

Two Models for Inferring Network Structure from Cascades

Dakan Wang and Yu Wu
Department of Computer Science
Stanford University
Stanford, CA 94305, USA
(vondrak, ywu2}@stanford.edu

Yu Zhang
Department of Computer Science
Trinity University
San Antonio, TX 77812, USA
yzhang@trinity.edu

ABSTRACT

In many real-world scenarios, the underlying network over which the diffusions and propagations spread is unobserved, i.e. the edges of the network are invisible. In such cases, we can only infer the network structure from underlying observations. The goal of this paper is to find a model that generates realistic cascades with observed data, so that it can help us with link prediction and outlier detection. For this purpose, we investigate two cascade models. The first model is a naive two-class cascades that includes one class of positive (infected) nodes and one class of negative (uninfected) nodes. In this model, we use the sparse logistic regression method to infer network edges. In the second model, we discard all negative training nodes and treat the whole network as a single class. In this model, we use the one-class Support Vector Machines to predict underlying edges. Experiments show that even if we discarded all negative training instances, we can still infer network edges accurately.

KEYWORDS

Social Network Structure, Information Cascades, Networks of Diffusion

1 INTRODUCTION

Processes that form cascades in a social network have been studied in a number of domains [12]. Examples include the spread of infectious diseases [6], the spread of new ideas and technologies [14], marketing [5], technology transfers [3, 4], computer virus transmission [2], and power systems [17]. Edges are very fundamental parts of a network to investigate the network structure. However, in many real world situations, the edges of a network are invisible and we can observe partial information about the edges. This information can help us infer the latent edges in the network. As a fundamental process in social network, cascade is a useful information source to infer edges. In general, a cascade describes how a contagion spreads or propagates across a network [16]. Research on cascade will enable us to gain insights about how people influence each other. However, sometimes it is very difficult to track a

cascade because of its complexity. Even worse, we may not know the structure of the network over which the cascade happens. In most cases, the available information is very limited. We may only know when the cascade propagates to a given node. For instance, in study of virus contagion, tracking the time when a person was infected may be much easier than tracking the source of the infection. In viral marketing, we can easily record the time when a customer made a purchase, but it is hard to know who had influenced him to make the decision. Therefore, the challenging task is to infer the latent network structure from the limited given information. Specially, we may only know the nodes, and sets of time data from a few cascades, but not the edges in the network. Our goal is to infer the hidden edges based on these limited information.

There are a few recent works focusing on inferring networks from cascades. Leskovec et al. [10] monitor one of the largest available collections of blog information and find that almost any metric they examined, such as size of cascades, size of blogs, and in- and out-degrees, follows a power law distribution.

Gomez-Rodriguez, Leskovec and Krause [11] convert the edge inference problem to combinatorial optimization problem [7]. They find that choosing the best set of k edges maximizing the likelihood of the data is NP-hard. They then develop NETINF, a greedy algorithm for inferring a near-optimal set of k directed edges. NETINF is able to scale to large real data sets by introducing some speed-up heuristics based on the submodularity of their evaluation function.

Myers and Leskovec [9] formulate the edge inference problem to an equivalent convex problem. This guarantees the optimality of the solution. In addition, they observe that social networks are sparse in a sense that on average nodes are connected to a constant number rather than a constant fraction of other nodes in the network. To enforce a sparse solution, they add the L_1 penalty regulation to the evaluation function.

Sadikov, Medina, Leskovec, and Garciamolina [13] address the problem of missing data in information cascades [8]. They propose a k -tree model of cascade (k is the number of parents for a node) to predict the missing data such as the size or depth of the tree. By given a fraction of a complete cascade C , they present a selection algorithm to select a *proxy* k -tree that best approximates C .

