

# Modeling User Attitude toward Controversial Topics in Online Social Media

**Huiji Gao**

Computer Science and Engineering  
Arizona State University  
Tempe, AZ 85281  
huiji.gao@asu.edu

**Jalal Mahmud, Jilin Chen, Jeffrey Nichols, Michelle Zhou**

IBM Almaden Research Center  
San Jose, CA 95120  
{jumahmud, jilinc, jwnichols, mzhou}@us.ibm.com

## Abstract

The increasing use of social media platforms like Twitter has attracted a large number of online users to express their attitude toward certain topics. Sentiment, opinion, and action, as three essential aspects of user attitude, have been studied separately in various existing research work. Investigating them together not only brings unique challenges but can also help better understand a user's online behavior and benefit a set of applications related to online campaign and recommender systems. In this paper, we present a computational model that estimates individual social media user's attitude toward controversial topics in terms of the three aspects and their relationships. Our model can simultaneously capture the three aspects so as to predict action and sentiment based on one's opinions. Experiments on multiple social media campaign datasets demonstrated that our attitude model can more effectively predict people's sentiment, opinion and action than approaches that treat these aspects separately.

## Introduction

Recent years micro-blogging platforms such as Twitter have seen a rapid growth, and the emergence of social media campaigns, where massive number of people express strong opinions and provide support for social causes of public interest. Social media users often show varied attitude toward such campaigns. Some people may support a campaign for one reason, while others may support the same campaign for some other reasons; some people may hold the opinion to themselves, while others may actively help propagating relevant information to other people.

As a concrete example, consider the social media campaign against "fracking". Fracking, or hydraulic fracturing, is the process of extracting natural gas from shale rock layers, and the process has been hotly debated in the public due to its potential impact on energy and environment. Two social media users in the campaign against fracking may hold the same negative sentiment toward fracking due to different opinions. For instance, one user Joe may support the opinion that fracking causes damage to environment, believing that fracking should be immediately stopped. Meanwhile,

another user Bill may also believe that fracking harms environment, but is meanwhile against the position of stopping fracking completely, believing that better regulation of fracking is called for. Due to their different opinions, Joe and Bill may have different tendency to spread a petition that calls for stopping fracking, despite their shared negative sentiment. Such nuanced relationships between sentiment, opinion, and action has not been captured well by traditional sentiment or opinion analysis work (Abu-Jbara, Hassan, and Radev 2012; Jiang et al. 2011; Tan et al. 2011; Somasundaran and Wiebe 2010). Meanwhile, prior behavior prediction work on social media (e.g., (Yang et al. 2010; Feng and Wang 2013) for predicting replies, retweets on Twitter) are agnostic on the underlying opinions for observable behaviors, thus missing the potential effect of opinions in their prediction efforts.

Motivated by this gap, we present a unified computational model that captures people's sentiment toward a topic, their specific opinion, and their likelihood of taking an action. Our model is inspired by an established theoretical framework in psychological and marketing research on *attitudes* and *attitude models*, where attitude is defined as a unified concept containing three aspects: "*feelings*", "*beliefs*", and "*actions*" (McGuire 1968; Eagly and Chaiken 1993; Schiffman and Kanuk 2010). According to the framework, beliefs are acquired on *attitude object* (e.g., a topic, product or person), which in turns influences the feelings on the object and the actions w.r.t. the attitude object. Our computational model operationalizes this framework mathematically, casting feelings, beliefs, and actions into users' sentiment, opinion, and action toward a topic on social media.

Figure 1 shows an illustrative example of user attitude toward a controversial topic (fracking) on Twitter. At sentiment level, it shows two sentiments toward fracking (support fracking vs. oppose fracking). Note that, a user may neither support nor oppose fracking. However, for clarity we do not include such neutral sentiment in the example. At opinion level, a user may have one or more opinions w.r.t. different facets of fracking. For example, "Fracking damages environment" is an opinion regarding to the "environment" facet of fracking and the example tweet on the left side of Figure 1 contains that opinion. Similarly, "Fracking is safe" is an opinion regarding to the "safety" facet, and the example tweet on the right side of Figure 1 contains that opinion.

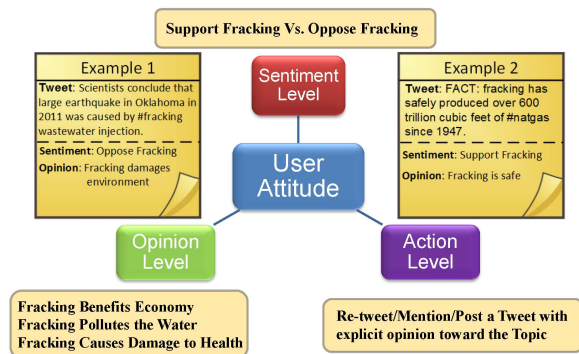


Figure 1: Illustrative Example of User Attitude toward a Controversial Topic (Fracking) on Twitter

Each opinion has a sentiment associated with it (the first opinion is negative toward “fracking” and the second opinion is positive toward “fracking”). A user who has multiple opinions may have an overall sentiment toward the topic (not shown in Figure 1). Finally, at action level, a user may retweet/mention/post a tweet containing such opinions.

The key to model a user’s attitude in terms of sentiment, opinion and action is to capture the relationships among them. In this work, we capture such relationships using a feature-based collaborative filtering method so as to predict action and sentiment based on one’s opinions. To demonstrate the effectiveness of our approach, we have conducted extensive experiments with two social media campaign datasets. Our main contributions are the following:

- We propose a feature-based collaborative filtering method to estimate a user’s attitude toward a controversial topic. Our method mathematically models a user’s opinion, sentiment and action together.
- We capture the relationships between a user’s opinions and his sentiment/actions. Our model is able to provide explicit explanations to a user’s action and sentiment toward a controversial topic through his opinions.
- We perform extensive experiments with multiple social media campaign datasets to demonstrate that our model outperforms a number of baselines in predicting sentiment, opinion and action.

## Related Work

Opinion and sentiment analysis is widely researched. There are prior works on detecting sentiment from various forms of textual data such as documents, blogs, tweets. Li et al. proposes a Topic-Level Opinion Influence Model (TOIM) that simultaneously incorporates topic factor and social influence in a two-stage probabilistic framework (Li et al. 2012). Lin et al. proposes a unsupervised probabilistic modeling framework based on Latent Dirichlet Allocation (LDA), which detects sentiment and topic simultaneously from text (Lin and He 2009). Li et al. described sentiment detection from micro-blogs using collaborative learning approach (Li et al. 2010). Hu et al. leverage social relations for sentiment analysis in both supervised and unsupervised ways (Hu et al.

2013a; 2013b). Guerra et al. measured the bias of social media users toward a topic, by solving a relational learning task over a network of users connected by endorsements (e.g., retweets in Twitter) (Calais Guerra et al. ).

