

ALLOWING FOR EQUAL OPPORTUNITIES FOR ARTISTS IN MUSIC RECOMMENDATION: A POSITION PAPER

Christine Bauer

Johannes Kepler University Linz
Institute of Computational Perception
christine.bauer@jku.at

ABSTRACT

Promoting diversity in the music sector is widely discussed on the media. While the major problem may lie deep in our society, music information retrieval contributes to promoting diversity or may create unequal opportunities for artists. For example, considering the known problem of popularity bias in music recommendation, it is important to investigate whether the short head of popular music artists and the long tail of less popular ones show similar patterns of diversity—in terms of, for example, age, gender, or ethnic origin—or the popularity bias amplifies a positive or negative effect.

I advocate for reasonable opportunities for artists—for (currently) popular artists and artists in the long-tail alike—in music recommender systems. In this work, I represent the position that we need to develop a deep understanding of the biases and inequalities because it is the essential basis to design approaches for music recommendation that provide reasonable opportunities. Thus, research needs to investigate the various reasons that hinder equal opportunity and diversity in music recommendation.

1. INTRODUCTION

Creating and maintaining diversity is an important and widely discussed topic in our society [27]. Thereby the debates on diversity are dominated by the challenges in promoting diversity as our society is prone to lay ground for unequal opportunities with respect to, for example, age, disability, gender, ethnic origin, religion, or sexual orientation throughout our society.

The issue of unequal opportunities is also relevant and a highly topical subject in the music sector. Some people voiced their concerns that there is a general discrimination of female artists [3, 16, 20, 24]. A similar inequality problem exists with respect to the little representation of black artists (especially black female artists) in high-popularity playlists on online music platforms [19, 20].

While the major problem may lie far beneath online music platforms or the music sector at large, the vast possibil-

ities of music information retrieval and recommendation may contribute tremendously in promoting diversity, inclusion, and equity—but may also be used to (intentionally or unintentionally) create unreasonable imbalances.

For instance, it is widely known that algorithms used for music recommendation are frequently prone to popularity bias [12]. This is a burden to inclusion, as such algorithms prioritize popular items and almost disregard the long tail of less popular items. In other words, the spectrum of suggested items is limited to a proportionally small set of items. As popularity bias is a common phenomenon in algorithmic filtering, research came up with diversity measures [15, 23] and there are various attempts to introduce diversity to recommendation algorithms [5, 6]. Studies (e.g., [9]) have shown that an increase in diversity has a positive effect on user experience, while the ideal degree of diversity may depend on user characteristics [10, 13, 21].

I postulate that we need to develop a deep understanding of the biases and inequalities because it is the essential basis to design approaches for music recommendation that are free from undesired biases and inequalities.

When we take a human-centric approach to music information retrieval (MIR), we need to consider all kinds of roles involved in MIR—not just the user. In this work, I put the—previously neglected—artists’ perspective in the loop. With the goal to provide reasonable opportunities for artists—for (currently) popular artists and artists in the long-tail alike—in music recommender systems, I take the position that research needs to investigate the various reasons that hinder equal opportunity and diversity in music recommendation.

This position paper is structured as follows: Section 2 presents the complexity of bias in music recommendation. Section 3 puts the artists’ perspective into the loop. Section 4 presents the fundamental research questions that have to be addressed to allow for equal opportunity for artists and promoting diversity in music consumption.

2. THE COMPLEXITY OF BIAS

Music recommendation relies on algorithmic decision-making. And an emerging body of literature has shown that algorithmic decision-making can go wrong in multiple ways [25], due to algorithmic problems, data sparsity, or actors gaming the system (e.g., via click manipulation). Typical problems include popularity bias, cold start prob-



lem, shilling attacks, grey-sheep problem, synonymy, as well as scalability and latency problems [14]. This leads to severe problems for society—from filter bubbles [18] to the reproduction and amplification of stereotypes and discrimination [22] to cognitive bias and humans’ overconfidence in algorithmic results [11]. Addressing these problems, there is a growing body of literature on fairness, accountability, and transparency in machine learning and artificial intelligence [2, 7, 17].

