Node Classification in Signed Social Networks

Jiliang Tang* Charu Aggarwal† Huan Liu‡

Abstract

Node classification in social networks has been proven to be useful in many real-world applications. The vast majority of existing algorithms focus on unsigned social networks (or social networks with only positive links), while little work exists for signed social networks. It is evident from recent developments in signed social network analysis that negative links have added value over positive links. Therefore, the incorporation of negative links has the potential to benefit various analytical tasks. In this paper, we study the novel problem of node classification in signed social networks. We provide a principled way to mathematically model positive and negative links simultaneously and propose a novel framework NCSSN for node classification in signed social networks. Experimental results on real-world signed social network datasets demonstrate the effectiveness of the proposed framework NCSSN. Further experiments are conducted to gain a deeper understanding of the importance of negative links for NCSSN.

1 Introduction

User information such as demographic values, interest, beliefs or other characteristics is crucial for social media service providers to customize their services to the users in many applications such as recommendations and personalized search. The more they know about users, the better they can serve them. However, social media sites can only collect very limited user information because most social media users do not share their information [33]. For example, more than 90% of users in Facebook[1] do not reveal their political views [1]. One way of bridging this knowledge gap is to infer missing user information by leveraging the pervasively available network structures in social media [29, 1]. An example of such inference is that of node classification in social networks. Since nodes represent users in social networks, we will use the terms "node" and "user" interchangeably.

The problem of node classification has been extensively studied [10] in the literature. Existing node classification algorithms can be mainly grouped into local classifier based methods and random walk based methods [5]. The vast majority of these algorithms have overwhelmingly focused on unsigned social networks (or social networks with only positive links) [21, 30, 17, 36, 34, 22]. However, social networks in social media can contain both positive and negative links. Examples of these signed social networks include Epinions[2] with trust and distrust links, and Slashdot[3] with friend and foe links. The work on node classification in signed social networks is rather limited.

The recent availability of signed social networks has encouraged increasing attention on signed social network analysis. Recent work shows that negative links have added value over positive links [14] and can benefit various analytical tasks. For example, negative links can significantly improve social recommendation performance [31, 18], and a small number of negative links can improve the performance of positive link prediction remarkably [11, 16]. Evidence from recent achievements in signed social networks suggests that negative links in signed social networks may be also potentially helpful in the task of node classification.

In this paper, we study the novel problem of node classification in signed social networks. This problem has not been previously studied. In essence, we investigate how to mathematically model positive and negative links in signed social networks for the problem of node classification. We design a novel framework, NCSSN, for node classification in signed social networks. Our main contributions are summarized below:

- We provide a principled approach to mathematically model positive and negative links for the problem of node classification in signed social networks;
- We propose a node classification framework NCSSN, which captures positive and negative links in signed social networks into a coherent model; and
- We evaluate the proposed framework in real-world signed social networks to understand the effec-

*Yahoo! Labs, jlt@yahoo-inc.com
†IBM T.J. Watson Research Center, charu@us.ibm.com
‡Arizona State University, huan.liu@asu.edu
[1]`https://www.facebook.com/`

[2]`http://www.epinions.com/`
[3]`http://slashdot.org/`

tiveness and mechanisms of the proposed NCSSN framework.

The remainder of this paper is organized as follows. In Section 2, we formally define the problem of node classification in signed social networks. In Section 3, we provide a way to model positive and negative links and introduce details about the proposed NCSSN framework with an optimization algorithm. Section 4 presents experimental results with discussions. Section 5 briefly reviews related work. Finally, Section 6 concludes with future work.

2 Problem Statement

Let $\mathcal{U} = \{u_1, u_2, \ldots, u_N\}$ be the set of N nodes in the social network. A signed network $\mathcal{G}$ can be decomposed into a positive component $\mathcal{G}_p$ and a negative component $\mathcal{G}_n$. Let $\mathbf{A}^p \in \mathbb{R}^{N \times N}$ be the adjacency matrix of $\mathcal{G}_p$ where $\mathbf{A}^p_{ij} = 1$ if u_i has a positive link to u_j, and $\mathbf{A}^p_{ij} = 0$, otherwise. Similarly, $\mathbf{A}^n \in \mathbb{R}^{N \times N}$ denotes the adjacency matrix of $\mathcal{G}_n$ where $\mathbf{A}^n_{ij} = 1$ if u_i has a negative link to u_j, $\mathbf{A}^n_{ij} = 0$ otherwise. Note that we only consider a binary weight $\{0, 1\}$ for links in this paper although the generalization of the proposed framework to continuous weights is straightforward.

Let $\mathcal{C} = \{c_1, c_2, \ldots, c_m\}$ be the set of m label classes. Assume that $\mathcal{U}^L = \{u_1, u_2, \ldots, u_n\}$ is the set of n labeled nodes where $n < N$ and $\mathcal{U}^U = \mathcal{U} \backslash \mathcal{U}^L$ is the set of $N - n$ unlabeled users. We use $\mathbf{Y} \in \mathbb{R}^{n \times m}$ to denote the label indicator matrix for $\mathcal{U}^L$ where $\mathbf{Y}_{ik} = 1$ if u_i is labeled as c_k, and $\mathbf{Y}_{ik} = 0$ otherwise.

With the aforementioned notations and definitions, the problem of node classification in a signed social network can be formally stated as follows:

Given a signed social network $\mathcal{G}$ with positive links $\mathbf{A}^p$, negative links $\mathbf{A}^n$, and labels $\mathbf{Y}$ for some nodes $\mathcal{U}^L$, the problem of node classification in a signed social network aims to infer labels for the unlabeled nodes $\mathcal{U}^U$.