Yang and Leskovec [15] develop a Linear Influence Model to infer the influence of a node from cascades. Different from predicting which node will influence which other nodes in the network, they focus on modeling the global influence of a node on the rate of diffusion through the (implicit) network. For each node, they estimate an influence function of how many subsequent infections can be attributed to the influence of that node over time. Therefore the number of newly infected nodes is a function of which other nodes got infected in the past.

In this paper, we model the edge inference problem to the traditional classification problem. We investigate two models for inferring edges from cascades. The first model is a two-class cascades that includes one class of positive (infected) nodes and one class of negative (uninfected) nodes. Here, we use the sparse logistic regression method to infer network edges. The second model discards all negative training nodes and treats the whole network as a single class. We use the one-class Support Vector Machines to predict underlying edges. Experiments show that the second model (the single class model) can still infer network edges accurately.

The paper is structured as follows: Section 2 formulates our problem and defines our assumptions. Section 3 presents the two models in theoretic approach. Section 4 presents and discusses the analysis results. Section 5 concludes the paper.

2 PROBLEM FORMULATION

Given a graph of n nodes $G = (V, E)$. The edge set E is unknown. Also given a set of cascades, each cascade is represented by a vector $\{\tau_1, \dots, \tau_n\}$ where τ_i indicates that node i is infected at time τ_i . If $\tau_i = \infty$, node i is not infected in the cascade. The problem is to infer the edges E from these cascades. Our idea is based on two assumptions: 1) earlier infected nodes may have edges to nodes infected later, and 2) infected nodes are very unlikely to have edges to uninfected nodes.

3 MODELS

3.1 Model A

Before proposing the model, let's present our assumptions first. [9] assumes the time period each

infected node takes to transmit the disease follows a random distribution. However, we notice there is some inconsistency between the model assumption and the objective function. Inspired by the maximum likelihood estimation in [9], we propose another assumption that an uninfected person i is not always susceptible to infection, and will only be infected when he shows up in the network. For instance, for cascades in blog sphere, one will not be influenced by other persons' blog-posts if he constantly stays off-line. The already infected person, say j , will decide whether or not to infect person i with probability A_{ij} . If j chooses to infect i , he will wait until person i shows up and then try to infect him. The probability of i being infected when he shows up can be defined in various ways to finish model assumption.

There are two points about our assumption:

- 1) Before getting infected at time τ_i , i can also show up at $\tau_i' < \tau_i$ (e.g., access the Internet in the blog case). It is not clear whether he was infected at that time. Either infected or not can be possible. It may be the case that he was infected, but we did not observe his infection at that time. Therefore to avoid ambiguity, we do not take a person's history before his infection into account.
- 2) For node i , which is never infected, we assume that he will not get infected whenever he shows up. To satisfy this, other infected node must not choose to infect node i , otherwise there will be a probability of i being infected when shows up. Hence the probability of i never infected is

$$\prod_{j \text{ infected}} (1 - A_{ij}) \quad (1)$$

Moreover, instead of assuming infected people transmit the virus to a susceptible one independently, we assume they impose a collective influence. This assumption is very common in real world. For instance, when one wants to purchase a product online, he may first check its online reviews and then make the decision. In other words, the other people who have already bought the product influence the person collectively through reviews. We still maintain the assumption that a susceptible person will only get infected when he shows up. Different from the above model, we do not keep the restriction that A_{ij} ranges between 0 and 1, but assume it can take any real value. Large positive value for A_{ij} means that the person j is very likely to infect person i , while negative value indicates that the person j may prevent the person i from being infected (can be interpreted as i distrusts j very much). When i shows up at time τ_i , the influence of j on i is $w(\tau_i - \tau_j)A_{ij}$, where $w(t)$ is the distribution defined in [9]. The collective influence on person i at τ_i is

$$\sum_{\tau_j < \tau_i} (w(\tau_i - \tau_j) A_{ij}) \quad (2)$$

Finally, we assume that the collective influence on person i who has never been infected is

$$\prod_{j \text{ infected}} A_{ij} \quad (3)$$

as explained in the above section.

