Poisson Trust Factorization for Incorporating Social Networks into Personalized Item Recommendation

Allison J.B. Chaney
Department of Computer Science
Princeton University
achaney@cs.princeton.edu

Prem Gopalan
Department of Computer Science
Princeton University
pgopalan@cs.princeton.edu

David M. Blei
Department of Computer Science
Princeton University
dblei@cs.princeton.edu

Abstract

Many web users are faced with the problem of selecting which books to read and movies to watch. Traditionally, we ask our trusted friends for recommendations, but algorithmic recommendation models make those choices even easier, saving us time and effort by steering us towards media we are more likely to enjoy. The downside to most algorithmic recommendations is that, for some people, part of the appeal of reading or consuming other media is in creating shared experiences with friends. In this work, we aim to bridge this gap. We present a model that incorporates social network information into recommendation models, reintroducing the social aspect to recommendation and having the potential to improve overall recommendations. Further, our model discovers the latent trust that exists between users in a network and allows us to consider which users are more trustworthy than others, providing us insight into the social network's dynamics.

1 Introduction

Recommendation systems are based on collaborative filtering, algorithms that harnessing recurring patterns of consumption across users and items. In particular, matrix factorization for collaborative filtering has provided a scalable and extensible solution to the recommendation problem [5].

Matrix factorization operates on data about which users purchased (or clicked or viewed) which items. These data hide usage patterns that can be uncovered by the algorithms and exploited to predict future purchases.

Many modern sources of user purchase data, however, also contain a social network. We will incorporate this network into algorithmic recommendation, positing that knowing the items that our trusted friends have purchased can improve traditional recommendation systems. With the latent patterns as a seed, our model asserts that a user is more likely to enjoy things enjoyed by her trusted friends. We note that the idea of "trust" is learned in our model—each user may trust some friends more than others.

There is previous work on incorporating social networks into recommendation. Yang, et al. [7] outline several matrix factorization and nearest neighbor approaches, though most of the methods require that the trust between users is known in advance. Of the matrix factorization approaches (which our approach builds on), only SoRec [6] does not require trust already present as weights on network connections. SoRec factorizes two matrices at once: the user-item matrix and the user-user social network, tying an individual's preferences for items to their preferences for friends. While this is an interesting approach, it does not directly address the idea of friends influencing purchases.

In this work, we studied a data set from Goodreads[1], a website that combines the ratings of books with social network information. With this website, a user can rate and review books, maintain lists of books called "bookshelves," and befriend other users. Books in a users list can be rated using a five-star system, or remain unrated since one of the default bookshelves is books *to read*. We fit our model using data that contains 964,234 ratings with 1 to 5 stars. This set contains ratings of over 208,319 items by 14,468 users; there are 31,499 network connections among those users.

In the following, we describe our model, algorithm, and report on a preliminary study of our data.

2 Poisson Trust Factorization

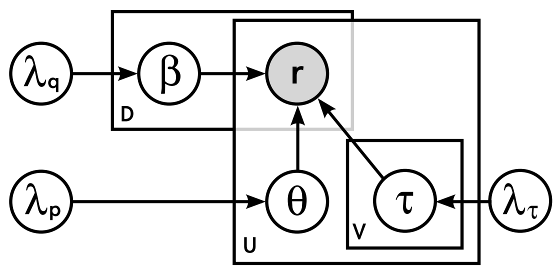

Figure 1: A directed graphical model of Poisson trust factorization.

Our model extends Poisson matrix factorization [2, 1], which replaces the traditional Gaussian approximation of ratings with a Poisson. In addition to latent user and item embeddings, our model discovers the trust between connected users, and this trust is then incorporated into the generative process of the ratings. The intuition behind this model is that there are multiple explanations for a positive rating. Some ratings can be explained from the user's latent preferences, and some can be explained by a the user's friends enjoying the book. We call our model Poisson trust factorization (PTF). Figure 2 shows the graphical model

We now describe the generative process behind PTF. First, user preferences come from a Gamma distribution:

$$\theta_u \sim \text{Gamma}(a, b) \tag{1}$$

(Recall that the Gamma is an exponential family distribution on positive scalars.) Item characteristics also come from a Gamma:

$$\beta_i \sim \text{Gamma}(c, d) \tag{2}$$

For each set of connected users, the trust variable is generated in a similar way:

$$\tau_{uv} \sim \text{Gamma}(f, g) \tag{3}$$

Note this is a directed relationship. User u might trust user v, but not vice-versa.

