Case Studies of Signed Networks

Christopher Wang

December 10, 2014

Abstract

Many studies on signed social networks focus on predicting the different relationships between users. However this prediction usually separate the links into positive or negative. Neutral links are either ignored or considered as part of negative links. This project is focused on the idea that if neutral links can provide information which in turn enhances the prediction of other links.

1 Introduction

There has been social networks online involving both positive and negative interactions. Those networks often include mechanisms for users to evaluate one another, or of content they create. These evaluations can be binary: Positive links can be formed if one user trusts another user, supports another user's opinions, and so on, while negative links suggest distrust, or disapproval, such as Wikipedia, and Stack Overflow [1]. Other networks may support rating systems that allow users to review products or other user's content by different scores. One of the examples is Epinions [3].

One can imagine that such networks can behave very differently from those with only plain interactions with no "emotions". A number of papers have begun to investigate some fundamental questions about signed networks. One of the fundamental questions to ask is that what are the properties or features of the networks (either local or not) that have huge impact on the sign of the given link. This question can help us to understand some of the underlying principles. Answers to this questions also lead us to proper solutions of a common task in online communities: to suggest or recommend new relationships to a given user. A substantial amount of on-going research is along this line of direction. In particular, researchers are focused on how to predict the sign of an unkown link in signed networks.

However, there are also a lot of networks allowing people to express "feelings" more than just positive or negative. For example, online rating systems like Epinions, Amazon/ebay rating systems, etc. The possible ratings are usually more than binary, people can go extreme, and also can be neutral. Previous research often treat the neutral links as the same role as the negative links. In this project, the author tries to address the problem that how the neutral links can be used together with the positive and negative relationships in the signed network. It is not hard to imagine that, when we attempt to suggest new relationship to a user or recommend a new commodity, to isolate those neutral relationships would be of huge help.

2 Previous Works

Social networks with both positive and negative interactions can be found in different applications, such as E-marketplaces, opinions and activity sharing sites, social sites, etc. The corresponding network can be expressed as *signed network*, where the sign of a edge being positive or negative depending on the attitude of the link generator to the recipient. There has been extensive research focusing on exploring how to predict the sign of an edge in such social networks using different metrics [1, 4].

2.1 Predicting the Edges by the Similarity Between Users

The paper [1] studies how the similarity in the characteristics of two users is able to significantly strengthen the analysis of user-to-user evaluations. It shows that the probability of a positive evaluation ($P(+)$) grows as the similarity between the evaluator and the target increases. Next the paper study the interaction between similarity and status, and discover that the evaluations are less status-driven as the similarity between users increase. Additionally, it studies the relationship between the evaluators similarity and status for each dataset, which provides an insight on evaluators.

Previous studies assume the relationship of the P(+) as a function of status difference to be monotonically decreasing. However [5] has shown otherwise, dip curve, which suggests that users are tough evaluator when they have approximately the same status. This phenomenon is the opposite of what was found in the previous paper, that users who are similar tends to be positive toward each other. But the paper shows with insight to user that evaluate, it can model the mechanism behind the dip phenomenon.

Finally the paper introduce **ballot-blind prediction**, an amazing method that predict the outcome with the identities of the first k evaluator, but without knowing their votes. This method shows without knowing the actual votes, it is possible to predict the outcome with small sampling. Also the performance gradually increases if the actual votes are known.

2.2 Edge Prediction by Logistic Regression

In the paper [4], the authors consider the sign prediction of an edge using logistic regression classifier, taking the form

$$P(+|x) = \frac{1}{1 + e^{-(b_0 + \sum_i^n b_i x_i)}},$$

where x_i's are features, and b_i's are the coefficients to estimate based on the data.

The features that are used in the model have two classes. The first class only contains the positive/negative in/out degrees of the two end points of the edge we want to predict. The second class considers the *triad* involving the edge to be predicted, consisting of edges that are neighbors of either endpoint.

The datasets with signed edges are often overwhelmingly of positive edges. To ensure that the random guess yields 50% correctness, the methodology of Guha et al. is used to create a balanced dataset for prediction. The logistic regression shows an amazing result of predicting the edges.

While this paper operates the prediction in the local level, rather than the global one, it is seen that the balance is operating more strongly at local level, since it was able to find an approximate global status ordering but no not a global organization of these networks into opposing fractions.