We now try to understand the above definition in another way. Each node i is associated with a set of cascades. In some cascades, node i was infected at τ_i , while in others i never got infected. For a cascade in which an node i infected at τ_i , the vector contains non-zero entry $w(\tau_i - \tau_j)$ for every j such that $\tau_j < \tau_i$. For a cascade in which node i never infected, the vector contains entry 1 for every already infected node. The i^{th} row of A , denoted A_i , is the weight vector characterizing how much node i trust other nodes. The probability of node i infected at τ_i or never infected is a function of the dot product (\bullet) between feature and weight vector. Then we reduce the cascade problem to a classification problem: we have some positive and negative instances, which are respectively different cascades in which node i got infected or not. For each instance we have a feature vector. And we would like to infer the weight vector from the provided training data. Then we may adopt some efficient discriminate model such as logistic regression, SVM (Support Vector Machines). If we want to guarantee sparsity of $|A|$, we may use sparse logistic regression. One possible advantage of this model over [9] is that we do not have an inconvenient product of terms like

$$1 - \prod_{\tau_j < \tau_i} (1 - w(\tau_i - \tau_j) A_{ij}) \quad (4)$$

To summarize, we now mathematically define the above idea. For a certain node i , we have a set of n_i instances $\{x^{(j)}, y^{(j)}\}_{j=1}^{n_i}$. Here $y^{(j)} \in \{0, 1\}$ represents whether node i is infected in cascade j . And $x^{(j)}$ is a vector $(x_0^j, \dots, x_k^j, \dots, x_n^j)$ with entry x_k^j to be

$$x_k^j = \begin{cases} w(\tau_i^j - \tau_k^j) & \tau_k^j < \tau_i^j \\ 0 & \text{else} \end{cases} \quad (5)$$

for $y_i = 1$, and

$$x_k^j = \begin{cases} 1 & j \text{ is infected} \\ 0 & \text{else} \end{cases} \quad (6)$$

for $y_i = 0$. We would like to learn from the training set our transmission weight $A_i = (A_{i1}, \dots, A_{in})$ and a hypothesis $h_{A_i}: x \rightarrow y$. We use logistic function to

$$h_A(x) = \frac{1}{1 + e^{-(A_i)^T x}} \quad (7)$$

And assume that

$$\begin{aligned} p(y = 1 / x, A_i) &= h_A(x) \\ p(y = 0 / x, A_i) &= 1 - h_A(x) \end{aligned} \quad (8)$$

The objective function we would like to maximize is

$$F = \sum_{i=1}^{n_i} y^i \log(h_A(x^i)) + (1 - y^i) \log(1 - h_A(x^i)) + \gamma \|A_i\|_1 \quad (9)$$

Finally we use the large scale sparse logistic regression package [1] to solve the above objective function.

3.2 Model B

In the above section, we cast the network inference problem to be a binary classification problem and treat the cascades containing node i to be positive instances and those not containing i to be negative instances. We implicitly assume that the cascades not containing i indicate that there are no edges between i and previously infected edges. However, we claim that second assumption is presumptuous. There are lots of examples against this assumption. We take hashtag propagation as an example. Here each hashtag is a cascade. There are many cascades that one user does not adopt a hashtag and some of the users he follows adopt that hashtag. These cascades are not evidences that there are no edges between node i and already infected nodes. Thus these negative instances are not truly negative. Treating them to be negative maybe harmful for inferring edges. On the other side, those positive instances (cascades containing node i) do indicate that there may be some edges between node i and some previously infected nodes. Those we are facing a problem of only having confident positive instances. It is lucky that there have been many works dealing with problems which only contain only positive instances. Such problems are called one-class classification problem. By definition, one class classification aims distinguish one class of objects from all other possible objects, by learning from a training set containing only the objects of that class. This is exactly the problem we want to solve. It also worth mentions that there are many provided easy-to-use software packages for one-class classification.

