EXPLORING POPULARITY BIAS IN MUSIC STEAMING SERVICES

Vera Crabtree
Ithaca College
vcrabtree@ithaca.edu

Sean McQuillan
Ithaca College
smcquillan@ithaca.edu

Douglas Turnbull
Ithaca College
dturnbull@ithaca.edu

Motivation: Is there popularity bias on commercial music streaming sites?

We have long known that music consumption follows a long-tail distribution in which a small group of artists receive the vast majority of the attention from listeners. Is this due to popularity bias? Or is there another reason for this inequality found on music streaming services?

Simulated User Experiment

- We created a set U of twelve simulated users $u \in U$ based on real user data from the **LFM-1B-Subset**
- We randomly-selected four users within each of the three subgroups
  - Low mainstream users
  - Medium mainstream users
  - High mainstream users
- For each of our three streaming music services, we created a new simulated user account
- From each of these accounts, we "follow" or "like" the top-ten artists for the user
- We then play each of the artists' top song once all the way through
- The next day, in a round robin style, we record a set of 10 top recommended artists for each of the generated mixes
- We then calculate the difference in popularity between the inputted artist and the outputted recommendation

Group Average Popularity (ΔGAP)

To measure the change of the average popularity between the user profile inputs and the recommendation outputs, we compute ΔGAP, which is defined as the popularity lift in artists recommended over artists in the user profiles.

We would expect ΔGAP to be 0 when the average popularity of the artists recommended by the music service is equal to the average popularity of artists the users listen to on that service.

$$\Delta GAP(U) = \frac{GAP(U)_R - GAP(U)_P}{GAP(U)_P}.$$

Results

- We had expected to find evidence to support the hypothesis that personalized music recommendation plays a role in the accelerating rich-get-richer phenomenon for music consumption
- However, our results did not find a evidence of popularity bias in music recommendations from three popular commercial streaming services
- Popularity bias on commercial streaming services may result from the inclusion (or exclusion) of songs on popular human curiated playlists

ϕ_S - Spotify Popularity			
	Spotify	Amazon	YouTube
Overall ΔGAP	0.00	-0.13	0.06
Low MS ΔGAP	0.00	-0.29	0.10
Medium MS ΔGAP	0.02	-0.07	0.11
High MS ΔGAP	-0.01	-0.05	-0.01
ϕ_{LFM} - LFM-1B-Subset Popularity			
	Spotify	Amazon	YouTube
Overall ΔGAP	-0.22	-0.32	-0.12
Low MS ΔGAP	-0.37	-0.74	0.10
Medium MS ΔGAP	-0.33	-0.21	-0.14
High MS ΔGAP	-0.10	-0.26	-0.19

Acknowledgements

This research was supported by NSF grant IIS-1901330/1901168 and NEH grant HAA-280975-21.

We would also like to thank the anonymous ISMIR reviewers for their highly constructive feedback.