using the following formula.

$$p(y_{u,n} = k | \mathbf{Y}_{-(u,n)}, \mathbf{Z}, \mathbf{S}, \mathbf{X}, \alpha, \eta, \delta)$$

$$\propto \frac{p(y_{u,n} = k, \mathbf{Y}_{-(u,n)} | \alpha)}{p(\mathbf{Y}_{-(u,n)} | \alpha)} \cdot \frac{p(\mathbf{Z} | y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \mathbf{X}, \eta)}{p(\mathbf{Z}_{-(u,n)} | \mathbf{Y}_{-(u,n)}, \mathbf{X}_{-(u,n)}, \eta)}$$

$$\cdot p(\mathbf{S} | y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \delta)$$

$$= \frac{C_{u,-n}^k + \alpha}{C_{u,-n}^{(\cdot)} + Y\alpha} \cdot \frac{\prod_{t=1}^T \prod_{a=0}^{C_{k,-(u,n)}^t - 1} (C_{k,-(u,n)}^t + \eta + a)}{\prod_{b=0}^{C_{u,-n}^{(\cdot)} - 1} (C_{k,-(u,n)}^{(\cdot)} + T\eta + b)}$$

$$\cdot p(\mathbf{S} | y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \delta). \quad (2)$$

Here all C 's are counters. $C_{u,-n}^k$ is the number of times we observe the viewpoint k from u 's posts, excluding the n -th post, based on $\mathbf{Y}_{-(u,n)}$. $C_{u,n}^t$ is

the number of times we observe topic t from user u 's n -th post, based on $\mathbf{Z}_{u,n}$. And $C_{k,-(u,n)}^t$ is the number of times we observe topic t associated with viewpoint k , excluding user u 's n -th post. Note that we need \mathbf{X} to know which words are assigned to the background topic so we can exclude them for $C_{u,n}^t$ and $C_{k,-(u,n)}^t$. $C_{u,-n}^{(\cdot)}$ is the number of times we observe any viewpoint from u 's posts, excluding the n -th post. $C_{u,n}^{(\cdot)}$ and $C_{k,-(u,n)}^{(\cdot)}$ are defined similarly.

The last term is further expanded as follows:

$$p(\mathbf{S}|y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \delta) = p(s_{u,n}|y_{u,n} = k, \mathcal{Y}_{u,n}, \delta) \cdot p(\mathbf{S}_{-(u,n)}|y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \delta). \quad (3)$$

Here $p(s_{u,n}|y_{u,n} = k, \mathcal{Y}_{u,n}, \delta)$ is computed according to Eqn. (1). For the latter term, we need to consider posts which reply to user u 's n -th post because the value of $y_{u,n}$ affects these posts.

$$p(\mathbf{S}_{-(u,n)}|y_{u,n} = k, \mathbf{Y}_{-(u,n)}, \delta) \propto \prod_{(u',n'): y_{u,n} \in \mathcal{Y}_{u',n'}} p(s_{u',n'}|y_{u',n'}, \mathcal{Y}_{u',n'}, \delta). \quad (4)$$

Next, we show how we jointly sample $x_{u,n,l}$ and $z_{u,n,l}$. We jointly sample them because when $x_{u,n,l} = 0$, $z_{u,n,l}$ does not need a value. We have the following formulas:

$$p(x_{u,n,l} = 1, z_{u,n,l} = t | \mathbf{X}_{-(u,n,l)}, \mathbf{Z}_{-(u,n,l)}, \mathbf{Y}, \mathbf{W}, \gamma, \eta, \beta, \beta^B) \propto \frac{C_{-(u,n,l)}^1 + \gamma}{C_{-(u,n,l)}^{(\cdot)} + 2\gamma} \cdot \frac{C_{y_{u,n,l}, -(u,n,l)}^t + \eta}{C_{y_{u,n,l}, -(u,n,l)}^{(\cdot)} + T\eta} \cdot \frac{C_{t, -(u,n,l)}^{w_{u,n,l}} + \beta}{C_{t, -(u,n,l)}^{(\cdot)} + V\beta}, \quad (5)$$

$$p(x_{u,n,l} = 0 | \mathbf{X}_{-(u,n,l)}, \mathbf{Z}_{-(u,n,l)}, \mathbf{Y}, \mathbf{W}, \gamma, \eta, \beta, \beta^B) \propto \frac{C_{-(u,n,l)}^0 + \gamma}{C_{-(u,n,l)}^{(\cdot)} + 2\gamma} \cdot \frac{C_{B, -(u,n,l)}^{w_{u,n,l}} + \beta^B}{C_{B, -(u,n,l)}^{(\cdot)} + V\beta^B}. \quad (6)$$

Here again the C s are counters defined in similar ways as before. For example, $C_{-(u,n,l)}^1$ is the number of times we observe 1 assigned to an x variable, excluding $x_{u,n,l}$.