3 The Proposed Node Classification Framework in Signed Social Networks

There are two understandings about positive and negative links in signed social networks - (1) independent understanding, which considers positive and negative links as two distinct features (or dimensions) about users [24]; and (2) dependent understanding, which treats positive and negative links as tightly related features in a single structure [16]. Correspondingly we can model independent and dependent information of positive and negative links for the problem of node classification in signed social networks. In this section, we first give details about models of independent and dependent information of signed social networks, and then introduce the proposed NCSSN framework with an optimization method.

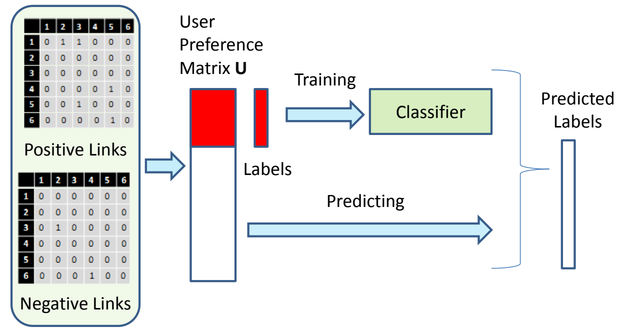

Figure 1: Modeling Independent Information of Positive and Negative Links.

3.1 Modeling Independent Information of Positive and Negative Links

Positive and negative links have been investigated separately (or independent understanding) in [24, 27] and the research results suggest that (1) we cannot extend properties of positive links to negative links and negative links present distinct properties; and (2) the formations of both positive and negative links are related to user preferences. Let $\mathbf{U}_i \in \mathbb{R}_+^{1 \times K}$ be the latent user preference of u_i and $\mathbf{U} = [\mathbf{U}_1; \mathbf{U}_2; \ldots; \mathbf{U}_N] \in \mathbb{R}^{N \times K}$ be the user preference matrix. Both positive and negative links are related to user preferences, which indicates that we can model independent information of positive and negative links by learning the user preference $\mathbf{U}$. The basic idea of capturing independent information for node classification is demonstrated in Figure 1 - we learn user preference matrix $\mathbf{U}$ from $\mathbf{A}^p$ and $\mathbf{A}^n$, and based on $\mathbf{U}$ and labeled nodes $\mathcal{U}^L$, we train the classifier, which is used to predict labels of unlabeled nodes $\mathcal{U}^U$. We formulate the basic idea in Figure 1 via the following optimization problem:

$$\min_{\mathbf{H}^p \geq 0, \mathbf{H}^n \geq 0, \mathbf{U} \geq 0, \mathbf{W}} \sum_{u_i, u_j \in \mathcal{U}} \|\mathbf{A}^p_{ij} - \mathbf{U}_i \mathbf{H}^p \mathbf{U}_j^\top\|_2^2 + \beta \sum_{u_i, u_j \in \mathcal{U}} \|\mathbf{A}^n_{ij} - \mathbf{U}_i \mathbf{H}^n \mathbf{U}_j^\top\|_2^2 + \alpha \sum_{u_i \in \mathcal{U}^L} \|\mathbf{U}_i \mathbf{W} - \mathbf{Y}_i\|_2^2 + \lambda(\|\mathbf{U}\|_F^2 + \|\mathbf{H}^p\|_F^2 + \|\mathbf{H}^n\|_F^2 + \|\mathbf{W}\|_F^2) \quad (3.1)$$

Next we give details about Eq. (3.1) as:

- A positive link from u_i to u_j can be modeled as the correlation between their preferences as $\mathbf{A}^p_{ij} = \mathbf{U}_i \mathbf{H}^p \mathbf{U}_j^\top$ where $\mathbf{H}^p \in \mathbb{R}_+^{K \times K}$ is the correlation matrix for positive links [26].

- Since negative links are also related to user preferences, similarly we model a negative link from u_i to u_j as $\mathbf{A}^n_{ij} = \mathbf{U}_i\mathbf{H}^n\mathbf{U}_j^\top$. Since positive and negative links present distinct properties, we introduce a different correlation matrix $\mathbf{H}^n \in \mathbb{R}^{K\times K}$ for negative links.
- We train a linear classifier $\mathbf{W} \in \mathbb{R}^{K\times m}$ based on user preferences and label information of labeled users $\mathcal{U}^L$ as $\mathbf{Y}_i = \mathbf{U}_i\mathbf{W}$.

The parameter β controls contributions of independent information from negative links. The term $\lambda(\|\mathbf{H}^p\|_F^2 + \|\mathbf{H}^n\|_F^2 + \|\mathbf{U}\|_F^2 + \|\mathbf{W}\|_F^2)$ is introduced to avoid overfitting.

Eq. (3.1) can be rewritten into the matrix form as follows:

$$\begin{aligned}\min_{\mathbf{H}^p\geq 0,\mathbf{H}^n\geq 0,\mathbf{U}\geq 0,\mathbf{W}} &\|\mathbf{A}^p - \mathbf{U}\mathbf{H}^p\mathbf{U}^\top\|_F^2 \\ &+ \beta\|\mathbf{A}^n - \mathbf{U}\mathbf{H}^n\mathbf{U}^\top\|_F^2 + \alpha\|\mathbf{C}(\mathbf{U}\mathbf{W} - \mathbf{Y})\|_F^2 \\ &+ \lambda(\|\mathbf{U}\|_F^2 + \|\mathbf{H}^p\|_F^2 + \|\mathbf{H}^n\|_F^2 + \|\mathbf{W}\|_F^2)\end{aligned} \tag{3.2}$$

where $\mathbf{C} \in \mathbb{R}^{N\times N}$ is a diagonal matrix, where $\mathbf{C}_{ii} = 1$ if $u_i \in \mathcal{U}^L$ and $\mathbf{C}_{ii} = 0$, otherwise.

