Algorithmic Collective Action in Recommender Systems: Promoting Songs by Reordering Playlists

Joachim Baumann[1,*], Celestine Mendler-Dünner[2]

[1]University of Zurich, Zurich University of Applied Sciences

[2]ELLIS Institute, Tübingen, Max Planck Institute for Intelligent Systems, Tübingen and Tübingen AI Center

Abstract

We investigate algorithmic collective action in transformer-based recommender systems. Our use case is a collective of fans aiming to promote the visibility of an artist by strategically placing one of their songs in the existing playlists they control. The success of the collective is measured by the increase in test-time recommendations of the targeted song. We introduce an easily implementable strategy towards this goal and test its efficacy on a publicly available recommender system model used in production by a major music streaming platform. Our findings reveal that even small collectives (controlling less than 0.01% of the training data) can achieve up to $25\times$ amplification of recommendations by strategically choosing the position at which to insert the song. Further, we find that the strategy only minimally impairs user experience; recommendations of other songs are largely preserved, and newly gained recommendations are taken from diverse songs of varying popularity levels. Taken together, our findings demonstrate how algorithmic collective action can be effective while not necessarily being adversarial, raising new questions around fairness, incentives, and social dynamics in recommender systems.

Keywords

music recommendation, collective action, power dynamics, transformer models, participatory AI

1. Motivation

In the ever-evolving landscape of music discovery, the challenge of sifting through the overwhelming number of tracks released daily has become increasingly difficult for both platforms and streamers. This has resulted in a strong dependence on platforms like Spotify, Deezer, or Apple Music, which distribute and promote music through song recommendations. These systems rely on historical data to learn user preferences and predict future content consumption [1, 2, 3, 4, 5].

It has been widely documented that music recommendation systems suffer from popularity bias as they tend to concentrate recommendation exposure on a limited fraction of artists, often overlooking new and emerging talent [6, 7, 8, 9, 10, 11]. As the success and visibility of artists are deeply influenced by the algorithms of these platforms, this can lead to a considerable imbalance in the music industry [12, 13] and reinforce existing inequalities [14, 15]. As a result, artists have started to fight for more transparency and fairer payments for online streaming

EWAF'24: European Workshop on Algorithmic Fairness, July 01–03, 2024, Mainz, Germany

[*]Work conducted during a research internship at the Max-Planck Institute for Intelligent Systems, Tübingen.

baumann@ifi.uzh.ch (J. Baumann); celestine@tue.ellis.eu (C. Mendler-Dünner)

0000-0002-0877-7063 (J. Baumann); 0000-0002-9880-7173 (C. Mendler-Dünner)

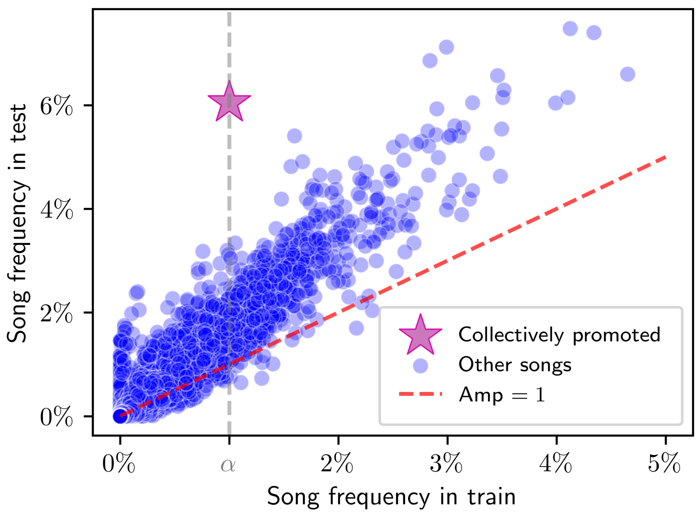

Figure 1: By strategically choosing the position at which to insert a song in a playlist, collectives can achieve a disproportional amplification in recommendation frequency relative to training set occurrences, compared to other songs.

services. The "Justice at Spotify" campaign, launched by the Union of Musicians and Allied Workers [16], has been signed by more than 28,000 artists. At the same time, the International Society for Music Information Retrieval has been arguing for promoting the discovery of less popular artists by recommending 'long-tail' items [17], as have other researchers [18, 19, 20].

