Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

Recommending Positive Links in Signed Social Networks by Optimizing a Generalized AUC

Dongjin Song[*] **and David A. Meyer**[†]
[*]Department of Electrical and Computer Engineering,University of California, San Diego
La Jolla, USA, 92093-0409. Email: dosong@ucsd.edu
[†]Department of Mathematics,University of California, San Diego
La Jolla, USA, 92093-0112. Email: dmeyer@math.ucsd.edu

Abstract

With the rapid development of signed social networks in which the relationships between two nodes can be either positive (indicating relations such as like) or negative (indicating relations such as dislike), producing a personalized ranking list with positive links on the top and negative links at the bottom is becoming an increasingly important task. To accomplish it, we propose a generalized AUC (GAUC) to quantify the ranking performance of potential links (including positive, negative, and unknown status links) in partially observed signed social networks. In addition, we develop a novel link recommendation algorithm by directly optimizing the GAUC loss. We conduct experimental studies based upon Wikipedia, MovieLens, and Slashdot; our results demonstrate the effectiveness and the efficiency of the proposed approach.

Introduction

Signed social networks, in which the relationship between two nodes can be either positive (indicating relations such as trust) or negative (indicating relations such as distrust), are increasingly common in recent years. For instance, Slashdot (Lampe, Johnston, and Resnick 2007; Brzozowski, Hogg, and Szabo 2008; Kunegis, Lommatzsch, and Bauckhage 2009) is a social website of technology related news; in Slashdot Zoo, users can tag each other as friends (like) or foes (dislike) based upon their comments on different articles. Another example is Epinions (Guha et al. 2004), which is a product review website with an active user community. Users can express whether they trust or distrust other users based upon their reviews. Unlike traditional unsigned social networks which can be represented as binary adjacency matrices (with 1 indicating the existence of a link and 0 indicating the unknown status of a link), a signed network can be modeled as an adjacency matrix in which an entry is 1 (or -1) if the relationship is positive (or negative) and 0 if the relationship is unknown.

One fundamental problem in signed social networks is link recommendation. In unsigned social networks, link recommendation aims to suggest to each user a list of people (items) to whom the user probably will create new connections. In signed social networks, however, given a user, the aim could be to rank people (items) this user is interested in (i.e., positive links) on the top and people (items) this user is not interested in (i.e., negative links) at the bottom. To achieve this goal, it is important to quantify effectively how good such a ranking list is.

In the past few years, the area under the receiver operating characteristic (ROC) curve (AUC) (Hanley and Mcneil 1982) has shown its effectiveness for measuring the performance of binary classification or ranking. For instance, Cortes and Mohri (2004) showed that the average AUC is monotonically increasing as a function of the classification accuracy, while the standard deviation for uneven distributions and higher error rates was observed; Clauset et al. (2008) as well as Menon and Elkan (2011) used AUC to measure the effectiveness of link prediction; Rendle et al. (2009) utilized AUC to quantify personalized ranking performance within social networks.

The AUC, however, is not an appropriate way to quantify the ranking performance in signed networks because it only applies to the binary case, rather than to the triplet (positive, negative, and unknown) that we consider. In particular, given a signed social network, if AUC treats positive links as positive samples and the other links as negative samples, it will be impossible to quantify the ranking quality of negative links since unknown status links can either be positive or negative. Similarly, if AUC treats negative links as negative samples and the other links as positive samples, it will not measure the ranking performance of positive links correctly.

A plausible model for user behavior in signed social networks is that more extreme positive and negative relationships will be revealed before less extreme ones. Such a model implies that a personalized ranking list of potential links should place positive links at the top, negative links at the bottom, and unknown status links in the middle. Therefore, we propose a novel criterion—a generalized AUC (GAUC)—to quantify the ranking performance in signed social networks, and develop a new algorithm to perform link recommendation. The contributions of our work include:

- we propose a generalized AUC which will be maximized only if all positive links are ranked on top, all negative links are ranked at the bottom, and all unknown status links are in the middle.

- we develop a link recommendation approach by directly minimizing the loss of the proposed GAUC; the proposed model enjoys both low computational complexity and high memory efficiency.
- we conduct experimental studies based upon three real world datasets, i.e., Wikipedia, MovieLens, and Slashdot; the results demonstrate the effectiveness of the generalized AUC for quantifying link recommendation in signed social networks; the results also demonstrate the efficiency and the effectiveness of the proposed approach.

Related Work

In the past few years, various approaches (Heider 1946; Doreian and Mrvar 2009; Leskovec, Huttenlocher, and Kleinberg 2010a; 2010b; Dong et al. 2012; Yang et al. 2012; Ye et al. 2013; Song and Meyer 2014) have been developed to investigate signed social networks. Among these approaches, while most of them concentrate on edge sign prediction and community detection, few have focussed on link recommendation. Although the link recommendation problem in signed social networks is more difficult than in conventional unsigned social networks, we review related work in the conventional setting to motivate our proposed approach. In general, there are two types of link recommendation approaches in unsigned social networks. One uses network topology and another is model-based.

