Encoding temporal and structural information in machine learning models for recommendation

Tiphaine Viard[1] and Raphaël Fournier-S'niehotta[2]

[1] Riken Advanced Intelligence Project, Tokyo, Japan
tiphaine.viard@riken.jp
[2] CEDRIC, CNAM, Paris, France
fournier@cnam.fr

Abstract. Recommender systems focus on the task of selecting relevant items for a user, according to their taste: movies to watch, books to read, etc. In this paper, we provide a proof-of-concept that graphs and link streams can be relevant for the recommendation task.
We design simple features modelling the structure and its evolution with time of recommender systems in order to validate this claim. We learn these features alongside classical content-based features, using a gradient boosting machine model (*XgBoost*) and perform rating prediction on a movie rating dataset, MovieLens20M, and a book rating dataset, from GoodReads. We obtain comparable performance to the state-of-the-art, and some interesting leads in terms of explicability.

Keywords: Recommender systems · machine learning · graphs · link streams

1 Introduction

Collaborative filtering methods, among others, rely on a matrix modelisation of the user-item rating data. The matrix can be seen as a graph, leveraging the structural and social information between users and items, see [3] for a review. However, that data can also naturally be seen as an interaction stream, where a user u rates an item i at a time t. Past works have demonstrated the relevance of graphs for recommenders [5]. The *link streams* model [1] is a recent proposition to extend graph theory to jointly model the temporal and structural aspects of such data and this work is the first attempt at using a time-varying graph model for recommendation.

In this paper, we show that link stream modelling produces relevant descriptors for a large-scale recommendation task. We test our approach on interactions from the well-known Movielens 20M dataset[3], along with a more recent one, Goodreads[4]. We devise relevant content-based features for both datasets and generate link stream-based features, adapted to the datasets' parameters.

[3] Grouplens: `grouplens.org/datasets/movielens/20m/`.
[4] UCSD Book Graph: `sites.google.com/eng.ucsd.edu/ucsdbookgraph/home`.

We then feed a state-of-the-art machine learning algorithm for recommendation (*XgBoost*) to learn the recommendation task.

We evaluate the relevance of our link-stream features by comparing their performance to a content-based-only baseline, and to state-of-the-art results. We obtain promising results in terms of RMSE, even though our primary objective is not to be directly competitive with state-of-the-art results, but to demonstrate the point of incorporating link stream features into more complex models.

Our analysis of graph and link stream-based features shows that they provide an accurate and fine-grained description of the datasets at hand. This makes an interesting point towards more explainability in recommender systems.

2 Problem setting and proposed model

Given a user u and a movie i, we focus on the task of predicting the rating assigned by u to i on a scale from 0 to 5, *i.e.* a *regression* task, since ratings are ordered, non-independent classes. Let U be a set of users, and I a set of movies, with $|U| = n$ the number of users, $|I| = m$ the number of movies, and F a set of features such that $|F| = f$. Our model inputs a $(n \cdot m) \times f$ matrix, and outputs a $n \cdot m$ matrix P, where each element r of P is the predicted rating, between 0 and 5. The feature set F is composed of numerical and categorical variables divided into three sets: $F(u)$, a set of user-based features, $F(m)$, with movie-based features, and $F(u,i)$, the set of interaction-based features (see Section 3).

Modelling a recommender system as a *bipartite graph* $G = (U, I, E)$ is rather natural, with U the set of nodes for the users, I for the items, and $e \in E \subseteq U \otimes I$ representing an interactions between a user and an item. A *bipartite link stream* $L = (T, U, I, E_L)$ is defined by a time span T, a set of users U, a set of items I, and a set of links $E_L \subseteq T \times U \otimes I$ [1]. Nodes u and i are linked at time t if $(t, ui) \in E$. We say that (b, e, ui) is a link of L if $[b, e]$ is a maximal interval of T such that u and i are linked at all t in $[b, e]$.

The usual properties of graphs have been generalized to link streams [1], enabling the study of interaction streams with a single modelling structure, and without resorting to snapshots.

3 Data and feature engineering

For our evaluation, we focus on a movie rating dataset (MovieLens) and a book rating dataset (Goodreads), with similar characteristics (time-based ratings, easily usable content-based information, size magnitude) which makes them comparable. The MovieLens 20M dataset comprises of 20,000,263 interactions, involving 138,493 users and 27,278 movies over the course of 20 years.

The Children Goodreads dataset contains 10,080,558 interactions between 462,164 users and 122,741 books, over the course of 11 years.

Let us now explain the features we design. From the available information in the MovieLens dataset, we obtained 39 content-based features.We extract

Dataset	Metric	With stream features		Without stream features	
		train	test	train	test
MovieLens 20M	MAE	**0.60234**	**0.63427**	0.63421	0.64427
	RMSE	**0.76349**	**0.80954**	0.82682	0.83961
	NDCG@10	**0.99212**	**0.97209**	0.98212	0.94863
Goodreads Children	MAE	**0.5299**	**0.5899**	0.63343	0.64986
	RMSE	**0.6818**	**0.7561**	0.7986	0.82652
	NDCG@10	1.0	**0.9901**	1.0	0.9772

Table 1. MAE, $NDCG$@10 and $RMSE$ values for the train and test datasets of both Movielens and Goodreads, with and without link stream or graph features.

the same features from the GoodReads dataset where available, ending with 10 content-based features.