3.2 Modeling Dependent Information of Positive and Negative Links Users are likely to be more similar to their friends (or users with positive links) than their foes (or users with negative links), i.e., users should sit closer to their friends than their foes [28]. This observation paves a way for us to model dependent information of positive and negative links.

For $\langle i,j,k\rangle$ where u_i has a positive link to u_j and a negative link to u_k, we force u_i closer to her friend u_j than her foe u_k in terms of their user preferences to capture dependent information of positive and negative links. To achieve this goal, we consider the following two cases:

- *Case 1:* If a user u_i sits closer to her friend u_j than her foe u_k, i.e., $\|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2 < 0$, we should not penalize this case since it is consistent with the finding.
- *Case 2:* If a user u_i sits closer to her foe u_k than her friend u_j, i.e., $\|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2 > 0$, this case contradicts with the finding, and we should add a penalty to pull u_i to u_j from u_k.

Based on the aforementioned analysis, we propose the following formulation to capture dependent information from positive and negative links as:

$$\min \sum_{\langle i,j,k\rangle\in\mathcal{S}} max(0, \|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2) \tag{3.3}$$

Next, we give details on the inner workings of Eq. (3.3):

- When a user u_i sits closer to her friend u_j than her foe u_k, i.e., *Case 1*, the minimizing term in Eq. (3.3) is 0 since $\|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2 < 0$. Therefore We do not add any penalty for *Case 1*;
- When a user u_i sits closer to her foe u_k than her friend u_j, i.e., *Case 2*, the minimizing term in Eq. (3.3) is $\|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2$ since $\|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2 > 0$. Eq. (3.3) will pull u_i to u_j from u_k for *Case 2*.

Eq. (3.3) can be rewritten in the matrix form as follows:

$$\begin{aligned}&\sum_{\langle i,j,k\rangle\in\mathcal{S}} max(0, \|\mathbf{U}_i - \mathbf{U}_j\|_2^2 - \|\mathbf{U}_i - \mathbf{U}_k\|_2^2) \\ &= \sum_{\langle i,j,k\rangle\in\mathcal{S}} f_{ijk} Tr(\mathbf{M}^{ijk}\mathbf{U}\mathbf{U}^\top)\end{aligned}$$

where $\mathbf{M}^{ijk}$ has $\mathbf{M}^{ijk}_{ij} = \mathbf{M}^{ijk}_{ji} = \mathbf{M}^{ijk}_{kk} = -1$ and $\mathbf{M}^{ijk}_{ik} = \mathbf{M}^{ijk}_{ki} = \mathbf{M}^{ijk}_{jj} = 1$ with other entries equal to zero. The term f_{ijk} is defined as follows:

$$f_{ijk} = \begin{cases} 1 & \text{if } Tr(\mathbf{M}^{ijk}\mathbf{U}\mathbf{U}^\top) > 0 \\ 0 & \text{otherwise} \end{cases}.$$

3.3 The Optimization Problem With model components for independent and dependent information of positive and negative links, the proposed node classification framework in signed social networks NCSSN is to solve the following optimization problem:

$$\begin{aligned}\min_{\mathbf{H}^p\geq 0,\mathbf{H}^n\geq 0,\mathbf{U}\geq 0,\mathbf{W}} &\|\mathbf{A}^p - \mathbf{U}\mathbf{H}^p\mathbf{U}^\top\|_F^2 + \alpha\|\mathbf{C}(\mathbf{U}\mathbf{W} - \mathbf{Y})\|_F^2 \\ &+ \beta\|\mathbf{A}^n - \mathbf{U}\mathbf{H}^n\mathbf{U}^\top\|_F^2 + \gamma \sum_{\langle i,j,k\rangle\in\mathcal{S}} f_{ijk} Tr(\mathbf{M}^{ijk}\mathbf{U}\mathbf{U}^\top) \\ &+ \lambda(\|\mathbf{H}^p\|_F^2 + \|\mathbf{H}^n\|_F^2 + \|\mathbf{U}\|_F^2 + \|\mathbf{W}\|_F^2)\end{aligned} \tag{3.4}$$

where β and γ control independent and dependent information from negative links, respectively. The relationships among $\mathbf{U}$, $\mathbf{H}^p$, $\mathbf{H}^n$ and $\mathbf{W}$ make the problem of finding optimal solutions for all parameters in Eq. (3.4) difficult to determine simultaneously. In this work, we adopt an alternate optimization scheme [9] for Eq. (3.4) where we optimize one component while fixing others.

Let $\mathcal{L}$ be the Lagrangian function as:

$$\begin{aligned}\mathcal{L} = &f(\mathbf{H}^p, \mathbf{H}^n, \mathbf{U}, \mathbf{W}) - Tr((\Lambda^U)^\top U) \\ &- Tr((\Lambda^{H^p})^\top\mathbf{H}^p) - Tr((\Lambda^{H^n})^\top\mathbf{H}^n)\end{aligned} \tag{3.5}$$

where $f(\mathbf{H}^p, \mathbf{H}^n, \mathbf{U}, \mathbf{W})$ is the objective function of Eq. (3.4). The notations Λ^U, Λ^{H^p} and Λ^{H^n} are the Lagrangian multipliers for non-negativity of $\mathbf{U}$, $\mathbf{H}^p$ and $\mathbf{H}^n$.