Network topological approaches recommend links based upon the network topological structure. For instance, Liben-Nowell and Kleinberg (2007) showed that common neighbors, Jaccard's coefficient, and Adamic/Adar (Adamic and Adar 2003) can be used for link prediction and recommendation; Katz (1953), which measures the ensemble of all paths, also demonstrated its effectiveness for link prediction and recommendation. These approaches, however, cannot perform link recommendation effectively when little network topological information is available, which often happens in real world sparse social networks.

To tackle this issue, model-based approaches have been developed. Specifically, there are two types of model-based methods. One is pointwise approaches, such as singular value decomposition (SVD), non-negative matrix factorization (NMF) (Lee and Seung 1999; Song, Meyer, and Min 2014), probabilistic matrix factorization (Salakhutdinov and Mnih 2007), and matrix factorization (MF) (Koren, Bell, and Volinsky 2009), which aim to reconstruct the network adjacency matrix with a low rank approximation. Another is pairwise approaches, such as maximum margin matrix factorization (Weimer, Karatzoglou, and Smola 2008) and Bayesian personalized ranking (BPR) (Rendle et al. 2009), which aim to provide a personalized ranking list based upon pairwise comparisons.

Although most of these approaches can be applied directly for link recommendation in signed social networks, they may not perform well because their objectives are inconsistent with that of link recommendation in signed networks. Therefore, we introduce a generalized AUC (GAUC) to quantify the personalized ranking performance in signed networks and develop a model to perform link recommendation by directly minimizing the loss of GAUC.

A Generalized AUC

In this section, we first introduce the area under the ROC curve (AUC) and then present a generalized AUC to measure the ranking performance in signed networks.

AUC

Given a binary classifier f and a training set $(a_i, b_i)_{i=1}^n$ with $a_i \in \mathbb{R}^d$ and $b_i \in \{-1, 1\}$, let $\mathcal{P} = \{a_i \mid b_i = 1\}$ be the set of positive samples and $\mathcal{N} = \{a_i \mid b_i = -1\}$ the set of negative samples. Then the AUC is defined by:

$$AUC = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{a_i \in \mathcal{P}} \sum_{a_j \in \mathcal{N}} I\big(f(a_i) > f(a_j)\big) \quad (1)$$

where $I(\cdot)$ is an indicator function which is 1 if the condition in the parenthesis is satisfied and 0 otherwise; $|\mathcal{P}|$ and $|\mathcal{N}|$ are the numbers of positive samples and negative samples, respectively. AUC is the value of the Wilcoxon-Mann-Whitney statistic (Hanley and Mcneil 1982) which is essentially the probability that a random element of one set $f(a_i)$ is larger than a random element of another $f(a_j)$. With an ideal ranking list, AUC should be 1 representing each positive sample is ranked higher than all the negative samples. For a random ranking, AUC will be 0.5.

A Generalized AUC

Given a signed social network which contains positive, negative, and unknown status links, an ideal personalized ranking list tends to rank positive links (indicating a relationship such as trust) on the top, negative links (indicating a relationship such as distrust) at the bottom, and unknown status links in the middle. Traditional AUC, however, cannot quantify such a ranking list appropriately because it considers only the binary case. For instance, if AUC treats positive links as positive examples and the other links as negative examples, it cannot quantify the ranking quality of negative links since unknown status links could be either positive or negative. Similarly, if AUC treats negative links as negative samples and the others as positive samples, it cannot measure the ranking quality of positive links.

Note that although mean average precision (MAP) and normalized discounted cumulative gain (NDCG) can be used to measure ranking performance in signed networks, they may not perform well because they tend to overestimate the positive links on the top and cannot quantify the negative links appropriately, because we aim to rank them at the bottom of the ranking list and there are much more unknown status links than negative links in real world applications.

To resolve these issues, we develop a novel criterion, named generalized AUC (GAUC), to measure the ranking performance in signed networks based upon the assumption that more extreme positive and negative relationships will be revealed before less extreme ones. Specifically, suppose we are given a classifier f and a training set $(a_i, b_i)_{i=1}^n$ with $a_i \in \mathbb{R}^d$ the ith sample, $b_i \in \{-1, 0, 1\}$ its label.

