

Unsupervised Sentiment Analysis with Signed Social Networks

Kewei Cheng,¹ Jundong Li,¹ Jiliang Tang,² Huan Liu¹

1. Computer Science and Engineering, Arizona State University, Tempe, 85281, USA

2. Computer Science and Engineering, Michigan State University, East Lansing, 48824, USA
{kcheng18, jundongl, huan.liu}@asu.edu, tangjili@msu.edu

Abstract

Huge volumes of opinion-rich data is user-generated in social media at an unprecedented rate, easing the analysis of individual and public sentiments. Sentiment analysis has shown to be useful in probing and understanding emotions, expressions and attitudes in the text. However, the distinct characteristics of social media data present challenges to traditional sentiment analysis. First, social media data is often noisy, incomplete and fast-evolved which necessitates the design of a sophisticated learning model. Second, sentiment labels are hard to collect which further exacerbates the problem by not being able to discriminate sentiment polarities. Meanwhile, opportunities are also unequivocally presented. Social media contains rich sources of sentiment signals in textual terms and user interactions, which could be helpful in sentiment analysis. While there are some attempts to leverage implicit sentiment signals in positive user interactions, little attention is paid on signed social networks with both positive and negative links. The availability of signed social networks motivates us to investigate if negative links also contain useful sentiment signals. In this paper, we study a novel problem of unsupervised sentiment analysis with signed social networks. In particular, we incorporate explicit sentiment signals in textual terms and implicit sentiment signals from signed social networks into a coherent model SignedSenti for unsupervised sentiment analysis. Empirical experiments on two real-world datasets corroborate its effectiveness.

Introduction

The popularity of social media services greatly diversifies the way people communicate and socialize, enabling users to share and exchange opinions in different aspects. The sheer volume of opinion-rich data provides rich sources in understanding individual and public opinions. For example, unveiling the opinions of customers is valuable for business advertisers in devising better targeted marketing tactics (Liu 2012); politicians could also adjust their campaign strategies according to the aggregated sentiments of tweets about election (O'Connor et al. 2010). As a traditional way to identify subjective information from source materials, sentiment analysis has received increasingly attention (Taboada et al. 2011; Kamvar and Harris 2011; Bollen, Mao, and Pepe

2011; Hu et al. 2013b; Tang et al. 2015b). Also, understanding sentiments can naturally advance a variety of real-world applications, such as recommendations, marketing and disaster relief (Ding and Liu 2007; Pang and Lee 2008; Zhang et al. 2014).

Traditional sentiment analysis methods either work in a supervised way to build classifiers from manually annotated sentiment labels (Pang and Lee 2004; Pang, Lee, and Vaithyanathan 2002) or are performed in an unsupervised scenario with a pre-defined sentiment lexicon (O'Connor et al. 2010; Wiebe, Wilson, and Cardie 2005; Wilson, Wiebe, and Hoffmann 2005). More often than not, social media data is distinct from traditional i.i.d. text data – they are not independently created but are inherently linked by user interactions. Another unique property is that social media data is often unlabeled, while sentiment labels are costly and labor-intensive to obtain. Motivated by sentiment consistency (Abelson 1983) and emotional contagion (Hatfield, Cacioppo, and Rapson 1994) in social science theories, rich sources of sentimental signals may exist among user interactions, and there are a surge of research (Hu et al. 2013b; Wang et al. 2015a; Tang et al. 2015b) attempting to exploit user interactions in understanding and predicting sentiment polarity of social media data. Nonetheless, most of them are supervised or semi-supervised by employing feature selection techniques (Li et al. 2016).

Aforementioned approaches predominantly focus on unsigned social networks where only positive user interactions are observed. In addition to positive links, many real-world social media platforms also consist of negative links, such as distrust relations in Epinions¹ and foes in Slashdot². The availability of negative links (Leskovec, Huttenlocher, and Kleinberg 2010b) brings about richer source of information and recent advances in signed network mining show that negative links have some added value over positive interactions. Furthermore, many learning tasks (Leskovec, Huttenlocher, and Kleinberg 2010a; Tang, Aggarwal, and Liu 2016) are enhanced by the modeling of negative links. Recent advances in signed social network analysis motivate us to investigate if negative links could also help us perform sentiment analysis, especially when the sentiment labels are

¹<http://www.epinions.com/>

²<https://slashdot.org/>

scarce or even unavailable.

Despite the potential opportunities from negative links, the development of a principled learning model for unsupervised sentiment analysis with signed social networks is still in its infancy. The reason can mainly be attributed as follows: (1) Different from positive links, negative links carry out different sentiment information. For example, trust information is often a good indicator of positive emotions such as joy and altruism; while distrust relations may be indicators of negative emotions like anger and pessimism. Hence, sentiment analysis with signed social networks can not simply be extended in a straightforward way; (2) Majority of existing sentiment analysis methods with unsigned social networks are based on some social theories, assuming that sentiment may spread along positive interactions and individuals tend to share similar opinions when they are connected. Nonetheless, these theories may not be directly applicable to signed social networks where individuals with negative links may show contrastive opinions. Hence, performing sentiment analysis with signed social networks is not a trivial problem.