3.4 Interaction polarity prediction

The problem of detecting agreement and disagreement from forum posts is relatively new. One possible solution is to use supervised learning, which requires training data (Galley et al., 2004; Abbott et al., 2011; Andreas et al., 2012). However, training

data are also likely domain and language dependent, which makes them hard for re-use. For our task, we take a simpler approach and use a sentiment lexicon together with some heuristics to predict the polarity of interaction expressions. Specifically, we first identify interaction sentences following the strategies from Hassan et al. (2012). We assume sentences containing mentions of the recipient of a post are interaction sentences. Next, we consider words within a text window of 8 words surrounding these mentions. We then use a subjectivity lexicon to label these words. To form an English lexicon, we combine three popular lexicons: the sentiment lexicon used by Hu and Liu (2004), Multi-Perspective Question Answering Subjectivity Lexicon by Wilson et al. (2005) and SentiWordNet by Baccianella et al. (2010). Since we also work with a Chinese data set, to form the Chinese sentiment lexicon, we use opinion words from HowNet² and NTUSD by Ku et al. (2007). To predict the polarity of an interaction expression, we simply check whether there are more positive sentiment words or more negative sentiment words in the expression, and label the interaction expression accordingly.

We would like to stress that since this interaction classification step is independent of the latent variable model, we can always apply a more accurate method, but this is not the focus of this work.

4 Models for Comparison

In our experiments, we compare our model, Joint Viewpoint-Topic Model with User Interaction (JVTM-UI), with the following baseline models.

JVTM: The model is shown in Figure 3(a), a variant of JVTM-UI that does not consider user interaction. Through comparison with it, we can check the effect of modeling user interactions.

JVTM-G: We consider JVTM-G in Figure 3(b), a variant of JVTM which assumes a global viewpoint distribution. Comparison with it allows us to check the usefulness of user identity in the task.

UIM: The third model we consider is a User Interaction Model (UIM) in Figure 3(c), where we rely on only the users' interactions to infer the viewpoints. We use it to check how well viewpoints can be discovered from only user interaction expressions.

²http://www.keenage.com/html/e_index.html

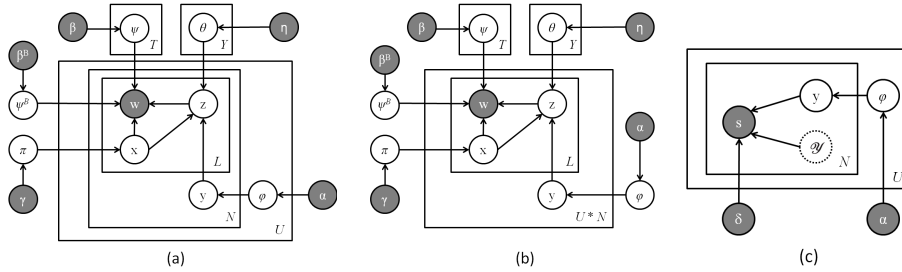


Figure 3: (a) JVTM: Joint Viewpoint-Topic Model. (b) JVTM-G: JVTM with a global viewpoint distribution. (c) UIM: User-Interaction Model.

TAM: The last model we consider is the one by Paul et al. (2010). As TAM is applied at document collections, we first concatenate all the posts by the same user into a pseudo document and then apply TAM.

5 Experiments and Analysis

In this section, we evaluate our model with a set of baseline models using two data sets.

| Name | Issue | #Posts | #Users |
|------|-------------------------------|--------|--------|
| EDS1 | Vote for Obama | 2599 | 197 |
| EDS2 | Arizona Immigration Law | 738 | 59 |
| EDS3 | Tax Cuts | 276 | 26 |
| CDS1 | Tencent and Qihoo dispute | 30137 | 2507 |
| CDS2 | Fang Zhouzi questions Han Han | 76934 | 1769 |
| CDS3 | Liu Xiang in London Olympics | 29486 | 2774 |

Table 2: Some statistics of the data set.