The other advantage of modeling network inference to be a one-class classification problem is that training the model will be more efficient. Indeed, in many situations, the proportion of positive instances in which node i is infected is very small. Thus our training data size will be significantly reduced if we only consider the positive instances.

The one-class classification problem can be formulated as follows:

$$\begin{aligned} \text{Min } & \frac{1}{2} |w|^2 + \frac{1}{vl} \sum_{i=1}^n \epsilon_i - \rho \\ \text{subject to } & w \cdot x_i \geq \rho - \epsilon_i \end{aligned} \quad (10)$$




Fig. 1: PR curves for Erdos Random Graph Different Transmission Time Configurations

4 EXPERIMENT

4.1 Datasets

As in [11, 9], there are two categories of datasets we are considering using: synthetic datasets and real world datasets. The synthetic dataset is mainly used for debugging programs. For this milestone, we only use synthetic datasets. The dataset we used is similar to the one in [9]. That is, we generate an Erdos random graph. For each generated edge, we randomly generate the transmission weight A_{ij} . Then we try to generate a set of cascades. For each cascade, we randomly generate a start node and take it as the newly infected node. Then we iteratively simulate the cascading process. Pick the node i with the earliest infection time from the set of newly infected node. For the picked node, decide whether to infect the other uninfected node based on the transmission weight A_{ij} . If node i choose to infect j , then randomly generate a time t_{ij} according to a specific distribution. If node j is not in the newly infected node set, add it to the set. If it is already in the set and t_{ij} is smaller than its infection time, update the infection time. Finally we remove node i from the newly infected node. In this way we can generate a cascade. As in [9], we also generate enough cascades for an informational training dataset.

4.2 Model A

We test whether a latent link can be inferred correctly. We compare the links inferred from the cascades with the links in real network and get the precision and recall of the model. For different choice of the sparsity parameter, precision is different. Thus we can plot a precision-recall curve. Note that we have connected the first point on the curve to (0; 1) and the last point to (1, 0).

4.2.1 Results on Synthetical Datasets

We tested our algorithm on synthetic datasets. As in [9], we generate an Erdos random graph consisting of

512 nodes and 1024 edges. In the graph, the transmission probability A_{ij} is sampled between 0.05 and 1.

For the time generation function $w(t)$, we also tried three probability distributions detailed as follows:

- Power law distribution $(\alpha-1)t^{-\alpha}$, where $\alpha = 9.5$.
- Exponential distribution $\frac{1}{\alpha} e^{-\frac{1}{\alpha}t}$, where $\alpha = 9.5$.
- Weillbull distribution $\frac{k}{\alpha} (\frac{x}{\alpha})^{k-1} e^{-(\frac{x}{\alpha})^k}$, where $\alpha = 9.5$ and $k = 2.3$.

Each cascade is generated as specified in 5.1. We generate enough cascades such that at least 95% edges are used to transmit a disease.

We show the experimental results in Figure 1. We plot three precision-recall curves for the three generated networks. Each point on the curve corresponds the model trained by setting the sparsity parameter to a certain value. Different points on the curve reflect models with different sparsity parameters. When sparsity parameter decreases, the recall increases and the precision increases.

Our model achieves very good performance comparable with, and even better than that obtained in [9]. However, it is quite weird that our model's performance is bad on the graph whose remission time is generated from a power law distribution. The reason is probably related to that our value of α is not the same as that set in [9]. We will try to find the reason in future work.

Apart from the Erdos random graph, we also test our model on another random graph, which is generated from the preferential attachment mechanism as taught in class. We report our results in Figure 2. As in Erdos random graph, we also tried three different probability density functions for sampling transmission time.