There are also prior works to aggregate message level sentiments to infer a social media user’s overall sentiment toward a target, or predicting a user’s opinion/sentiment toward a topic/target. Kim et al. described user-level sentiment prediction using collaborative filtering approach (Kim et al. 2013). Tan et al. described user-level sentiment analysis using social network information (Tan et al. 2011).

However, none of the existing research on opinion and sentiment analysis predicts the likelihood of taking an action based on the current sentiment/opinion toward a topic. In addition, they did not describe the relationship of user actions, opinions and sentiments. In our work, we mathematically model such relationships, and predict the likelihood of an action and sentiment based on user’s opinions.

Meanwhile, a number of prior works exist on predicting user actions on social media, such as predicting replies (Mahmud et al. 2013), retweets (Yang et al. 2010; Feng and Wang 2013) or follow behaviors (Hannon, Bennett, and Smyth 2010). However, such works do not address the problem of predicting actions based on a current opinion of a user toward a target. In contrast to existing works on action prediction, our work predicts action which is the result of an opinion about a topic.

## Methodology

In this section, we introduce our methodology for attitude modeling. We focus on Twitter users, although our core technology can be easily applied to other social media platforms. On Twitter, people often express their attitude toward a topic by retweeting a tweet containing an opinion. For concreteness, we only consider retweet actions which are result of different opinions. A user’s retweeting action toward a target (a tweet) is driven by two factors: the user who performs the action on the target, and the target which is acted on. Different users may have different preferences on different targets (tweets), resulting in various retweeting actions. Therefore, it is essential to capture a user’s preferences toward a tweet when inferring his/her retweeting behavior.

Collaborative filtering (CF) methods are commonly used to infer a user’s preferences toward a target in recommender systems (Schafer et al. 2007). Among various CF approaches, model-based methods such as matrix factorization (MF) have been shown to be an efficient approach which infers a user’s preferences based on latent dimensions while addressing the data sparseness problem (Koren 2008). This inspires us to adopt MF approach for retweeting action inference. i.e., inferring a user’s probability of taking an action on the target (retweeting a tweet). On the other hand, it also provides opportunities to model a user’s sentiment and opinion (to be discussed later). In the following sections, we first introduce the basic matrix factorization model and discuss how it provides opportunities to infer a user’s retweeting action together with opinions and sentiment, then discuss the challenges of modeling user attitude with basic model, and finally propose our model **ATMiner**.

Table 1: Opinions in the Fracking Dataset

+/-	Opinion	Description	Tweet Example
+	Economy and Energy	Fracking benefits economy and energy	Fracking saves us money; fracking creates jobs; fracking reduces greenhouse gas emissions. @Ameratex
	Safety	Fracking is safe	FACT: fracking has safely produced over 600 trillion cubic feet of #natgas since 1947.
-	Oil Spill	Fracking causes oil spill	Lives in a pineapple under the sea. BP oil spill.
	Environment	Fracking damages environment	Scientists conclude that large earthquake in Oklahoma in 2011 was caused by #fracking wastewater injection.
	Health	Fracking causes health problems	To anyone speaking of the economic "benefits" of fracking: what use is that money if your food and water are full of poison.
	Economy	Fracking does not help economy	The amount of money BP lost from the oil spill could buy about 30 ice cream sandwiches for everyone on earth.
	General	Fracking is bad	Yoko Ono took a tour of gas drilling sites in PA to protest fracking. Suddenly she's against breaking up rock groups.
	Call for Action	Fracking should be stopped	#DontFrackNY @NYGovCuomo Protect our kids and families from #fracking. Please RT!

### Basic Model for Attitude Modeling

We first introduce the basic low-rank matrix factorization (**MF**) model for retweeting action inference. Let  $\mathbf{u} = \{u_1, u_2, \dots, u_m\}$  be the set of users, and  $\mathbf{v} = \{v_1, v_2, \dots, v_n\}$  be the set of tweets, where  $m$  and  $n$  denote the number of users and tweets, respectively.  $\mathbf{R} \in \mathbb{R}^{m \times n}$  is a user-tweet matrix with each element  $\mathbf{R}_{ij}$  representing the retweeting action made by user  $u_i$  toward a tweet  $v_j$ .  $\mathbf{R}_{ij} = 1$  indicates that there is a retweeting action, and 0 otherwise.

Let  $\mathbf{U} \in \mathbb{R}^{m \times d}$  be the user latent preferences and  $\mathbf{V} \in \mathbb{R}^{n \times d}$  be the tweet latent profile, with  $d \ll \min(m, n)$  being the number of latent factors. The basic **MF** model approximates  $u_i$ 's retweeting preference on a tweet  $v_j$  via solving the following optimization problem:

$$\min_{\mathbf{U}, \mathbf{V}} \|\mathbf{R} - \mathbf{UV}^T\| + \mathcal{L}(\mathbf{U}, \mathbf{V}) \quad (1)$$

Where  $\mathcal{L}(\mathbf{U}, \mathbf{V})$  represents appropriate regularization terms w.r.t. the latent factors.

The basic **MF** approach has been proven to be effective in various recommender systems. It provides opportunities for modeling a user's attitude due to the following properties:

- It captures users' latent preferences and tweets' latent profiles, therefore is able to estimate a user's essential retweeting preference and predict his/her retweet action toward a tweet.
- A user's opinions indicate his preferences on different opinions regarding to the topic, while his latent preferences obtained from **MF** indicate his preferences on latent factors of the topic. Here, an opinion corresponds to an explicit factor of the topic. For example, "fracking is safe" is an opinion of "fracking", as expressed through the tweet "FACT: fracking has safely produced over 600 trillion cubic feet of #natgas since 1947". In many applications (e.g., clustering), latent factors obtained from **MF** can be summarized to explicit factors. Thus, we expect that a user's latent preferences are able to provide a potential indication to his opinions toward the topic through

the relationship between latent factors and explicit factors, indicating the opportunity for modeling his opinions.

- The basic **MF** model is flexible and allows us to include prior knowledge such as observed sentiment information, introduced in the next section.

Thus, we start from the basic **MF** model for modeling user attitude. However, it cannot be directly applied to attitude modeling with online campaign data, due to the specific data properties and challenges.

#### • Cold-start users in online campaign

In online campaign, a user may only take one or few actions to express his sentiment/opinions toward the topic, e.g., retweeting an opinion-oriented tweet related to the topic, resulting in an extremely sparse behavior history, which causes the cold-start problem regarding to users who have none or very few observed retweeting actions, widely known in recommender systems. As pointed by (Schein et al. ), cold-start user behavior modeling presents significant challenges in capturing user preferences due to the lack of sufficient historical observations.