To compute $\mathbf{U}$*, we fix* $\mathbf{H}^p$*,* $\mathbf{H}^n$ *and* $\mathbf{W}$*.* The derivative of $\mathcal{L}$ with respect to $\mathbf{U}$ is as follows:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{U}} = \mathbf{E} - \mathbf{B} - \Lambda^U$$

where the nonnegative matrices $\mathbf{B}$ and $\mathbf{E}$ are defined as follows:

$$\begin{aligned}
\mathbf{B} &= (\mathbf{A}^p)^\top \mathbf{U}\mathbf{H}^p + \mathbf{A}^p\mathbf{U}(\mathbf{H}^p)^\top + \alpha\big((\mathbf{C}\mathbf{Y}\mathbf{W}^\top)^+ \\
&+ (\mathbf{C}\mathbf{U}\mathbf{W}\mathbf{W}^\top)^-\big) + \beta((\mathbf{A}^n)^\top \mathbf{U}\mathbf{H}^n + \mathbf{A}^n\mathbf{U}(\mathbf{H}^n)^\top) \\
&+ \gamma\big(\sum_{\langle i,j,k\rangle \in \mathcal{S}} f_{ijk}(\mathbf{M}^{ijk}\mathbf{U} + \mathbf{U}(\mathbf{M}^{ijk})^\top)\big)^- \\
\mathbf{E} &= \mathbf{U}(\mathbf{H}^p)^\top \mathbf{U}^\top \mathbf{U}\mathbf{H}^p + \mathbf{U}\mathbf{H}^p\mathbf{U}^\top\mathbf{U}(\mathbf{H}^p)^\top \\
&+ \alpha\big((\mathbf{C}\mathbf{U}\mathbf{W}\mathbf{W}^\top)^+ + (\mathbf{C}\mathbf{Y}\mathbf{W}^\top)^-\big) \\
&+ \beta(\mathbf{U}(\mathbf{H}^n)^\top\mathbf{U}^\top\mathbf{U}\mathbf{H}^n + \mathbf{U}\mathbf{H}_n\mathbf{U}^\top\mathbf{U}(\mathbf{H}^n)) \\
&+ \gamma\big(\sum_{\langle i,j,k\rangle \in \mathcal{S}} f_{ijk}(\mathbf{M}^{ijk}\mathbf{U} + \mathbf{U}(\mathbf{M}^{ijk})^\top)\big)^+ + \lambda\mathbf{U}
\end{aligned} \tag{3.6}$$

where for any matrix $\mathbf{X}$, $(\mathbf{X})^+$ and $(\mathbf{X})^-$ denote the positive and negative parts of $\mathbf{X}$, respectively.

Setting $\frac{\partial \mathcal{L}}{\partial \mathbf{U}} = 0$ and using the KKT complementary condition $\mathbf{U}_{ij}\Lambda^U_{ij} = 0$, we can derive the update rule for $\mathbf{U}$ as follows:

$$\mathbf{U}_{ij} \leftarrow \mathbf{U}_{ij}\sqrt{\frac{\mathbf{B}_{ij}}{\mathbf{E}_{ij}}} \tag{3.7}$$

To compute $\mathbf{H}^p$*, we fix* $\mathbf{U}$*,* $\mathbf{H}^n$ *and* $\mathbf{W}$*.* The derivative of $\mathcal{L}$ with respect to $\mathbf{H}^p$ is as follows:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{H}^p} = \mathbf{U}^\top\mathbf{U}\mathbf{H}^p\mathbf{U}^\top\mathbf{U} + \lambda\mathbf{H}^p - \mathbf{U}^\top\mathbf{A}^p\mathbf{U} - \Lambda^{H^p} \tag{3.8}$$

Setting $\frac{\partial \mathcal{L}}{\partial \mathbf{H}^p} = 0$ and using the KKT complementary condition $\mathbf{H}^p_{ij}\Lambda^{H^p}_{ij} = 0$, we can get the update rule for $\mathbf{H}^p$ as follows:

$$\mathbf{H}^p_{ij} \leftarrow \mathbf{H}^p_{ij}\sqrt{\frac{[\mathbf{U}^\top\mathbf{A}^p\mathbf{U}]_{ij}}{[\mathbf{U}^\top\mathbf{U}\mathbf{H}^p\mathbf{U}^\top\mathbf{U} + \lambda\mathbf{H}^p]_{ij}}} \tag{3.9}$$

To compute $\mathbf{H}^n$*, we fix* $\mathbf{U}$*,* $\mathbf{H}^p$ *and* $\mathbf{W}$*.* The derivative of $\mathcal{L}$ with respect to $\mathbf{H}^n$ is as follows:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{H}^n} = \mathbf{U}^\top\mathbf{U}\mathbf{H}^n\mathbf{U}^\top\mathbf{U} + \lambda\mathbf{H}^n - \mathbf{U}^\top\mathbf{A}^n\mathbf{U} - \Lambda^{H^n} \tag{3.10}$$

Setting $\frac{\partial \mathcal{L}}{\partial \mathbf{H}^n} = 0$ and using the KKT complementary condition $\mathbf{H}^n_{ij}\Lambda^{H^n}_{ij} = 0$, we can get the update rule for $\mathbf{H}^n$ as:

$$\mathbf{H}^n_{ij} \leftarrow \mathbf{H}^n_{ij}\sqrt{\frac{[\mathbf{U}^\top\mathbf{A}^n\mathbf{U}]_{ij}}{[\mathbf{U}^\top\mathbf{U}\mathbf{H}^n\mathbf{U}^\top\mathbf{U} + \lambda\mathbf{H}^n]_{ij}}} \tag{3.11}$$

To compute $\mathbf{W}$*, we fix* $\mathbf{U}$*,* $\mathbf{H}^p$ *and* $\mathbf{H}^p$*.* Setting $\frac{\partial \mathcal{L}}{\partial \mathbf{W}} = 0$, we obtain the following:

$$\mathbf{W} = (\mathbf{U}^\top\mathbf{C}\mathbf{U} + \lambda\mathbf{I})^{-1}\mathbf{U}^\top\mathbf{C}\mathbf{Y} \tag{3.12}$$

After we learn $\mathbf{W}$ and $\mathbf{U}$, the label of a node in the signed social network $u_i \in \mathcal{U}^U$ can be predicted as follows:

$$c^* = \arg\max_{c_j \in \mathcal{C}}([\mathbf{U}_i\mathbf{W}]_j) \tag{3.13}$$

4 Experiments

In this section, we present experiments, which evaluate the effectiveness of the proposed NCSSN framework. We begin by introducing experimental settings. Then we will present performance comparisons and then the impact of negative links. Finally, we present sensitivity analysis of NCSSN.