2. Proposed strategy

We explore ***algorithmic collective action*** as a means for emerging artists to gain exposure in ML-powered recommender systems by mobilizing their fan base. Algorithmic collective action [21] describes the coordinated effort of a group of individuals to strategically report the part of the training data they control in order to impact the outcomes of a learning algorithm.

In music recommender systems the training data consists of user-generated playlists p. Each playlist is composed of an ordered list of songs. We want collective action to preserve user experience. Thus, we design collective action strategies under an ***authenticity constraint***:

$$\text{A strategy } h : p \to p' \text{ is authentic iff:} \quad \mathrm{Lev}(p, h(p)) \leq 1$$

where Lev denotes the Levenshtein distance [22], also known as edit distance in information theory, counting the number of operations needed to transform one sequence into another.

We propose a concrete strategy that satisfies this constraint. Our strategy consists of inserting an agreed-upon target song s^* at a specific position in the playlist, as shown in Figure 2. The position i to insert the song s^* is chosen by identifying the least likely song s_0 among the songs in the playlist p and placing s^* right after s_0.

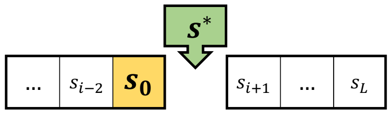

Figure 2: Song insertion strategy: placing the target song s^* right after s_0 in a playlist of length L.

Intuition for the strategy. Sequential recommenders are trained to approximate the conditional distributions of songs. For a given context window, the model then recommends one of the top K most likely songs to follow this context. Our strategy aims to exploit contexts that are overrepresented in the data the collective controls to increase the chance of meeting the top K threshold. To this end, it selects contexts that end on a low-frequency song (for small collectives, these typically appear only once in the controlled playlists). To find these low-frequency songs participants of the collective can share information about the playlists they own, gather stream counts by scraping Spotify playlists, or use public APIs to gather external song statistics.

Notice that the probabilistic assumption on the sequence model is not specific to the model architecture or the training algorithm used. This makes the strategy robust and easy to implement in practice.

3. Success of collective action

We empirically test the success of our strategy against a recent transformer-based automatic playlist continuation model [4]. The model has been deployed and made publicly available by Deezer—one of the biggest streaming platforms in the world. To train the model we use the Spotify million playlist dataset [23], treating each playlist as a user, and randomly sampling a small α-fraction to compose the collective. We find that by strategically placing the target song, small collectives can achieve disproportional representation at test time, see Figure 1. The star shows that a collective of size 1% can achieve that the target song is recommended in 6% of the playlist continuations at test time. This corresponds to a factor 6 amplification comparing training time and test time occurrences. In contrast, placing the song at the end of every playlist is largely ineffective. Also interesting to observe in Figure 1 is that a similar strategy does not seem to be implemented by any artist in the investigated data.

We further experiment with collective sizes from 0.03% to 3%, and show amplification in Figure 3. Interestingly, even tiny user collectives, controlling as few as 60 out of 1 million playlists can achieve an amplification of $25\times$. This is $40\times$ more than an average song occurring at the same frequency in the training data. Notably, this can be achieved by choosing the position at which to insert one song strategically, while leaving the rest of the playlist untouched.

