Recommender systems increase exposure diversity. Or do they ? A complex networks approach.

Augustin Godinot and Fabien Tarissan

Université Paris-Saclay, CNRS, ISP, Ecole Normale Supérieure Paris-Saclay, 4 avenue des Sciences, 91190 Gif-sur-Yvette,

1 Introduction

While the ever-increasing emergence of online services has led to a growing interest in the development of recommender systems, the algorithms underpinning such systems have begun to be criticized for their role in limiting the variety of content exposed to users [1]. In this context, the notion of diversity has been proposed as a way of mitigating the side effects resulting from the specialization of recommender systems [2]. However, we still know little about how classic recommendation paradigms affect users' behavior in terms of diversity. In this work, using a recent development on measuring diversity in *Heterogeneous Networks* [5], we show the relevance of a complex network approach to recommender systems evaluation: by conducting random walks on a tripartite graph modeling how users are exposed to different musical categories via recommendations, we were able to analyze the effect of a well-known recommendation model [3] on the diversity of users' attention.

The results show that the musical selections of a large proportion of users are diversified as a result of the recommendations. However, our study also reveals that this observation only holds for diversity in the sense of *variety* [7] and that recommendations usually fail to provide a *balanced exposure* to the different categories.

2 Networks, recommendations and diversity

Evaluating the diversity. To measure the diversity of the recommendations made to users, we rely on the recent work on measuring diversity in *Heterogeneous Networks* [5,6]. Let U be a set of users, I a set of items (or songs) and T a set of tags (or categories). We define the tripartite graph $\mathbb{T} = (U, I, T, E_{UI}, E_{IT})$ of users' activity where $E_{UI} \subseteq U \times I$ is the set of (weigthed) links relating users to their past musical records and $E_{IT} \subseteq I \times T$ is the set of links relating songs to their musical categories.

After normalizing the weights (out edges weights of a node must sum to 1), we conduct a two stages random walk to obtain the probability $p_{u \to t}$ of reaching each tag $t \in T$ from each user $u \in U$. From the distribution $(p_{u \to t})_{t \in T}$, we derive in Equation 1 the α-diversity $D_\alpha(u)$ of the user u (see [5] for a detailed explanation and properties). Intuitively, such an index measures the extent to which the probabilities are well distributed over the categories and defines the **organic diversity** of each user, that is the diversity of its past musical habits.

$$D_\alpha(u) = \left(\sum_{t \in T} p_{u \to t}^{\alpha} \right)^{\frac{1}{1-\alpha}} \text{ if } \alpha \neq 1 \quad \text{and} \quad D_1(u) = \left(\prod_{t \in T} (p_{u \to t})^{p_{u \to t}} \right)^{-1} \text{ if } \alpha \to 1 \quad (1)$$

Suppose now that we have information on how songs are exposed to users via recommendations, represented as a new set of links $E_{UI}^r \subseteq U \times I$. This defines a second tripartite graph $\mathbb{T}^r = (U, I, T, E_{UI}^r, E_{IT})$ on which one can conduct the same computations, leading to the measure of the **post-recommendation exposure diversity**, that is the diversity of the categories exposed to the users via recommendations. By comparing the organic and the post-recommendation exposure diversity, one can then study the effect of recommendations to the diversity of users' attention. In particular it becomes possible to measure whether or not the recommendations increase diversity.

Generating recommendations For its wide usage, its relevance to music recommendation and its relative simplicity, we choose to study the *Matrix Factorization for Implicit Datasets* model [3]. This model was trained on the *Million Song Dataset* [4], which was cleaned and randomly down-sampled to 100,000 users. Let r_{ui} be the number of times a user $u \in U$ has played the item $i \in I$. From this quantity, we derive two indicators: p_{ui} indicates whether a user u previously interacted with an item i and c_{ui} stands for the confidence we have in the binary proposition "user u likes item i".

$$c_{ui} = 1 + \mu r_{ui} \quad \text{and} \quad p_{ui} = \begin{cases} 1 \text{ if } r_{ui} > 0 \\ 0 \text{ otherwise} \end{cases} \quad (2)$$

$$x^*, y^* = \min_{x,y} \sum_{u,i} c_{ui} \left(p_{ui} - x_u^T y_i \right)^2 + \lambda \left(\|x\|_2^2 + \|y\|_2^2 \right) \quad (3)$$

After training the model (see Equation 3), we obtain estimated values $\hat{p}_{ui} = (x_u^*)^T y_i^*$ of whether a user u is likely to interact with an item i. Finally, we recommend the k items ($k = 10$, 50 and 500) with the best score $\hat{p}_{ui}$ to the user. Guided by cross validation, we choose the following parameter values: 3000 latent factors (ie x_u and y_u are vectors with 3000 coordinates), $\mu = 40$, $\lambda = 10^6$.