4.1 Experimental Settings For this study, we collected two datasets from Epinions and Slashdot that explicitly allow users to establish both positive and negative links.

Slashdot is a technology news platform in which users can create friend (positive) and foe (negative) links to other users. Users in Slashdot can join in some interest groups. For users, we collect their positive and negative links as well as groups that they join. These group identifiers are treated as the class labels.

Epinions is a popular product review site. Users can create both positive (trust) and negative (distrust) links to other users. They can write reviews for products from various categories. We chose these categories as the class labels of users. In particular, for users, we collected their positive and negative links, and also the categories of the products for which they write reviews. For a user who writes reviews for products from multiple categories, we chose the one with most products she writes reviews to as her label.

Some additional preprocessing was performed on these two datasets by filtering users without any positive and negative links, and class labels with a limited number of users. A number of key statistics of these datasets are illustrated in Table 2[4]. It is evident from

[4] We will make these two datasets publicly available via `http://jiliang.xyz/signed.html`

Table 1: Performance Comparison of Node Classification in Slashdot and Epinions.

Algorithms	Slashdot				Epinions			
	5%	10%	15%	20%	5%	10%	15%	20%
ICA	19.99	20.86	21.25	21.34	10.57	11.00	11.45	11.99
GReg	19.57	20.91	21.23	22.03	9.71	11.02	11.51	12.10
SocDim	20.07	21.49	21.93	22.32	10.91	11.22	11.71	12.31
NCSSN	23.62	24.74	25.41	25.87	12.29	12.82	13.69	14.03
Random	8.33	8.31	8.34	8.34	4.98	5.00	5.00	5.02

Table 2: Statistics of Datasets

Datasets	Slashdot	Epinions
# Of Users	14,799	23,280
# of Positive Links	189,642	291,422
Density of Positive Links	1.7e-3	1.1e-3
# of Negative Links	42,829	40,792
Density of Positive Links	3.9e-4	1.5e-4
# of Classes	10	20

these statistics that (a) positive links are denser in signed social networks; and (b) the signed social network in Slashdot is denser than that in Epinions.

In each case, we randomly choose $x\%$ of nodes labeled and the remaining $1 - x\%$ as unlabeled nodes for testing. For each x, we repeat the experiments 10 times and report the average performance. Since it is very common for labels to be sparsely specified, we chose relatively small values of x, which were $\{5, 10, 15, 20\}$. One commonly used measure, referred to as Micro-F1, is adopted to assess the classification performance. This measure is defined as $Micro\text{−}F1 = \frac{2pr}{p+r}$, where p and r denote the micro-average of precision and recall, respectively. These measures are formally defined as follows:

$$p = \frac{\sum_{j=1}^{m} TP_j}{\sum_{j=1}^{m}(TP_j + FP_j)},\ r = \frac{\sum_{j=1}^{m} TP_j}{\sum_{j=1}^{m}(TP_j + FN_j)}$$

where TP_j, FP_j, and FN_j are true positive, false positive and false negative of the j-th class label, respectively.

4.2 Comparison of Classification Performance

In this subsection, we evaluate the effectiveness of the proposed framework NCSSN in term of node classification performance. The comparison results are demonstrated in Table 1. The algorithms in the table are defined as follows:

- *ICA*: This algorithm uses local neighborhood information to learn local classifiers [17].
- *GReg*: This algorithm propagates labels from labeled nodes to unlabeled nodes by performing random walks on the social network [35].
- *SocDim*: This algorithm first extracts social dimensions for users and then utilizes them as features for discriminative learning [29].
- *Random*: this algorithm randomly assigns labels to unlabeled nodes.

Note that *ICA*, *GReg* and *SocDim* are developed for unsigned social networks and we apply them to signed social networks by ignoring negative links. For methods with parameters, we used cross-validation to determine their values. The parameters for *NCSSN* are set as $\{\alpha = 1, \beta = 0.5, \gamma = 0.3, K = 500\}$ in Slashdot and $\{\alpha = 1, \beta = 1, \gamma = 0.7, K = 2000\}$ in Epinions. A detailed sensitivity analysis of *NCSSN* will be discussed in the following subsections. We make some key observations from Table 1:

- In general, with the increase in the number of labeled nodes, the classification performance consistently increases for all methods in Table 1.
- *ICA*, *GReg* and *SocDim* obtains much better performance than *Random*. For example, *SocDim* gains 140.94% and 118.20% relative improvement with 5% labeled nodes in Slashdot and Epinions, respectively. These results support the importance of positive links in node classification.
- The proposed framework *NCSSN* consistently outperforms all baseline methods (e.g., *ICA*, *GReg* and *SocDim*), which only exploit positive links, on both datasets. Compared to the best performance of *ICA*, *GReg* and *SocDim*, the relative performance improvement of the proposed framework *NCSSN* is shown in Table 3. In addition to positive links, signed social networks also provide negative links and *NCSSN* captures positive and negative links into a coherent model. These observations suggest the importance of negative links and we will investigate the effects of negative links on the proposed framework in the following subsection. .

Table 3: Relative Performance Improvement of *NCSSN* Compared to the Best Performance of *ICA*, *GReg* and *SocDim*.