4. Externalities

As we have seen, collective action offers an effective lever for platform participants to promote their interests on algorithm-driven platforms. However, strategies can only be effective if they are not creating equally strong incentives for other players in the system to counter them. In

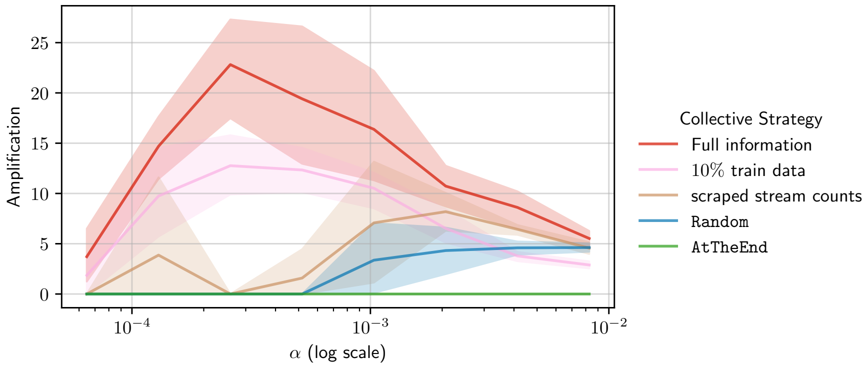

Figure 3: Empirical amplification of our proposed strategy with different levels of information (compared to two baselines `Random` and `AtTheEnd`).

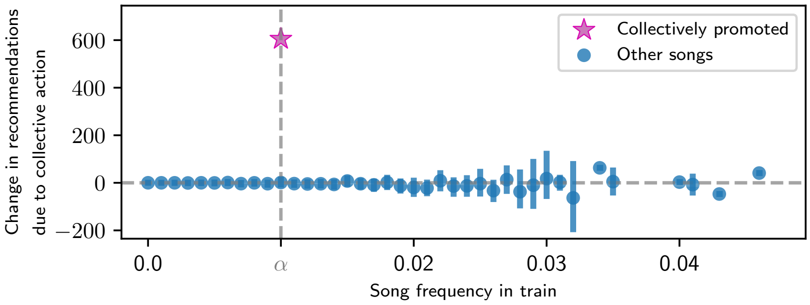

Figure 4: Impact of collective action on other songs when controlling 1% of training data ($\alpha = 1\%$). Songs are sorted by their training set frequency and aggregated into 50 evenly spaced bins with 95% CI.

Table 1

Effect on recommendation performance—measured using standard metrics [24]: Normalized Discounted Cumulative Gain (NDCG), R-precision, and number of clicks to find relevant song (#C).

	Performance	Performance loss Strategy		Performance loss
	No strategy	all	only participants	replace relevant
NDCG	0.29 ± 0.15	0.01 ± 0.01	0.00 ± 0.0	0.03 ± 0.04
R-precision	0.22 ± 0.11	0.01 ± 0.01	0.00 ± 0.0	0.03 ± 0.04
#C	2.52 ± 1.34	0.01 ± 0.01	-0.01 ± 0.1	0.03 ± 0.04

the following we study the externalities of our strategy, choosing $\alpha = 1\%$.

First, we focus on the effect of our strategy on other artists. In Figure 4, we visualize the change in total recommendation counts for individual songs, binned according to their training set frequency. The purple star indicates the song promoted by the collective. We find that the gained recommendations are taken from songs of varying popularity levels, and no artist appears to be affected disproportionally.

Second, we focus on the effect of our strategy on the recommendation performance. Table 4 shows performance along multiple metrics. The loss seems to be very small for the platform. In comparison, the last column shows an alternative strategy with the same success, but replacing a relevant song every time s^* is recommended. This shows that the gained recommendations often replace irrelevant songs causing relatively little harm. Similarly, we see little performance drop for the platform participants, suggesting that their recommendations are also widely preserved at test time. If the song were to be actually relevant for individuals in the collective such a strategy could even help increase recommendation performance.