In addition, the paper also studied the problem that how the negative links change the edge sign prediction result. To elaborate it, the authors in [4] compares the results of predicting the positive edge using two different feature sets. One is only derived from the positive edges of the given network, while the other used the information of negative links. As a result, there is a significant improvement by using negative links.

3 Dataset Description

Epinions is an online reviewing site where users can write reviews about products and assign them a numeric rating from 1 to 5. The dataset has 132K users, 1.5M reviews, and 13.6M ratings of those reviews. There is no doubt that 5-star rating represents positive and 1-star rating represent negative. I

Table 1: Rating Distribution of Epinion Data

Rating	Ratio
5	76.95%
4	15.10%
3	5.61%
2	2.32%
1	.017%

will figure out what is the reasonable range for neutral ratings depend on experiments. The major concern is that the data is overwhelmed by rating "5". Rating Distribution However, according to the above overall fractions of rating that users on Epinion gives to review articles. As mentioned in [1], they categorized 5-star rating as "positive" evaluation and the rest as negative. This is one of the difficulties that I have ran into when trying to figure out if netrual evaluation can have a role in predicting evaluation. The original idea was to categorized 5-star as positive, 1-star as negative, and the rest as neutral. As shown on the table above there are hardly any 1-star rating, more than 92% of the ratings are either 4-star or 5-star. It is obvious that on Epinion people almost never gives a bad review. Around 43.6% of the users have only given 5-star to others' reviews and around 99.5% has never given 1-star to any reviews. However an interesting result is that around 17% of the users has never given a 5-star rating. If we can figured out these users' property, it might be able to improve predicting "positive" evaluation.

The dataset that I used are given in the following format: Right now only the first three column was used in my research. ReviewId is the Id of the review, userId is the user that is evaluating the review, and rating is the score the user evaluated. The author of each review can be found with the following dataset:

Table 3: Data Format of Review Author

ReviewId	UserId	SubjectId
1445594	718357	149002425217
1445595	220568	149003604865
1445596	717325	5303145344

4 Problem Formulation and Learning Methods

Notice that there are substantial online rating systems that allow users to evaluate certain products or other user's articles. In such networks, not only positive and negative feedbacks are given. They often comes with a rating score from 1 to 5, such as Epinions, and Amazon product review, etc. For those networks, there exists "neutral" evaluations.

The problem we will be solving in this project is predicting neutral evaluation of a given edge in social network. Earlier works [1] have studied what might affect the types of evaluation that one user gives to another, whether it is relative status between two users or similarity in the characteristics of two users. In most of the analysis neutral votes are either ignored or grouped into negative votes. Is there a pattern in which the evaluator gives a neutral vote to the target?

If there is a pattern, by removing the neutral votes the network would be left with rating 1 and 5, namely, the "extreme" attitudes. If one can further predict rating 5 we might be able to build a recommendation system that is highly likeable by the user. What kind of features can be useful in predicting neutral votes?

4.1 Predicting Edge Class

The network contains two types of nodes, user and review article. Each article contains only in-degree from other users. The goal is to predict the rating a user gives to a review based on various features. Different from [4], the edge here is from user to item,

Table 2: Data Format of User Ratings

ReviewId	UserId	Rating	Status	Date	Type	VerticalId
139431556	591156	5	0	2001/01/10	1	2518365
139431556	204358	5	0	2001/01/10	1	2518365
139431556	368725	5	0	2001/01/10	1	2518365

not from user to user. By connecting each review to the author we can build a reputation system for each user, summing up the ratings of the reviews the user had written. Obviously, past reputations are useful in predicting the future reviews written by the user.