Fig. 2: PR curves for Preferential Attachment Graph of Different Transmission Time

Fig. 2: PR curves for Preferential Attachment Graph of Different Transmission Time Configurations

4.2.2 Results on Real World Datasets

Apart from synthetic dataset, we also tested our model on two real world datasets, specified as below:

- An email network in a research institution consisted of 437 nodes and 2805 edges.
- A collaboration network. We extract the largest component of the network. The number of nodes and edges are 379 and 608 respectively.

For the collaboration network, we sampled the edge transmission probability A_{ij} uniformly. For the email

network, we set $A_{ij} = 1 - (1 - \phi)(1 - \epsilon)^{m_{ij}}$, where $\phi = 0:05$, $\epsilon = 0:0001$ and m_{ij} is the number of emails from i to j . This parameter set was suggested in [9]. For each network, we also adopted three different probability density functions for transmission time as in synthetic datasets. The parameters for those probability density functions are:

- Power law distribution $\alpha = 2.5$.
- Exponential distribution $\mu = 9.5$.
- Weillbull distribution $\alpha = 3.6$ and $k = 2.5$.




Fig. 3: PR Curves for Email Network of Different Transmission Time Configurations




Fig. 4: PR Curves for Collaboration Network of Different Transmission Time

We reported our results in Figure 3 and Figure 4. Our model can also handle edge prediction for real world datasets. It seems that that our model performance is worse than that in [9]. However, we want to point out that our model has large improvement space because we do not use enough training data. Indeed, cascades we sampled only covered 90% edges, while those in [9] covered 99% edges. The reason that we did not sample as many cascades as in [9] is that in order to sample cascades to cover 95% edges, we need to sample much more cascades since for some uncovered edges, the transmission probability is very small. Since each cascade will be one additional instance for each node. Our laptop memory cannot hold that many instances. However, sampling more cascades for getting more training instances can improve the performance significantly. We tried sampling cascades to cover 85% edges and 90% edges and tested each setting's performance. And the PR curve in the later setting lifted by 15 points. Thus we expect our model to be much better than its current state if our laptop memory can hold enough training instances.

4.3 Model B

In this section, we conducted experiments to test our model B. The data were generated as in the above experiments. However, we discarded all negative instances and only kept those positive instances for each node. We used one-class SVM component in libSVM to infer the edge weight. The libSVM tool does not provide the functionality of learning a sparse weight vector. Therefore we at this stage cannot get a PR curve with respect to the sparsity parameter as in the experiments in the above section. Thus in this experiment we evaluated our experiment from another perspective: we want to check whether the weights inferred by our model satisfy that the weights of the edges are bigger than the weights of the non-edges. For instance, consider a graph with n nodes.

We now want to investigate what nodes have links to node s . Without loss of generality, we suppose that nodes $i_1^s, i_{21}^s, \dots, i_m^s$ have links to node s . Then we rank all the nodes by our inferred weight vector w_s . Suppose the top ranked m nodes are $j_1^s, j_2^s, \dots, j_m^s$. Then we count how many of our top ranked nodes have edges to node i , or we calculate the cardinality of the set $\{i_1^s, i_{21}^s, \dots, i_m^s\} \cap \{j_1^s, j_2^s, \dots, j_m^s\}$. Suppose that:

$$\begin{aligned} tt_s &= \# \{i_1^s, i_{21}^s, \dots, i_m^s\} \cap \{j_1^s, j_2^s, \dots, j_m^s\}, \\ tf_s &= \{i_1^s, i_{21}^s, \dots, i_m^s\} - \{j_1^s, j_2^s, \dots, j_m^s\}, \\ ft_s &= \{j_1^s, j_2^s, \dots, j_m^s\} - \{i_1^s, i_{21}^s, \dots, i_m^s\}. \end{aligned} \quad (11)$$

We use the following precision and recall to evaluate our model

$$\begin{aligned} P &= \frac{tt_s}{tt_s + ft_s} \\ R &= \frac{tt_s}{tt_s + tf_s} \end{aligned} \quad (12)$$