Let $\mathcal{P} = \{a_i \mid b_i = 1\}$ be the set of positive samples, $\mathcal{N} = \{a_i \mid b_i = -1\}$ be the set of negative samples, and $\mathcal{O} = \{a_i \mid b_i = 0\}$. Then GAUC can be defined as:

$$\begin{aligned} \textit{GAUC} = &\frac{\eta}{|\mathcal{P}|(|\mathcal{O}| + |\mathcal{N}|)} \sum_{a_i \in \mathcal{P}} \sum_{a_s \in \mathcal{O} \bigcup \mathcal{N}} I\big(f(a_i) > f(a_s)\big) \\ &+ \frac{1-\eta}{|\mathcal{N}|(|\mathcal{O}| + |\mathcal{P}|)} \sum_{a_t \in \mathcal{O} \bigcup \mathcal{P}} \sum_{a_j \in \mathcal{N}} I\big(f(a_j) < f(a_t)\big) \end{aligned} \tag{2}$$

where $|\mathcal{P}|$ denotes the number of positive samples, $|\mathcal{N}|$ represents the number of negative samples, and $|\mathcal{O}|$ is the number of unknown status samples. The first and second terms quantify the ranking performance of positive links and negative links, respectively. $0 \leq \eta \leq 1$ is a parameter which controls the tradeoff between these two terms. In this work, we set $\eta = \frac{|\mathcal{P}|}{|\mathcal{P}|+|\mathcal{N}|}$, i.e., the tradeoff is controlled by the relative fraction of the number of positive links and that of negative links. After substituting η in Eqn. (2), we can obtain

$$\begin{aligned} \textit{GAUC} = &\frac{1}{(|\mathcal{P}| + |\mathcal{N}|)} \Big(\frac{1}{(|\mathcal{O}| + |\mathcal{N}|)} \sum_{a_i \in \mathcal{P}} \sum_{a_s \in \mathcal{O} \bigcup \mathcal{N}} I\big(f(a_i) \\ &> f(a_s)\big) + \frac{1}{(|\mathcal{O}| + |\mathcal{P}|)} \sum_{a_j \in \mathcal{N}} \sum_{a_t \in \mathcal{O} \bigcup \mathcal{P}} I\big(f(a_j) < f(a_t)\big) \Big). \end{aligned} \tag{3}$$

Like the AUC, GAUC will be 1 if we have a perfect ranking list and will be 0.5 if we have a random ranking list. The main difference between them is that GAUC can jointly quantify the ranking quality over positive links and negative links, in the presence of unknown links. We are aware that different variants of AUC (Nakas and Yiannoutsos 2004; Li 2009) have been developed in the past. GAUC differs from these variants by focusing on the head and tail of a ranking list.

Optimizing the Generalized AUC

In this section, we first state the problem we aim to study. Then, we propose a novel model for link recommendation by directly minimizing the loss of GAUC. Finally, we introduce an optimization procedure for our proposed approach.

Problem Statement

Suppose we are given a partially observed signed network $X \in \mathbb{R}^{n \times n}$ with $X_{ij} \in \{1, -1, 0, ?\}$, where 1 denotes a positive link, -1 represents a negative link, 0 is an unknown status link, and ? denotes a potential positive or negative link. In the training stage, we treat both 0 and ? as zero (i.e., potential links) and study the underlying mechanism for ranking observed positive links on the top and negative links at the bottom. In the test phase, we evaluate how these potential links are ranked based upon the relative positions of potential positive and negative links.

We aim to learn a mapping function f such that a ranking score for the link at i-th row and j-th column of X can be produced as

$$f(i, j, X) = \widehat{X_{ij}}. \tag{4}$$

Since many real world signed social networks are sparse graphs with low rank structure, a low rank model which is memory efficient can be employed to obtain the ranking score as follows:

$$f(i, j, U, V) = f(U_i, V_j) = U_i^T V_j \tag{5}$$

where $U_i \in \mathbb{R}^r$, $V_j \in \mathbb{R}^r$, and r is the rank ($r \ll n$).

Link Recommendation Model

Since GAUC is a reasonable way to quantify the ranking performance in signed social networks, an ideal link recommendation model would be expected to optimize GAUC directly. Based upon the low rank model in Eqn. (5), the loss of GAUC can be defined as:

$$\begin{aligned} &1 - \textit{GAUC}(U, V) = \\ &\frac{1}{(|\mathcal{P}| + |\mathcal{N}|)} \Big(\frac{1}{(|\mathcal{O}| + |\mathcal{N}|)} \sum_{X_{ij} \in \mathcal{P}} \sum_{X_{is} \in \mathcal{O} \bigcup \mathcal{N}} I(U_i^T V_j \\ &\leq U_i^T V_s) + \frac{1}{(|\mathcal{O}| + |\mathcal{P}|)} \sum_{X_{ij} \in \mathcal{N}} \sum_{X_{is} \in \mathcal{O} \bigcup \mathcal{P}} I(U_i^T V_j \geq U_i^T V_s) \Big). \end{aligned} \tag{6}$$