#### • Implicit opinions

In the basic **MF** model, a user  $u_i$ 's latent preferences indicate his preferences on latent factors of the topic. Although such preferences are conceptually similar as his opinions, they cannot be explicitly described. Thus, strategies need to be introduced to bridge the gap between the latent preferences factorized by **MF** and explicit opinions expected to be returned as output.

#### • Unknown overall sentiment

The basic **MF** approach models a user's actions through latent user preferences, while the overall sentiment is not considered. However, a user may present multiple opinions containing both positive and negative sentiment, which raises challenges to infer his/her overall sentiment.

According to the above challenges, we propose **ATMiner** to model user attitude in online social media in terms of *feature*

selection for preference approximation, opinion regularization, and sentiment regularization.

### Feature Selection for Preference Approximation

According to the cold-start problem, techniques relying on various observations of historical retweeting actions may fail due to the lack of information. However, social media sites provide abundant user-related information, such as posts, friendships, etc. Such information could be potentially utilized as user-related features to approximate a user’s preferences without observing sufficient user actions. Thus, we introduce user-related features to estimate a user’s latent preferences as shown below,

$$\min_{\mathbf{W}, \mathbf{V}} \|\mathbf{R} - \mathbf{F}(\mathbf{W}, \mathbf{X})\mathbf{V}^\top\| + \varphi \|\mathbf{W}\|_1 + \mathcal{L}(\mathbf{W}, \mathbf{V}) \quad (2)$$

where  $\mathbf{F}(\mathbf{W}, \mathbf{X}) = \mathbf{X}\mathbf{W}^\top$  is a linear function with  $\mathbf{X} \in \mathbb{R}^{m \times f}$  as user-related features and  $\mathbf{W} \in \mathbb{R}^{d \times f}$  as the feature coefficient,  $f$  denotes the dimension of user feature space. Considering the large amount of user-generated content in social media, the feature space is usually in large dimension. Therefore, we introduce Lasso (Hastie, Tibshirani, and Friedman 2001) (a widely used sparse regularization in many data mining applications) into our model for simultaneously feature selection purpose. Here,  $\|\mathbf{W}\|_1$  is the corresponding sparse regularization term, where  $\|\cdot\|_1$  represents 1-norm of a matrix with  $\|\mathbf{W}\|_1 = \sum_i \sum_j \mathbf{W}_{i,j}$ ,  $\varphi$  is the parameter to control the feature sparsity.

### Opinion Regularization

The basic MF model factorizes the action matrix into user preferences and tweet profiles. As we are interested in discovering a user’s opinions, a proper strategy is necessary to constrain the user preferences on latent (implicit) factors into explicit opinions. Thus, we propose an opinion regularization to constrain the user-related latent factor with the context of observed opinions, as shown below,

$$\min_{\mathbf{W}} \mathcal{L}_O = \|\mathbf{F}(\mathbf{W}, \mathbf{X}) - \mathbf{O}\| \quad (3)$$

$s.t. \mathbf{W} \geq 0$

Here,  $\mathbf{O} \in \mathbb{R}^{m \times d}$  denotes user-opinion distribution observed from training data. Each element  $\mathbf{O}_{i,j}$  is a categorical value representing user  $i$ ’s preferences on opinion  $j$ . By minimizing the distance between factorized user latent preferences and observed user-opinion distribution, we could force the factorized latent user preferences bounded in the opinion space, therefore making the implicit latent preferences explicitly representing user opinions. It is worth noting that a non-negative constraint on the opinion space has been introduced in Eq. (3), since the opinion strength in real-world is commonly non-negative.

### Sentiment Regularization

Since a user may hold various opinion containing both positive and negative aspects, determining the overall sentiment from a user’s opinions becomes difficult as the relationships

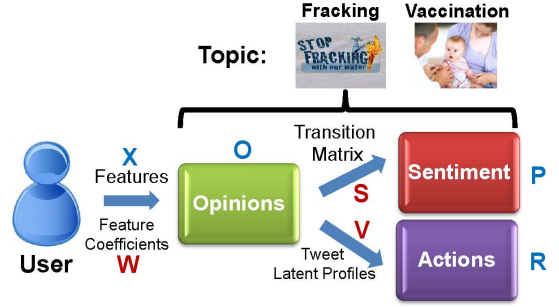


Figure 2: Modeling Relationships among Sentiment, Opinions and Actions.

among them are unknown. In this paper, we introduce a transition matrix to capture such relationships under the following sentiment constraint.

$$\min_{\mathbf{W}} \mathcal{L}_P = \|\mathbf{F}(\mathbf{W}, \mathbf{X})\mathbf{S} - \mathbf{P}\| \quad (4)$$

$s.t. \mathbf{W} \geq 0, \mathbf{S} \geq 0$

$\mathbf{S} \in \mathbb{R}^{d \times k}$  denotes a opinion-sentiment transition matrix with  $k$  the number of sentiment polarities (In this work,  $k$  is set to 2 representing positive and negative).  $\mathbf{P} \in \mathbb{R}^{m \times k}$  denotes user-sentiment distribution observed from training data. The non-negative constraint of the transition matrix is also introduced to capture the non-negativeness of sentiment strength.

### ATMiner: Modeling User Attitude in Online Social Media

We have discussed how to extend the basic MF methodology for modeling attitudes of users. Now we introduce our proposed model, **ATMiner**, for modeling user attitude

$$\min_{\mathbf{W} \geq 0, \mathbf{S} \geq 0, \mathbf{V} \geq 0} \|\mathbf{R} - \mathbf{F}(\mathbf{W}, \mathbf{X})\mathbf{V}^\top\|_F^2 + \lambda \|\mathbf{F}(\mathbf{W}, \mathbf{X}) - \mathbf{O}\|_F^2 \quad (5)$$

$$+ \eta \|\mathbf{F}(\mathbf{W}, \mathbf{X})\mathbf{S} - \mathbf{P}\|_F^2 + \varphi \|\mathbf{W}\|_1$$

$$+ \alpha (\|\mathbf{W}\|_F^2 + \|\mathbf{V}\|_F^2 + \|\mathbf{S}\|_F^2)$$

$s.t. \mathbf{F}(\mathbf{W}, \mathbf{X}) = \mathbf{X}\mathbf{W}^\top.$

where  $\lambda$  and  $\eta$  controls the opinion regularization and sentiment regularization, respectively. A small  $\lambda$  ( $\eta$ ) indicates a weak relationship between the factorized latent factor and the observed opinions (sentiment), while a large  $\lambda$  indicates that we force them to be as close as possible.  $\alpha$  is a parameter introduced to avoid over-fitting.

