Network-Based Extension of Multi-Relational Factorization Models

Fatemeh Vahedian, Robin Burke & Bamshad Mobasher
Center for Web Intelligence, Depaul University, Chicago, IL 60604
{fvahedia, rburke, mobasher}@cs.depaul.edu

ABSTRACT

Complex heterogeneous networks contain many types of relations, both local to a particular entity and distant in the network. Multi-relational factorization schemes that incorporate multiple local relations have shown improved recommendation accuracy. This paper extends this prior work on multi-relational factorization to include extended relations derived from network data and demonstrates improved accuracy for these extended hybrids.

1. INTRODUCTION

Multi-relational factorization models have emerged as a state-of-the-art approach to recommendation in areas such as social networks, where both items and users are characterized by relations of multiple types [4, 3]. In such formulations, there is a main "target" relation for which predictions will be generated and multiple "auxiliary" relations that contribute information. For example, in a movie recommendation setting, the relation "user-movie" is the target relation and other relations, such as "movie-genre", would be considered auxiliary. In this paper, we extend the multi-relational approach in [3] to include multi-step relations.

Figure 1 shows an example of a network containing movie preference data. We can view each relation as a typed edge in such a heterogeneous network. For example, there is a "genre" edge connecting each movie with the genre by which it is labeled, and we can create a movie-genre relation by collecting all such edges. More distant relations can be created by considering multi-step typed paths, called *meta-paths* [7]. For example, the user-actor relation, which does not appear directly in the data, can be composed by following all user-movie/movie-actor meta-paths.

2. HETEROGENEOUS NETWORKS

A heterogeneous network is a directed graph in which both nodes and edges have types. Two edges of the same type, by definition, share the same object types at their originating and end points. A meta-path is a sequence of edge types, a composition of relations created by the edges that exist between two object types. Since there may be multiple edges that correspond to a given edge type, traversing a meta-path yields not a single node, but rather a set of destination nodes. For example, consider in the Movie schema, a meta-path consisting of just the movie-genre edge. If we start with "Whiplash" from Figure 1 and follow this meta-path, we would arrive at a set of destination nodes: {"Drama", "Music"}. In our examples, we will typically denote an edge type with the initials of the beginning and ending node types.

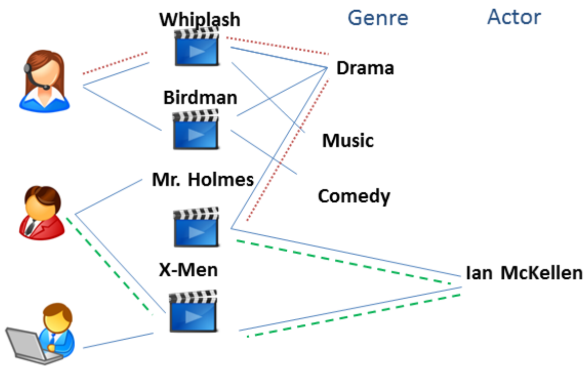

Figure 1: Heterogeneous network example

Prior work has demonstrated that meta-paths of various lengths could contribute to a multi-component weighted hybrid in a variety of network settings [1, 2, 6, 5]. In this work we use meta-paths to build multiple relations for a multi-relational factorized model.

3. MULTI-RELATIONAL FACTORIZATION

Multi-relational factorization models have improved prediction performance and are currently considered a state-of-the-art models in recommender systems and relational prediction research [3]. To date, this work has used direct relations, such as user-movie or movie-actor in our example. Our contribution here is to extend this work with relations generated using extended meta-paths.

In multi-relational matrix factorization models, one *target* relation is predicted and the remaining *auxiliary* relations are used as side information. For example, if the task is to recommend movies to users, the user-movie relation is the target relation and the other links between nodes such as movie-genre and movie-actor are auxiliary. In the multi-relational matrix factorization model DMF described in [3], different latent feature models are defined for each relation. Parameters are learned from the factorization process in such a way that they are optimized for the best performance on each relation, associating one latent feature vector model with each relation.

The CATSMF model is proposed in [3] to improve the efficiency of the DMF model when applied to multiple prediction targets. Since the DMF model must learn parameters for each relation individually, the number of parameters to be learned grows by a factor of number of relations in the network. In order to deal with this problem, CATSMF limits the parameters needed for the auxiliary relations by coupling them together. It also enables the learning of interactions between the different auxiliary relations.

