Graph-Based Recommendations: Make the Most Out of Social Data

Amit Tiroshi[1,2], Shlomo Berkovsky[1,3], Mohamed Ali Kaafar[1], David Vallet[1], and Tsvi Kuflik[2]

[1] NICTA, Australia
[2] University of Haifa, Israel
[3] CSIRO, Australia

Abstract. Recommender systems use nowadays more and more data about users and items as part of the recommendation process. The availability of auxiliary data, going beyond the mere user/item data, has the potential to improve recommendations. In this work we examine the contribution of two types of social auxiliary data – namely, tags and friendship links – to the accuracy of a graph-based recommender. We measure the impact of the availability of auxiliary data on the recommendations using features extracted from both the auxiliary and the original data. The evaluation shows that the social auxiliary data improves the accuracy of the recommendations, and that the greatest improvement is achieved when graph features mirroring the nature of the auxiliary data are extracted by the recommender.

Keywords: Graph-based recommendations, feature extraction, social data, music recommendations.

1 Introduction

The popularity of Web-based recommender system has led to the development of a spectrum of recommendation techniques. Most of them exploit, in a canonic form, three information entities: users, items, and feedback of users for items. Numerous prior works have shown that the accuracy of the generated recommendations improves when the representation of these entities is enriched by *auxiliary* external data, such as user's demographic data, item domain knowledge, or information on the recommendation constraints [2]. This finding was validated across a variety of recommendation techniques, application domains, and types of auxiliary data that can be used by the recommender.

Here, we investigate the exploitation of an auxiliary data originating from an online social networking system in a graph-based recommender. The choice of the social auxiliary data is driven by the abundance and ease of access to social data. Nowadays, it is common for users to have accounts on a social network (often, on more than one), to express their opinions, stay in touch with contacts, and share content of interest [5]. All this information can be captured and mined, and potentially serve as the source of a rich auxiliary user information for the

V. Dimitrova et al. (Eds.): UMAP 2014, LNCS 8538, pp. 447–458, 2014.

recommender [7,10]. Specifically, we leverage two types of social auxiliary data: (i) free-text tags assigned by users to content items, and (ii) online friendship links established between the social network users.

The focus on graph-based recommendations is also natural. In previous works, graph-based representation of the recommender data was shown to successfully encapsulate the relationships between the entities and to facilitate the generation of accurate recommendations [15,16]. It also allows for automatic extraction and population of graph-based features, which further improve the recommendation accuracy. Hence, our goal in this work is to study how the inclusion of auxiliary tags and friendship link data, along with the graph features extracted from this data, affects the accuracy of the graph-based recommender.

To answer this question, we use a publicly available data extracted from the music Web-site Last.fm.[1] The dataset consists of 1,892 users and 17,632 artists, whom the users tagged or listened to. That is, every user-artist pair is accompanied by the by the set of tags assigned by the user to the artist and by the number of times the user listened to the artist. Every user tagged or listened to, on average, 98.56 and 49.06 artists, respectively - and, vice versa, every artist was tagged or listened by 14.89 and 5.26 users, respectively. The dataset also contains information regarding 12,717 friendship links established between users on Last.fm. We represent this data as a graph, where the users, artists, and tags are the nodes, and 'listens' and 'friend' relationships are the edges. The assigned tags are reflected in the graph by the user-tag and tag-artist edges (see schematic diagrams in Figure 1).

We experiment with four graph schemes, where the tag, friendship, or both tag-and-friendship auxiliary data is included. For each schema, we extract and populate two sets of graph-based features. The first refers to a subset of *basic* features that can be extracted only from the original user, item, and feedback entities, disregarding the inclusion of the auxiliary data. The second is the *extended* set of features, where the basic features are augmented by a set of new features that mirror the nature of the included auxiliary data. We feed these features into a Gradient Boosted Decision Tree classifier [8] to predict withheld listening numbers, and recommend the top-ranked artists to users.