Figure 2 gives an illustration of how the model works. A user holds opinions toward the topic, which is approximated through user-related features. The opinions and tweet latent profile together results in an action (retweeting). Opinions also transfer to sentiment through a transition matrix. Input to the model is a set of training users with their corresponding sentiment distribution  $\mathbf{P}$ , opinions distribution  $\mathbf{O}$ , action history  $\mathbf{R}$ , and user features  $\mathbf{X}$ . Outputs are Feature Coefficients  $\mathbf{W}$  (Opinion Level), Transition Matrix  $\mathbf{S}$  (Sentiment Level), and Tweet Latent Profile  $\mathbf{V}$  (Action Level).

Alternative algorithm is commonly used as it is difficult to provide a direct closed-form solution for the above optimization problem. Thus, we apply an alternative algorithm

Table 2: Opinions in the Vaccination Dataset

+/-	Opinion	Description	Tweet Example
+	General_Positive	Positive Information (Opinion) about vaccination	Vaccination campaign launches with hope of halting measles outbreak <a href="http://t.co/H2B6ujFx22">http://t.co/H2B6ujFx22</a>
	Call for Action	Vaccination should be continued	To not vaccinate is like manslaughter. Vaccinate!
	CounterNegative	Counter negative information about vaccination	Six vaccination myths - and why they're wrong. <a href="http://t.co/BX7kq0SOjz">http://t.co/BX7kq0SOjz</a>
-	General_Negative	Negative Information (Opinion) about vaccination	Vaccination has never been proven to have saved one single life.
	Sideeffect	Vaccination causes disease	Until the #Vaccination was introduced RT @trutherbot: Cancer was a rarity more than 200 years ago.
	No-enforcement	Criticize forced vaccination	Police State? Registry System Being Set Up to Track Your Vaccination Status - <a href="http://t.co/fkSWDbYAbB">http://t.co/fkSWDbYAbB</a>

Table 3: Statistical Information of the Datasets

	Fracking	Vaccination
No. of Users	5,387	2,593
No. of Positive Users	1,562	1617
No. of Negative Users	3,822	976
Duration	1/13-3/13	5/13-10/13
No. of Historical Tweets	458,987	226,541
No. of Opinions	8	6
No. of Action Tweets	162	105
No. of Features	10,907	4,803

to find optimal solutions for the three variables  $\mathbf{W}$ ,  $\mathbf{S}$ , and  $\mathbf{V}$ . The key idea is to minimize the objective function w.r.t. one variable while fixing the other variables, as similar to (Ding, Li, and Jordan 2008). The algorithm will keep updating the variables until convergence or reaching the number of maximum iterations. Due to the space limit, we will not present the detailed inference but list the final updating rules below,

$$\begin{aligned}
\mathbf{W}(i, j) &\leftarrow \mathbf{W}(i, j) \sqrt{\frac{[\mathbf{V}^T \mathbf{R}^T \mathbf{X} + \lambda \mathbf{O}^T \mathbf{X} + \eta \mathbf{S} \mathbf{P}^T \mathbf{X}](i, j)}{A(i, j)}} \\
\mathbf{S}(i, j) &\leftarrow \mathbf{S}(i, j) \sqrt{\frac{\eta(\mathbf{W} \mathbf{X}^T \mathbf{P})(i, j)}{[\eta(\mathbf{W} \mathbf{X}^T \mathbf{X} \mathbf{W}^T \mathbf{S}) + \alpha \mathbf{S}](i, j)}} \\
\mathbf{V}(i, j) &\leftarrow \mathbf{V}(i, j) \sqrt{\frac{(\mathbf{R}^T \mathbf{X} \mathbf{W}^T)(i, j)}{[\mathbf{V} \mathbf{W} \mathbf{X}^T \mathbf{X} \mathbf{W}^T + \alpha \mathbf{V}](i, j)}} \quad (6)
\end{aligned}$$

where  $A$  is defined as

$$A = \mathbf{V}^T \mathbf{V} \mathbf{W} \mathbf{X}^T \mathbf{X} + \lambda \mathbf{W} \mathbf{X}^T \mathbf{X} + \eta \mathbf{S} \mathbf{S}^T \mathbf{W} \mathbf{X}^T \mathbf{X} + \varphi e_o e_x^T + \alpha \mathbf{W} \quad (7)$$

The time complexity of the above learning algorithm is  $O(mnf)$ , where  $m$  is the number of users,  $n$  the number of tweets, and  $f$  the number of features.

### Predicting User Attitude

Once we have trained our attitude model, we can apply the model to predict attitude of a new user  $u$  in a test set. In particular, we do the following predictions:

- **Opinion Prediction**

For a user  $u$  in test set, by obtaining the corresponding features  $X_u$ , the model can predict his/her opinion  $O_u$  through

$O_u = F(X_u, W) = X_u W^T$ , where  $W$  is the feature coefficient that learned from the model. Note that in our experiment, a user can hold more than one opinions, thus, this task corresponds to a multi-label classification problem.

- **Sentiment Prediction**

The sentiment of a test user  $u$  is estimated through the user-related features  $X_u$ , and the opinion-sentiment transition matrix  $S$  learned from our model, i.e.,  $P_u = F(X_u, W)S = X_u W^T S$ .

- **Retweeting Action Inference**

The probability of a test user  $u$  taking an action on a target tweet  $t$  is estimated through the user-related features  $X_u$ , and the tweets latent profile  $V_j$ , i.e.,  $R_{i,j} = X_u W^T V_j^T$ . Similar to opinion prediction, a user is allowed to retweet more than one tweet, therefore this task also corresponds to a multi-label classification problem.

## Experiments

In this section, we evaluate our attitude model **ATMiner** on two real-world datasets. In particular, we evaluate the model in terms of three tasks, i.e., sentiment prediction, opinion prediction, and action inference. Before we delve into experimental details, we first discuss our online campaign datasets, experiment setup, and baseline models.

### Datasets

We select ‘‘fracking’’ and ‘‘vaccination’’ as the controversial topics for investigating user attitude. We use Twitter’s streaming API to obtain 1.6 million tweets related to fracking topic from Jan, 2013 to March, 2013, with a set of fracking-related keywords (such as fracking, frac flowback, frac fluid, frackivist, fracturing, frac water, shale, hydrofracturing). For vaccination dataset, we obtained 1.1 million tweets related to vaccination topic from May, 2013 to Oct, 2013, with a set of vaccination-related keywords (such as vaccination, vaccine, vaccines, vaccinate). For each dataset, we ranked all the crawled tweets based on their retweeted times, and select those which are retweeted for more than 100 times as our action tweets. There were 162 action tweets in ‘‘fracking’’ dataset and 105 action tweets in ‘‘vaccination’’ dataset. From these action tweets, we obtained users who retweeted them. There were 5387 users who retweeted

