

# ARTIST AND STYLE EXPOSURE BIAS IN COLLABORATIVE FILTERING BASED MUSIC RECOMMENDATIONS

Andres Ferraro, Dmitry Bogdanov, Xavier Serra

Music Technology Group -  
Universitat Pompeu Fabra  
first.lastname@upf.edu

Jason Yoon  
Kakao Corp.

jason.yoon@kakaocorp.com

## ABSTRACT

Algorithms have an increasing influence on the music that we consume and understanding their behavior is fundamental to make sure they give a fair exposure to all artists across different styles. In this on-going work we contribute to this research direction analyzing the impact of collaborative filtering recommendations from the perspective of artist and music style exposure given by the system. We first analyze the distribution of the recommendations considering the exposure of different styles or genres and compare it to the users' listening behavior. This comparison suggests that the system is reinforcing the popularity of the items. Then, we simulate the effect of the system in the long term with a feedback loop. From this simulation we can see how the system gives less opportunity to the majority of artists, concentrating the users on fewer items. The results of our analysis demonstrate the need for a better evaluation methodology for current music recommendation algorithms, not only limited only to user-focused relevance metrics.

## 1. INTRODUCTION

There are multiple factors that make design decisions for a music recommendation system a complex problem. Some decisions can be related to theoretical aspects of music, while others may have ideological or social connotations, may be subjective, not possible to quantify, or be changing depending on time and context [13].

Collaborative filtering methods are typically used to generate a recommendation by identifying patterns in what people listen from historical information. The drawback of these methods is that since they do not consider any other than information about interactions between users and items, it is not possible to generate recommendations for new items (the cold-start problem). Also the recommendations tend to follow the distribution of popularity of the music [8] with the most popular items being recommended more (the long-tail recommendation problem).

Celma and Cano [3] show this by analysing navigation, clustering and connectivity in artist similarity networks built with collaborative filtering data.

Some solutions had been proposed for long-tail and cold-start recommendations using audio information and metadata [12, 15]. By combining the best of these solutions we can create hybrid systems that can reduce these problems, but integrating multiple systems can be complicated [14]. With the advances in deep learning, new methods are able to automatically learn a representation from the data without the need for manually selecting the features [12, 15]

Still, these solutions have some issues, in particular related to the fact that they work as black-boxes. For example, it is difficult to explain the results and it is hard to know if different musical styles are well-represented. Also, previous works do not show how robust these methods are to biased datasets and if it is possible to generate recommendations for new styles or genres that are less present in the user-item interactions.

The growth of music streaming services in the last years has increased the importance of music recommender systems, and reducing the choice overload is commonly referred to as one of the advantages of these systems. But is important to understand the increasing impact that these systems have to what people listen. They define which song will be the next hit, how much will an artist earn or even which music genres might receive almost zero promotion. This raises some ethical issues that had been discussed in previous works, Holzapfel et al. [6] raise the question if a group of artists that are never recommended by a system can be considered a case of discrimination. As researchers, we have to think about the implications of the systems we develop and the importance of assuring every artist has a fair chance to reach the public [5].

Recently, there have been studies trying to address these issues. Cramer et al. [4] summarizes possible algorithmic biases and highlights that music recommendations for "balanced" not-biased consumption may not necessarily lead to optimal experience for many users. McInerney et al. [10] propose a bandits approach to balance exploration and exploitation in the recommendations for the users, but they do not address the impact on the exposure of different artists or music styles. Mehrotra et al. [11] proposes a way to understand the trade-off between relevance, satisfaction and fairness in music recommendations. In this



case, fairness measures the diversity of the level of popularity of recommendations, but it does not capture the overall exposure of the artists or the different musical styles.

Following these studies, we demonstrate preliminary results of our on-going research that gives a better understanding of the influence of music recommendation systems on users’ behavior that could affect artists’ exposure. We show that the distribution of the recommendations in terms of their artists, styles or genres is different from what the users had listened before. Also, we show that with time the system tends to recommend fewer items, therefore, focusing user interactions on fewer artists, which is not the desired behavior of the system.