Since the indicator function $I(\cdot)$ is non-convex, hinge loss can be used as a convex surrogate in Eqn. (6). Therefore, an upper bound of GAUC loss can be derived based upon the following two inequalities (as shown in Figure 1):

$$I(U_i^T V_j \leq U_i^T V_s) \leq \max\Big(0, U_i^T (V_s - V_j) + 1\Big), \tag{7}$$

and

$$I(U_i^T V_j \geq U_i^T V_s) \leq \max\Big(0, U_i^T (V_j - V_s) + 1\Big). \tag{8}$$

Assuming $|\mathcal{O}| \gg |\mathcal{P}|$ and $|\mathcal{O}| \gg |\mathcal{N}|$ which holds for real world sparse signed graphs, the upper bound of GAUC loss can be written as the following objective:

$$\begin{aligned} \mathcal{Q}(U, V) = &\sum_{i=1}^n \sum_{j=1}^n \sum_{s=1}^n \max\Big(0, U_i^T (V_s - V_j) + 1\Big) \cdot \\ &I(X_{ij} = 1, X_{is} \neq 1) + \lambda_U \sum_i U_i^T U_i + \\ &\sum_{i=1}^n \sum_{j=1}^n \sum_{s=1}^n \max\Big(0, U_i^T (V_j - V_s) + 1\Big) \cdot \\ &I(X_{ij} = -1, X_{is} \neq -1) + \lambda_V \sum_j V_j^T V_j \end{aligned} \tag{9}$$

where the second and fourth terms are regularization terms used for preventing over-fitting. λ_U and λ_V are two hyperparameters for controlling the scale of regularization terms.

Optimization

Although $\mathcal{Q}(U, V)$ is non-convex and non-smooth with respect to U and V, a sub-gradient method can be employed to minimize the objective over U and V alternately. Specifically, the partial derivative of $\mathcal{Q}(U, V)$ given U_i is:

$$\begin{aligned} &\frac{\partial \mathcal{Q}(U, V)}{\partial U_i} = \\ &\begin{cases} \sum_{j=1}^n \sum_{s=1}^n (V_s - V_j) + \lambda_U U_i, & \text{if } X_{ij} = 1, X_{is} \neq 1, \\ & \text{and } U_i^T (V_s - V_j) > -1; \\ \sum_{j=1}^n \sum_{s=1}^n (V_j - V_s) + \lambda_U U_i, & \text{if } X_{ij} = -1, X_{is} \neq -1, \\ & \text{and } U_i^T (V_j - V_s) > -1; \\ \lambda_U U_i, & \text{otherwise} \end{cases} \end{aligned} \tag{10}$$

Algorithm 1 Optimization of $\mathcal{Q}(U, V)$

Input: X, U, V, α, number of batches b, number of iterations t, threshold ς, maximum iteration T.
Initialize: set $t = 0$, initialize U_0 and V_0 randomly
repeat
  $t = t + 1$;
  Calculate $\frac{\partial \mathcal{Q}(U_t, V_t)}{\partial U_t}$ based upon Eqn. (10);
  $U_{t+1} = U_t - \alpha \frac{\partial \mathcal{Q}(U_t, V_t)}{\partial U_t}$;
  Calculate $\frac{\partial \mathcal{Q}(U_t, V_t)}{\partial V_t}$ based upon Eqn. (11) and Eqn. (12);
  $V_{t+1} = V_t - \alpha \frac{\partial \mathcal{Q}(U_t, V_t)}{\partial V_t}$;
until $|\mathcal{Q}(U_{t+1}, V_{t+1}) - \mathcal{Q}(U_t, V_t)| < \varsigma$ or $t > T$

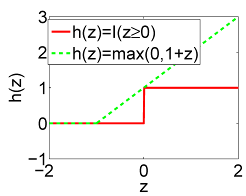

Figure 1: Convex surrogate loss of $I(z \geq 0)$.

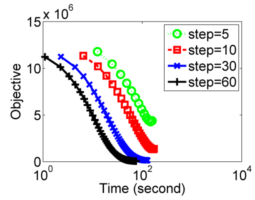

Figure 2: The efficiency of the proposed approach.

the partial derivative of $\mathcal{Q}(U, V)$ given V_j is:

$$\frac{\partial \mathcal{Q}(U, V)}{\partial V_j} = \begin{cases} -\sum_{i=1}^{n}\sum_{s=1}^{n} U_i + \lambda_V V_j, & \text{if } X_{ij} = 1, X_{is} \neq 1, \\ & \text{ and } U_i^T(V_s - V_j) > -1; \\ \sum_{i=1}^{n}\sum_{s=1}^{n} U_i + \lambda_V V_j, & \text{if } X_{ij} = -1, X_{is} \neq -1, \\ & \text{ and } U_i^T(V_j - V_s) > -1; \\ \lambda_V V_j, & \text{otherwise} \end{cases} \quad (11)$$

and the partial derivative of $\mathcal{Q}(U, V)$ given V_s is:

$$\frac{\partial \mathcal{Q}(U, V)}{\partial V_s} = \begin{cases} \sum_{i=1}^{n}\sum_{j=1}^{n} U_i + \lambda_V V_s, & \text{if } X_{ij} = 1, X_{is} \neq 1, \\ & \text{ and } U_i^T(V_s - V_j) > -1; \\ -\sum_{i=1}^{n}\sum_{j=1}^{n} U_i + \lambda_V V_s, & \text{if } X_{ij} = -1, X_{is} \neq -1, \\ & \text{ and } U_i^T(V_j - V_s) > -1; \\ \lambda_V V_s, & \text{otherwise.} \end{cases} \quad (12)$$