5.1 Data Sets and Experimental Settings

We focus our work on finding users’ viewpoints on a controversial issue, where we assume that there are two contradictory viewpoints. We use two data sets on controversial issues. The first data set comes from Abu-Jbara et al. (2012) and Hassan et al. (2012). This data set originally was used for finding subgroups of users, so the annotations were done at user level, i.e. for each user there is a label indicating which subgroup he/she belongs to. We use the top-3 mostly discussed threads with two subgroups for our study.

In reality, controversial issues are often discussed across threads. We thus constructed another large data set which contains more than one thread for each issue. We chose three hot issues from one of the most popular Chinese online forums — TianYa

Club³. The three issues are “Fang Zhouzi questions Han Han”⁴, “Tencent and Qihoo dispute”⁵, and “Liu Xiang in London Olympics”⁶. All these issues triggered heated discussions on the forum and we found that most of the users were divided into two different groups.

We crawled the data set using the TianYa API⁷. The API allows users to issue queries and get threads most related to the queries. For each issue, we used entities involved in the event as queries and obtained 750 threads for each query. We then extracted all the posts in the threads. As there are users who posted irrelevant posts in the forum, we then filtered out those users who did not mention the entities or had fewer than 4 posts.

We refer to the first set of data in English as EDS1, EDS2 and EDS3, and the second set of data in Chinese as CDS1, CDS2 and CDS3. Some statistics of the resulting data set are shown in Table 2.

For all the models, we set $Y = 2$. We set $T = 10$ for the English data sets and $T = 40$ for the Chinese data sets. We run 400 iterations of Gibbs sampling as burn-in iterations and then take 100 samples with a gap of 5 to obtain our final results. We empirically set $\beta = 0.01$, $\beta^B = 0.1$, $\gamma = 10$ and $\delta = 0.1$ for our model on all the data sets. α and η are set through grid search where they take values in $\{0.01, 0.001\}$. For each data set, we choose the best setting for each model and report the corresponding results.

³http://en.wikipedia.org/wiki/Tianya_Club

⁴http://en.wikipedia.org/wiki/Fang_Zhouzi

⁵http://en.wikipedia.org/wiki/360_v._Tencent

⁶http://en.wikipedia.org/wiki/Liu_Xiang

⁷<http://open.tianya.cn/index.php>

5.2 Identification of viewpoints

We first evaluate the models on the task of identifying viewpoints. For fair comparison, each model will output a viewpoint label for each post. For JVTM-UI, JVTM, JVTM-G and UIM, after we learn the model, each post will directly have a viewpoint assignment. For TAM we cannot directly get each post’s viewpoint as the model assumes a document-level viewpoint distribution. To estimate each post’s viewpoint in this model, we use viewpoint assignment at the word level learnt from the model. Then for each post, we label its viewpoint as the viewpoint that has the majority count in the post.

Ideally, we would like to manually label all the posts to obtain the ground truth for evaluation. Since there are too many posts, we only labeled a sample of them. For each issue, we randomly selected 150 posts to label their viewpoints. For each post, we asked two different annotators to label its viewpoint. We made sure that the annotators understand the issue and the two major viewpoints before they annotated the posts. Specifically, as the Chinese data sets are about some controversial issues around the entities involved, we then defined two major viewpoints as *support* and *not support* the entity who initiated the event. The entities of data set CDS1, CDS2 and CDS3 are *Fang Zhouzi*, *Tencent* and *Liu Xiang* respectively. For each given post, the annotators were asked to judge whether the post has expressed viewpoints and if so, what is its corresponding viewpoint. We measure the agreement score using Cohen’s kappa coefficient. The lowest agreement score for an issue is 0.61 in the data set, showing good agreement. We then used the set of posts that were labeled with the same viewpoint by the two annotators as our evaluation data for all the models.

Since our task is essentially a clustering problem, we use *purity* and *entropy* to measure the performance (Manning et al., 2008). Furthermore, we also use *accuracy* where we choose the better alignment of clusters with ground truth class labels and compute the percentage of posts that are “classified” correctly. For purity and accuracy, the higher the measure is the better the performance. For entropy, the lower the measure is the better the performance.