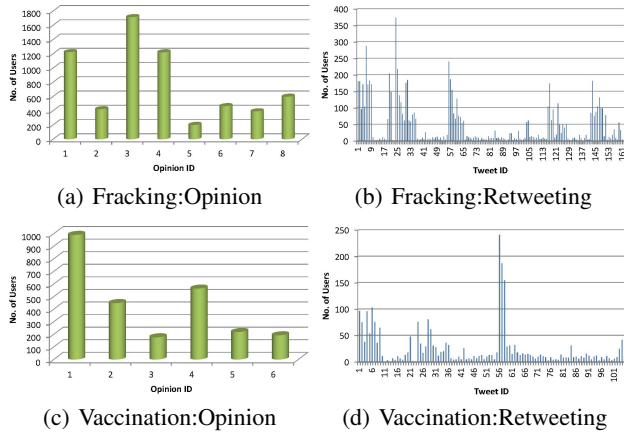


Figure 3: Opinion Distribution and Retweeting Distribution among Users in Fracking and Vaccination Datasets

Table 4: Performance Comparison of Sentiment Prediction

Method	$F_{Micro}$ (Fracking)	$F_{Micro}$ (Vaccination)
Random	0.5204	0.5014
Least Squares	0.6702	0.7552
Lasso	0.7686	0.7840
ATMiner	<b>0.7896</b>	<b>0.8130</b>

an action tweet in “fracking” dataset and 2593 users who retweeted an action tweet in “vaccination” dataset.

Training our model requires ground truth of sentiment, opinion and action of these users. Ground truth of action for each such user is readily available from the retweet of action tweets (for each action tweet in our dataset, the user either retweeted the tweet or not). To generate the ground truth of sentiment and opinion of these users, we followed a supervised approach. However, instead of manually labelling each user in our datasets, we manually labeled only the action tweets (162 for “fracking” and 105 for “vaccination”). Specifically, we summarized eight fracking-related opinions and six vaccination related opinions listed in Table 1 and Table 2, where “+” and “-” denote positive and negative sentiment associated with the corresponding opinion. We then assign these opinions to users based on their corresponding retweeting actions. The assignment follows the traditional assumption of studying user retweeting behavior, that when a user retweets a tweet, we assume he endorses the tweet content (Boyd, Golder, and Lotan 2010; Conover et al. 2011; Welch et al. 2011; Calais Guerra et al. ). The assignment of each user on different opinions are considered as opinion ground truth. We then label each user’s sentiment based on his opinion assignment, i.e., if the majority opinions assigned to this user are positive, the user is labeled as positive, otherwise negative. For each user in our dataset, we also crawled historical tweets (100 max) that were posted before the time when the user first retweeted an action tweet. These historical tweets are used to generate user features.

Table 3 shows the statistical information of our datasets. Figure 3 plots the opinion distribution and retweeting distribution among all the users in our datasets.

## Experiment Setup

We construct user features based on users’ historical tweets. We employ unigram model while removing stop-words to construct the feature space, and use term frequency as feature value. To generate the training and testing set, we split users into 10 folds, with each fold containing 10% of total users. We keep seven folds (70%) for training and the remaining (30%) for testing. To train the model,  $O$  is obtained from the ground truth of users’ opinion distribution in training data. Similarly,  $P$  is obtained from ground truth of users’ sentiment distribution in training data and  $R$  is obtained from ground truth of users’ action history in training data.

All the parameters of our model are set through cross-validation. Specifically, we set  $\lambda = 0.5$ ,  $\eta = 0.5$ ,  $\varphi = 2$ , and  $\alpha = 0.1$ . We will discuss the effectiveness of the parameters in later part of the paper.

For each user in the training data, we obtain his ground-truth sentiment, opinion, and retweeting actions to train the model. For each user in our test data, we perform three prediction tasks described in the previous section.

Since our prediction model relies on user features, we compare our **ATMiner** with supervised learning approaches. Three baseline methods are introduced as below.

### • Random Prediction

We apply a random predictor on the dataset to evaluate the prediction difficulty of our problem. The random predictor randomly predicts a user’s sentiment to be positive or negative. For opinion prediction and retweeting action inference, due to their multi-label property, it randomly picks one or more opinions/retweets for prediction according to the multi-label evaluation metrics (to be discussed later).

### • Least Squares

Least Squares (**LS**) is a state-of-the-art approach for supervised learning (Bishop and Nasrabadi 2006). For each prediction task, we use user features and corresponding labels to train LS, which is then applied on the test data.

### • Lasso

Due to the large number of user features, we also introduce “Lasso” (Hastie, Tibshirani, and Friedman 2001), one of the most popular sparse learning methods, as our baseline method for comparison.

We do not consider the basic matrix factorization model as a baseline since **MF** immediately fails due to the large number of cold-start users.

## Evaluation Metrics

Since the opinion prediction and action inference involve a multi-label classification problem, we adopt micro-averaged F-measure to evaluate the prediction performance, as defined below:

$$F_{Micro} = \frac{(1 + \beta^2)P \cdot R}{\beta^2 P + R} \quad (8)$$

where  $\beta$  is a parameter to control the weight between precision and recall.  $\beta = 1$  is commonly used, and we consider

Table 5: Performance Comparison of Opinion Prediction and Action Inference

Dataset	Methods	Opinion Prediction				Action Inference			
		$P@1$	$R@1$	$F@1$	$F_{Overall}$	$P@1$	$R@1$	$F@1$	$F_{Overall}$
Fracking	Random	0.1646	0.1425	0.1527	0.1646	0.0087	0.0061	0.0071	0.0121
	Least Squares	0.3401	0.3058	0.3221	0.3589	0.1082	0.0762	0.0894	0.1148
	Lasso	0.4302	0.3910	0.4097	0.4488	0.1280	0.0903	0.1059	0.1221
	ATMiner	<b>0.4681</b>	<b>0.4210</b>	<b>0.4433</b>	<b>0.4869</b>	<b>0.1527</b>	<b>0.1077</b>	<b>0.1263</b>	<b>0.1451</b>
Vaccination	Random	0.1862	0.1702	0.1778	0.1832	0.0121	0.0132	0.0126	0.0267
	Least Squares	0.2401	0.3078	0.2697	0.3243	0.1071	0.1092	0.1081	0.1197
	Lasso	0.2848	0.3575	0.3170	0.3685	0.1170	0.1103	0.1135	0.1521
	ATMiner	<b>0.3359</b>	<b>0.4752</b>	<b>0.3935</b>	<b>0.4272</b>	<b>0.1329</b>	<b>0.1273</b>	<b>0.1300</b>	<b>0.1676</b>

Table 6: Prediction Performance ( $F_{Micro}/F_{Overall}$ ) w.r.t. Various Training Data Size