For practical applications, the size (or the number of the non-zero entries) of sparse signed networks maybe too large to handle. In this case, stochastic sub-gradient descent can be employed. In particular, assuming $|\mathcal{O}| \gg |\mathcal{P}|$ and $|\mathcal{O}| \gg |\mathcal{N}|$, let $q = |\mathcal{P}| + |\mathcal{N}|$ and b be the number of batches. The computational complexity of $\frac{\partial \mathcal{Q}(U,V)}{\partial U}$ (or $\frac{\partial \mathcal{Q}(U,V)}{\partial V}$) will be $O(\frac{n(q+r)}{b})$. The specific optimization procedure is shown in Algorithm 1.

Experiment

Datasets

We consider two well-known signed directed social networks, i.e., Wikipedia (Burke and Kraul 2008) and Slashdot (Lampe, Johnston, and Resnick 2007) [1]. The Wikipedia data comprise a voting network for promoting candidates to the role of admin. Slashdot is a social website focusing on technology related news. In Slashdot Zoo, users can tag each other as friends (like) or foes (dislike) based upon comments on articles.

Table 1: The statistics of two datasets.

Datasets	Wikipedia	MovieLens	Slashdot
Nodes	7,118	6040/3952	82,144
Edges	103,747	739,012	549,202
+edges	78.78%	77.84%	77.4%
−edges	21.21%	22.16%	22.6%
Density	0.0020	0.0309	0.000081

We also consider Movielens 1M dataset[2] which contains 6040 users as well as 3952 items. Although this dataset is mainly used for collaborative filtering, we preprocess it so that ratings of 4 and 5 are treated as positive links, ratings of 1 and 2 are treated as negative links, and other ones are treated as unknown status links.

The detailed statistics of these three datasets are provided in Table 1.

Evaluation Metrics

Given a fully observed signed social network $X \in \mathbb{R}^{n \times n}$ with $X_{ij} \in \{1, -1, 0\}$, we randomly remove a fraction (80%, 60% and 40%) of positive and negative links and use the rest to form a partially observed network for training (as X_{Train}). The zero entries in X_{Train} are called potential links because they could either be positive or negative in the future. The removed links form a test set X_{Test}.

To evaluate the effectiveness of the proposed model for link recommendation, we utilize GAUC (over 1, -1, and 0) in Eqn. (3), AUC (over 1 and -1) in Eqn. (1), and mean average precision (MAP) (over 1 and -1) to quantify the ranking performance over X_{Test}. To evaluate the effectiveness of top-k link recommendation, we also report its associated average Recall@k and average Precision@k for reference. Specifically, Recall@k is defined as:

$$\textit{Recall@}k = \frac{\#\textit{positive links in the top } k}{\#\textit{positive links}},$$

and Precision@k is given as

$$\textit{Precision@}k = \frac{\#\textit{positive links in the top } k}{\#\textit{positive links and negative links in the top } k}.$$

Parameter Setting

There are three hyper-parameters in our model, i.e., λ_U, λ_V, and k. We set $\lambda_U = \lambda_V$ for simplicity and search over the grid of $\{1, 5, 10, 20, 50, 100, 200\}$ to find the optimal setting for λ_U and λ_V. We also search over the grid of $\{10, 30, 50, 70, 90\}$ to find the optimal setting for k. Specifically, we conduct 5 fold cross-validation on X_{train} and

[1] These datasets are available online at http://snap.stanford.edu/data/.

[2] This dataset is available online at http://grouplens.org/datasets/movielens/.

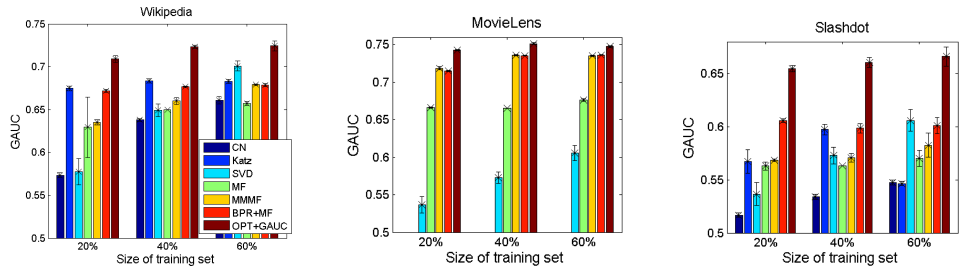

(a) GAUC on Wikipedia. (b) GAUC on MovieLens. (c) GAUC on Slashdot

Figure 3: GAUC on Wikipedia, MovieLens, and Slashdot. Error bars represent the standard deviations of uncertainty.