	5%	10%	15%	20%
Slashdot	+17.68%	+15.12%	+15.87%	+15.91%
Epinions	+12.65%	+14.26%	+16.91%	+13.97%

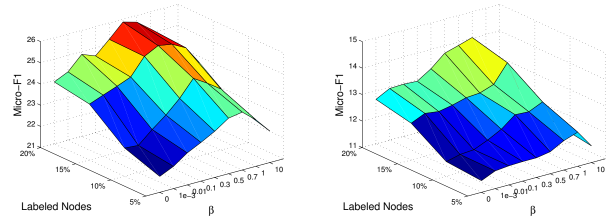

(a) Slashdot (b) Epinions

Figure 3: β vs Percentage of Labeled Nodes.

In summary, the aforementioned analysis provides the insights that (a) both positive and negative links in signed social networks can help node classification; and (b) the proposed *NCSSN* framework obtains significant performance improvement compared to other methods.

4.3 Impact of Negative Links Based on the performance comparison in the previous subsection, we observe that the proposed *NCSSN* framework improves the classification performance significantly. In addition to positive links, the proposed framework *NCSSN* also captures negative links in signed social networks. In this subsection, we investigate the impact of negative links on *NCSSN*.

The optimization problem for *NCSSN* in Eq. (3.4) provides two terms to model negative links - (1) $\|\mathbf{A}^n - \mathbf{U}\mathbf{H}^n\mathbf{U}^\top\|_F^2$ modeling independent information from negative links ; and (2) $\sum_{\langle i,j,k\rangle\in\mathcal{S}} f_{ijk}Tr(\mathbf{M}^{ijk}\mathbf{U}\mathbf{U}^\top)$ modeling dependent information from negative links. To study the impact of negative links on the proposed framework, we systematically eliminate the effects of these two components by defining the following variants of *NCSSN*:

- NCSSN\II - We eliminate the effect of the independent component on *NCSSN*. In particular, we remove the term $\|\mathbf{A}^n - \mathbf{U}\mathbf{H}^n\mathbf{U}^\top\|_F^2$ from the optimization problem in Eq. (3.4) by setting $\beta = 0$.

- NCSSN\DI - We eliminate the effect of the dependent component on *NCSSN* by removing the term $\sum_{\langle i,j,k\rangle\in\mathcal{S}} f_{ijk}Tr(\mathbf{M}^{ijk}\mathbf{U}\mathbf{U}^\top)$ from the optimization problem in Eq. (3.4) by setting $\gamma = 0$.

- NCSSN\DII - We eliminate the effects of both independent and dependent components on *NCSSN* by setting $\beta = 0$ and $\gamma = 0$ in Eq. (3.4).

The parameters in all variants are determined with cross-validation and the performance comparison of *NCSSN* and its variants are demonstrated in Figure 2. We make following observations:

- When we eliminate the effect of the independent component, the performance of NCSSN\II degrades in comparison with *NCSSN*. For example, the performance reduces 8.03% and 6.95% with 5% of labeled users in Slashdot and Epinions, respectively. These results suggest that the independent component in *NCSSN* is important.

- We make similar observations for NCSSN\DI when we eliminate the effect of the component modeling dependent information.

- When we eliminate the effects of both components, the performance of NCSSN\DII further reduces compared to NCSSN\II and NCSSN\DI. These results suggest that these components are complementary to each other.

In summary, via the component analysis of *NCSSN*, we conclude that (a) both components can contribute to the performance improvement of *NCSSN*; (b) it is necessary to model both because they contain complementary information.

4.4 Parameter Analysis The proposed *NCSSN* framework has three important parameters. The first is β, which controls the contribution from independent information from negative links, the second is γ, which controls the contribution from dependent information from negative links, and the third is the number of dimensions K of latent user factors. Hence we study the effect of each of the three parameters by fixing the other 2 to see how the performance of *NCSSN* varies with the percentage of labeled nodes.

We fix $\{\gamma = 0.3, K = 500\}$ and $\{\gamma = 0.7, K = 2000\}$ for Slashdot and Epinions, respectively, and vary β as $\{0, 0.001, 0.1, 0.3, 0.5, 0.7, 1, 10\}$. The performance variations with respect to β and the percentage of labeled nodes are depicted in Figures 3(a) and 3(b) for Slashdot and Epinions, respectively. The following observations are evident:

- When β increases from 0, eliminating the impact of independent information from negative links, to 0.1, the performance increases dramatically in both datasets. These results further support the

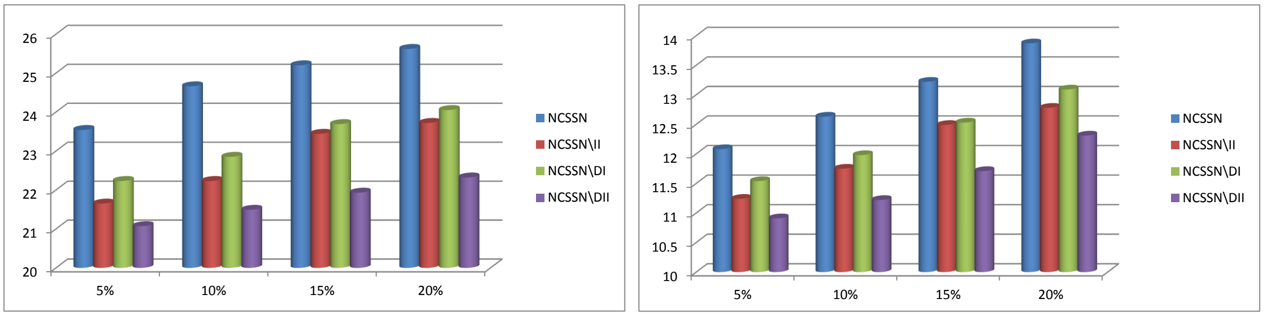

(a) Slashdot (b) Epinions

Figure 2: The Impact of Negative Links on The Proposed Framework.