3 Results and conclusions

In Figure 1, we plot the diversity increase (difference between the post-recommendation and organic diversities) with respect to the organic diversity of each user and different number items recommended ($k = 10$ left, 50 middle and 500 right). We chose to focus on two values of diversity order: $\alpha = 0$ (top) captures the *variety* of exposed content (it is maximized when the number of tags reached is maximum) whereas $\alpha = 2$ (bottom) involves a notion of *balanced exposure* (it is maximized when the tags distribution is uniform).

The plot clearly shows that the musical selections of a large proportion of users are diversified as a result of the recommendations and that diversity also increases with the number of items recommended (from left to right). However, the plot also reveals

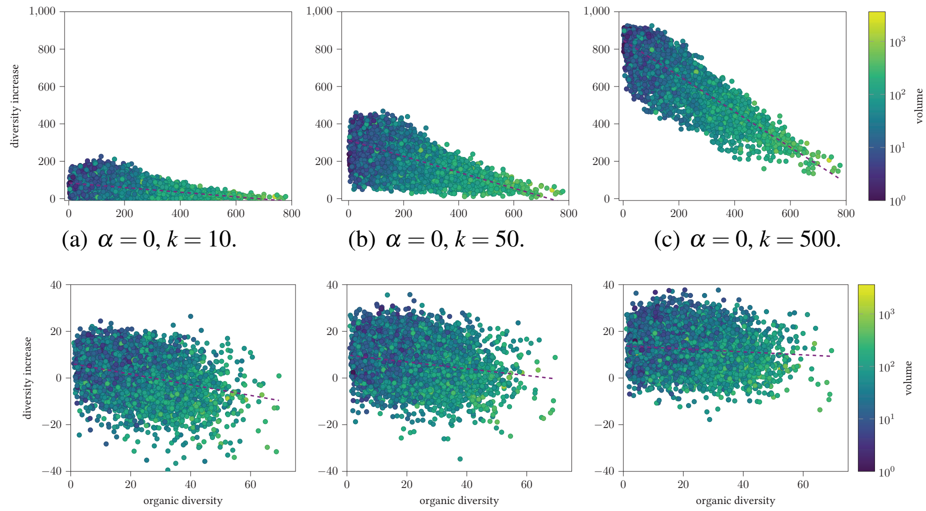

(a) $\alpha = 0, k = 10.$ (b) $\alpha = 0, k = 50.$ (c) $\alpha = 0, k = 500.$

(d) $\alpha = 2, k = 10.$ (e) $\alpha = 2, k = 50.$ (f) $\alpha = 2, k = 500.$

Fig. 1. Diversity increase with respect to the organic diversity

that the recommendations do not have the same effect when diversity is considered in terms of variety (top) or balance (bottom) and suggests that the main effect of the recommendations is to introduce new categories into users' musical habits. As soon as balance is taken into account in the measure, the recommendations are less effective in increasing diversity, which is even negative after the recommendations for a significant fraction users.

We believe that the complex network approach proposed in this work, as well as the practical investigation conducted on a collaborative filtering approach, applied to a real musical dataset, shed new light on how researchers working on recommender systems could examine the ethical effects of algorithmic recommendations.

This work is funded in part by the ANR grant ANR-20-CE38-0011.

References

1. Helberger, N.: On the Democratic Role of News Recommenders. Digital Journalism 7(8), 993–1012 (jun 2019)
2. Helberger, N., Karppinen, K., D'Acunto, L.: Exposure diversity as a design principle for recommender systems. Information, Communication & Society 21(2), 191–207 (dec 2016)
3. Hu, Y., Koren, Y., Volinsky, C.: Collaborative Filtering for Implicit Feedback Datasets. In: 2008 Eighth IEEE International Conference on Data Mining. IEEE (dec 2008)
4. McFee, B., Bertin-Mahieux, T., Ellis, D.P., Lanckriet, G.R.: The million song dataset challenge. In: Proceedings of the 21st international conference companion on World Wide Web - WWW '12 Companion. p. 909. ACM Press, Lyon, France (2012)
5. Morales, P.R., Lamarche-Perrin, R., Fournier-S'niehotta, R., Poulain, R., Tabourier, L., Tarissan, F.: Measuring diversity in heterogeneous information networks. Theoretical Computer Science 859, 80–115 (mar 2021)
6. Poulain, R., Tarissan, F.: Investigating the lack of diversity in user behavior: The case of musical content on online platforms. Information Processing & Management 57(2), 102169 (mar 2020)
7. Stirling, A.: A general framework for analysing diversity in science, technology and society. Journal of The Royal Society Interface 4(15), 707–719 (feb 2007)