In this paper, we study the problem of sentiment analysis with signed social networks under an unsupervised scenario. In essence, we aim to answer the following two questions: (1) Do the positive and negative interactions among users reveal different sentiment polarities in the text? (2) How to explicitly model positive and negative interactions among users for sentiment analysis in an unsupervised way? To answer these two questions, we propose an unsupervised sentiment analysis framework - SignedSenti. The main contributions are summarized as follows:

- We verify that positive and negative interactions among users help unveil different sentiment polarities in the text;
- We propose a novel framework SignedSenti to leverage explicit sentiment signals in textual terms and implicit sentiment signals in positive (negative) user interactions for unsupervised sentiment analysis;
- We evaluate the effectiveness of the proposed SignedSenti framework on real-world signed social networks.

Problem Statement

We use bold uppercase characters for matrices (e.g., $\mathbf{A}$), bold lowercase characters for vectors (e.g., $\mathbf{a}$), normal lowercase characters for scalars (e.g., a). Also, We represent i -th row of matrix $\mathbf{A}$ as $\mathbf{A}_{i*}$, j -th column as $\mathbf{A}_{*j}$, (i, j) -th entry as $\mathbf{A}_{ij}$, transpose as $\mathbf{A}'$, trace as $\text{tr}(\mathbf{A})$ if $\mathbf{A}$ is a square matrix. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, its Frobenius norm is defined as $\|\mathbf{A}\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^d \mathbf{A}_{ij}^2}$. $\mathbf{I}_n$ denotes the identity matrix of size n -by- n .

Let $\mathcal{T} = \{t_1, t_2, \dots, t_m\}$ be a set of m text posts and $\mathcal{F} = \{f_1, f_2, \dots, f_d\}$ be a set of d textual terms. As shown in Figure 1, the matrix representation of $\mathcal{T}$ is $\mathbf{X} \in \mathbb{R}^{m \times d}$. Each text post may be a review or a comment for a product or an article, respectively. Assume these m text posts are describing a set of l items $\mathcal{O} = \{o_1, \dots, o_l\}$ (e.g., $\{o_1, \dots, o_4\}$ in Figure 1). Their relations are encoded in a text-item relation matrix $\mathbf{O} \in \{0, 1\}^{m \times l}$ where $\mathbf{O}_{ij} = 1$ if text post t_i is about




Figure 1: An illustration of unsupervised sentiment analysis with signed social networks.

item o_j , otherwise $\mathbf{O}_{ij} = 0$. Also, we assume that these m text posts are generated by n distinct social media users $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$. Matrix $\mathbf{T} \in \{0, 1\}^{n \times m}$ shows the authorship between users and text posts such that $\mathbf{T}_{i,j} = 1$ if text post t_j is posted by user u_i , $\mathbf{T}_{i,j} = 0$ otherwise. In addition to positive user interactions, social media users can also be negatively connected, we use $\mathbf{A} \in \mathbb{R}^{n \times n}$ to denote the signed adjacency matrix where $\mathbf{A}_{ij} = 1$, $\mathbf{A}_{ij} = -1$ and $\mathbf{A}_{ij} = 0$ represent positive, negative and missing links from user u_i to u_j , respectively. The relations among posts $\mathcal{T}$, items $\mathcal{O}$ and users $\mathcal{U}$ are shown in the middle of Figure 1; while an illustration of matrices $\mathbf{O}$, $\mathbf{T}$ and $\mathbf{A}$ are demonstrated at the bottom of Figure 1.

With above notations, we now define the concepts of positive linked set, negative linked set and make the signed link based partial order assumption. They act as preliminaries in understanding the proposed framework SignedSenti.

Definition 1 Positive Linked Set:

For a specific text post t_i on the item o_r posted by user u_a , its positive linked set $\mathcal{P}(t_i)$ is defined as the whole set of text posts t_j on the same item o_r that are posted by user u_b , where u_b is positively connected from u_a , i.e., $\mathcal{P}(t_i) = \{t_j | \forall (j, r, a, b) \text{ s.t. } \mathbf{O}_{ir} = 1, \mathbf{O}_{jr} = 1, \mathbf{T}_{ai} = 1, \mathbf{T}_{bj} = 1, \mathbf{A}_{ab} = 1\}$.

Definition 2 Negative Linked Set:

For a specific text post t_i on the item o_r posted by user u_a , its negative linked set $\mathcal{N}(t_i)$ is defined as the whole set of text posts t_k on the same item o_r that are posted by user u_b , where u_b is negatively connected from u_a , i.e., $\mathcal{N}(t_i) = \{t_k | \forall (k, r, a, b) \text{ s.t. } \mathbf{O}_{ir} = 1, \mathbf{O}_{kr} = 1, \mathbf{T}_{ai} = 1, \mathbf{T}_{bk} = 1, \mathbf{A}_{ab} = -1\}$.