5. Conclusion

We designed an easy-to-implement collective action strategy under a natural authenticity constraint. We demonstrated that it can be effective in promoting a target song even for tiny collectives, while minimally impairing overall user experience. This suggests a widely unexplored design space for effective collective action strategies that differ from typical adversarial data poisoning attacks [c.f. 25, 26, 27, 28]. They offer a powerful data lever [29, 30], and an approach to participatory AI [31]. Thus, understanding the role of economic power [32, 33], formalizing incentives [34], as well as quantifying long-term payoffs, dynamics and equilibria under collective action promises to be a fruitful direction for future work.[1]

Acknowledgments

We would like to thank Moritz Hardt for many insightful and formative discussions throughout the course of this work. We would also like to thank Mila Gorecki, Ricardo Dominguez-Olmedo, Ana-Andreea Stoica, and André Cruz for invaluable feedback on the manuscript, and Olawale Salaudeen, Florian Dorner, Stefania Ionescu, and Tijana Zrnic for helpful feedback on earlier versions of this work. We would also like to thank the anonymous reviewers for their feedback.

References

[1] C. Hansen, C. Hansen, L. Maystre, R. Mehrotra, B. Brost, F. Tomasi, M. Lalmas, Contextual and sequential user embeddings for large-scale music recommendation, in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020, p. 53–62.

[2] F. Tomasi, J. Cauteruccio, S. Kanoria, K. Ciosek, M. Rinaldi, Z. Dai, Automatic music playlist generation via simulation-based reinforcement learning, in: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '23, 2023, p. 4948–4957.

[3] D. Moor, Y. Yuan, R. Mehrotra, Z. Dai, M. Lalmas, Exploiting sequential music preferences via optimisation-based sequencing, in: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, 2023, p. 4759–4765.

[4] W. Bendada, G. Salha-Galvan, T. Bouabça, T. Cazenave, A Scalable Framework for Automatic Playlist Continuation on Music Streaming Services, in: Proceedings of the 46th

[1] An extended version of the paper can be found on arXiv [35].

International ACM SIGIR Conference on Research and Development in Information Retrieval, 2023, pp. 464–474.
[5] W. Bendada, T. Bontempelli, M. Morlon, B. Chapus, T. Cador, T. Bouabça, G. Salha-Galvan, Track Mix Generation on Music Streaming Services Using Transformers, in: Proceedings of the 17th ACM Conference on Recommender Systems, 2023, pp. 112–115.
[6] P. M. Napoli, Requiem for the long tail: Towards a political economy of content aggregation and fragmentation, International Journal of Media & Cultural Politics 12 (2016) 341–356.
[7] C. Bauer, M. Kholodylo, C. Strauss, Music recommender systems: challenges and opportunities for non-superstar artists, in: 30th Bled eConference, 2017, pp. 21–32.
[8] C. Bauer, Allowing for equal opportunities for artists in music recommendation, in: Proceedings of the 1st Workshop on Human-Centric Music Information Research Systems, 2019, pp. 16–18.
[9] M. P. Coelho, J. Z. Mendes, Digital music and the "death of the long tail", Journal of Business Research 101 (2019) 454–460.
[10] E. Blake, Data shows 90 percent of streams go to the top 1 percent of artists, 2020. https://www.rollingstone.com/pro/news/top-1-percent-streaming-1055005.
[11] D. Jannach, L. Lerche, F. Gedikli, G. Bonnin, What recommenders recommend – an analysis of accuracy, popularity, and sales diversity effects, in: User Modeling, Adaptation, and Personalization, Springer Berlin Heidelberg, 2013, pp. 25–37.
[12] L. Aguiar, J. Waldfogel, Platforms, power, and promotion: Evidence from spotify playlists, The Journal of Industrial Economics 69 (2021) 653–691.
[13] R. Prey, M. Esteve Del Valle, L. Zwerwer, Platform pop: disentangling spotify's intermediary role in the music industry, Information, Communication & Society 25 (2022) 74–92.
[14] L. Porcaro, C. Castillo, E. Gómez, Diversity by design in music recommender systems, Transactions of the International Society for Music Information Retrieval (2021).
[15] T. Tofalvy, J. Koltai, "Splendid Isolation": The reproduction of music industry inequalities in Spotify's recommendation system, New Media & Society 25 (2023) 1580–1604.
[16] Union of Musicians and Allied Workers, Justice at Spotify, 2021. URL: https://weareumaw.org/justice-at-spotify.
[17] C. Bauer, Report on the ISMIR 2020 special session: how do we help artists?, ACM SIGIR Forum 54 (2020).
[18] Ò. Celma, Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space, Springer Berlin, Heidelberg, 2010.
[19] D. Turnbull, L. Barrington, G. Lanckriet, Five approaches to collecting tags for music, in: ISMIR, 2008.
[20] S. Craw, B. Horsburgh, S. Massie, Music recommendation: Audio neighbourhoods to discover music in the long tail, in: Case-Based Reasoning Research and Development, Springer International Publishing, 2015, pp. 73–87.
[21] M. Hardt, E. Mazumdar, C. Mendler-Dünner, T. Zrnic, Algorithmic Collective Action in Machine Learning, in: Proceedings of the 40th International Conference on Machine Learning, volume 202, PMLR, 2023, pp. 12570–12586.
[22] V. I. Levenshtein, Binary codes capable of correcting deletions, insertions and reversals, Soviet Physics Doklady 10 (1966) 707.
[23] C.-W. Chen, P. Lamere, M. Schedl, H. Zamani, Recsys Challenge 2018: Automatic Music

