Abstract: Recommendation using Extended Paths in Complex Networks

Robin Burke and Fatemeh Vahedian
Center for Web Intelligence
DePaul University
Chicago, IL USA
{rburke, fvahedia}@cdm.depaul.edu

1 Introduction

A heterogeneous network is defined by a diversity of objects and relations. For example, the Yelp social network can be viewed as a heterogeneous network with users, businesses, categories and locations as nodes and the many types of relations between these objects can be considered edges. A sample of this network structure can be seen in Figure 1.

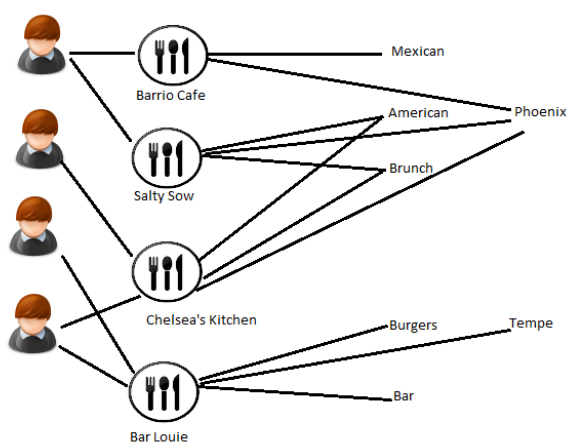

Figure 1: Yelp meta-path example

One of the key challenges for recommender systems in the social web and other complex domains is the effective integration of the many dimensions of available data about users. There are anticipated benefits to accuracy, diversity and personalization in taking more information into account, but also computational drawbacks in dealing with multi-dimensional data.

We have found that extended meta-paths, extracting relations beyond those local to an individual or to the recommended entity, can be effective for recommendation in a variety of data sets and in a variety of algorithmic settings. In particular, we have experimented with recommendation algorithms based on a multi-component hybrid model and based on multi-relational matrix factorization. We have shown that our meta-path-based approach to recommendation in heterogeneous networks yields improvements in both accuracy and diversity in social tagging systems and other complex heterogeneous networks.

2 Weighted Hybrid

A weighted hybrid recommender is a system comprised of multiple recommendation components, each of which returns a real-valued score for a combination of user and item. The scores from all the components are combined in a weighted sum [Burke, 2002]. More formally,

$$s(u, i) = \sum_j \alpha_j s_j(u, i)$$

where $s(u, i)$ is the overall score computed for a user-item combination, $s_j(u, i)$ is the score computed by the jth component, and α_j is the weight associated with the jth component. In our experiments, the weights are learned through an optimization procedure.

In our multi-component decomposition of heterogeneous network recommendation, the components themselves are built from two-dimensional matrices familiar to researchers in collaborative recommendation [Desrosiers and Karypis, 2011]. A user-based matrix is one in which the rows are users and the columns are "profiles" of the user along various dimensions. Users are compared on the basis of their profiles, and peer users form a neighborhood from which a target user's preferences for unknown items can be extrapolated. Item-based recommendation can be configured similarly.

The profile for a given user in a given component is created by following meta-paths from the user. For example , in Yelp, one profile might be defined by the set of businesses that a user has rated; another might be the set of terms that appear in reviews the user has written. Any meta-path can be used for this purpose and there is no requirement that meta-paths be simple: node and edge types can be revisited. For example, in the Yelp example, one component from a hybrid might represent users in terms of the types of businesses found in the locations that they frequent. This component would be built using the UBLB meta-path, which loops from businesses to locations and back to businesses again. Component generation could in theory continue indefinitely. However, there are significant computational costs in generating components and in optimizing a hybrid with a large number of components.

Our work on weighted hybrids has yielded a number of results. One is that there are a number of non-obvious tradeoffs in creating larger hybrids from extended network metapaths. Depending on the dataset, components built from longer meta-paths may perform better in terms of recall/precision performance than the corresponding component with a shorter prefix of that path. Predictive power is therefore not a simple decreasing function of the length of the path. Recommendation diversity is generally enhanced through the use of longer meta-paths, but this effect is dataset dependent. We also find, as expected, significantly longer weight optimization times when working with larger collections of components. See [Burke and Vahedian, 2013; Burke *et al.*, 2014; Vahedian and Burke, 2014; Vahedian, 2014]. for additional details about this work.

Figure 2 shows a typical result from this research. Hybrids incorporating longer meta-paths, such as HM5, demonstate improved performance over those with components drawn from direct edges in the networks, such as H1.

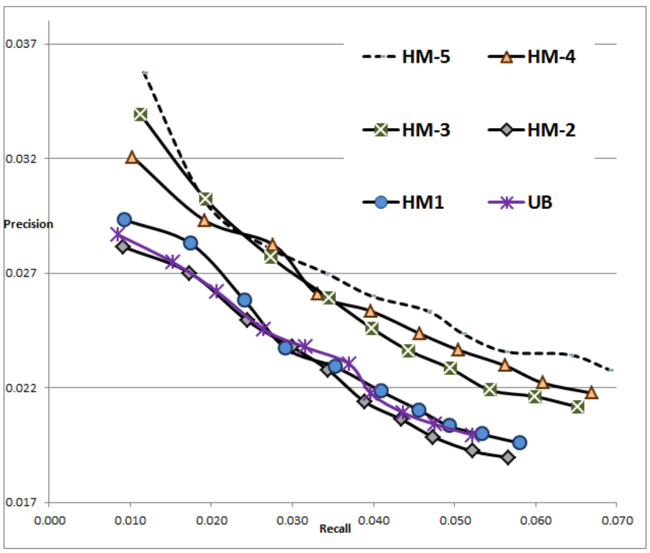

Figure 2: Recall vs. precision for Yelp dataset

Because of the non-linear nature of component utility, one important question is whether the hybrid weights can be predicted or at least estimated from the characteristics of the data. Such an estimate is made even more essential by the unbounded nature of the set of metapath-based components, since limiting the number of components is key to making weight learning efficient. We have experimented with entropy-based measures of the contribution of each component, with the aim of finding a metric with which to discriminate between components and filter out those unlikely to be useful, prior to the weight learning step. In [Vahedian and Burke, 2014], we show that it is possible to predict the information contribution of different meta-path-based components and thereby limit the size and complexity of the resulting hybrid, without significant sacrifice of accuracy.