Finally, ratings come from a combination of base matrix factorization and a trust component:

$$r_{ui} \sim \text{Poisson}\left(\theta_u^\top \beta_i + \sum_{v \in N(u,i)} \tau_{uv} r_{vi}\right) \tag{4}$$

In inference, it will be useful to employ auxiliary latent variables $r_{ui} = \sum_n z_{uin}$ where $n = \{n_1^{MF}, \cdots, n_K^{MF}, n_1^{trust}, \cdots, n_V^{trust}\}$:

$$z_{uik}^{MF} \sim \text{Poisson}(\theta_{uk}\beta_{ik}) \tag{5}$$

$$z_{uiv}^{trust} \sim \text{Poisson}(\tau_{uv} r_{vi}) \tag{6}$$

[1]www.goodreads.com

3 Posterior Inference

We have defined a model of trust and factors. In inference, we condition on a social network and ratings data to approximate the conditional distribution of the factors and user trusts.

To handle large data sets, we developed a variational inference algorithm to approximate the posterior [4, 2]. Our derivation is based on the complete conditionals, i.e., the conditional distribution of each latent variable given the other latent variables and observations [3]. With the auxillary variables, we can derive the complete conditionals of PTF. For the Gamma variables (user representation, item representation, and trust), they are:

$$\theta_{uk} \mid \beta, \tau, z, m, r \sim \text{Gamma}\left(a + \sum_i z_{uik},\ b + \sum_i \beta_{ik}\right) \tag{7}$$

$$\beta_{ik} \mid \theta, \tau, z, m, r \sim \text{Gamma}\left(c + \sum_u z_{uik},\ d + \sum_u \theta_{uk}\right) \tag{8}$$

$$\tau_{uv} \mid \theta, \beta, z, m, r \sim \text{Gamma}\left(f + \sum_{i \in D(u,v)} z_{uiv},\ g + \sum_{i \in D(v)} r_{vi}\right) \tag{9}$$

For the auxillary variables, the complete conditional is

$$z_{ui} \mid \theta, \beta, \tau, m, r \sim \text{Mult}\left(r_{ui}, \phi\right) \tag{10}$$

where

$$\phi \propto \langle \theta_{u1}\beta_{i1}, \cdots, \theta_{uK}\beta_{iK}, \tau_{u1}r_{1i}, \cdots, \tau_{uV}r_{Vi} \rangle \tag{11}$$

We embed these distributions into a coordinate ascent variational inference algorithm by first setting up a mean-field variational distribution with exponential family factors in the same families as above and then iteratively optimizing these factors to the expected natural parameter of the complete conditional. This algorithm goes uphill in the KL divergence between the variational distribution and the posterior.

4 Results

We fit our model to 1M ratings from Goodreads and achieved the average user precision values shown in Table 1. We find that our model outperforms both the trust model if matrix factorization is excluded, as well as baseline matrix factorization.

Table 1: Average per-user precision on held out data.

threshold	1	25	100
trust only	1.85%	1.03%	0.56%
MF100 only	8.47%	2.85%	1.39%
trust + MF100	8.81%	2.90%	1.42%

In addition to looking at precision, we can inspect individual predictions to better understand the relationships between users; looking at books in common and fitted trust parameters can shed insight on the explanations behind recommendations, which can lead to a more interpretable recommendation.

5 Discussion and Future Work

While the precision performance of our model only beats baseline matrix factorization by a small amount, it is our hope that this model improves interpretability and greater personalization for individuals. Recommendations are able to be explained as coming from a friend's rating, which can lead to improved interfaces.

There are many directions for future work. Users studies can be done to validate the interpretability of the results; this can provide insight into the areas in which the model performs well or poorly, and provides a metric of success in addition to precision. It is also important to study the impact of improved recommendation explanations as well as the differences in the recommendations themselves.

Further, the ratings model of Equaton 4 implicitly weights the factor side by the number of components and the trust side by the number of connections. We suspect that learning this relative weighting will give further improvements.

The model can also be extended to include hierarchical and topical trust—trust can be explained as a function of the user's credulity and the friend's credibility, and this can be modeled across the K latent topics or dimensions discovered by matrix factorization.

References

[1] J. Canny. GAP: A Factor Model for Discrete Data. *ACM Conference on Information Retrieval (SIGIR)*, 2004.

[2] P. Gopalan, J.M. Hofman, and D.M. Blei. Scalable recommendation with poisson factorization. *Submitted*, 2013.

[3] M. Hoffman, D. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1303–1347), 2013.

[4] Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Mach. Learn.*, 37(2):183–233, November 1999.

[5] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems, 2009.

[6] Hao Ma, Haixuan Yang, Michael R. Lyu, and Irwin King. Sorec: Social recommendation using probabilistic matrix factorization. In *17th ACM Conference on Information and Knowledge Management*, CIKM '08, 2008.

[7] Xiwang Yang, Harald Steck, Yang Guo, and Yong Liu. On top-k recommendation using social networks. In *Proceedings of the sixth ACM conference on Recommender systems*, RecSys '12, pages 67–74, New York, NY, USA, 2012. ACM.