The evaluation highlights two key findings. Firstly, we show that the inclusion of auxiliary tag and friendship data improves the accuracy of the generated recommendations, whereas the inclusion of both achieves the greatest improvement. We also observe that the information encapsulated in the user friendship links turns to have more influence on the recommendations than the artist tagging information. Secondly, we show that features, which were conceived in a way that reflects the very nature of the auxiliary data being used, yield substantially more accurate recommendations.

2 Methodology

The effect of social auxiliary data on a graph-based recommender entails two questions: (i) what data is included in the graph representation, and (ii) which

[1] `http://www.last.fm`, Last.fm – Listen to free music and watch videos.

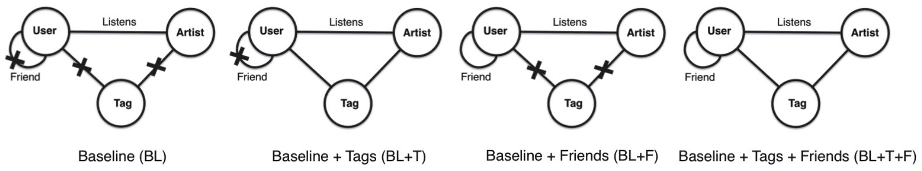

Fig. 1. Four schematic graph model

graph features are extracted and populated. We describe the methodology for graph-based data representation and feature extraction in the sections below.

2.1 Graph Models

The canonic Last.fm data used by graph-based recommender uses three entities: users, artists, and feedback (in this case - number of listens). The users and items are represented as the graph nodes, whereas the number of listens is expressed through the label of the edge between the two nodes. We use four different graph models, which use different types of social auxiliary data: no auxiliary data (denoted as baseline, BL), user-artists tags data (BL+T), user-to-user friendship links (BL+T), and both the tag and friendship data (BL+T+F). The graph models are illustrated in Figure 1.

When the auxiliary tag data is included, the free-text tags are also represented as the graph nodes. This way, a user-artist tag is converted into two edges: between the user and the tag, and between the tag and the artist. When the auxiliary friendship data is included, bi-directional edges between the nodes of users who friended each other are established. Finally, the inclusion of the tag and friendship data at the same time augments the baseline graph with the tag nodes, user-tag edges, tag-artists edges, and user-to-user edges.

2.2 Feature Subsets

Given the above graph-based representation of the data, we extract and populate a set of graph features. Some of these features can be populated directly from the data, e.g., number of artists listened by a user, number of users who listened to an artist, average number of listens for all artists listened by a user, and so forth. Other features inherently rely on the graph-based representation, e.g., node degree centrality, average neighborhood degree, PageRank score, and clustering coefficient. Intuitively, these features are not populated from the data, but rather quantify the importance of nodes in the graph-based representation of the data. Note that these features can be populated both for the user and artists nodes, and we refer the reader to [16] for an elaborate discussion of the graph features that can be extracted and populated.

We differentiate in our work between two groups of features. The *basic* features are extracted from the graph-based data representation only and disregard the unique nature of the social auxiliary data. For example, consider the PageRank score feature. The value of this feature for a given user node can be populated regardless of the auxiliary data. Including the friendship edges will clearly affect the value of this feature, but it can still be extracted and populated for the four graph models. We denote by F the set of basic features, e.g., F_{BL+T} is the set of basic features extracted from the graph-based representation of data augmented by user-artist tags.

The set of *extended* features includes, on top of the above-mentioned basic features, also new features that mirror the nature of the included auxiliary data. For example, consider a feature defined as "the ratio of user's friends, who also listened to a certain artist." This feature leverages the very notion of friendship between two users, on top of only considering the presence of a new edge between the two nodes in the graph-based representation of the data. We denote the set of extended features by $\hat{F}$, e.g., $\hat{F}_{BL+T+F}$.