Recent advances in signed network analysis (Tang et al. 2015a) show that users are likely to be more similar to their

Table 1: Statistics of datasets

Statistics	Epinions	Slashdot
# of posts	1,559,803	133,335
# of items	200,952	72,241
# of users	326,978	7,897
# of positive links	717,667	52,639
# of negative links	123,705	17,535

friends then their foes. Hence, it motivates us to investigate if friends are more likely to exhibit similar sentiments than foes on the same item, leading to the following signed link based partial order assumption:

Assumption 1 Signed Link Based Partial Order:

For text post t_j in the positive linked set of t_i and text post t_k in the negative linked set of t_i , sentiment polarity of t_i is usually more similar to the sentiment polarity of t_j than t_k . We denote such property as signed link partial order which can be formulated as follows:

$$\text{sim}(t_i, t_j) > \text{sim}(t_i, t_k), t_j \in \mathcal{P}(t_i), t_k \in \mathcal{N}(t_i) \quad (1)$$

Then the problem of unsupervised sentiment analysis with signed social networks can be stated as follows:

Given: a set of social media posts $\mathcal{T}$, a set of items $\mathcal{O}$, a set of social media users $\mathcal{U}$, and available relations including the user-text relation $\mathbf{T}$, user-user relation $\mathbf{A}$ (either positive or negative) and text-item relation $\mathbf{O}$;

Infer: the sentiment polarities of all posts in $\mathcal{T}$.

Data Analysis

Datasets

We used two real-world datasets from Epinions and Slashdot which include both positive and negative links to perform unsupervised sentiment analysis. Detailed statistics of these two datasets are shown in Table 1.

Epinions: Epinions is a product review website where users share their reviews about products. Users can build either trust or distrust relations to other users. We crawled a set of reviews, products and users as well as their interactions. The unigram model is employed on product reviews to construct the feature space, and term frequency is used as feature weight. For the evaluation purpose, we take the rating scores of reviews as ground truth of sentiment labels. In particular, the ratings of 4, 5 and 6 are considered as positive labels while the ratings of 1, 2 and 3 are taken as negative labels.

Slashdot: Slashdot is a technology news website for users to share and comment new articles on science and technology. Users can tag others as friends or foes. Likewise, we crawled and collect comments, articles, users and their relations. The feature space is also built with unigram model and the ratings of comments are employed to establish ground truth in the same way as Epinions.

Signed Link Based Partial Order Assumption

We would like to validate whether the signed link based partial order assumption holds for text posts in real-world signed networks.

First, we define the sentiment similarity between two text posts t_i and t_j as $\text{sim}(t_i, t_j) = \|\mathbf{y}_i - \mathbf{y}_j\|_2$, where $\mathbf{y}_i \in \mathbb{R}^{1 \times k}$ and $\mathbf{y}_j \in \mathbb{R}^{1 \times k}$ are the ground truth of sentiment labels for text posts x_i and x_j , respectively. k denotes the number of sentiment labels. With the definition of text post sentiment similarity, to verify if the signed link based partial order assumption holds, we construct two vectors $\mathbf{s}_p$ and $\mathbf{s}_n$ of the same length. Elements in $\mathbf{s}_p$ denote the sentiment similarity of two text posts t_i and t_j , where t_j is from the positive linked set of t_i . Elements in $\mathbf{s}_n$ indicate the sentiment similarity between two text posts t_i and t_k where t_k is from the negative linked set of t_i . To validate the assumption, we first sample 500 pairs in each group to construct $\mathbf{s}_p$ and $\mathbf{s}_n$, and then conduct two sample t-test on these two vectors. The null hypothesis is $H_0 : \mathbf{c}_p \geq \mathbf{c}_n$ while the alternative hypothesis is $H_1 : \mathbf{c}_p < \mathbf{c}_n$. In the formulations, $\mathbf{c}_p$ and $\mathbf{c}_n$ represent the sample means in these two groups $\mathbf{s}_p$ and $\mathbf{s}_n$, respectively. The null hypothesis is rejected at the significant level $\alpha = 0.01$ with p -values of $4.3e^{-7}$ and $7.2e^{-4}$ in Epinions and Slashdot, respectively. It indicates that the signed link based partial order assumption indeed holds in real-world signed social networks. In other words, it suggests the existence of implicit sentiment signals among positive and negative user interactions, which paves way for unsupervised sentiment analysis.

Proposed Framework-SignedSenti

In this section, we discuss how to model both positive and negative user interactions in understanding and predicting sentiment polarities in an unsupervised scenario.

Basic Model for Unsupervised Sentiment Analysis

Unsupervised sentiment analysis is naturally a clustering problem. Specifically, we would like to cluster text posts into k different sentiment groups. Let $\mathbf{U} \in \mathbb{R}^{m \times k}$ be the text-sentiment cluster matrix such that $\mathbf{U}_{ij} = 1$ if text post t_i belongs to class c_j , and $\mathbf{U}_{ij} = 0$ otherwise. In essence, it can be modeled by solving the following nonnegative matrix factorization problem:

$$\begin{aligned} \min_{\mathbf{U}, \mathbf{V}} & \|\mathbf{X} - \mathbf{UV}'\|_F^2 + \gamma(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2) \\ \text{s.t. } & \mathbf{U} \geq 0, \mathbf{V} \geq 0, \mathbf{U} \in \{0, 1\}^{m \times k}, \mathbf{U}'\mathbf{1} = \mathbf{1}, \end{aligned} \quad (2)$$

where $\mathbf{V} \in \mathbb{R}^{d \times k}$ is a term-sentiment matrix, and each row of $\mathbf{V}$ shows the distribution of each term in these k sentiment groups. $\gamma(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2)$ is introduced to avoid overfitting.