(a) AUC on Wikipedia. (b) AUC on MovieLens (c) AUC on Slashdot

Figure 4: AUC on Wikipedia, MovieLens, and Slashdot. Error bars represent the standard deviations of uncertainty.

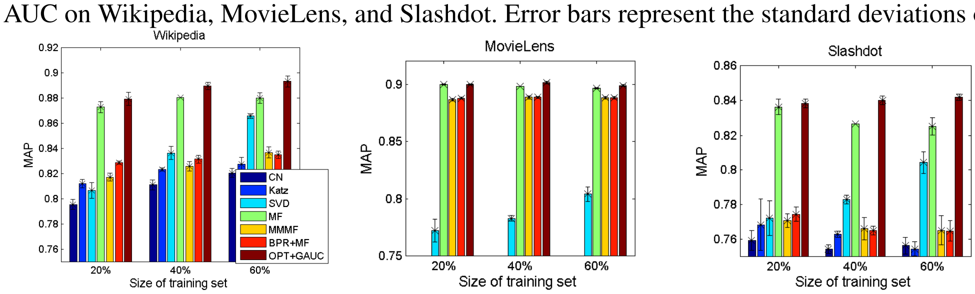

(a) MAP on Wikipedia. (b) MAP on MovieLens. (c) MAP on Slashdot

Figure 5: MAP on Wikipedia, MovieLens, and Slashdot. Error bars represent the standard deviations of uncertainty.

the parameter combination which achieves the best average GAUC is employed for test.

Results

Baslines: We compare the proposed approach (OPT+GAUC) with various baseline approaches to demonstrate its effectiveness. Among these baselines, common neighbor (CN) (Liben-Nowell and Kleinberg 2007) and Katz (Katz 1953) are obtained based upon the network topological structure; singular value decomposition (SVD) and matrix factorization (MF) (Koren, Bell, and Volinsky 2009) are two representative pointwise approaches for collaborative filtering; maximum margin matrix factorization (MMMF) (Weimer, Karatzoglou, and Smola 2008) and Bayesian personalized ranking based upon matrix factorization (BPR+MF) (Rendle et al. 2009) are two popular pairwise approaches for personalized ranking. For fair comparison, β in Katz is set as 0.005 because it achieves the best performance; we select the hyper-parameters of MF, MMMF, and BRP+MF in a similar way as we do for OPT+GAUC. To ensure our results are reliable, we conduct each experiment 5 times; the average GAUC/AUC/MAP/Recall@k/Precision@k and their associated standard deviations are reported for comparison.

Efficiency: We study the efficiency of the proposed approach by employing stochastic sub-gradient descent over 20% of the Wikipedia dataset. In particular, we partition X_{Train} into b batches and perform sub-gradient descent over U and V iteratively. Figure 2 shows the efficiency of the proposed approach with different numbers of batches ($b \in \{5, 10, 30, 60\}$). We observe that when b varies from 5 to 60, the objective function of our proposed approach converges faster. This is because the computational complexity of the sub-gradients linearly depends on b. Moreover, we observe that the GAUC/AUC/MAP does not decay very much while b is increasing.

Link recommendation: Figures 3, 4, and 5 show the GAUC/AUC/MAP and their associated standard deviations of various approaches on three datasets when the size of training set varies from 20% to 60%. We observe that CN is generally outperformed by other approaches because it only considers the neighborhood structure of the network; since Katz not only encodes neighborhood structure but also considers the high-order relationships, it can outperform CN in most cases. Note that CN and Katz cannot work on MovieLens because it is a bipartite network. For GAUC in Fig-

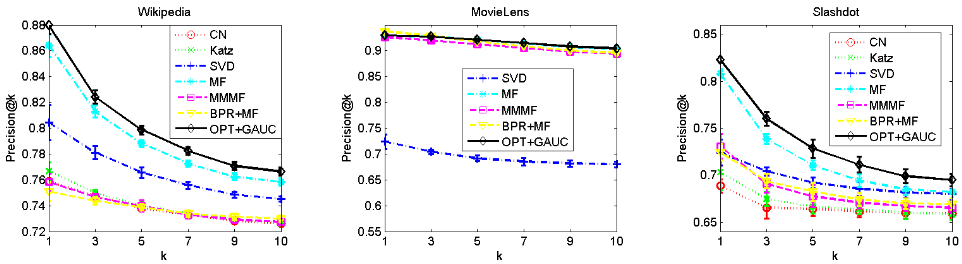

(a) Precision@k on Wikipedia. (b) Precision@k on MovieLens. (c) Precision@k on Slashdot

Figure 6: Precision@k for Wikipedia, MovieLens, and Slashdot. Error bars represent the standard deviations of uncertainty.