We would like to highlight the dual impact of the *new* features $\hat{F} \setminus F$. These features can be populated using the auxiliary data only, as this data does not exist in the BL graph model. In the first instance, the mere presence of these features in the recommendation process may affect the accuracy of the graph-based recommender. However, the inclusion of the auxiliary data and the augmentation of the graph-based representation with the new nodes and edges, may affect the values of the basic features in F and, indirectly, also affect the recommendation generation process. We will investigate this duality in the experimental part of the work.

3 Experimental Setting

3.1 Dataset

We use a publicly available dataset extracted from the Last.fm music Web-site and used in [3]. The dataset contains 1,892 users and 17,632 artists, whom the users tagged and/or listened to. The dataset also contains social information regarding the friendship links established between Last.fm users. There are overall 12,717 such bidirectional friendship links. There are in total 11,946 unique tags in the dataset, which were assigned by users to artists 186,479 times. Each user assigned on average 98.56 tags, 18.93 of them being distinct. Each artist was assigned with 14.89 tags on average, 8.76 of them being distinct.

A brief characterisation of the dataset is shown in Figure 2. We first observe in Figure2(a) the distribution of the number of friends per user. We note that the average number of user-to-user edges in the BL+F model is low, which is illustrated by the vast majority of users having less than 10 friends, and about half of users having less than 4 friends. Nevertheless, friendship-based features prove to be important for the recommendations, as per the next section. Intuitively, the existence of the friendship edge between two users can be a good indicator of similar tastes, and as such, friendship-based features are expected to affect the recommendation process.

In Figure2(b), we show the distribution of the average number of tags each user shares with other users. This is presumably an indicator of the connect-

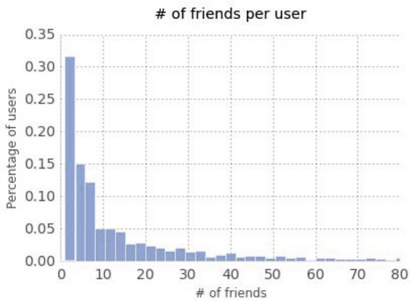

(a) Distribution of the number of friends per user

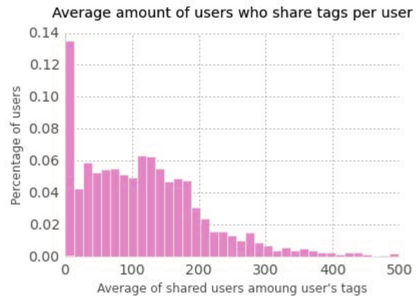

(b) Distribution of the average number of tags shared with other users

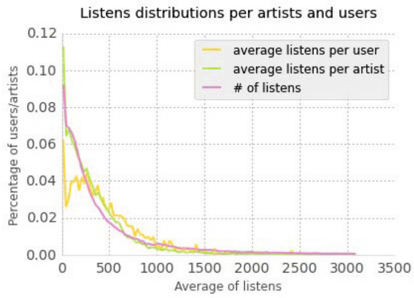

(c) Distribution of the average number of listens per user/artist and overall

Fig. 2. Data characteristics

edness of users in the BL+T graph. We observe that, on average, a user tag is found redundant across a non-negligible set of other users. In essence, many tags are used by a significant portion of users, while only 13.5% of tags do not propagate in the graph. In Section 4, we show that this feature is amongst the most important features extracted from the auxiliary tag data.

Lastly, we measure the distributions of the number of listens per artist, user, and in total. We observe that the overall and per artist distribution are highly similar. The user-based distribution resembles the same behaviour, but, as expected drops faster. This aligns with the intuition that the number of users who listen to several hundreds of artists is smaller than the number of artists who are listened by several hundreds of users [11].

3.2 Prediction and Metrics

We perform a 5-fold cross validation in which 5 different training sets are created. We prune users with less than 5 ratings to ensure that every user has at least one rating in the test set and four training ratings. For each training fold, we create a graph for each graph model outlined in Figure 1. Each graph generates the set of basic features, as well as a different set of extended features associated with the auxiliary data being included. The generated features are used as the input for a linear regressor, trained to predict the number of listens for a given user-artist pair. We use the Gradient Boost Decision Tree regressor, which is often used for general classification and prediction problems [8].