Tasks	Methods	Fracking				Vaccination			
		$T_1$	$T_3$	$T_5$	$T_7$	$T_1$	$T_3$	$T_5$	$T_7$
Sentiment Prediction ( $F_{Micro}$ )	Random	0.5204	0.5204	0.5204	0.5204	0.5014	0.5014	0.5014	0.5014
	Least Squares	0.6200	0.6337	0.6411	0.6702	0.6334	0.6813	0.7326	0.7552
	Lasso	0.7246	0.7321	0.7611	0.7686	0.6847	0.7435	0.7714	0.7840
	ATMiner	<b>0.7444</b>	<b>0.7525</b>	<b>0.7772</b>	<b>0.7896</b>	<b>0.7341</b>	<b>0.7628</b>	<b>0.8078</b>	<b>0.8130</b>
Opinion Prediction ( $F_{Overall}$ )	Random	0.1646	0.1646	0.1646	0.1646	0.1832	0.1832	0.1832	0.1832
	Least Squares	0.3155	0.3305	0.3332	0.3589	0.2555	0.2807	0.3130	0.3243
	Lasso	0.3974	0.4039	0.4462	0.4488	0.3062	0.3339	0.3462	0.3685
	ATMiner	<b>0.4264</b>	<b>0.4349</b>	<b>0.4788</b>	<b>0.4869</b>	<b>0.3564</b>	<b>0.3941</b>	<b>0.4058</b>	<b>0.4272</b>
Action Inference ( $F_{Overall}$ )	Random	0.0121	0.0121	0.0121	0.0121	0.0267	0.0267	0.0267	0.0267
	Least Squares	0.1005	0.0992	0.1083	0.1148	0.0905	0.1092	0.1103	0.1197
	Lasso	0.1009	0.1152	0.1232	0.1221	0.1019	0.1232	0.1331	0.1521
	ATMiner	<b>0.1161</b>	<b>0.1260</b>	<b>0.1416</b>	<b>0.1451</b>	<b>0.1137</b>	<b>0.1430</b>	<b>0.1523</b>	<b>0.1676</b>

that value of  $\beta$  in our experiments.  $P$ ,  $R$  are micro-precision and micro-recall respectively, with each one defined as,

$$P = \frac{\sum_i^c (TP)_i}{\sum_i^c (TP)_i + \sum_i^c (FP)_i}$$

$$R = \frac{\sum_i^c (TP)_i}{\sum_i^c (TP)_i + \sum_i^c (FN)_i} \quad (9)$$

where  $TP_i$  represents the true positive number of class  $i$ ,  $FP_i$  represents the false positive number of class  $i$ ,  $FN_i$  represents the false negative number of class  $i$ , and  $c$  is the total number of classes.

We use the above metrics to evaluate the sentiment classification performance. Since opinion prediction and action inference correspond to multi-label classification, we introduce “@1” and “Overall” for the above metrics to evaluate the two tasks. “@1” indicates that a user can only hold one opinion (or retweet one tweet). Thus, we select one opinion (tweet) with the highest probability as predicted label for the user. “Overall” indicates that we predict  $n$  opinions (retweet actions) for each user where  $n$  is the total number of opinions (retweet actions) that user has in the testing data.

### Comparison of Prediction Performance with Various Approaches

We compare **ATMiner** with the baseline methods on three prediction tasks, under the experimental setting described above. Table 4 and Table 5 show the prediction performance. Note that  $P$  and  $R$  are not presented in Table 4 as they are the

same as  $F_{Micro}$  for binary sentiment classification. Similarly,  $P_{Overall}$  and  $R_{Overall}$  are not presented in Table 5 as they are the same as  $F_{Overall}$  in multi-label classification. The results present a set of key observations:

- The three prediction tasks exhibit different level of prediction difficulty. The performance of random prediction indicates the prediction difficulty of each task. The sentiment prediction has  $F_{Micro}$  around 50% under a random prediction, opinion prediction achieves around 15% to 18%, while the action prediction is the most difficult one among three, achieving only 1% approximately. This is due to the different number of class labels of each task, with the sentiment classification containing two classes, opinion classification containing eight classes for fracking and six classes for vaccination, and the action inference containing the most labels (162 for fracking and 105 for vaccination).
- **Lasso** always performs better than **LS** among all the three tasks, indicating that feature selection is necessary to reduce the large user feature space. Compared to **Lasso**, **LS** does not perform well due to the “curse of dimensionality” from the large number of user features.
- Among all the approaches, **ATMiner** performs the best. It takes advantage of user features to approximate his/her opinion, and simultaneously captures the relationship among a user’s sentiment, opinion, and retweeting actions, resulting in a good performance on all three tasks. The improvement of **ATMiner** over **LS** and **Lasso** indi-

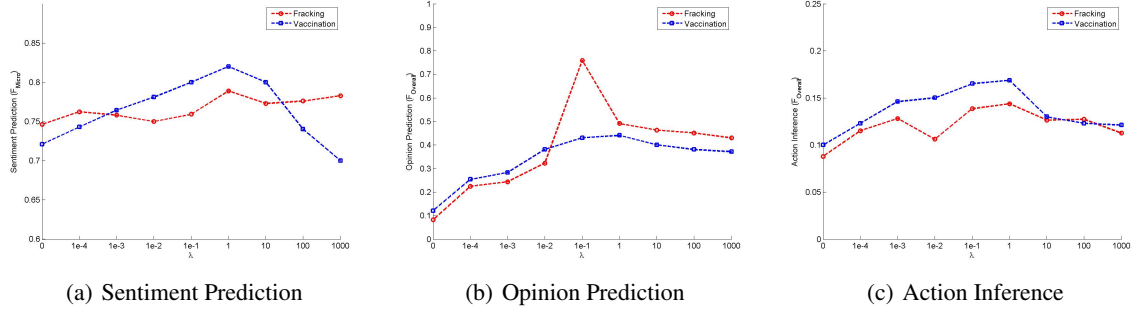


Figure 4: The performance of **ATMiner** w.r.t.  $\lambda$

cates that the relationships among sentiment, opinion, and actions are essential in capturing a user’s attitude, while **ATMiner** is capable to capture such relationships and utilized them for prediction.

### Performance on Various Training Data Sizes

Here, we investigate the importance of the size of training data for each of the prediction tasks. We equally split the dataset into 10 folds and 7 of them are used as training data. To study the performance w.r.t. different size of training data, we randomly select one, three, five, and seven folds from the training data, denoted as  $T_1$ ,  $T_3$ ,  $T_5$ , and  $T_7$ , while keep the testing data fixed. We use  $T_1$ ,  $T_3$ ,  $T_5$ , and  $T_7$  to train the models, respectively, and then apply on the testing data. Table 6 shows the corresponding prediction performance. Due to the space limit, we only present  $F_{Overall}$  for opinion prediction and retweeting action inference, as similar performance can be observed on other metrics. From the results, we summarize the key observations:

- The size of training data affects the prediction performance on all the three tasks. For example, from  $T_1$  to  $T_7$ , **ATMiner** has 6.07%, 14.19%, and 24.98% relative improvement on sentiment prediction, opinion prediction, and action inference respectively with fracking data. Similar performance can be observed with vaccination data.
- For all three tasks, the performance of **ATMiner** is sensitive to the size of training data. However, the performance at  $T_1$  and  $T_3$  indicates that **ATMiner** can afford to work well with a relatively low amount of training data. The reason is that **ATMiner** can capture the relationships among sentiment, opinion and action, thus even with a small amount of labeled data, the information within each category could enhance each other and improve the prediction performance of each task.
- The performance of **ATMiner** on all the three tasks at different sizes of training data is better than baselines, indicating the stability and effectiveness of our model.