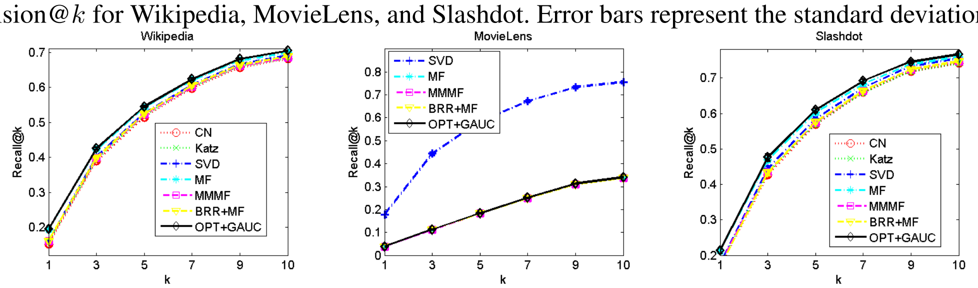

(a) Recall@k on Wikipedia. (b) Recall@k on MovieLens. (c) Recall@k on Slashdot

Figure 7: Recall@k for Wikipedia, MovieLens, and Slashdot. Error bars represent the standard deviations of uncertainty.

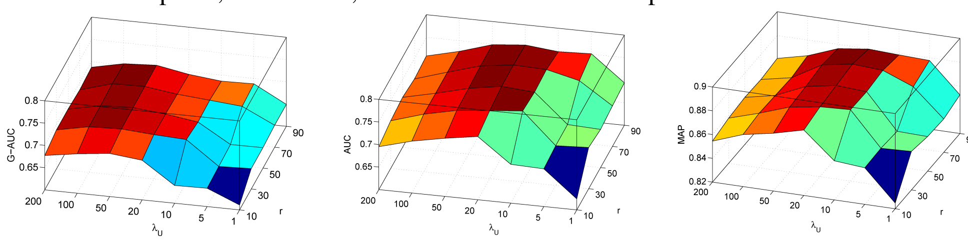

(a) GAUC on 20%Wikipedia. (b) AUC on 20%Wikipedia. (c) MAP on 20%Wikipedia.

Figure 8: Parameter sensitivity study for Wikipedia (20% training) with respect to GAUC, AUC, and MAP.

ure 3, we observe that, pairwise approaches, i.e., MMMF and BPR+MF, outperform pointwise approaches (i.e., SVD and MF) in most cases. This may be because SVD tends to over-fit the data and MF only reconstructs the partially observed network based upon observed positive as well as negative links (it neglects unknown status links). For AUC and MAP, SVD cannot perform as well as MF because SVD tends to over-fit the data (especially for smaller datasets); MMMF and BPR+MF are outperformed by MF as they do not directly model negative links. On the contrary, they treat both negative and unknown status links as 0. OPT+GAUC generally outperforms all baseline algorithms regarding GAUC/AUC/MAP, this is because OPT+GAUC models the triplet in signed networks more reasonably than other approaches.

Top-k link recommendation: We further investigate the effectiveness of the proposed approach by comparing its Precision@k and Recall@k with baseline methods (the size of training set is 40% for Wikipedia, MovieLens, and Slashdot) in Figure 6 and 7, respectively. We observe that while OPT+GAUC is among the best approaches for Recall@k (except in Figure 7(b)), it consistently achieves the best Precision@k in Figure 6(a) and 6(c). This indicates that modeling unknown status links helps to increase the performance of top-k link recommendation. We also observe that while SVD achieves the best Recall@k in Figure 7(b), it gets the worst Precision@k in Figure 6(b). This may be due to over-fitting. Note that all the other approaches except SVD achieve similar performance in Figure 6(b) and 7(b), this may be because MovieLens is a relatively dense graph and it contains sufficient observed links to estimate the ranking of potential links.

Parameter sensitivity: We study the sensitivity of OPT+GAUC with respect to the regularization parameters $\lambda_U \in \{1, 5, 10, 20, 50, 100, 200\}$ and $r \in \{10, 30, 50, 70, 90\}$. When we vary the value of λ_U or rank r, we keep the other parameters fixed. We plot the GAUC/AUC/MAP with respect to λ_U and r in Figure 8. We observe that OPT+GAUC is very stable as it achieves good GAUC/AUC/MAP when λ_U varies from 10 to 100 and k varies from 30 to 90. The results of other settings are similar and are omitted here due to space limitations.