For graph and link stream features, we devised 21 features, mixing older graph models ideas and new additions. **Neighbourhood-based features** explore the relations between users and items over time. We use graph and link stream degree statistics (mean, maximum, std. dev., assortativity). **Inter-contact time features** model the dynamics of the interactions. **Clique-based features** are original link stream features similar to clustering. Given the size of the data, we resort to sampling maximal balanced bipartite cliques, *i.e.* cliques corresponding to dense subgroups of users all rating a substantial number of items. Our features describe this sampling, using, for any given node, the fraction of cliques containing it, the balancedness and average duration of those cliques.

4 Evaluation setting and recommendation results

We perform the recommendation task by feeding our features into *XgBoost*.We tuned *XgBoost* using Bayesian optimization on the hyper parameter space, and obtain optimal results with deep trees and a small learning rate. We report our recommendation results in terms of MAE, RMSE and NDCG@k (*Normalized Discounted Cumulative Gain*, for a ranking of $k = 10$ elements). Table 1 presents our results for both datasets, comparing two settings: one with link stream features and content-based features, and the baseline, using only the content-based features. We see that adding link stream features lead to significantly better learning performance than using only content-based features (bold values). For the MovieLens dataset, these results are below the best results in the state-of-the-art (RMSE: 0.7652 in [2], with a deep learning model). There is no similar baseline on the GoodReads dataset to compare ourselves. However, our model is significantly simpler, and as we discuss in Section 5, has interesting implications in terms of explicability.

In addition to the classic performance metrics presented above, we evaluate the descriptiveness of the link stream features we devise. Figure 1 shows the relative importance of features as selected by *XgBoost*.We can see that the introduced link stream and graph features are commonly used as split points by

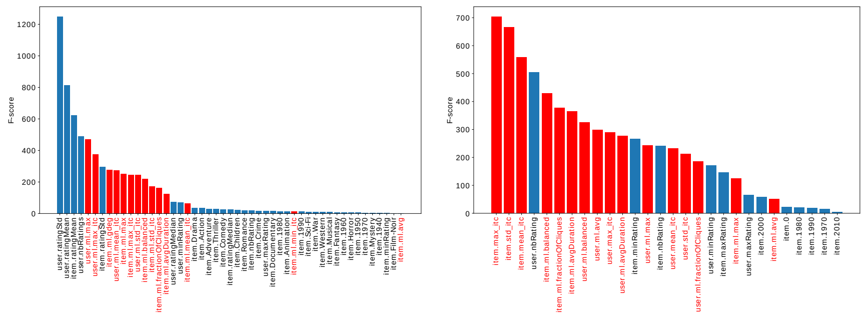

Fig. 1. Feature importance as selected by *XgBoost*, with link stream features in red. Left: MovieLens, right: GoodReads Children.

the boosting algorithm, which supports the claim that such features are subtle descriptors of the structure and dynamics of the datasets. See Figure 1 for the details. Moreover, graphs and link streams are easily visualizable, which is likely to improve explicability of recommendations substantially.

5 Conclusion and perspectives

Our contribution is a proof-of-concept of incorporating link streams features in a classic recommender system environment, validated on two large-scale datasets. We perform slightly below the state-of-the-art, but our model focuses on remaining simple and explicable, with link streams. For more details on this work, see [4]. The source code is available at `https://bitbucket.org/tiph_viard/social_recommendation`.

As perspectives, since many recommender systems rely on finding similar-minded users, establishing a robust model of communities with link streams may prove very efficient. Recent advances in the design of graph neural networks also brings promising models to close the performance gap. We hope our experiment prolong the fruitful exchanges between the "recommender system" and the "social network analysis" communities.

References

1. Latapy, M., Viard, T., Magnien, C.: Stream graphs and link streams for the modeling of interactions over time. arXiv preprint arXiv:1710.04073 (2017)
2. Strub, F., Gaudel, R., Mary, J.: Hybrid recommender system based on autoencoders. In: Workshop on Deep Learning for Recommender Systems. pp. 11–16. ACM (2016)
3. Tang, J., Hu, X., Liu, H.: Social recommendation: a review. Social Network Analysis and Mining **3**(4), 1113–1133 (2013)
4. Viard, T., Fournier-S'niehotta, R.: Augmenting content-based rating prediction with link stream features. Computer Networks **150**, 127–133 (2019)
5. Zhou, T., Ren, J., Medo, M., Zhang, Y.C.: Bipartite network projection and personal recommendation. Phys. Rev. E **76**(4) (2007)