In order to evaluate each graph model and its associated features, we measure the performance of the regressor in predicting the number of listens and in ranking the artists. For each user, we create a candidate set of artists by picking 10 artists out of the true set of top-50 artists listened by the user and complementing these by randomly chosen artists. For example, candidate set of 100 artists includes 10 artists listened by the user and 90 random artists.

Then we use the regressor to predict the number of listens for each artist in the candidate set, rank the set accordingly, and compute precision at 10 (P@10) for every user, as the performance metric [13]. That is, we compute the intersection between the top-10 artists in the predicted ranked set and the real 10 artists listened by the user that were included in the candidate set. Finally, we average P@10 across all the test set users. This evaluation method is known as *top-n recommendations* [6], and is applied to evaluate recommenders that use implicit factors, e.g., number of views or behavior logs.

4 Results

As outlined in Section 2, social auxiliary data allows for different representations of the graph model, which, in turn, offer new possibilities for the extraction and population of graph features. In this section, we first study the impact of such auxiliary data and the corresponding features on the accuracy of graph-based recommendations. Then, we analyze the difference between the recommendations generated using the basic and the extended set of graph features.

4.1 Auxiliary Data as a Source of Graph Features

In this experiment, we observe the performance of the recommender when using the extended set of features, i.e., both basic and new features, as extracted and populated from the four graph models. Figure 3 shows the P@10 values averaged across all test set users, as obtained for the BL, BL+T, BL+F, and BL+T+F graph models, while the candidate set size varies from 50 to 150 artists. The boundaries of the boxes denote the 25th and 75th percentile of P@10, and the average P@10 is marked by the dots inside the boxes.

First, we clearly observe that the inclusion of the social auxiliary data of either the artist tags or friendships links substantially improves P@10. When both the tags and friendship links are included in the BL+T+F model, we observe the highest average P@10 across all the candidate set sizes. The improvement with respect to the BL model that includes no auxiliary data is statistically significant, ranging from 70% for candidate set size of 50 to 84% for candidate set size of 150. Note that a modest improvement (2.9% for candidate set size of 50 to 9%

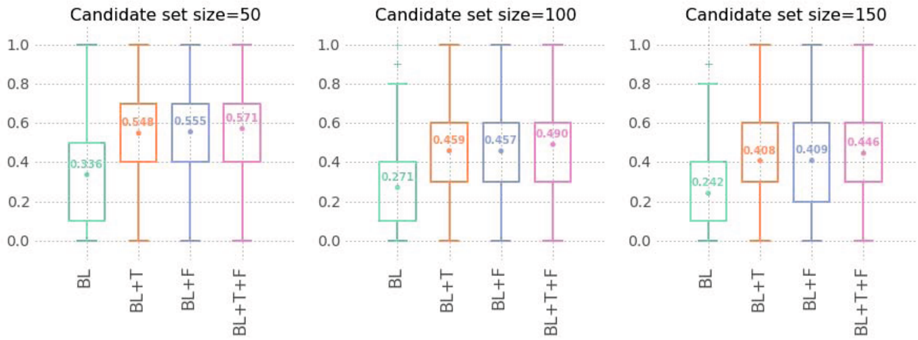

Fig. 3. Average precision for the four graph models and varying candidate set size

for candidate set size of 150) of the BL+T+F model with respect to the best-performing model that included either the tag or the friendship data. This shows that including both types of social auxiliary data further improves the accuracy of the recommendations.

Notably, the BL+T and BL+F models obtain very similar P@10 scores across all three candidate set sizes, showing the positive effect of the inclusion of the auxiliary data. However, as noted in Section 3.1, the tag data includes more than 186K user-artist tags, whereas the friendship data consists of only 12K user-to-user links. The observation that the obtained precision is similar indicates that a single user friendship link is more influential than a single artist tag and yields a greater improvement in the recommendation accuracy.