## 2. PROPOSED ANALYSIS

In this work, we use a basic Matrix Factorization [7] algorithm and *Echo Nest Profile Subset* to build a user-track matrix and generate 10 track recommendations for each user. We use the associated tags from *Last.fm Dataset* to analyze how recommendations are distributed across the different musical styles in comparison with listening statistics from our dataset representing the initial preferences of users. We also show a simulation of how these recommendations can affect user behavior in the long term. For this we take the recommendation of the system for each user and increase the counter in the original user-track matrix, simulating that the users listened to all recommendations by the system. We then retrain the model and generate new recommendations. We repeat this process 30 times.

### 2.1 Datasets


The Million Song Dataset (MSD) [9] is a large dataset of audio features and metadata expanded by the Music Information Retrieval community with additional information including tags, lyrics and other annotations. The Echo Nest Taste Profile Subset [2] provides play counts by 1,019,318 users covering 384,546 songs from MSD. For this work we only consider users and items with more than 30 interactions (128,374 tracks and 445,067 users). Additionally, the Last.fm Dataset [2] provides song-level tags extracted from *Last.fm* for a subset of MSD. These tags are crowd-sourced and cover genre, instrumentation, moods and eras. One track can have multiple tags.

### 2.2 Metrics

For a better understanding of system behavior, we need to define metrics that can assess how probable is for new or less popular artists to be recommended and compare those across different styles. It is also valuable to know to how many different users each artist is recommended.

In this work, we use the *Gini index* to measure the distribution of how many users each artist gets recommended to, but in future works other metrics should be also considered (for example, proposed for multistakeholder recommendation approaches [1]).

We also use *Coverage* to measure the percentage of different artists globally recommended. With this metric we



**Figure 1:** Distribution of recommendations and users listening. Values are average percentages per music style.

can have an idea of the amount of artists that the system gives zero promotion.

## 3. RESULTS

### 3.1 Distribution of recommendations

Figure 1 shows the global tag distribution of all user-track recommendations pairs (10 tracks per user) compared to such a distribution for initial user listening behavior for the top 20 tags.<sup>1</sup> For the rest of the tags, the system is recommending 9.4% less compared with what the users listened to. Table 1 similarly reports an average percentage of recommendations and initial user preferences in three tag and artists categories grouped by their popularity in terms of the original play counts.

We can see a clear popularity bias in what users listen to, and this bias is further reinforced by recommendations, which may be not the desired behavior. The system is recommending more top tags and less long-tail tags than what people listened to.

	Tags 1-5	Tags 5-2k	Tag 2k-50k
Recommended	<b>4.7807</b>	<b>0.0347</b>	0.0001
Listened	4.1195	0.0327	<b>0.0003</b>

(a) Tags

	Artists 1-5	Artists 5-2k	Artists 2k-18k
Recommended	<b>1.5672</b>	<b>0.0433</b>	0.0003
Listened	0.6182	0.0370	<b>0.0014</b>


(b) Artists

**Table 1:** Average percentage of recommendations and user play counts for (a) tags and (b) artists with different popularity.

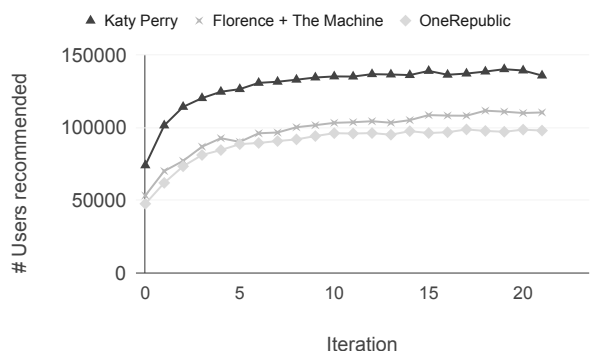
### 3.2 Simulating feedback loops

Figure 2 shows the results of simulating the feedback loop of the recommendations. We can see how the *Gini index*

<sup>1</sup> Note that there are 51,699 tags and therefore it is not possible to show all of them.