### Investigate the Effectiveness of Parameters

Besides  $\alpha$  which controls the over-fitting that is commonly adopted in many learning models, **ATMiner** has three essential parameters:  $\lambda$  to control the opinion regularization,

$\eta$  to control the sentiment regularization, and  $\varphi$  to control the sparseness of feature space. To investigate the parameter effect, we study each parameter by evaluating the model performance when varying the value of one parameter and keeping the other parameter values fixed as their optimal values. Figure 4, Figure 5, and Figure 6 plot the model performance w.r.t.  $\lambda$ ,  $\eta$ ,  $\varphi$ , respectively. The values of three parameters are set as  $\{0, 1 \times e^{-4}, 1 \times e^{-3}, 1 \times e^{-2}, 1 \times e^{-1}, 1, 10, 100, 1000\}$ . From the figures, we observe

- $\lambda$  is a critical parameter to all the three tasks. When  $\lambda$  increases from 0 to  $10^3$ , the prediction performance on three tasks exhibits a similar trend on both datasets, i.e., first increases, reaches its peak, then decreases, indicating the sensitivity of  $\lambda$  to the model performance. The importance of  $\lambda$  can be interpreted through **U** in Figure 2.  $\lambda$  controls the formation of **U**, which is directly related to the opinion. **U** also affects the sentiment through the transition matrix **S**, and affects the retweeting action together with **V**. When  $\lambda$  is taking a small value, the user latent preference **U** can not obtain sufficient information from observed opinion information **O**, resulting in a poor prediction performance in three tasks. On the other hand, when  $\lambda$  is taking a large value, **U** has to be forced to be extremely close to **O**, therefore making the model severely over-fitted on opinion preferences, which affects the other two tasks as well.
- Changing the value of  $\eta$  immediately affects the sentiment prediction performance on both datasets. It increases very fast when  $\eta$  increases from 0 to  $1 \times e^{-4}$ . On the other hand, the performance of opinion prediction and action inference are not quite affected by  $\eta$  when  $\eta$  takes a small value (There is a trend of increase in opinion prediction with vaccination data, while its magnitude is substantially lower than the change in sentiment prediction). Although a small  $\eta$  would make the sentiment regularization worse and result in an inappropriate value of **U** and **S**, **U** can still be correctly learned through the setting of  $\lambda$ , therefore the opinion prediction is not quite affected. Same argument holds for the action inference. When  $\eta$  increases to a large value, performance of both opinion prediction and action inference become poor. This is because a large  $\eta$  bonds the value **U** and **S** into a very small range, therefore even the setting of  $\lambda$  can not help find a optimal value of  $U$  in that small range. This also leads to the poor perfor-



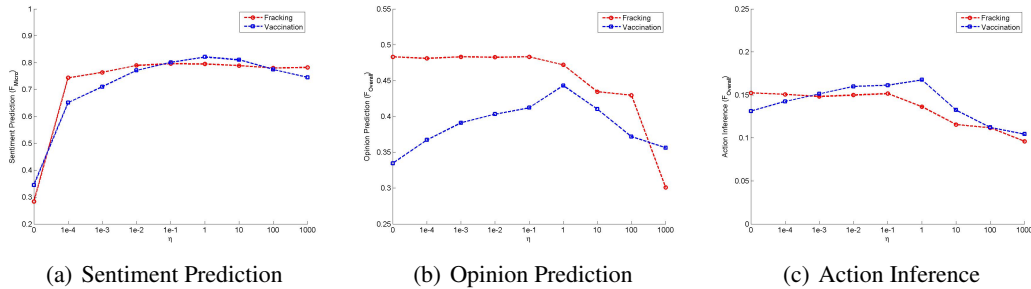


Figure 5: The performance of **ATMiner** w.r.t.  $\eta$

mance of action inference. The performance of three tasks w.r.t  $\eta$  indicates one advantage of **ATMiner**. When there are insufficient observations of users’ sentiment, we could just remove the sentiment regularization while the model could still perform opinion prediction and action inference, which demonstrates the flexibility of **ATMiner**.

- The model performance w.r.t  $\varphi$  on both datasets indicates that opinion prediction and action inference are more sensitive to the feature space than sentiment prediction. Although sentiment prediction may fail due to large feature space without sparseness, the additional information learned from opinion and action could make up this deficiency. The results demonstrate the effectiveness of **ATMiner**, indicating that sparseness does improve the performance and learning sentiment, opinion and action simultaneously could capture the inner relationships and improve the performance as well.

## Discussion

Here we discuss generality of our findings and real world application of our work.

### Generalizability

Since our model is topic-oriented and based on supervised learning, the requirement for topic-dependent training data could be a limitation when applying to a new topic. For each new topic, our method requires pre-defined action tweets (e.g., tweets which are retweeted a certain number of times) and then labelling those action tweets for opinions and sentiment for each opinion. Labels of action (e.g., retweeting) of training users are observable in the training data, hence action labelling does not require any manual effort. Since each such action is on an action tweet (e.g., retweeting an action tweet) and action tweets are manually labelled for opinions, construction of opinion ground truth (i.e. user opinion matrix) for training users does not require any manual effort. Similarly, sentiment labelling of training users is done based on his opinion assignment, and does not need any additional manual effort.

In this work, we have manually labelled the action tweets. In future, we plan to leverage crowd workers from popular crowd sourcing platforms such as Amazon Mechanical Turk for labelling effort. Furthermore, we plan to investigate whether sentiment-LDA (Li, Huang, and Zhu 2010) may be

useful to discover opinions with their sentiment from action tweets.

For the sake of clarity and concreteness, we have only considered retweeting actions in this work. However, our model should be generalizable to other types of actions such as creation of hashtags, mentions or tweeting. Furthermore, our method of attitude modeling should be widely applicable to different social media platforms (e.g., Facebook). In future, we plan to incorporate different actions and study usage of our model in different social media platforms.

### Applicability

Our work can benefit applications related to online campaigns and recommender systems. For example, a government campaign on Twitter supporting vaccination can engage with the followers who are more likely to take certain action (e.g., spreading a campaign message) based on their opinions. A marketing campaign can engage with people who have higher positive sentiment toward the campaign based on their specific opinions. As another example, when anti-government messages are spread in social media, government would want to spread counter messages to balance that effort and hence identify people who are more likely to spread such counter messages based on their opinions. In future, we will integrate our model with different social media campaign applications and study its usage in the real world.