We observe a drop in the obtained P@10 scores when the size of the candidate set increases. This is expected, as the selection of top-10 listened artists out of a set of 150 candidates is inherently harder than out of 50 candidates only. Nevertheless, the drop in accuracy is smaller than one may expect. Specifically, P@10 of the BL+T+F model drops by 14% when the size of the candidate set is doubled from 50 to 150, and by 22% when it is tripled to 150.

To assess the fluctuations in the accuracy of the graph-based recommender, we measure the precision scores obtained for various users. In Figure 4-left, we plot the average P@10 obtained for all the users having a certain number of friends. The regression line shows that P@10 increases with the number of friends, i.e., users who established many friendship links with other users get more accurate recommendations than those who established a few. A strong positive correlation of 0.82 is observed between between the number of friends of a user and P@10 achieved for the user.

No dependency is observed is a similar experiment referring to the number of tags a user assigned. Hence, we turn to the popularity of the assigned tags. Figure 4-right shows the P@10 of a user as a function of the average popularity of tags used by the user, i.e., average number of other users who used these tags. As shown by the regression line, P@10 increases with the popularity of the used tags, although in this case the correlation is weaker, 0.29. Note that

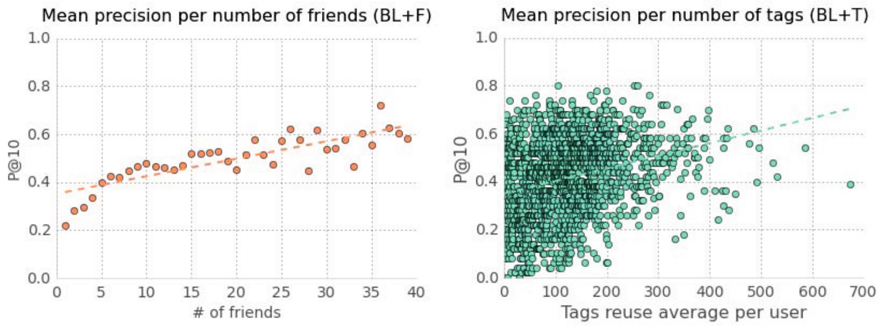

Fig. 4. Precision vs. number of friends (left), and average popularity of tags (right)

the precision obtained for users with non-popular tags fluctuates all over the precision spectrum. This suggests that the dependency between the number and popularity of tags, and the accuracy of recommendations is harder to establish.

4.2 Basic and Extended Feature Set

We turn now to the analysis of the performance of the basic and extended features sets extracted from various graph models. For each graph model, we compare the average P@10 of the predictions generated using the extended feature set $\hat{F}$ with the one generated using the basic feature set F. The candidate set size is fixed in this experiment to 50. The results of the comparison are shown in Figure 5, where the solid boxes denote the extended feature set $\hat{F}$ and the dotted boxes - the basic set F. Note that the comparison is impossible for the BL graph model, since no auxiliary data is included there and $\hat{F} = F$.

It can be observed that the extended feature sets outperform the basic ones across the boards. The improvement contributed by the new features $\hat{F} \setminus F$ is

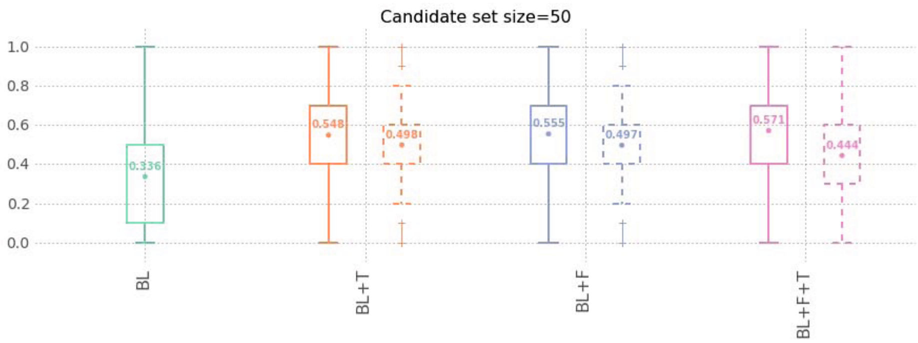

Fig. 5. Precision of the extended (solid boxes) vs the basic (dotted boxes) feature set for the four graph models