We give an overview of the all the averaged model results on the data sets in Figure 4. We observed that




Figure 4: Averaged results of the models in identification of viewpoints.

UIM performs relatively better than other methods except our model. This shows user interactions are important features to identify post viewpoints. Overall, our model has a better performance as it is with higher purity and accuracy, and lower entropy.

| | JVTM-UI | UIM | JVTM | TAM | JVTM-G | |
|------|----------|-------------|------|-------------|--------|------|
| EDS1 | <i>P</i> | 0.77 | 0.74 | 0.64 | 0.65 | 0.63 |
| | <i>E</i> | 0.72 | 0.76 | 0.90 | 0.92 | 0.94 |
| | <i>A</i> | 0.77 | 0.74 | 0.61 | 0.60 | 0.57 |
| EDS2 | <i>P</i> | 0.82 | 0.78 | 0.68 | 0.65 | 0.64 |
| | <i>E</i> | 0.69 | 0.73 | 0.79 | 0.86 | 0.90 |
| | <i>A</i> | 0.81 | 0.78 | 0.68 | 0.68 | 0.65 |
| EDS3 | <i>P</i> | 0.79 | 0.73 | 0.65 | 0.64 | 0.62 |
| | <i>E</i> | 0.67 | 0.79 | 0.88 | 0.89 | 0.87 |
| | <i>A</i> | 0.79 | 0.73 | 0.65 | 0.64 | 0.62 |
| CDS1 | <i>P</i> | 0.87 | 0.83 | 0.83 | 0.82 | 0.82 |
| | <i>E</i> | 0.61 | 0.64 | 0.65 | 0.66 | 0.64 |
| | <i>A</i> | 0.60 | 0.58 | 0.59 | 0.58 | 0.57 |
| CDS2 | <i>P</i> | 0.71 | 0.65 | 0.61 | 0.63 | 0.60 |
| | <i>E</i> | 0.80 | 0.85 | 0.92 | 0.95 | 0.96 |
| | <i>A</i> | 0.71 | 0.65 | 0.61 | 0.61 | 0.59 |
| CDS3 | <i>P</i> | 0.78 | 0.78 | 0.78 | 0.78 | 0.78 |
| | <i>E</i> | 0.73 | 0.75 | 0.70 | 0.72 | 0.73 |
| | <i>A</i> | 0.67 | 0.59 | 0.67 | 0.66 | 0.63 |

Table 3: Results on viewpoint identification on the all data sets.

Table 3 shows the detailed results on the data sets. We perform the 2-tailed paired t-test as used by Abu-Jbara et al. (2012) on the results. All the result differences are at 10% significance level if not with further clarification. First, JVTM has a better performance over JVTM-G, which shows it is important to consider user identity in the task. Second, JVTM and TAM have similar performance on

EDS1 and CDS2, but JVTM has a relatively better performance on EDS2, EDS3, CDS1 and CDS3. This shows it is helpful to consider each viewpoint’s topic preference. Although as studied by Paul et al. (2010), by only using unigram features, TAM may not be able to cluster viewpoints accurately, our study shows that the results can be improved when adding each viewpoint’s topic focus. Third, UIM has relatively better performance than the other models, which demonstrates that user interactions alone can do a decent job in inferring viewpoints. Finally, our proposed model has the best performance across the board in terms of all three evaluation metrics. Note that, our proposed model significantly outperforms other methods at 5% significance level except at 10% significance level over JVTM model. This shows by jointly modeling topics, viewpoints and user interactions, our model can better identify posts with different viewpoints.

5.3 Identification of user groups

We also use another task to evaluate our model. The task here is finding each user’s viewpoint and subsequently grouping users by their viewpoints. This task has been studied by Abu-Jbara and Radev (2012), Dasigi et al. (2012), Abu-Jbara et al. (2012) and Hassan et al. (2012). For the English data set, the user-level group labels are provided by the original data set. For the Chinese data set, we randomly selected 150 users for each issue and manually labeled them according to their viewpoints as reflected by their posts. If a user’s posts do not clearly suggest a viewpoint, we label her as neutral. Again we asked two human judges to do annotation. The agreement scores are above 0.70 for all issues, showing substantial agreement. This score is higher than viewpoint identification, which suggests that it is easier to judge a user’s viewpoint than a single post’s viewpoint. We use the set of users who have got the same labels by the two human judges for our experiments. Similarly we compute *purity*, *entropy* and *accuracy* to evaluate the clustering results.