Sentiment Signals from Textual Terms

It has been widely studied in literature (Wang, Lu, and Zhai 2010) that the overall sentiment of a text post is strongly correlated with sentiment of terms in the post. In other words, some terms may contain strong sentiment signals in identifying sentiment polarities. For example, the words of “wonderful” and “appealing” in a text post may express positive emotions while the words of “terrible” and “disappointed” could express negative emotions. The rich sentiment signals in terms help to bridge the gap between the difficulties in

obtaining sentiment labels and the necessity of label supervision in sentiment analysis. To leverage sentiment signals in rich textual information, we employ a widely used sentiment lexicon SentiWordNet (Esuli and Sebastiani 2006) to obtain sentiment polarities of terms. SentiWordNet contains positive, negative and objective scores between 0 and 1 for all synsets in WordNet. In WordNet, there are a total of 117,659 words and phrases. Let $\mathbf{P} \in \mathbb{R}^{d \times k}$ be a term-sentiment indication matrix which encodes sentiment signals of words. Since our task is polarity sentiment analysis, we set $k = 2$ and let $\mathbf{P}_{i1}$ denote the positive score of term f_i while $\mathbf{P}_{i2}$ represents the negative score of term f_i . To take advantage of the textual sentiment signal, we force the above term-sentiment matrix $\mathbf{V}$ in the base model to be consistent with the term-sentiment indication matrix $\mathbf{P}$ by minimizing:

$$\min_{\mathbf{V}} \|\mathbf{V} - \mathbf{P}\|_F^2. \quad (3)$$

It should be noted that the number of sentiment signals, i.e., k should be adapted according to the needs whether to perform binary or multi-class sentiment polarity analysis.

Exploiting Positive and Negative Interactions

The signed link based partial order assumption suggests that for each text post, its sentiment is more similar to posts in its positive linked set than posts in its negative linked set. In other words, it indicates that friends are more likely to reveal similar sentiments than foes on the same item. As $\mathbf{U} \in \mathbb{R}^{m \times k}$ denotes the sentiment polarity hard assignment matrix, we use $\|\mathbf{U}_{i*} - \mathbf{U}_{j*}\|_2^2$ to represent the sentiment similarity between two text posts t_i and t_j . To model the signed link based partial order assumption, there are two cases that we need to discuss. For each text post t_i , (1) if another text post t_j in its positive linked set is more closer to the text post t_k in its negative linked set, i.e., $\|\mathbf{U}_{i*} - \mathbf{U}_{j*}\|_2^2 - \|\mathbf{U}_{i*} - \mathbf{U}_{k*}\|_2^2 < 0$, we do not need to penalize it; (2) if its negative linked set is more closer to its positive linked set, i.e., $\|\mathbf{U}_{i*} - \mathbf{U}_{j*}\|_2^2 - \|\mathbf{U}_{i*} - \mathbf{U}_{k*}\|_2^2 > 0$, we should add a penalty to pull the sentiment of t_i be more closer to t_j than to t_k . Mathematically, it can be formulated by solving the following objective function:

$$\min \sum_{(i,j,k) \in \Omega} \max(0, \|\mathbf{U}_{i*} - \mathbf{U}_{j*}\|_2^2 - \|\mathbf{U}_{i*} - \mathbf{U}_{k*}\|_2^2), \quad (4)$$

where Ω denotes all triplets that satisfies the signed link based partial order assumption, i.e., $\Omega = \{(i, j, k) | i \in \mathcal{T}, j \in \mathcal{P}(t_i), k \in \mathcal{N}(t_i)\}$. The above penalty term can be further reformulated as:

$$\begin{aligned} & \sum_{(i,j,k) \in \Omega} \max(0, \|\mathbf{U}_{i*} - \mathbf{U}_{j*}\|_2^2 - \|\mathbf{U}_{i*} - \mathbf{U}_{k*}\|_2^2) \\ &= \sum_{(i,j,k) \in \Omega} w_{ij}^k \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}'), \end{aligned} \quad (5)$$

where $\mathbf{M}$ is a sparse matrix with all entries equal to zero except that $\mathbf{M}_{ij} = \mathbf{M}_{ji} = \mathbf{M}_{kk} = -1$ and $\mathbf{M}_{ik} = \mathbf{M}_{ki} = \mathbf{M}_{jj} = 1$. $\mathbf{M}_{ij}^k$ is the matrix $\mathbf{M}$ with elements associated with triplet (i, j, k) and w_{ij}^k is defined as follows:

$$w_{ij}^k = \begin{cases} 1 & \text{if } \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}') > 0 \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

Objective Function of SignedSenti

With the model components of sentiment signals from terms and the signed link based partial order assumption, the final objective function of unsupervised sentiment analysis with signed social network can be formulated as follows:

$$\begin{aligned} & \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{X} - \mathbf{U} \mathbf{V}'\|_F^2 + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}') \\ & + \beta \|\mathbf{V} - \mathbf{P}\|_F^2 + \gamma (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2) \\ & \text{s.t. } \mathbf{U} \geq 0, \mathbf{V} \geq 0, \mathbf{U} \in \{0, 1\}^{m \times k}, \mathbf{U}' \mathbf{1} = \mathbf{1}. \end{aligned} \quad (7)$$

Parameters α and β control the contribution of sentiment signals from terms and signed social networks, respectively.