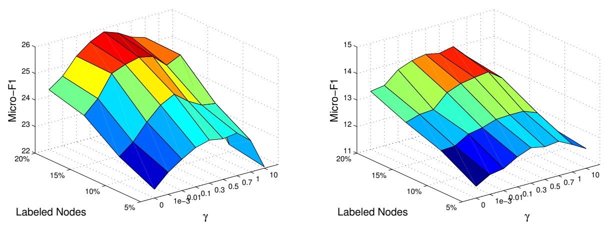

(a) Slashdot (b) Epinions

Figure 4: γ vs Percentage of Labeled Nodes.

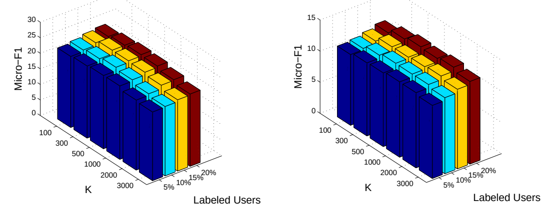

(a) Slashdot (b) Epinions

Figure 5: Number of Latent Dimensions K vs Percentage of Labeled Nodes.

importance of independent information of negative links in the classification performance.

- The performance is relatively stable in fairly large ranges such as $\beta \in [0.3, 0.7]$ for Slashdot, and $\beta \in [0.5, 1]$ for Epinions. This property eases the parameter setting process.
- When β increases from 1 to 10, the performance decreases. When β is large, independent information of negative links dominate the learning process. This causes the estimates of $\mathbf{U}$ and $\mathbf{W}$ overfitting.

We vary γ as $\{0, 0.001, 0.1, 0.3, 0.5, 0.7, 1, 10\}$ while setting $\{\beta = 0.5, K = 500\}$ and $\{\beta = 1, K = 2000\}$ for Slashdot and Epinions, respectively. The performance variations in terms of γ and the percentage of labeled nodes are demonstrated in Figures 4(a) and 4(b) for Slashdot and Epinions, respectively. We have very similar observations for γ as β. With the increase of γ, the performance first increases, reaches its peak value and degrades. This pattern can be used to determine the optimal value of γ for *NCSSN*.

Next, we varied K as $\{100, 300, 500, 1000, 2000, 3000\}$ while we fixed $\{\beta = 0.5, \gamma = 0.3\}$ and $\{\beta = 1, \gamma = 0.7\}$ for Slashdot and Epinions, respectively. The performance variations in terms of the number of latent dimensions K and the percentage of labeled users are illustrated in Figures 5(a) and 5(b) for Slashdot and Epinions, respectively. We make the following observations:

- In general, with the increase of the number of latent dimensions, the performance tends to increase first, reaches its best performance, and then decreases.
- We note that *NCSSN* achieves its best performance when $K = 500$ in Slashdot while $K = 2000$ in Epinions. There are two possible reasons why *NCSSN* achieves its best performance with relative smaller number of latent dimensions in Slashdot than Epinions - (a) the signed social network in Epinions has more users than that in Slashdot; and (b) users in Slashdot are better connected than these in Epinions.

Among the these three parameters, the performance of *NCSSN* is most sensitive to β than γ and the number of latent dimensions K.

5 Related Work

Our work is related to node classification in unsigned social networks and mining signed social networks. In this section, we will give a brief review.

5.1 Node Classification in Unsigned Social Networks

Many real-world problems can be modeled as the node classification problem in unsigned social networks, which aims to predict labels of unlabeled nodes in an unsigned social network by giving the network with some labeled users [22]. According to [5], node classification algorithms in unsigned social networks can be mainly divided into local classifier methods and random walk based methods.

Local classifier based methods use local neighborhood information to learn local classifiers. Iterative classification methods (or ICA) [21] and its variants [17] construct feature vectors for nodes from the information known about them and their neighborhood. These feature vectors are then used along with labeled users to build a local classifier. Weighted-vote relational neighbor (wvRN) takes a weighted average of the class probabilities in the neighborhood for classification [19].

Random walk based methods propagate the labels by performing random walks on the network such as label propagation [36, 3], graph regularization [34, 35] and adsorption [4]. Other related approaches include inference using graphical models [30], metric labeling [13], and spectral partitioning [20].

5.2 Mining Signed Social Networks

Although mining signed networks is still in the early stages of development, some research tasks have been relatively well studied such as positive and negative link prediction and community detection. A comprehensive overview about mining signed social networks can be found in [28].

Positive and negative link prediction infers new positive and negative links by giving a snapshot of a signed network. In [16], local topology-based features based on balance theory are extracted to improve the performance of a logistic regression classifier in signed relation prediction. Features derived from longer cycles in signed networks can be used to improve the positive and negative link prediction performance [7]. In [12], a low-rank matrix factorization approach with generalized loss functions is proposed to predict trust and distrust relations.

The problem of community detection in signed networks determines groups of users that are densely connected with positive links within the same group and negative links between groups [23, 15]. In [32], an agent-based random walk model is proposed to identify communities. In addition to modularity maximization, frustration minimization is defined as another optimized object for signed social networks, which tries to minimize the number of positive links among different groups and the number of negative links inside the same group [2]. Several approaches are discussed in [8] about how to extend normalized cut from unsigned to signed social networks for community detection. In [6], a signed probabilistic mixture (SPM) model is proposed to detect overlapping communities in undirected signed networks.