Figure 5 shows the averaged results of all the models. Similar to previous experiment, our model has a better performance compared to the competing models.

The results on the each data set are shown in Table 4. The tables show that similar trends can be

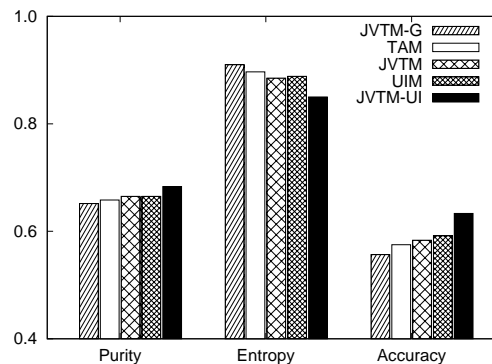


Figure 5: Averaged results of the models in identification of user groups.

| | | JVTM-UI | UIM | JVTM | TAM | JVTM-G |
|------|----------|-------------|------|------|------|--------|
| EDS1 | <i>P</i> | 0.67 | 0.67 | 0.67 | 0.67 | 0.67 |
| | <i>E</i> | 0.85 | 0.88 | 0.89 | 0.89 | 0.91 |
| | <i>A</i> | 0.63 | 0.59 | 0.58 | 0.59 | 0.57 |
| EDS2 | <i>P</i> | 0.77 | 0.77 | 0.77 | 0.77 | 0.77 |
| | <i>E</i> | 0.72 | 0.76 | 0.74 | 0.75 | 0.76 |
| | <i>A</i> | 0.62 | 0.59 | 0.60 | 0.58 | 0.59 |
| EDS3 | <i>P</i> | 0.68 | 0.63 | 0.61 | 0.61 | 0.58 |
| | <i>E</i> | 0.90 | 0.92 | 0.95 | 0.96 | 0.97 |
| | <i>A</i> | 0.68 | 0.63 | 0.61 | 0.58 | 0.57 |
| CDS1 | <i>P</i> | 0.64 | 0.60 | 0.61 | 0.61 | 0.60 |
| | <i>E</i> | 0.91 | 0.97 | 0.96 | 0.96 | 0.97 |
| | <i>A</i> | 0.61 | 0.55 | 0.55 | 0.56 | 0.53 |
| CDS2 | <i>P</i> | 0.69 | 0.69 | 0.69 | 0.69 | 0.69 |
| | <i>E</i> | 0.83 | 0.89 | 0.85 | 0.89 | 0.89 |
| | <i>A</i> | 0.62 | 0.57 | 0.56 | 0.58 | 0.54 |
| CDS3 | <i>P</i> | 0.67 | 0.63 | 0.64 | 0.60 | 0.60 |
| | <i>E</i> | 0.89 | 0.91 | 0.92 | 0.93 | 0.96 |
| | <i>A</i> | 0.64 | 0.62 | 0.60 | 0.56 | 0.54 |

Table 4: Results on identification of user groups on the all the data sets.

observed for the task of user group identification. We also perform the 2-tailed paired t-test on the results. We find our model significantly outperforms other models in terms of accuracy at 5% significance level, and purity and entropy at 10% significance level. Overall speaking, our joint model performed the best among all the models for this task for all three metrics. This shows that it is important to consider the topical preference of individual viewpoint, user’s identify as well as the interactions between users.




Figure 6: The user interaction network in a discussion thread about “will you vote obama.” Green (left) and white (right) nodes represent users with two different viewpoints. Red (thin) and blue(thick) edges represent negative and positive interactions.

5.4 User interaction network

To gain some direct insight into our results, we show the user interaction network from one thread in Figure 6. Here each node denotes a user, and its color denotes the predicted viewpoint of that user. A link between a pair of users means these users have interactions and the interaction types have a dominant polarity. The polarities of these links are predicted using the interaction expressions and a sentiment lexicon, whereas the viewpoints of different users are learned by JVTM-UI, making use of the interaction polarities. The figure shows that clearly there are mostly positive interactions between users with the same viewpoint and mostly negative interactions between users with different viewpoints. Note that, our method to identify user interaction polarity is rule-based. As this step serves as a preprocessing step for our latent variable model, we can always use a more accurate method to improve the performances.