3 Multi-relational Matrix Factorization

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 [Gantner *et al.*, 2010; Drumond *et al.*, 2014]. In such formulations, we have a main "target" relation that where predictions will be generated and multiple "auxiliary" relations that contribute information. Research in this area has concentrated on efficient use of direct relations such user-business or business-location from our Yelp example. However, as in our multi-component recommendation formulation, it is possible to construct auxiliary relations that go beyond the immediate connections in the network and to formalize these relations as meta-path expansions.

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 businesses to users, the user-business relation is the target relation and the other links between nodes such as business-location and busines-type are auxiliary. In the multi-relational matrix factorization model DMF described in [Gantner *et al.*, 2010], 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 individually.

CATSMF model is proposed in [Drumond *et al.*, 2014] to improve the efficiency of the DMF model when applied to multiple 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.

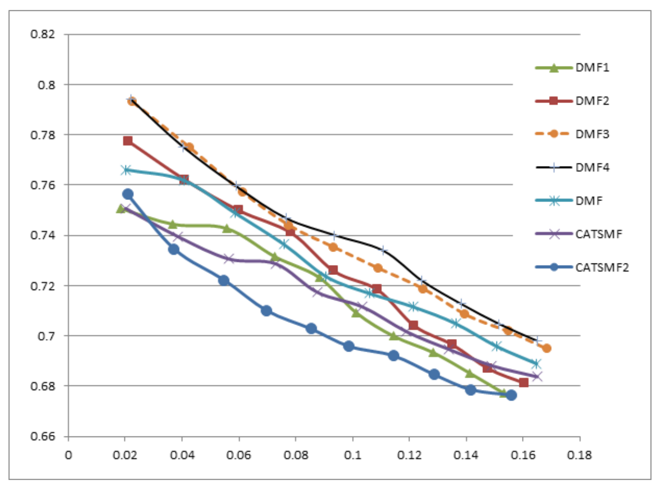

Figure 3: Recall vs. precision for MovieLens dataset

Our work in this area is preliminary. However, our results so far indicate that, as in the multi-component case, extended meta-paths generate useful relations for multi-relational factorization. Figure 3 shows precision/recall results for seven multi-relational factorization algorithm variants using a movie recommendation dataset. The versions of the DMF algorithm incorporating two-step paths, DMF1 and

DMF2, do not show better accuracy than the original algorithm, but those incorporating three-step extended representations of the user-movie relation (DMF3 and DMF4) do. The best performing variant is DMF4, which excludes 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, whereas the more extended relations are not.

4 Conclusion

The work described here demonstrates the utility of extended meta-path expansions for recommendation generation in multiple algorithmic contexts. Both in linear weighted hybrids and in multi-relational factorization, relations using lengthier meta-paths outperform those using only direct relations. In our linear weighted hybrids, we have found some diversity advantages as well.

Limiting the number of meta-paths considered is crucial to make these extended models tractable. We have explored information-theoretic techniques for predicting meta-path utility with some promising results. Future work will concentrate on the continued exploration of the tradeoff between recommendation performance and model complexity.

References

[Burke and Vahedian, 2013] Robin Burke and Fatemeh Vahedian. Social web recommendation using metapaths. In *RSWeb@RecSys*, 2013.

[Burke *et al.*, 2014] Robin D. Burke, Fatemeh Vahedian, and Bamshad Mobasher. Hybrid recommendation in heterogeneous networks. In *User Modeling, Adaptation, and Personalization - 22nd International Conference, UMAP 2014, Aalborg, Denmark, July 7-11, 2014. Proceedings*, pages 49–60, 2014.

[Burke, 2002] R. Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, 2002.

[Desrosiers and Karypis, 2011] Christian Desrosiers and George Karypis. A comprehensive survey of neighborhood-based recommendation methods. In Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor, editors, *Recommender Systems Handbook*, pages 107–144. Springer, 2011.

[Drumond *et al.*, 2014] Lucas Rego Drumond, Ernesto Diaz-Aviles, Lars Schmidt-Thieme, and Wolfgang Nejdl. Optimizing multi-relational factorization models for multiple target relations. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, CIKM '14, pages 191–200, New York, NY, USA, 2014. ACM.

[Gantner *et al.*, 2010] Zeno Gantner, Lucas Drumond, Christoph Freudenthaler, Steffen Rendle, and Lars Schmidt-Thieme. Learning attribute-to-feature mappings for cold-start recommendations. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 176–185. IEEE, 2010.

[Vahedian and Burke, 2014] Fatemeh Vahedian and Robin D. Burke. Predicting component utilities for linear-weighted hybrid recommendation. In *Proceedings of the 6th Workshop on Recommender Systems and the Social Web (RSWeb 2014) co-located with the 8th RecSys 2014*, 2014.

[Vahedian, 2014] Fatemeh Vahedian. Weighted hybrid recommendation for heterogeneous networks. In *Eighth ACM Conference on Recommender Systems, RecSys '14*, pages 429–432, 2014.