Conclusions

In this paper, we proposed a generalized AUC (GAUC) to quantify the ranking performance in signed networks. Based upon the GAUC loss, we derived a link recommendation model by directly minimizing this loss and introduced an optimization procedure. We conducted experimental studies based upon three real world networks, i.e., Wikipedia, MovieLens, and Slashdot. Our experiment results demonstrated the effectiveness and efficiency of the proposed approach.

Acknowledgement

This work was partially supported by the U.S. Department of Defenses Minerva Research Initiative under ARO grants W911NF-09-1-0081 and W911NF-12-1-0389.

References

Adamic, L. A., and Adar, E. 2003. Friends and neighbors on the web. *Social Networks* 25(3):211 – 230.

Brzozowski, M. J.; Hogg, T.; and Szabo, G. 2008. Friends and foes: ideological social networking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 817–820.

Burke, M., and Kraul, R. 2008. Mopping up: Modeling wikipedia promotion decisions. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, 27–36.

Clauset, A.; Moore, C.; and Newman, M. E. J. 2008. Hierarchical structure and the prediction of missing links in networks. *Nature* 453:98–101.

Cortes, C., and Mohri, M. 2004. Auc optimization vs. error rate minimization. In *Proceedings of Advances in Neural Information Processing Systems*, 537–546. MIT Press.

Dong, Y.; Tang, J.; Wu, S.; Tian, J.; Chawla, N. V.; Rao, J.; and Cao, H. 2012. Link prediction and recommendation across heterogeneous social networks. In *Proceedings of the IEEE 12th International Conference on Data Mining*, 181–190.

Doreian, P., and Mrvar, A. 2009. Partitioning signed social networks. *Social Networks* 1–11.

Guha, R. V.; Kumar, R.; Raghavan, P.; and Tomkin, A. 2004. Propogation of trust and distrust. In *Proceedings of the 13th International Conference on World Wide Web*, 403–412.

Hanley, J. A., and Mcneil, B. J. 1982. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*.

Heider, F. 1946. Atitudes and cognitive organization. *Journal of Psychology* 21:107–112.

Katz, L. 1953. A new status index dervied from sociometric analysis. *Psychometrika* 18(1):39–43.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *IEEE Computer* 42(8):30–37.

Kunegis, J.; Lommatzsch, A.; and Bauckhage, C. 2009. The slashdot zoo: Mining a social network with negative edges. In *Proceedings of the 19th International Conference on World Wide Web*, 541–570.

Lampe, C.; Johnston, E.; and Resnick, R. 2007. Follow the reader: filtering comments on slashdot. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1253–1262.

Lee, D., and Seung, H. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401:788–799.

Leskovec, J.; Huttenlocher, D. P.; and Kleinberg, J. M. 2010a. Predicting positive and negative links in online social networks. In *Proceedings of the 20th International Conference on World Wide Web*, 641–650.

Leskovec, J.; Huttenlocher, D. P.; and Kleinberg, J. M. 2010b. Signed networks in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1361–1370.

Li, Y. 2009. A generalization of auc to an ordered multiclass diagnosis and application to longitudinal data analysis on intellectual outcome in pediatric brain-tumor patients. *Ph.D Thesis of Geogia State University*.

Liben-Nowell, D., and Kleinberg, J. 2007. The link prediction for social networks. *Journal of American Society for Information Science and Technology* 58(7):1019–1031.

Menon, A. K., and Elkan, C. 2011. Link prediction via matrix factorization. In *In Proceedings of the European Conference on Machine Learning, Part II, LNAI2912*, 437–452.

Nakas, C. T., and Yiannoutsos, C. T. 2004. Ordered multiple-class roc analysis with continuous measurements. *Statistics in Medicine* 23:3437–3449.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 452–461.

Salakhutdinov, R., and Mnih, A. 2007. Probabilistic matrix factorization. In *Proceedings of the Neural Information Processing Systems*.

Song, D., and Meyer, D. A. 2014. A model of consistent node types in signed directed social networks. In *Proceedings IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 72–80.

Song, D.; Meyer, D. A.; and Min, M. R. 2014. Fast nonnegative matrix factorization with rank-one admm. In *NIPS 2014 Workshop on Optimization for Machine Learning (OPT2014)*.

Weimer, M.; Karatzoglou, A.; and Smola, A. 2008. Improving maximum margin matrix factorization. *Machine Leanring* 72(3):263–276.

Yang, S. H.; Smola, A. J.; Long, B.; Zha, H.; and Chang, Y. 2012. Friend or frenemy?: predicting signed ties in social networks. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*.

Ye, J.; Cheng, H.; Zhu, Z.; and Chen, M. 2013. Predicting positive and negative links in signed social networks by transfer learning. In *Proceedings of the 23th International Conference on World Wide Web*.