Notice that in this case, $tf_s = ft_s$, leading to $P = R$. We only use precision to evaluate our model. The goal of the experiments is to show that when we have observed enough cascades, the edges in the network can also be inferred even if we discard all the negative instances. We investigated how the performance of our model change with respect to edge cover ratio. Here edge cover ratio means the proportion of edges covered by the generated cascades. The larger the cover ratio, the more cascades we generated. We show in Figure 5 and 6 that the precision-cover ratio curve for two networks for different probability density function of transmission time. From the Figures, we can see that if the edge cover ratio is high enough, or number of cascades is large, then we may get satisfactory performance even if we do not consider negative instances.




Fig. 5: Precision-Cover Ratio of Preferential Attachment Network Configurations




Fig. 6: Precision-Cover Ratio of Email Network Configurations

5 CONCLUSION

We investigated two models for inferring network edges from cascades. The first model is based on sparse logistic regression. It is easy to understand and both efficient and effective. Our second model is based on one-class SVM. Experiments showed that even if we discarded all negative training instances, we can also infer network edges accurately.

In the future, we will test our model on larger scale data. One of the interesting data is the twitter dataset. Indeed, we have explored the twitter dataset. We once wanted to use hashtag propagation to infer the who following whom links. However, after a careful data analysis, we found that the hashtag propagation may not provide very useful information for inferring who following whom links. If possible, we may try some other evaluations. One possible way is to test whether our inferred latent network structure conveys information that is consistent with our intuition. For instance, from the inferred weight between pairs of nodes, we may discover several communities. Then we may check whether different communities have different interests (e.g., one community is interested in sports, while the other is interested in entertainment).

REFERENCES

- [1] <http://www.stanford.edu/~boyd/>.
- [2] Albert R, Jeong H and Barabasi A, Error and attack tolerance of complex networks, *Nature*, 406: 378-382, 2000.
- [3] Bass F, A new product growth model for consumer durables, *Management Science*, 15: 215-227, 1969.
- [4] Brown J and Reinegen P, Social ties and word-of-mouth referral behavior, *Journal of Consumer Research*, 14(3): 350-362, 1987.
- [5] Domingos P and Richardson M. Mining the network value of customers. *ACM*, 2001.
- [6] Gladwell M. *The Tipping Point*. Little Brown, 2000.
- [7] Kempe D, Kleinberg JM, and Tardos E. *Maximizing the spread of influence through a social network*. In the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Pp. 137–146, 2003.
- [8] Kossinets G. *Effects of missing data in social networks*. *Social Networks*, 28:247–268, 2006.
- [9] Myers S and Leskovec J. *On the convexity of latent social network inference*. In *Advances in Neural Information Processing Systems*, 23, edited by J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel and A. Culotta, 2010.
- [10] Leskovec J, Mcglohon M, Faloutsos C, Glance N, and Hurst M. *Cascading behavior in large blog graphs*. In *WSDM*, 2007.
- [11] Rodriguez MG, Leskovec J, and Krause A, *Inferring networks of diffusion and influence*. In *proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1019-1028, 2010.
- [12] Rogers EM. *Diffusion of Innovations*. Free Press, 1995.
- [13] Sadikov E, Medina M, Leskovec J, and Garciamolina H. *Correcting for missing data in information cascades*. In *Proceedings of WSDM'2011*, pp.55~64, 2010.
- [14] Surowiecki J. *The Wisdom of Crowds*. Random House, 2004.
- [15] Yang J and Leskovec J. *Modeling information diffusion in implicit networks*. In *IEEE International Conference on Data Mining*, Stanford InfoLab, 2010.
- [16] Watts DJ and Dodds PS. *Influentials, networks, and public opinion formation*. *Journal of Consumer Research*, 34(4):441–458, 2007.
- [17] Watts DJ, A simple model of global cascades in random networks, in *Proceedings of the National Academy of Sciences*, 99(9): 5766-5771, 2002.