

Towards Effective Exploration/Exploitation in Sequential Music Recommendation

Himan Abdollahpouri
DePaul University
USA
habdolla@depaul.edu

Steve Essinger
Pandora Media, Inc.
USA
sessinger@pandora.com

ABSTRACT

Music streaming companies collectively serve billions of songs per day. Radio-based music services may intersperse audio advertisements among the songs as a means to generate revenue, much like traditional FM radio. Regardless of the monetization approach, the recommender system should decide when to play content that the listener is known to enjoy (exploit) and content that is novel to the listener (explore). Recommender systems that rely on this explore/exploit type framework have been deployed in a wide variety of applications such as movies, books, music, shopping and more. In this work, we investigate the impact of different ad/song sequences on listener behavior. In particular, we focus on the impact of exploring new song content for the listener given the previous sequence of ads and songs in the listener’s session. Our results show that the prior sequence matters when considering song exploration and that this prior sequence has an impact on the listener’s tendency to interrupt their current session.

1 INTRODUCTION

Recommender systems (RS) have been deployed in numerous domains including music, movies, e-commerce and books. In music recommendation, one of the overarching goals of the RS is to find the best song to play for each listener, personalized to their specific taste(s) in music. In general, companies offering music recommendation services provide two different types of subscriptions: (1) Ad-supported membership where the music is free, but the listener is subject to advertisements and (2) premium membership where listener pays a monthly membership fee in exchange for ad-free listening. This paper focuses on the former, ad-supported listening. Unsurprisingly, listeners prefer hearing songs over ads. However, the business depends on the revenue that it makes from the ads and cannot operate without serving them. Therefore, playing ads is crucial to keep the business alive and should be considered as a content served to the listener along with music.

One of the fundamental concepts in RS is the idea of exploration and exploitation [9]. This paradigm results in a balance between recommending content the system has high certainty the user would like (exploitation) and the content for which there is less certainty (exploration). Without exploration, users would become stuck in a filter bubble and continue to see a narrow set of products. This is a missed opportunity to experience other products that could be of interest to them [5, 8]. Another reason for exploration is when the number of items matching a user’s interest is limited and the system should not recommend the same item again to the user. For example, in online dating [7], it is possible that the system

has already recommended all the available people who match the user’s interest and, so, exploring a wider range of people is needed in order to be able to generate new recommendations. Therefore, providing exploratory content to a user is a key component for discovery. We conducted an experiment on a music recommendation application and our results show that the previous sequence of events in a listener’s session is important in deciding whether the RS should provide subsequent exploratory types of content.

2 SONG/AD SEQUENCE ANALYSIS

We have compiled data from a large-scale music recommendation service for our analysis. To find the effect of different sequences of songs and ads on the probability of a user switching the station after listening to an exploratory song, we looked at one million sessions on mobile devices where the ad placement had been made completely at random. Note that the randomness of ad placement is important in order to make sure our analysis is not biased toward any particular ad placement algorithm. We compare the impact of explore songs versus exploit songs in the context of the previous three events. For example, given the prior three events Ad, Song, Song, where each song is an exploit, what is the probability of the listener changing the station if the next song spun for them is an explore song versus the probability of station change given an exploit song? Station change is used as a proxy for discontent with the current stream of music.

We calculated the probabilities of users changing the station when they are exposed to different sequences of ads and songs as

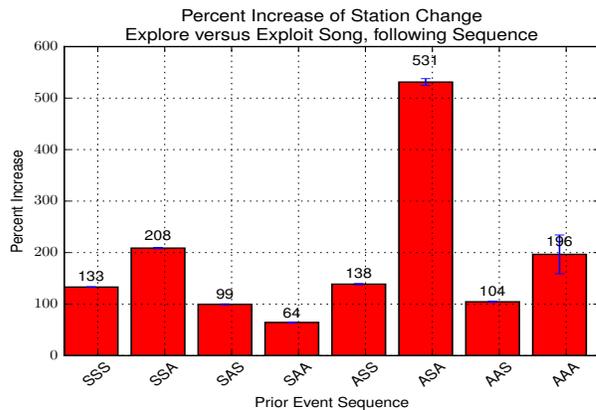


Figure 1: Percent increase of the probability of station change for an explore song vs. an exploit song, following different sequences of exploit songs and ads.

follows: there are a total of 8 possible event combinations for a set of three items as shown in figure 1. We denote explore song by S' and exploit song by S . *Station change* is represented by, C .

$$\text{Percentage Difference} = \frac{P(C | S') - P(C | S)}{P(C | S)} * 100 \quad (1)$$

$P(C | S')$ is the probability of a user changing the station given the last played content is an explore song. $P(C | S)$ is the probability of a user changing the station when the last played content is an exploit song. The lower and upper confidence bounds for the computed percentage increases, shown in figure 1 as vertical blue lines on top of the bars, are computed as follows,

$$\left[\frac{P(C | S') - P(C | S)}{P(C | S)} \pm 1.96 * SE \right] * 100 \quad (2)$$

where SE (i.e. the standard error) is calculated using equation 3,

$$\sqrt{\frac{P(C | S') * (1 - P(C | S'))}{N_{S'}} + \frac{P(C | S) * (1 - P(C | S))}{N_S}} \quad (3)$$

where $N_{S'}$ is the total number of times an explore song has been played. The total number of times an exploit song has been played is denoted by N_S . Figure 1 shows the percent increase of station changes after playing an explore versus an exploit song when a user has observed the respective prior sequence of exploit songs and ads. Due to the company's data privacy policy, we have not included the individual probabilities of switching the station for explore and exploit songs, but have provided the probability difference of change.