## Conclusions

In this paper, we have presented a computational model to estimate a user’s attitude in terms of sentiment, opinion and likelihood of taking an action toward controversial topics in social media. Through the model, we can capture the relationships among these aspects so as to predict action and sentiment based on one’s opinions. Our model extended traditional matrix factorization approach by usage of features, opinion and sentiment regularization. We have also presented an algorithm to learn the parameters of our model. Our experiments using two real world campaign datasets demonstrate that our model outperforms a number of baselines in predicting sentiment, opinion and action.

## Acknowledgments

Research was sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA) under the Social Media in Strategic Communication (SMISC) program, Agreement

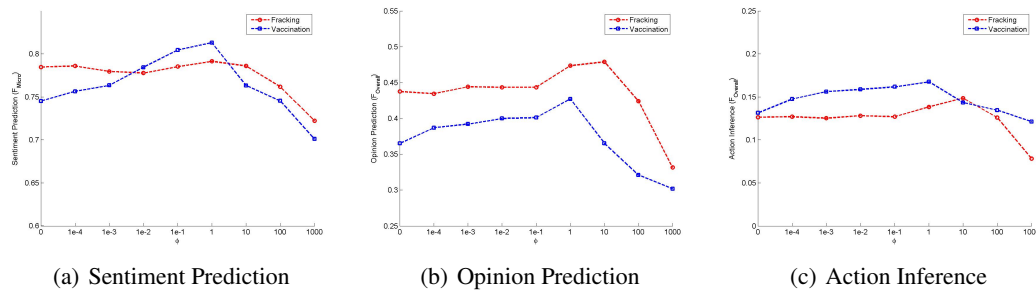


Figure 6: The performance of **ATMiner** w.r.t.  $\varphi$

Number W911NF-12-C-0028. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Defense Advanced Research Projects Agency or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

## References

- Abu-Jbara, A.; Hassan, A.; and Radev, D. 2012. Attitudeminer: mining attitude from online discussions. In *Proceedings of the NAACL-HLT 2012*, 33–36.
- Bishop, C. M., and Nasrabadi, N. M. 2006. *Pattern recognition and machine learning*, volume 1. Springer New York.
- Boyd, D.; Golder, S.; and Lotan, G. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *Proc. of the HICSS '10*, 1–10.
- Calais Guerra, P. H.; Veloso, A.; Meira, Jr., W.; and Almeida, V. From bias to opinion: a transfer-learning approach to real-time sentiment analysis. In *Proc. of the KDD '11*, 150–158.
- Conover, M.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Political polarization on twitter. In *Proc. of the ICWSM 2011*.
- Ding, C.; Li, T.; and Jordan, M. I. 2008. Nonnegative matrix factorization for combinatorial optimization: Spectral clustering, graph matching, and clique finding. In *Proc. of the ICDM 2008*, 183–192.
- Eagly, A., and Chaiken, S. 1993. *The Psychology of Attitudes, 10th edition*. Harcourt Brace Jovanovich.
- Feng, W., and Wang, J. 2013. Retweet or not?: personalized tweet re-ranking. In *Proc. of the WSDM 2013*, 577–586.
- Hannon, J.; Bennett, M.; and Smyth, B. 2010. Recommending twitter users to follow using content and collaborative filtering approaches. In *Proc. of the RecSys '10*, 199–206.
- Hastie, T.; Tibshirani, R.; and Friedman, J. J. H. 2001. *The elements of statistical learning*, volume 1. Springer.
- Hu, X.; Tang, J.; Gao, H.; and Liu, H. 2013a. Unsupervised sentiment analysis with emotional signals. In *Proc. of the WWW 2013*, 607–618.
- Hu, X.; Tang, L.; Tang, J.; and Liu, H. 2013b. Exploiting social relations for sentiment analysis in microblogging. In *Proc. of the WSDM 2013*, 537–546.
- Jiang, L.; Yu, M.; Zhou, M.; Liu, X.; and Zhao, T. 2011. Target-dependent twitter sentiment classification. In *Proc. of the HLT '11*, 151–160.
- Kim, J.; Yoo, J.; Lim, H.; Qiu, H.; Kozareva, Z.; and Galstyan, A. 2013. Sentiment prediction using collaborative filtering. In *Proc. of the ICWSM 2013*.
- Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the SIGKDD 2008*, 426–434. ACM.
- Li, G.; Hoi, S. C. H.; Chang, K.; and Jain, R. 2010. Microblogging sentiment detection by collaborative online learning. In *Proc. of the ICDM 2010*, 893–898.
- Li, D.; Shuai, X.; Sun, G.; Tang, J.; Ding, Y.; and Luo, Z. 2012. Mining topic-level opinion influence in microblog. In *Proc. of the CIKM 2012*, 1562–1566.
- Li, F.; Huang, M.; and Zhu, X. 2010. Sentiment analysis with global topics and local dependency. In *Proc. of the AAAI 2010*, 1371–1376.
- Lin, C., and He, Y. 2009. Joint sentiment/topic model for sentiment analysis. In *Proc. of the CIKM 2009*, 375–384.
- Mahmud, J.; Zhou, M. X.; Megiddo, N.; Nichols, J.; and Drews, C. 2013. Recommending targeted strangers from whom to solicit information on social media. In *Proc. of the IUI 2013*, 37–48.
- McGuire, W. 1968. *The nature of attitudes and attitude change*. Addison-Wesley.
- Schafer, J. B.; Frankowski, D.; Herlocker, J.; and Sen, S. 2007. The adaptive web. Springer-Verlag. chapter Collaborative filtering recommender systems, 291–324.
- Schein, A. I.; Popescul, A.; Ungar, L. H.; and Pennock, D. M. Methods and metrics for cold-start recommendations. In *Proc. of the SIGIR 2002*, 253–260.
- Schiffman, L., and Kanuk, L. 2010. *Consumer Behavior, 10th edition*. Prentice Hall.
- Somasundaran, S., and Wiebe, J. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop*, 116–124.
- Tan, C.; Lee, L.; Tang, J.; Jiang, L.; Zhou, M.; and Li, P. 2011. User-level sentiment analysis incorporating social networks. In *Proc. of the KDD '11*, 1397–1405.
- Welch, M. J.; Schonfeld, U.; He, D.; and Cho, J. 2011. Topical semantics of twitter links. In *Proc. of the WSDM '11*.
- Yang, Z.; Guo, J.; Cai, K.; Tang, J.; Li, J.; Zhang, L.; and Su, Z. 2010. Understanding retweeting behaviors in social networks. In *Proc. of the CIKM '10*, 1633–1636.