Table 1. Feature importance rankings for the four graph schemes

	F_{BL}		$\hat{F}_{BL}$	
#1	user_avg_nei_degree	(100%)	user_avg_nei_degree	(100%)
#2	user_pagerank	(56%)	user_pagerank	(56%)
#3	artist_pagerank	(50%)	artist_pagerank	(50%)
#4	user_clustering_coef	(49%)	user_clustering_coef	(49%)
#5	artist_degree_centrality	(47%)	artist_degree_centrality	(47%)
	F_{BL+T}		$\hat{F}_{BL+T}$	
#1	artist_pagerank	(100%)	fraction_of_shared_neighbours	(100%)
#2	user_avg_nei_degree	(96%)	fraction_of_shared_tags	(27%)
#3	user_pagerank	(87%)	user_avg_nei_degree	(18%)
#4	artist_degree_centrality	(79%)	user_clustering_coef	(17%)
#5	user_clustering_coef	(57%)	user_pagerank	(17%)
	F_{BL+F}		$\hat{F}_{BL+F}$	
#1	user_clustering_coef	(100%)	fraction_of_shared_neighbours	(100%)
#2	artist_pagerank	(73%)	fraction_of_shared_friends	(36%)
#3	user_pagerank	(65%)	user_pagerank	(23%)
#4	artist_degree_centrality	(58%)	user_clustering_coef	(20%)
#5	user_degree_centrality	(53%)	artist_degree_centrality	(16%)
	F_{BL+T+F}		$\hat{F}_{BL+T+F}$	
#1	artist_pagerank	(100%)	fraction_of_shared_neighbours	(100%)
#2	user_pagerank	(81%)	fraction_of_shared_friends	(41%)
#3	user_avg_nei_degree	(70%)	fraction_of_shared_tags	(18%)
#4	user_clustering_coef	(68%)	user_clustering_coef	(12%)
#5	user_degree_centrality	(46%)	user_avg_nei_degree	(11%)

comparable for BL+T and BT+F models (10% and 12%, respectively), while it is substantially higher for the BL+T+F model (29%). This clearly shows that extracting features that reflect the nature of the included auxiliary data and enriching the set of basic features is beneficial, as this improves the accuracy of the recommendations. It should be noted that some improvement is observed also for the basic feature set F, but this improvement can be leveraged if the new features $\hat{F} \setminus F$ are extracted and populated.

We note that out of the three graph models with the basic feature set, the lowest P@10 is achieved, somewhat surprisingly, by F_{BL+T+F}. That is, although $\hat{F}_{BL+T+F}$ is superior to $\hat{F}_{BL+T}$ and $\hat{F}_{BL+T}$ (see Figure 3), in the case of basic feature sets we observe that F_{BL+T} and F_{BL+T} both outperform F_{BL+T+F}. We conjecture that including both types of social auxiliary data but not extracting and populating the new features leads to redundancy in the graph and degrades the performance of the recommender.

Related to this is the analysis of most important features in each feature set. A by-product of the Gradient Boost Decision Tree is the feature importance ranking, which communicates the the contribution[2] of every feature in the set to correct predictions of the user-artist number of listens. Table 1 compares the 5 most important features along with their importance scores, for the four graph models. Like in Figure 5, the basic and the extended feature sets are identical for the BL model.