4. EXPERIMENTS AND RESULTS

We build on the DMF and CATSMF models by incorporating additional relations built from extended meta-paths. For this paper, we used a 33% subset of the MovieLens 1M dataset [1]. There are four relations directly available in this data, as indicated in our prior examples: user-movie, movie-actor, movie-director, and movie-genre. In addition to these four direct relations, we generated six meta-path relations starting from the user. Figure 2 shows the meta-paths generated and the different experimental conditions. We ran the DMF and CATSMF algorithms using only the direct relations, and then built augmented versions of each using additional relations derived from the meta-paths shown. The models were optimized using BPR as the optimization criterion (BPR-opt), as described in [3].

	um	ma	md	mg	uma	umd	umg	umam	umdm	umgm
DMF	■	■	■	■						
DMF2	■	■	■	■	■	■	■	■	■	■
DMF3	■	■	■	■				■	■	■
CATSMF	■	■	■	■						
CATSMF2	■	■	■	■	■	■	■			

Figure 2: Relations in each recommendation model

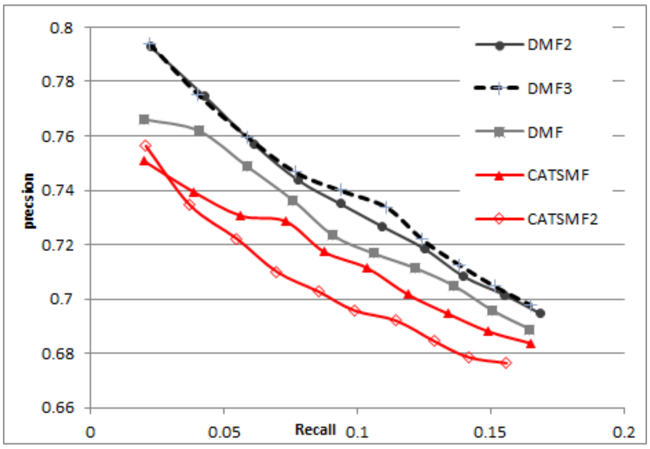

Figure 3: Recall vs. precision for MovieLens dataset

The user rating data was randomly partitioned into 80% training and 20% test data. Relations were generated from the training data and factorized. Figure 3 shows the results for recall and precision on recommendation lists of length 1-10 for the five algorithm variants shown in Figure 2.

The key finding is that the versions of the DMF algorithm that incorporate longer meta-paths demonstrate improvements in both precision and recall. The best performing variant is DMF3, which does not include the two-step relations UMA, UMD, UMG. This finding makes sense in that the movie-actor, movie-director, and movie-genre relations are already incorporated in the DMF model.

Unlike the results in [3] using different data, we did not find that CATSMF with its coupled approach offered better results than DMF algorithm. Part of the reason may be the increased level of personalization required by this recommendation task, as opposed to the relational completion tasks used in the CATSMF work. We are still exploring the reasons for these differences in performance.

5. CONCLUSION

We have shown that recommendation using multi-relational matrix factorization in networked data can be enhanced through in the inclusion of relations derived from meta-path expansions. Although the results here are only for a single data set, we have found benefits of such extended paths in other data sets in our work with linear weighted hybrids and are exploring the application of multi-relational factorization in these areas as well.

Previous work [2] has shown that the utility of relations based on extended meta-paths is not necessarily a decreasing function of path length – a finding confirmed here. Since the set of meta-path relations is by its nature unbounded, it is essential to find some means for limiting the set of components considered. We are exploring heuristics to enable the selection of the most useful relations.

Acknowledgments

This work was supported in part by the National Science Foundation under Grant No. IIS-1423368 (Multi-dimensional Recommendation in Complex Heterogeneous Networks).

6. REFERENCES

[1] R. Burke and F. Vahedian. Social web recommendation using metapaths. In *RSWeb@RecSys*, 2013.

[2] R. D. Burke, F. Vahedian, and B. Mobasher. Hybrid recommendation in heterogeneous networks. In *UMAP 2014*, pages 49–60, 2014.

[3] L. R. Drumond, E. Diaz-Aviles, L. Schmidt-Thieme, and W. Nejdl. Optimizing multi-relational factorization models for multiple target relations. In *CIKM 2014*, pages 191–200, New York, NY, USA, 2014. ACM.

[4] Z. Gantner, L. Drumond, C. Freudenthaler, S. Rendle, and L. Schmidt-Thieme. Learning attribute-to-feature mappings for cold-start recommendations. In *ICDM 2010*, pages 176–185. IEEE, 2010.

[5] F. Vahedian. Weighted hybrid recommendation for heterogeneous networks. In *RecSys '14*, pages 429–432, 2014.

[6] F. Vahedian and R. D. Burke. Predicting component utilities for linear-weighted hybrid recommendation. In *RSWeb 2014*, 2014.

[7] X. Yu, X. Ren, Y. Sun, Q. Gu, B. Sturt, U. Khandelwal, B. Norick, and J. Han. Personalized entity recommendation: A heterogeneous information network approach. In *WSDM 2014*, pages 283–292, 2014.

[1]http://grouplens.org/datasets/movielens/