The problem in Eq. (7) is difficult to solve due to the discrete constraint on $\mathbf{U}$. To tackle this issue, we relax the objective function by reformulating it as an orthogonal constraint. After the relaxation, Eq.(7) can be rewritten as:

$$\begin{aligned} & \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{X} - \mathbf{U} \mathbf{V}'\|_F^2 + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}') \\ & + \beta \|\mathbf{V} - \mathbf{P}\|_F^2 + \gamma (\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2) \\ & \text{s.t. } \mathbf{U} \geq 0, \mathbf{V} \geq 0, \mathbf{U}' \mathbf{U} = \mathbf{I}. \end{aligned} \quad (8)$$

Optimization Algorithm for SignedSenti

The objective function of the proposed SignedSenti framework is not convex w.r.t. both $\mathbf{U}$ and $\mathbf{V}$ simultaneously. Hence, we introduce an alternating algorithm to solving its optimization problem.

Update $\mathbf{U}$: First, we fix $\mathbf{V}$ to update $\mathbf{U}$. Specifically, when $\mathbf{V}$ is fixed, the objective function is convex w.r.t. the text-sentiment matrix $\mathbf{U}$. Thus, $\mathbf{U}$ can be obtained by solving the following optimization problem:

$$\begin{aligned} & \min_{\mathbf{U}} \mathcal{J}(\mathbf{U}) = \|\mathbf{X} - \mathbf{U} \mathbf{V}'\|_F^2 + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}') + \gamma \|\mathbf{U}\|_F^2 \\ & \text{s.t. } \mathbf{U} \geq 0, \mathbf{U}' \mathbf{U} = \mathbf{I}. \end{aligned} \quad (9)$$

The Lagrangian of Eq. (9) is:

$$\begin{aligned} & \min_{\mathbf{U}} \mathcal{L}(\mathbf{U}) = \|\mathbf{X} - \mathbf{U} \mathbf{V}'\|_F^2 + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k \text{tr}(\mathbf{M}_{ij}^k \mathbf{U} \mathbf{U}') \\ & + \gamma \|\mathbf{U}\|_F^2 + \text{tr}(\Gamma_u (\mathbf{U}' \mathbf{U} - \mathbf{I})) - \text{tr}(\Lambda_u \mathbf{U}'). \end{aligned} \quad (10)$$

where Γ_u and Λ_u are the Lagrange multipliers for constraints $\mathbf{U}' \mathbf{U} = \mathbf{I}$ and $\mathbf{U} \geq 0$, respectively. To compute $\mathbf{U}$, we take the partial derivative of Eq. (10) w.r.t. $\mathbf{U}$ and set it to be zero:

$$\Lambda_u = 2(\mathbf{U} \mathbf{V}' \mathbf{V} - \mathbf{X} \mathbf{V} + \gamma \mathbf{U} + \mathbf{U} \Gamma_u) + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k (\mathbf{M}_{ij}^k \mathbf{U} + \mathbf{M}_{ij}^{k'} \mathbf{U}). \quad (11)$$

With the KKT complementary condition for the nonnegativity constraint of $\mathbf{U}$, i.e., $(\Lambda_u)_{ij} \mathbf{U}_{ij} = 0$, we have:

$$\begin{aligned} & (\mathbf{U} \mathbf{V}' \mathbf{V} - \mathbf{X} \mathbf{V} + \gamma \mathbf{U} + \frac{\alpha}{2} \sum_{(i,j,k) \in \Omega} w_{ij}^k (\mathbf{M}_{ij}^k \mathbf{U} + \mathbf{M}_{ij}^{k'} \mathbf{U}) \\ & + \mathbf{U} \Gamma_u)_{ij} \mathbf{U}_{ij} = 0, \text{ where} \end{aligned} \quad (12)$$

$$\Gamma_u = -\frac{\alpha}{2} \sum_{(i,j,k) \in \Omega} w_{ij}^k (\mathbf{U}'(\mathbf{M}_{ij}^k \mathbf{U} + \mathbf{M}_{ij}^{k'} \mathbf{U})) - \mathbf{V}' \mathbf{V} + \mathbf{U}' \mathbf{X} \mathbf{V} - \gamma \mathbf{I}. \quad (13)$$

It leads to the following update rule for $\mathbf{U}$:

$$\mathbf{U}_{ij} \leftarrow \mathbf{U}_{ij} \sqrt{\frac{\mathbf{B}_{ij}}{\mathbf{E}_{ij}}}, \text{ where} \quad (14)$$

$$\mathbf{B} = 2\mathbf{X}\mathbf{V} + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k (\mathbf{M}_{ij}^k \mathbf{U} + \mathbf{M}_{ij}^{k'} \mathbf{U})^- + 2\mathbf{U}\Gamma_u^-, \quad (15)$$

$$\mathbf{E} = 2(\mathbf{U}\mathbf{V}'\mathbf{V} + \gamma \mathbf{U}) + \alpha \sum_{(i,j,k) \in \Omega} w_{ij}^k (\mathbf{M}_{ij}^k \mathbf{U} + \mathbf{M}_{ij}^{k'} \mathbf{U})^+ + 2\mathbf{U}\Gamma_u^+. \quad (16)$$