6 Conclusions

Most of the existing node classification algorithms focus on unsigned social networks; while little work exists for the case of signed social networks. It is evident from recent research that negative links have the potential to benefit various analytical tasks. In this paper, we study the novel problem of node classification in signed social networks. We first make a number of observations about positive and negative links, provide an approach to mathematically model independent and dependent information from positive and negative links and propose the NCSSN framework for node classification in signed social networks. Experiments are conducted on two real-world signed social network datasets, i.e., Slashdot and Epinions, and the results show the effectiveness of the proposed framework. Our detailed experiments also illustrate the importance of negative links in the proposed framework.

There are several interesting directions needing further investigation. First, since node classification for unsigned social networks has been extensively studied and this paper provides findings about positive and negative links, we would like to investigate how to transform algorithms from unsigned to signed social networks. Second, our recent work shows that negative links can be predicted for these social media sites which do not allow users to specify negative links such as Facebook and Twitter [25]; therefore we would like to apply the proposed framework to signed social networks with explicit positive links and predicted negative links from these social media sites. Finally mining signed social networks is still a very new topic and there are unprecedented opportunities; hence we will investigate more tasks in signed social networks.

Acknowledgments

This material is based upon work supported by, or in part by, the U.S. Army Research Office (ARO) under contract/grant number 025071, the Office of Naval Research(ONR) under grant number N000141010091, and the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

[1] M. A. Abbasi, J. Tang, and H. Liu. Scalable learning of users preferences using networked data. In *HT*, 2014.

[2] P. Anchuri and M. Magdon-Ismail. Communities and balance in signed networks: A spectral approach. In *ASONAM*, 2012.

[3] A. Azran. The rendezvous algorithm: Multiclass semi-supervised learning with markov random walks. In *ICML*, 2007.

[4] S. Baluja, R. Seth, D. Sivakumar, Y. Jing, J. Yagnik, S. Kumar, D. Ravichandran, and M. Aly. Video suggestion and discovery for youtube: taking random walks through the view graph. In *WWW*, 2008.

[5] S. Bhagat, G. Cormode, and S. Muthukrishnan. Node classification in social networks. In *Social network data analytics*,Springer, 2011.

[6] Y. Chen, X.-l. Wang, and B. Yuan. Overlapping community detection in signed networks. *arXiv preprint arXiv:1310.4023*, 2013.

[7] K.-Y. Chiang, N. Natarajan, A. Tewari, and I. S. Dhillon. Exploiting longer cycles for link prediction in signed networks. In *CIKM*, 2011.

[8] K.-Y. Chiang, J. J. Whang, and I. S. Dhillon. Scalable clustering of signed networks using balance normalized cut. In *CIKM*, 2012.

[9] C. Ding, T. Li, W. Peng, and H. Park. Orthogonal nonnegative matrix t-factorizations for clustering. In *KDD*, 2006.

[10] L. Getoor and C. P. Diehl. Link mining: a survey. *ACM SIGKDD Explorations Newsletter*, 7(2):3–12, 2005.

[11] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW*, 2004.

[12] C.-J. Hsieh, K.-Y. Chiang, and I. S. Dhillon. Low rank modeling of signed networks. In *KDD*, 2012.

[13] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. *JACM*, 2002.

[14] J. Kunegis, J. Preusse, and F. Schwagereit. What is the added value of negative links in online social networks? In *WWW*, 2013.

[15] J. Kunegis, S. Schmidt, A. Lommatzsch, J. Lerner, E. W. De Luca, and S. Albayrak. Spectral analysis of signed graphs for clustering, prediction and visualization. In *SDM*, 2010.

[16] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, 2010.

[17] Q. Lu and L. Getoor. Link-based classification. In *ICML*, 2003.

[18] H. Ma, M. R. Lyu, and I. King. Learning to recommend with trust and distrust relationships. In *Recsys*, 2009.

[19] S. A. Macskassy and F. Provost. A simple relational classifier. Technical report, DTIC Document, 2003.

[20] F. McSherry. Spectral partitioning of random graphs. In *FOCS* 2001.

[21] J. Neville and D. Jensen. Iterative classification in relational data. In *Proc. AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, 2000.

[22] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93, 2008.

[23] T. Sharma, A. Charls, and P. Singh. Community mining in signed social networks-an automated approach. *ICCEA09, Manila*, 9, 2009.

[24] M. Szell, R. Lambiotte, and S. Thurner. Multirelational organization of large-scale social networks in an online world. *PNAS*, 2010.

[25] J. Tang, S. Chang, C. Aggarwal, and H. Liu. Negative link prediction in social media. In *WSDM*, 2015.

[26] J. Tang, H. Gao, X. Hu, and H. Liu. Exploiting homophily effect for trust prediction. In *WSDM*, 2013.

[27] J. Tang, X. Hu, and H. Liu. Is distrust the negation of trust?: the value of distrust in social media. In *HT*, 2014.

[28] J. Tang, Y. Chang, C. Aggarwal, and H. Liu. A Survey of Signed Network Mining in Social Media. In *arXiv preprint arXiv:1511.07569*, 2015.

[29] L. Tang and H. Liu. Relational learning via latent social dimensions. In *KDD*, 2009.

[30] B. Taskar, P. Abbeel, and D. Koller. Discriminative probabilistic models for relational data. In *UAI*, 2002.

[31] P. Victor, C. Cornelis, M. De Cock, and A. Teredesai. Trust-and distrust-based recommendations for controversial reviews. In *WebSci'09*, 2009.

[32] B. Yang, W. K. Cheung, and J. Liu. Community mining from signed social networks. *TKDE*, 2007.

[33] E. Zheleva and L. Getoor. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. In *WWW*, 2009.

[34] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. *NIPS*, 2004.

[35] D. Zhou, J. Huang, and B. Schölkopf. Learning from labeled and unlabeled data on a directed graph. In *ICML*, 2005.

[36] X. Zhu, Z. Ghahramani, J. Lafferty, et al. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*, 2003.