**Figure 2:** Coverage and Gini index of the recommendations simulating feedback loops.



**Figure 3:** Number of users reached by a song recommendation on the example of three popular songs when simulating the feedback loop.

increases on each iteration, starting in 0.95 and going up to 0.98. A value of 1.0 indicates that the system is recommending the same songs to all users. In the same figure we see the evolution of the *Coverage* of the recommendations. For the first iteration the *Coverage* is 40 % but at the last iteration the *Coverage* is 20 % meaning that 80 % of the songs are not recommended by the system.

In Figure 3 we demonstrate how three very popular songs (the 3rd, 5th and 6th most played songs according to our initial user-track matrix) gather even more exposure from recommendations during the feedback loop iterations. These songs have been recommended to between 50,000 and 100,000 users at the first iteration, and ended up being recommended to 100,000 to 135,000 users after 10 iterations. These songs occupied top three positions among the most listened songs in our simulation only after one iteration. It is important to mention that in a real case there will be other interactions between users and items that are not considered here.

#### 4. CONCLUSIONS

In this work, we have considered how the popularity bias is affecting collaborative filtering recommendations based on Matrix Factorization. In our experiments, this algorithm is increasing the exposure of more popular musical styles,

while reducing the exposure in the long tail, which may be an undesired behaviour. The goal of our future research is to expand our analysis on state-of-the-art algorithms proposed for cold-start and long-tail music recommendation, which are still lacking such an evaluation.

#### 5. REFERENCES

- [1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. Beyond personalization: Research directions in multistakeholder recommendation. *arXiv preprint arXiv:1905.01986*, 2019.
- [2] Thierry Bertin-Mahieux, Daniel P.W. Ellis, Brian Whitman, and Paul Lamere. The million song dataset. In *Proc. of the 12th International Society for Music Information Retrieval Conf. (ISMIR)*, 2011.
- [3] Òscar Celma and Pedro Cano. From hits to niches?: or how popular artists can bias music recommendation and discovery. In *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*. ACM, 2008.
- [4] Henriette Cramer, Jean Garcia-Gathright, Sravana Reddy, Aaron Springer, and Romain Takeo Bouyer. Translation, tracks & data: an algorithmic bias effort in practice. In *Extended Abstracts of the 2019 CHI Conf. on Human Factors in Computing Systems*, 2019.
- [5] Andres Ferraro. Music cold-start and long-tail recommendation: bias in deep representations. In *Proceedings of the 13th ACM Conf. on Recommender Systems*, 2019.
- [6] Andre Holzapfel, Bob Sturm, and Mark Coeckelbergh. Ethical dimensions of music information retrieval technology. *Transactions of the International Society for Music Information Retrieval*, 2018.
- [7] Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE International Conf. on Data Mining*, 2008.
- [8] Peter Knees and Markus Schedl. A survey of music similarity and recommendation from music context data. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 2013.
- [9] Brian McFee, Thierry Bertin-Mahieux, Daniel PW Ellis, and Gert RG Lanckriet. The million song dataset challenge. In *Proc. of the 21st International Conf. on World Wide Web (WWW)*, 2012.
- [10] James McInerney, Benjamin Lacker, Samantha Hansen, Karl Higley, Hugues Bouchard, Alois Gruson, and Rishabh Mehrotra. Explore, exploit, and explain: personalizing explainable recommendations with bandits. In *Proc. of the 12th ACM Conf. on Recommender Systems (RecSys)*, 2018.
- [11] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proc. of the 27th ACM International Conf. on Information and Knowledge Management*, 2018.
- [12] Sergio Oramas, Oriol Nieto, Mohamed Sordo, and Xavier Serra. A deep multimodal approach for cold-start music recommendation. In *Proc. of the 2nd Workshop on Deep Learning for Recommender Systems*, 2017.
- [13] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskis. Music recommender systems. In *Recommender systems handbook*, pages 453–492. Springer, 2015.
- [14] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval (IJMIR)*, 2018.
- [15] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. Deep content-based music recommendation. In *Advances in neural information processing systems*, 2013.