[2] Note that the importance values of the features shown in Table 1 do not sum up to 100%. Instead, the importance of the top feature is marked as 100% and the importance of other features is scaled with respect to this top feature.

Note that the most important features in the extended feature sets of all three graph models are new features that are extracted from the auxiliary data. The $fraction\_of\_shared\_neighbours$ feature denotes the ratio between the number of nodes that are common neighbours (both tags and friends) of the user and artist and the overall number of neighbors. This feature is further broken down into $fraction\_of\_shared\_friends$ and $fraction\_of\_shared\_tags$. The first feature is steadily selected by the Gradient Boost Decision Tree as the most important feature, and it is accompanied by the ratio computed for the tags and friends in the BL+T and BL+F models, respectively. In the combined BL+T+F model, both features are selected, but the importance of the ratio for friends is higher than for tags, which aligns with our earlier observation that the auxiliary friendship data is more influential than the tag data.

As the new features from the $\hat{F} \setminus F$ set are selected as the most important features, other features from the basic set are pushed down the list. However, in all three models, the basic features included in top-5 list of the extended feature set were also present in the top-5 list of the basic set. This experiment highlights our observation regarding the importance of the extraction of the new features for the graph-based recommendation process.

5 Related Work

There are multiple works in the field that evaluate the contribution of social links to popular personalisation and recommendation techniques, such as collaborative and content-based filtering [12,4,10]. A hybrid recommender system that combines tags and social links was evaluated by Guy et al. in [10]. The authors compared the hybrid approach to stand-alone approaches that solely use only social links or tags, and it was found that the hybrid approach significantly outperformed others. A user study discovered that recommendations generated by the hybrid approach were regarded by users as the most interesting. Our work reaffirm their findings and shows the superiority the graph the exploits both the tag and friendship social auxiliary data.

Along similar lines, Freyne et al. developed a personalized model for recommending social network news updates [7]. The model combined in a linear manner the quantified strength of user-to-user online relationships with the observed importance of network activities for the user. Their strength of user-to-user relationship encompassed the activity of the two users, as well as their direct and indirect (through common friends) interaction. The model was evaluated in a user study and was found to accurately recommend relevant social network updates, to assist users in establishing and maintaining friendships, and to boost contribution of user-generated content. Recent reviews of additional works that leverage social data for the purpose of recommendation generation can be found in [9], [14], and [1].

In [12], Konstas et al. developed a graph-based approach based on random walks for generating recommendations over social datasets, among them Last.fm (this was a proprietary dataset, and not the publicly accessible one used in our work). The reported results showed an improvement in recommendations when using the random walk approach in comparison to the baseline standard collaborative filtering. The optimizations in that work surrounded a single graph algorithm, random walk with restarts, and its parameters, such as the walk restart. In our work, random graph walks with static parameters are represented by the PageRank score feature. We extend the work of Konstas et al. by considering a broader range of graph features and evaluating the basic and extended feature set on four graph models. We also provide insight on the most important individual features.

In our own previous works we studied the performance of graph-based recommendations and graph features in various scenarios [15,16]. We used graph features in two different domains: for business recommendation and for interest recommendation. In both use-cases, the graph features were found to have a positive impact on the prediction accuracy. The difference between the graphs and features in this work with respect to our previous works lays in the availability of social auxiliary data, which introduces new graph features, affects the existing features, and imposes a new type of relatedness between the graph nodes. Also, the friendship data connects nodes from the same type (user-to-user edges), thus, extending the bipartite user-item graph. In this work, we used a different machine learning technique and evaluation metric to generate the predictions, showing that the effect of graph features holds also for Gradient Boost Decision Tree and ranking-based tasks.

6 Conclusions

In this work we studied the effect of inclusion of social auxiliary data on graph-based recommendations. We discovered that the inclusion of both the tags assigned by users to items and of the friendship links established between users contributes to the accuracy of recommendation. The impact of the friendship data was found to be stronger than of the tags data, while the strongest impact was observed when both types of auxiliary data were included.