Update $\mathbf{V}$: Likewise, we fix $\mathbf{U}$ to update $\mathbf{V}$. When $\mathbf{U}$ is fixed, the objective function is convex w.r.t. the term-sentiment matrix $\mathbf{V}$. Hence, $\mathbf{V}$ can be obtained by solving:

$$\min_{\mathbf{V}} \mathcal{J}(\mathbf{V}) = \|\mathbf{X} - \mathbf{U}\mathbf{V}'\|_F^2 + \beta \|\mathbf{V} - \mathbf{P}\|_F^2 + \gamma \|\mathbf{V}\|_F^2 \quad (17)$$

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

The Lagrangian of Eq. (17) is:

$$\mathcal{L}(\mathbf{V}) = \|\mathbf{X} - \mathbf{U}\mathbf{V}'\|_F^2 + \beta \|\mathbf{V} - \mathbf{P}\|_F^2 + \gamma \|\mathbf{V}\|_F^2 - \text{tr}(\Lambda_v \mathbf{V}'), \quad (18)$$

where Λ_v is the Lagrange multipliers for the constraints $\mathbf{V} \geq 0$. We take the partial derivative of Eq. (18) w.r.t. $\mathbf{V}$ and set it to be zero:

$$\Lambda_v = 2(\mathbf{V}\mathbf{U}'\mathbf{U} - \mathbf{X}'\mathbf{U} + \beta(\mathbf{V} - \mathbf{P}) + \gamma \mathbf{V}). \quad (19)$$

Similarly, with the KKT complementary condition for the nonnegativity constraint of $\mathbf{V}$, i.e., $(\Lambda_v)_{ij} \mathbf{V}_{ij} = 0$, we have:

$$2(\mathbf{V}\mathbf{U}'\mathbf{U} - \mathbf{X}'\mathbf{U} + \beta(\mathbf{V} - \mathbf{P}) + \gamma \mathbf{V})_{ij} \mathbf{V}_{ij} = 0, \quad (20)$$

which leads to the following update rule for $\mathbf{V}$:

$$\mathbf{V}_{ij} \leftarrow \mathbf{V}_{ij} \sqrt{\frac{\mathbf{X}'\mathbf{U} + \beta \mathbf{P}}{\mathbf{V}\mathbf{U}'\mathbf{U} + (\beta + \gamma) \mathbf{V}}}. \quad (21)$$

With these update rules, the detailed algorithm of the proposed SignedSenti framework is illustrated in Algorithm 1. At the very beginning, we initialize $\mathbf{U}$, $\mathbf{V}$ randomly and calculate $\mathbf{M}$ from $\mathbf{T}$, $\mathbf{A}$ and $\mathbf{O}$. From line 3 to 7, we update $\mathbf{U}$ and $\mathbf{V}$ iteratively until converge. To update $\mathbf{U}$, we need to calculate w_{ij}^k and Γ_u at first. According to Eq.(6) and Eq.(13), the computation cost of obtaining w_{ij}^k and Γ_u are $\mathcal{O}(m^2k)$ and $\mathcal{O}(k^2d + kmd + m^2k + k^2m)$ respectively. With w_{ij}^k and Γ_u , we employ Eq.(14) to update $\mathbf{U}$, the computational cost of updating $\mathbf{U}$ is $\mathcal{O}(kmd + m^2k + k^2m)$. The total cost of computing $\mathbf{V}$ according to Eq.(21) is $\mathcal{O}(kmd)$. After we obtain $\mathbf{U}$, sentiment polarities of text texts can be obtained by performing K-Means on $\mathbf{U}$.

Experiments

Experimental Setting

Following a common way to assess performance of unsupervised sentiment analysis, we take clustering accuracy as the evaluation metric. Higher clustering accuracy often indicates better performance. SignedSenti is compared with the following baseline methods:

Algorithm 1: SignedSenti Algorithm

Input : $\{\mathbf{X}, \mathbf{T}, \mathbf{A}, \mathbf{O}, \mathbf{P}, k, \alpha, \beta, \gamma\}$

Output: sentiment polarity for each text post.

- 1 Initialize $\mathbf{U}$, $\mathbf{V}$ randomly;
 - 2 Compute $\mathbf{M}$ based on $\mathbf{T}$, $\mathbf{A}$ and $\mathbf{O}$;
 - 3 **while not converge do**
 - 4 Calculate w_{ij}^k according to Eq.(6) ;
 - 5 Compute Γ_u according to Eq.(13) ;
 - 6 Update $\mathbf{U}$ according to Eq.(14);
 - 7 Update $\mathbf{V}$ according to Eq.(21);
 - 8 **end**
 - 9 Employing $\mathbf{U}$ to predict sentiment polarity of text posts.
-

- **SentiStrength** (Thelwall, Buckley, and Paltoglou): SentiStrength is a lexicon-based unsupervised method that extracts sentiment strength from informal English with pre-defined sentiment lexicon.
- **MPQA** (Wiebe, Wilson, and Cardie 2005): It predicts sentiment polarity of text posts according to a manually labeled sentiment lexicon MPQA.
- **SentiWordNet** (Esuli and Sebastiani 2006): It determines sentiment scores of text posts via a widely used sentiment lexicon SentiWordNet.
- **K-Means**: As the one of the most representative clustering methods, it partitions the text posts into k sentiment polarities on the original textual terms.
- **NMF** (Pauca et al. 2004): Nonnegative matrix factorization is a popular method in text mining. It is also a variant of the proposed SignedSenti model by setting $\alpha = \beta = 0$.
- **SignedSenti-T**: It is a variant of the proposed SignedSenti that only employs the textual information for sentiment analysis. Specifically, we set $\alpha = 0$.
- **SignedSenti-L**: It is a variant of the proposed SignedSenti that does not explicitly leverage sentiment signals from textual terms. In particular, we set $\beta = 0$.