An exploit song is denoted by S and an ad is shown by A . As you can see, depending upon the previous sequence of songs and ads, the probability of a user switching the station when we show them an explore song is higher than the probability when we show an exploit song. This is true for all 8 different combinations of songs and ads. Moreover, some sequences are riskier than the others for placing an explore song. For example, the *ASA* sequence (which means playing an ad, then a song and then another ad) has the highest probability increase (+531.13%) of a user switching the station when given an explore song after that sequence. Clearly, this is not the best opportunity to explore new content. On the other hand, the *SAA* sequence has the lowest probability increase (+64.42%), but is still positive. While playing an explore song is still riskier than an exploit song in all cases, it is better to explore after particular sequences over others. Certainly, different sequences of songs and ads have different effects on station switching behavior and a recommender system should try to take these sequences into account when doing exploration and exploitation, as in our sequential music recommendation system. Overarching, instead of a blind explore-exploit platform, we advise taking an intelligent approach that accounts for a listener's state of listening (whether they are happy with the past couple of songs/ads or not) into account when deciding to exploit or when to explore.

3 RELATED WORK

The idea of explore-exploit has been studied in recommender systems by some researchers [9]. In particular, for single item recommendation, approaches like Multi-Armed Bandits have been

used to make a balance between exploration and exploitation [10]. Moreover, authors have previously proposed an approach for an effective balance between recommending popular and long-tail items [2]. A more similar idea to our work is done in [6] where authors investigated a proper timing for delivering the recommendation. However, in our work, we are not looking for a perfect timing for the recommendation in general as the user always should receive a content (song or ad) as recommendation. Our work is also novel as we look at the previous sequences of the recommendations as an indication for whether it is a good time for exploration or not.

4 CONCLUSION AND FUTURE WORK

In this work, we investigated the impact of different ad/song sequences on listener behavior. In particular, we focused on the impact of exploring new song content for the listener given the previous set of ads and songs in the listener's session. Our experimental results show that the previous sequence of ads/songs matters in deciding what the right time is for exploration versus exploitation. For our future work, we will launch an A/B experiment controlling for the placement of explore songs and see how different users behave when they observe different sequences of songs and ads. We will also investigate more sophisticated offline models, such as HMMs and RNNs in a reinforcement learning setting that could learn superior personalized playlist sequencing. This work is a starting point for a larger project in which we aim to optimize the stream of recommendations of mixed types of content (i.e. contents from different stakeholders) [1, 3, 4].

ACKNOWLEDGMENTS

We would like to thank Pandora Media, Inc. for access to their vastly rich dataset.

REFERENCES

- [1] Himan Abdollahpouri, Robin Burke, and Mobasher Bamshad. 2017. Recommender systems as multi-stakeholder environments. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP2017)*. ACM.
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling Popularity Bias in Learning to Rank Recommendation. In *Proceedings of the 11th ACM conference on Recommender systems*. ACM, To appear.
- [3] Himan Abdollahpouri and Steve Essinger. 2017. Multiple stakeholders in a music recommender system. In *1st International Workshop on Value-Aware and Multistakeholder Recommendation at RecSys 2017*.
- [4] Robin Burke and Himan Abdollahpouri. 2017. Patterns of Multistakeholder Recommendation. In *1st International Workshop on Value-Aware and Multistakeholder Recommendation at RecSys 2017*.
- [5] Oscar Celma. 2016. The Exploit-Explore Dilemma in Music Recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 377–377.
- [6] Nofar Dali Betzalel, Bracha Shapira, and Lior Rokach. 2015. Please, not now!: A model for timing recommendations. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 297–300.
- [7] Luiz Pizzato, Tomek Rej, Thomas Chung, Irena Koprinska, and Judy Kay. 2010. RECON: a reciprocal recommender for online dating. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 207–214.
- [8] Paul Resnick, R Kelly Garrett, Travis Kriplean, Sean A Munson, and Natalie Jomini Stroud. 2013. Bursting your (filter) bubble: strategies for promoting diverse exposure. In *Proceedings of the 2013 conference on Computer supported cooperative work companion*. ACM, 95–100.
- [9] Hastagiri P Vanchinathan, Isidor Nikolic, Fabio De Bona, and Andreas Krause. 2014. Explore-exploit in top-n recommender systems via gaussian processes. In *Proceedings of the 8th ACM Conference on Recommender systems*. ACM, 225–232.
- [10] Xinxi Wang, Yi Wang, David Hsu, and Ye Wang. 2014. Exploration in interactive personalized music recommendation: a reinforcement learning approach. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 11, 1 (2014), 7.