Following these observation, we thoroughly investigated the need for extracting new features, which mirror the nature of the included data. We assessed the contribution of these features and conclude that the greatest improvement in the accuracy of the recommendations is achieved when the inclusion of the auxiliary data is complemented by the extraction and population of the new features. Overall, our work shows highlights the benefits offered by the graph features to recommender systems.

This work raises several challenging questions that we leave for the future. One question pertains to leveraging the content of the tags for enriching the graph structure. For instance, multiple tags assigned by a user (or to an artist) may convey a similar message. We would like to analyse the textual content of the tags and establish graph links between similar tags. We will then study the impact of these links on the accuracy of the generated recommendations.

Another question deals with the importance of specific values of the features. In here, we measured the importance of the extracted features. It may turn out, for instance, that although the importance of a feature is low, the importance of specific values of the feature is high. For example, consider the importance of listening to a popular mainstream artist versus of listening to a niche punk band. We will develop models that identify these important values and exploit them for the recommendation purposes.

References

1. Abel, F., Herder, E., Houben, G.-J., Henze, N., Krause, D.: Cross-system user modeling and personalization on the social web. User Model. User-Adapt. Interact. 23(2-3), 169–209 (2013)
2. Berkovsky, S., Kuflik, T., Ricci, F.: Cross-technique mediation of user models. In: Wade, V.P., Ashman, H., Smyth, B. (eds.) AH 2006. LNCS, vol. 4018, pp. 21–30. Springer, Heidelberg (2006)
3. Cantador, I., Brusilovsky, P., Kuflik, T.: Second workshop on information heterogeneity and fusion in recommender systems. In: RecSys (2011)
4. Cantador, I., Konstas, I., Jose, J.M.: Categorising social tags to improve folksonomy-based recommendations. J. Web Sem. 9(1), 1–15 (2011)
5. Chen, T., Kaafar, M.A., Friedman, A., Boreli, R.: Is more always merrier?: A deep dive into online social footprints. In: WOSN, pp. 67–72. ACM (2012)
6. Cremonesi, P., Koren, Y., Turrin, R.: Performance of recommender algorithms on top-n recommendation tasks. In: RecSys, pp. 39–46 (2010)
7. Freyne, J., Berkovsky, S., Daly, E.M., Geyer, W.: Social networking feeds: recommending items of interest. In: RecSys, pp. 277–280 (2010)
8. Friedman, J.H.: Greedy function approximation: A gradient boosting machine. Annals of Statistics 29, 1189–1232 (2000)
9. Groh, G., Birnkammerer, S., Köllhofer, V.: Social recommender systems. In: Recommender Systems for the Social Web, pp. 3–42. Springer (2012)
10. Guy, I., Zwerdling, N., Ronen, I., Carmel, D., Uziel, E.: Social media recommendation based on people and tags. In: SIGIR, pp. 194–201 (2010)
11. Haupt, J.: Last.fm: People-powered online radio. Music Reference Services Quarterly 12(1-2), 23–24 (2009)
12. Konstas, I., Stathopoulos, V., Jose, J.M.: On social networks and collaborative recommendation. In: SIGIR, pp. 195–202 (2009)
13. G. Shani and A. Gunawardana. Evaluating recommendation systems. In *Recommender Systems Handbook*, pages 257–297. 2011.
14. Shapira, B., Rokach, L., Freilikhman, S.: Facebook single and cross domain data for recommendation systems. User Modeling and User-Adapted Interaction 23(2-3), 211–247 (2013)
15. Tiroshi, A., Berkovsky, S., Kaafar, M.A., Chen, T., Kuflik, T.: Cross social networks interests predictions based ongraph features. In: RecSys, pp. 319–322. ACM (2013)
16. Tiroshi, A., Berkovsky, S., Kaafar, M.A., Vallet, D., Chen, T., Kuflik, T.: Improving business rating predictions using graph based features. In: IUI, pp. 319–326. ACM (2014)