Sentiment Polarity Prediction Performance

In this subsection, we compare SignedSenti with other baseline algorithms. Noticed that in SigendSenti, we have three regularization parameters α, β, γ . We empirically set these parameters as $\{\alpha = 1, \beta = 0.5, \gamma = 0.7\}$ in Epinions and $\{\alpha = 1, \beta = 1, \gamma = 0.1\}$ in Slashdot. More discussions about the effectiveness of these parameters will be presented later. The comparison results of various unsupervised sentiment analysis algorithms on Epinions and Slashdot datasets are shown in Table 2. We make the following observations:

- SignedSenti consistently outperforms other baseline methods on both datasets with significant performance gain. We also perform pairwise Wilcoxon signed-rank test (Demšar 2006) between SignedSenti and these baseline methods, it shows SignedSenti is significantly better with a significance level of 0.05. The superiority of the proposed SignedSenti can be attributed to the utilization of external sources, including textual sentiment signals and positive (negative) user interactions.
- In general, traditional lexicon-based unsupervised methods such as SentiStrength, MPQA and SentiWordNet do

Table 2: Sentiment polarity prediction accuracy.

Method	Epinions	Slashdot
SentiStrength	0.521	0.628
MPQA	0.662	0.684
SentiWordNet	0.645	0.586
K-Means	0.644	0.677
NMF	0.637	0.648
SignedSenti-T	0.649	0.672
SignedSenti-L	0.714	0.700
SignedSenti	0.723	0.731




Figure 2: Parameter analysis of SignedSenti on Slashdot.

not perform well in the unsupervised case. This observations show the necessity to build a sophisticated learning model to automatically predict the sentiment polarities of text posts.

- SignedSenti also obtains better performance than traditional document clustering methods K-Means and NMF. The reason is that social media texts are often noisy and incomplete, hence without the guide of any sentiment signals or user interactions, it is difficult to discriminate the sentiment polarities of different text posts.
- The clustering accuracy of SignedSenti is higher than its variant SignedSenti-T. SignedSenti-T only leverages sentiment signals from terms and does not explicitly consider user interactions. Its inferiority to SignedSenti indicates that in addition to textual sentiment signals, positive and negative links also contain implicit rich sentiment signals that can boost the sentiment polarity prediction.

Parameter Analysis

The proposed SignedSenti has two important parameters α and β which controls the contribution of implicit sentiment signals from positive (negative) user interactions and textual terms respectively. We study the effect of each parameter by fixing the other to investigate how it affects the clustering performance. We only report the experimental result on Slashdot as we have similar observations on Epinions. In particular, we first fix $\{\beta = 1, \gamma = 0.1\}$ and vary α as $\{0, 0.01, 0.1, 0.3, 0.5, 0.7, 1, 10\}$. As shown in Figure 2(a), when α increase from 0 to 0.01 the performance increases dramatically which further validates the effectiveness of leveraging implicit sentiment signals in positive and negative interactions. If we continuously increase α , the performance is relatively stable in fairly large ranges $[0.01, 1]$, then it decreases when $\alpha > 1$. Similarly, to in-

vestigate how β affects the performance, we vary β as $\{0, 0.01, 0.1, 0.3, 0.5, 0.7, 1, 10\}$ by fixing $\{\alpha = 1, \gamma = 0.1\}$. The result is presented in Figure 2(b). Likewise, the performance increases significantly at the very beginning due to the increase of β from 0 to 0.01. After that, with the increase of β , the performance fluctuates in ranges of 71.5 and 73.5. To summary, the clustering performance is rather stable when we tune these two parameters in a wide range, which is very appealing in practice.

Related Work

In this section, we briefly review sentiment analysis in social media. Sentiment analysis in social media has been a surge of research recently. However, it faces some challenges mainly because of the bewildering combination of heterogeneous data sources and structures. Also, since labels of social media data are costly to obtain, unsupervised sentiment analysis is more desired. Recent years have witnessed some efforts in exploring external information for unsupervised sentiment analysis. As the most representative unsupervised sentiment analysis algorithms, lexicon-based methods (Taboada et al. 2011; O’Connor et al. 2010; Wilson, Wiebe, and Hoffmann 2005) determine sentiment polarity of texts by exploiting sentiment signals revealed by words or phrases. In addition to rich source of text information, abundant emotional signals are widely observed in social media. In (Hu et al. 2013a), the authors proposed a framework to incorporate two categories of emotional signals for unsupervised sentiment analysis. (Wang et al. 2015b) made one of the first attempt to leverage social media images for unsupervised sentiment analysis. Different from above mentioned approaches, we present the first study on unsupervised sentiment analysis with both positive and negative social interactions.