6 Conclusion

In this work, we proposed a novel latent variable model for viewpoint discovery from threaded forum posts. Our model is based on the three important factors: viewpoint specific topic preference, user identity and user interactions. Our proposed model captures these observations in a principled way. In particular, to incorporate the user interaction information, we proposed a novel generative process. Empirical evaluation on the real forum data sets showed that our model could cluster both posts and users with different viewpoints more accurately than the baseline models we consider. To the best of our

knowledge, our work is the first to incorporate user interaction polarity into a generative model to discover viewpoints.

In this work, we only considered unigrams. As some previous work has shown, more complex lexical units such as n -grams (Mukherjee and Liu, 2012) and dependency triplets (Paul et al., 2010) may improve the performance of topic models. We will consider these strategies in our future work. Currently we use a simple heuristic-based classifier to predict interaction polarity. In our further work, we plan to consider more accurate methods using deeper linguistic analysis. We did not study how to summarize the discovered viewpoints in this work, which is also something we will look into in our future work.

Acknowledgments

We thank the reviewers for their valuable comments on this work. We also thank Shuang Xia for his help on processing and labeling the data sets.

References

- Rob Abbott, Marilyn Walker, Pranav Anand, Jean E. Fox Tree, Robeson Bowmani, and Joseph King. 2011. How can you say such things?!?: Recognizing disagreement in informal political argument. In *Proceedings of the Workshop on Language in Social Media (LSM 2011)*, pages 2–11.
- Amjad Abu-Jbara and Dragomir R. Radev. 2012. Subgroup detector: A system for detecting subgroups in online discussions. In *ACL (System Demonstrations)*, pages 133–138.
- Amjad Abu-Jbara, Pradeep Dasigi, Mona Diab, and Dragomir R. Radev. 2012. Subgroup detection in ideological discussions. In *Proceedings of ACL 2012*, pages 399–409.
- Jacob Andreas, Sara Rosenthal, and Kathleen McKeown. 2012. Annotating agreement and disagreement in threaded discussion. In *Proceedings of LREC’12*.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Pradeep Dasigi, Weiwei Guo, and Mona T. Diab. 2012. Genre independent subgroup detection in online discussion threads: A study of implicit attitude using

- textual latent semantics. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 65–69.
- Yi Fang, Luo Si, Naveen Somasundaram, and Zhengtao Yu. 2012. Mining contrastive opinions on political texts using cross-perspective topic model. In *WSDM*, pages 63–72.
- Michel Galley, Kathleen McKeown, Julia Hirschberg, and Elizabeth Shriberg. 2004. Identifying agreement and disagreement in conversational speech: Use of bayesian networks to model pragmatic dependencies. In *Proceedings of ACL'04, Main Volume*, pages 669–676.
- Ahmed Hassan, Amjad Abu-Jbara, and Dragomir Radev. 2012. Detecting subgroups in online discussions by modeling positive and negative relations among participants. In *Proceedings of the 2012 EMNLP*, pages 59–70.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Lun-wei Ku, Yong-sheng Lo, and Hsin-hsi Chen. 2007. Using polarity scores of words for sentence-level opinion extraction. In *Proc. of the NTCIR-6 Workshop Meeting*, pages 316–322.
- Yue Lu, Hongning Wang, ChengXiang Zhai, and Dan Roth. 2012. Unsupervised discovery of opposing opinion networks from forum discussions. In *Proceedings of the 21st ACM international conference on Information and knowledge management, CIKM '12*, pages 1642–1646, New York, NY, USA. ACM.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schtze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, July.
- Arjun Mukherjee and Bing Liu. 2012. Modeling review comments. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 320–329.
- Michael J. Paul and Roxana Girju. 2010. A two-dimensional topic-aspect model for discovering multifaceted topics. In *AAAI*.
- Michael J. Paul, ChengXiang Zhai, and Roxana Girju. 2010. Summarizing contrastive viewpoints in opinionated text. In *EMNLP*, pages 66–76.
- Swapna Somasundaran and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 226–234.
- Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 116–124.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*.