Playlist Continuation, in: Proceedings of the 12th ACM Conference on Recommender Systems, RecSys ’18, 2018, pp. 527–528.
[24] H. Zamani, M. Schedl, P. Lamere, C.-W. Chen, An Analysis of Approaches Taken in the ACM RecSys Challenge 2018 for Automatic Music Playlist Continuation, ACM Trans. Intell. Syst. Technol. 10 (2019).
[25] Z. Tian, L. Cui, J. Liang, S. Yu, A comprehensive survey on poisoning attacks and countermeasures in machine learning, ACM Comput. Surv. 55 (2022).
[26] H. Zhang, Y. Li, B. Ding, J. Gao, Practical data poisoning attack against next-item recommendation, in: Proceedings of The Web Conference 2020, 2020, p. 2458–2464.
[27] Z. Yue, Z. He, H. Zeng, J. McAuley, Black-box attacks on sequential recommenders via data-free model extraction, in: Proceedings of the 15th ACM Conference on Recommender Systems, 2021, p. 44–54.
[28] Z. Yue, H. Zeng, Z. Kou, L. Shang, D. Wang, Defending substitution-based profile pollution attacks on sequential recommenders, in: Proceedings of the 16th ACM Conference on Recommender Systems, 2022, p. 59–70.
[29] N. Vincent, B. Hecht, S. Sen, “Data Strikes”: Evaluating the Effectiveness of a New Form of Collective Action Against Technology Companies, in: The World Wide Web Conference, 2019, pp. 1931–1943.
[30] N. Vincent, H. Li, N. Tilly, S. Chancellor, B. Hecht, Data Leverage: A Framework for Empowering the Public in Its Relationship with Technology Companies, in: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’21, 2021, pp. 215–227.
[31] A. Birhane, W. Isaac, V. Prabhakaran, M. Diaz, M. C. Elish, I. Gabriel, S. Mohamed, Power to the people? opportunities and challenges for participatory ai, in: Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, 2022.
[32] M. Hardt, M. Jagadeesan, C. Mendler-Dünner, Performative Power, in: Advances in Neural Information Processing Systems, 2022.
[33] M. Hardt, C. Mendler-Dünner, Performative prediction: Past and future, arXiv preprint arXiv:2310.16608 (2023).
[34] M. Olson, The logic of collective action: public goods and the theory of groups., Harvard University Press, 1965.
[35] J. Baumann, C. Mendler-Dünner, Algorithmic collective action in recommender systems: Promoting songs by reordering playlists, ArXiv preprint arXiv:2404.04269 (2024).