Conclusion

Due to vast opinion-rich resources brought by social media services, sentiment analysis for social media data has received increasing attention in recent years. As it is costly to obtain sentiment labels for social media data, unsupervised methods are more appealing in practice. Traditional unsupervised sentiment analysis method are either lexicon-based or employ sentiment signals from textual terms to determine sentiment polarity. However, social media data is not independent but are correlated by user interactions. And in many cases, users may also be negatively connected such as distrust relations and foes. The availability of both positive and negative links could be another rich source in deriving implicit sentiment signals for sentiment analysis. In this paper, we study a novel problem of unsupervised sentiment analysis with signed social networks. Methodologically, we propose to incorporate the signed social relations and sentimental signals from terms into a unified framework when we are lack of sentiment labels. We also conduct experiments on two real world signed social networks Epinions and Slashdot. The results show that the proposed SignedSenti has significantly better performance than state-of-the-art methods.

Acknowledgments

This material is, in part, supported by National Science Foundation (NSF) under grant number 1614576.

References

- Abelson, R. P. 1983. Whatever became of consistency theory? *Personality and Social Psychology Bulletin*.
- Bollen, J.; Mao, H.; and Pepe, A. 2011. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the 5th International Conference on Web and Social Media*, 450–453.
- Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research* 7(Jan):1–30.
- Ding, X., and Liu, B. 2007. The utility of linguistic rules in opinion mining. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 811–812. ACM.
- Esuli, A., and Sebastiani, F. 2006. Sentiwordnet: A publicly available lexical resource for opinion mining. In *Proceedings of LREC*, volume 6, 417–422. Citeseer.
- Hatfield, E.; Cacioppo, J. T.; and Rapson, R. L. 1994. *Emotional contagion*. Cambridge university press.
- Hu, X.; Tang, J.; Gao, H.; and Liu, H. 2013a. Unsupervised sentiment analysis with emotional signals. In *Proceedings of the 22nd international conference on World Wide Web*, 607–618. ACM.
- Hu, X.; Tang, L.; Tang, J.; and Liu, H. 2013b. Exploiting social relations for sentiment analysis in microblogging. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 537–546. ACM.
- Kamvar, S. D., and Harris, J. 2011. We feel fine and searching the emotional web. In *Proceedings of the 4th ACM international conference on Web search and data mining*, 117–126. ACM.
- Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010a. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*, 641–650. ACM.
- Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010b. Signed networks in social media. In *Proceedings of the the 28th SIGCHI conference on human factors in computing systems*, 1361–1370. ACM.
- Li, J.; Cheng, K.; Wang, S.; Morstatter, F.; Robert, T.; Tang, J.; and Liu, H. 2016. Feature selection: A data perspective. *arXiv:1601.07996*.
- Liu, B. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies* 5(1):1–167.
- O’Connor, B.; Balasubramanyan, R.; Routledge, B. R.; and Smith, N. A. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the 4th International Conference on Web and Social Media*, 122–129.
- Pang, B., and Lee, L. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd annual meeting on Association for Computational Linguistics*, 271. ACL.
- Pang, B., and Lee, L. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval* 2(1-2):1–135.
- Pang, B.; Lee, L.; and Vaithyanathan, S. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing*, 79–86. ACL.
- Pauca, V. P.; Shahnaz, F.; Berry, M. W.; and Plemmons, R. J. 2004. Text mining using non-negative matrix factorizations. In *Proceedings of the 4th SIAM International Conference on Data Mining*, 452–456. SIAM.
- Taboada, M.; Brooke, J.; Tofiloski, M.; Voll, K.; and Stede, M. 2011. Lexicon-based methods for sentiment analysis. *Computational linguistics* 37(2):267–307.
- Tang, J.; Aggarwal, C.; and Liu, H. 2016. Node classification in signed social networks. In *Proceedings of the 2016 SIAM International Conference on Data Mining*. SIAM.
- Tang, J.; Chang, Y.; Aggarwal, C.; and Liu, H. 2015a. A survey of signed network mining in social media. *arXiv preprint arXiv:1511.07569*.
- Tang, J.; Nobata, C.; Dong, A.; Chang, Y.; and Liu, H. 2015b. Propagation-based sentiment analysis for microblogging data. In *Proceedings of 2015 SIAM International Conference on Data Mining*. SIAM.
- Thelwall, M.; Buckley, K.; and Paltoglou, G. Sentistrength, 2011. available online.
- Wang, Y.; Hu, Y.; Kambhampati, S.; and Li, B. 2015a. Inferring sentiment from web images with joint inference on visual and social cues: A regulated matrix factorization approach. In *Proceedings of 9th International AAAI Conference on Web and Social Media*.
- Wang, Y.; Wang, S.; Tang, J.; Liu, H.; and Li, B. 2015b. Unsupervised sentiment analysis for social media images. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, 2378–2379.
- Wang, H.; Lu, Y.; and Zhai, C. 2010. Latent aspect rating analysis on review text data: a rating regression approach. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 783–792.
- Wiebe, J.; Wilson, T.; and Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation* 39(2-3):165–210.
- Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, 347–354. ACL.
- Zhang, Y.; Lai, G.; Zhang, M.; Zhang, Y.; Liu, Y.; and Ma, S. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, 83–92. ACM.