4.1.1 Features

The features I used are separated into two classes. The first class is based on the degrees of users and reviews. As we are interested in predicting the sign of the link from a user u to a review r, the features we considered are essentially the out-degree of the user u and in-degree of the review r. In more detail, they are divided into three categories: the incoming edges (all the ratings) to the review r, denoted by $d^p_{in}(r)$, $d^e_{in}(r)$, and $d^n_{in}(r)$; The incoming edges to each user (all the ratings to the reviews the user u wrote). We can represent them as $\sum_{r \in A(u)} d^p_{in}(r)$, $\sum_{r \in A(u)} d^n_{in}(r)$ and $\sum_{r \in A(u)} d^n_{in}(r)$, where $A(u)$ denotes the set of reviews that written by the user u; The outgoing edges from each user (all the reviews that the user u gave), denoted by $d^p_{out}(u)$, $d^n_{out}(u)$ and $d^n_{out}(u)$. Each rating edge is grouped as the following: rating $\{1\}$ (negative: n), rating $\{2, 3, 4\}$ (neutral: e), and rating $\{5\}$ (positive: p).

The second class is the similarity and status different between users. On Epinion, reviews are written based on products. These can be used to provide a description of the user's interest by considering the product's reviews a user had evaluated. The user's interest is represented as a list of products the reviews were based on. The similarity between users u and v is then the cosine similarity between their list of products: $s(u, v) = \frac{u_a v_a}{|u_a||v_a|}$. In case users u or v 's lists are empty, the score was smoothed by adding 1 to the equation $s(u, v) = \frac{u_a v_a + 1}{(|u_a|+1)(|v_a|+1)}$. As mention in [1]. The higher the similarity between users u and v, the probability of u given v positive votes is higher. We also consider the status difference between users, which is the reputation of u subtracting reputation of v.

4.1.2 Learning Methodologies

- **Support Vector Machine (SVM)**:

  The SVM model separate the classes by choosing a hyperplane. There are many hyperplanes that can separate two classes. SVM consider the one as the best hyperplane if it has the largest margin that separated two classes. Assuming the data is linearly seperatable the goal is to have the maximum margin between two hyperplane,

  $$wx - b = 1 \ and \ wx - b = -1$$

  From geometry we know that the distance between the two plane is $\frac{2}{||w||}$.

  I used weka package to build the SVM model. Since the dataset is enormous, the model was first built on a small training set and updated by providing the rest of the data.

- **Multinomial Model to Predict 3 Classes of the Edge**:

  We can also use a multinomial model to build a classifier to predict 3 classes of the edge sign: positive, negative and neutral. Here the response variable has 3 levels $\mathcal{S} = \{0, 1, 2\}$. Multinomial regression learns a model of the form

$$\Pr(S = s|X = x) = \frac{e^{\beta_{0s}+\beta_s^T x}}{\sum_{l=0}^{2} e^{\beta_{0l}+\beta_l^T x}},$$

where *beta* is a $p \times 3$ matrix of coefficients, p is the number of features that we are going to use. β_s referes to the s-th column for outcome category s.

I used R package *glmnet* to build the multinomial model. The lasso penalties and partial Newton algorithm are used [2].

4.2 Model Training and Validation

We used 10-fold method. After the feature extraction, the data is equally divided into 10 sets. At each time, we leave one set out as the test data, and use the rest of 9 sets to train the model.

5 Results and Discussions

5.1 Effects of Considering the Neutral Class

Table 4: Grouping Comparison

Grouping	Results
{1}{2,3,4}{5}	81.9269%
{1,2,3,4}{5}	81.8658%
{1}{2}{3}{4}{5}	81.6259%

We first investigate the role of neutral links here as well. In particular we consider the question of whether information about neutral links can be helpful in predicting the sign of a given link. In other words, how useful is it to know what are the items or reviews that a user neither like nor dislike very much, if we want to predict the presence of additional items or reviews that a user strongly likes?

In order to see the effect of considering neutral edges, we use the machine learning framework developed in previous sections to build classifiers that predict whether there exists a positive edge from a user to a review or not. Here we only consider the degree features, three types of grouping are implemented when predicting the positive edges for simplicity. Grouping 1, where neutral edge is by itself a feature, $\{1\}\{2, 3, 4\}\{5\}$, grouping 2, where neutral edge is consider as negative, $\{1, 2, 3, 4\}\{5\}$, and grouping 3 where each rating is a feature $\{1\}\{2\}\{3\}\{4\}\{5\}$. As shown in table 4, grouping into three gives slightly better result the reason might be that grouping into two creates an underfitting model where as the grouping into five intends to overfit. These results clearly demonstrate that there is an improvement to be gained by using information about neutral edges, even to predict the edge being positive or not.

5.2 Accuracy Results of predictions

5.2.1 Predicting the edge being positive or non-positive

Using this result I tried to improve the prediction by either adding similarity score, status difference or standardlization of the degree features. The standardlization means that instead of using the exact positive, negative, and neutral in/out degrees, the fraction of links being positive, negative or neutral out of the total in/out degree is used. The accuracy results are summarized in Table 5.

In my experiment I found that when the degree features are not standardlized, including similarity and status difference gives a slightly better results. In general including similarity produces a better results than including status difference. This validates the results in [1].

When the degree features are standardlized the results are increased by a huge amount. However the impacts of the similarity and status difference can be hardly noticed. The standardlized in-degree of the review articles play a much more significant role in the prediction prediction, compared to similarity and

Table 5: Accuracy of predicting positive $\{p\}$ vs. neutral and negative $\{n, e\}$ edges on Epinion data

$d_{in}(r)$	similarity	status difference	standardlized	accuracy
✓	×	×	×	81.9269%
✓	✓	×	×	82.5812%
✓	×	✓	×	82.3743%
✓	✓	✓	×	82.7144%
×	×	×	✓	84.3129%
×	✓	×	✓	84.3400%
×	×	✓	✓	84.3450%
×	✓	✓	✓	84.3593%
✓	×	×	✓	86.5184%
✓	✓	×	✓	86.5307%
✓	×	✓	✓	86.5844%
✓	✓	✓	✓	86.5837%

Table 6: Accuracy of predicting neutral $\{e\}$ vs. neutral and extreme $\{n, p\}$ edges on Epinion data

$d_{in}(r)$	similarity	status difference	standardlized	accuracy
✓	×	×	×	81.9312%
✓	✓	×	×	82.5922%
✓	×	✓	×	82.3837%
✓	✓	✓	×	82.7245%
×	×	×	✓	84.3061%
×	✓	×	✓	84.3336%
×	×	✓	✓	84.3381%
×	✓	✓	✓	84.3528%
✓	×	×	✓	86.5068%
✓	✓	×	✓	86.5201%
✓	×	✓	✓	86.5691%
✓	✓	✓	✓	86.5882%

status difference. This is reasonable since the standardlization of the in-degree of the review articles is a rough distribution estimate of the three classes.

Notice that some of the classification models were built excluding the in-degree of the review articles. Although the performance of the models without the in-degree of the articles ($\approx 84\%$) is better than that of the models without the standardlization ($\approx 82\%$), suprisingly the effect of similarity and status difference is minimal, compared to the degree features. This shows that the in/out degree of the review article has a more significant impact than similarity and status difference.

5.2.2 Predicting the edge being neutral or not neutral

Previous tests were focused only on predicting positive edges. As mention earlier if we can predict neutral edges, we might be able to create a better recommendation system. The performance is summarized in Table 6. Using the same features the model can predict neutral edges very well, the result is about the same as that of predicting positive edges. If one is interested in neutral edges the model can perform fairly well using the same set of features.

5.2.3 Multinomial Prediction Results

The multinomial prediction yields almost the same performance as SVM. Table 7 lists the accuracies of two models using all the derived features.

Table 7: Camparison of SVM and Multinomial Regression

all features	SVM	Multinomial Regression
	86.5802%	86.5906%

6 Conclusion

We have investigated an underlying properties that determines the sign of a given link in social networks that have positive, negative and neutral links. We have shown that neutral links play a very important role in learning the social networks. A rating system that is more than binary provides a slightly advantages in terms of edge prediction. In addition, we have provided a set of features that can predict the link sign fairly well. Evidence has shown that the quality of the review article is much more important than the social similarities of the authors or the reputation difference between the author and the user to rate the article.

References

[1] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. Effects of user similarity in social media. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 703–712. ACM, 2012.

[2] Jerome Friedman, Trevor Hastie, and Rob Tibshirani. glmnet: Lasso and elastic-net regularized generalized linear models. *R package version*, 1, 2009.

[3] Ramanthan Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM, 2004.

[4] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 641–650. ACM, 2010.

[5] Jure Leskovec, Daniel P Huttenlocher, and Jon M Kleinberg. Governance in social media:

A case study of the wikipedia promotion process.
In *ICWSM*, 2010.