Still, while some aspects of bias in data and algorithms are subject of interest in research and draw attention on the media (e.g., filter bubble and popularity bias), other biases are not addressed or may even not have been identified yet.

3. THE NEGLECTED ARTIST IN MUSIC INFORMATION RETRIEVAL

In the music information retrieval (MIR) community (and related communities), research on diversity typically takes the perspective of the system—for instance, to mitigate the cold-start problem [4]—or the user (here: music consumer)—to better meet user preferences [13, 26]. The perspective of the item suppliers is considered only occasionally. For instance, Reference [1] raise awareness that recommender systems in multi-stakeholder environments may be fair for one stakeholder while being unfair for other stakeholders. Reference [8] proposes an approach with the goal to provide all artists in a collection with the opportunity of being listened in recommendations. Taking a human-centric approach to MIR systems, the goal is to include the artists’ perspective in MIR research.

4. RESEARCH DIRECTIONS

Taking a human-centric perspective with the aim to allow for equal opportunity for artists and promoting diversity in music consumption, requires to address fundamental research questions concerning potential bias in current systems and, generally, in music consumption. For instance:

Research Question 1. *How is diversity in terms of, for example, age, disability, gender, ethnic origin, religion, or sexual orientation of artists represented in the long tail of the popularity distribution?*

How is diversity represented in the short head of popular artists?

How does the diversity in the long tail and the short head relate to each other, and to the entire population?

Research Question 2. *How does the popularity of music items reflect inherent user taste?*

How is the popularity of music items affected by what is offered on online music platforms, on playlist, in recommendations, in advertising, etc.?

Understanding bias is a prerequisite to address its various facets and mitigate them. One concrete research question could be formulated as follows:

Research Question 3. *What is the influence of using timbre of the singing voice for music recommendation on*

the artist gender distribution in recommended items?

If recommendations allow for little diversity in timbre, items will likely be sung by same-gender singers.

Overall, the goal of future work is to investigate the various facets reasons that hinder equal opportunity and diversity in music recommendation. A deep understanding of the biases and inequalities is the essential basis to design approaches for music recommendation that provide reasonable opportunities for artists—for (currently) popular artists and artists in the long-tail alike.

5. ACKNOWLEDGMENTS

This research is supported by the Austrian Science Fund (FWF): V579.

6. REFERENCES

- [1] Himan Abdollahpouri and Robin Burke. Multi-stakeholder recommendation and its connection to multi-sided fairness. In Robin Burke, Himan Abdollahpouri, Edward Malthouse, KP Thai, and Yongfeng Zhang, editors, *Proceedings of the Workshop on Recommendation in Multistakeholder Environments (RMSE 2019)*, number 2440 in CEUR Workshop Proceedings, Aachen, Germany, 2019. <http://ceur-ws.org/Vol-2440/paper3.pdf>.
- [2] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda. In *2018 CHI Conference on Human Factors in Computing Systems*, CHI’18, New York, NY, USA, 2018. ACM.
- [3] Pauwke Berkers and Julian Schaap. Gender inequality in metal music production. *Emerald Studies in Metal Music and Culture*, pages 145–149, 2018.
- [4] K. R. Bindu, Rhama Lalgudi Visweswaran, P. C. Sachin, Kundavai Devi Solai, and Soundarya Gunasekaran. Reducing the cold-user and cold-item problem in recommender system by reducing the sparsity of the sparse matrix and addressing the diversity-accuracy problem. In Nilesh Modi, Pramode Verma, and Bhushan Trivedi, editors, *Proceedings of International Conference on Communication and Networks*, pages 561–570, Singapore, 2017. Springer.
- [5] Keith Bradley and Barry Smyth. Improving recommendation diversity. In *Proceedings of the 12th National Conference in Artificial Intelligence and Cognitive Science*, AICS’01, pages 75–84, 2001.
- [6] Jinpeng Chen, Yu Liu, Jun Hu, Wei He, and Deyi Li. A novel framework for improving recommender diversity. In *Behavior and Social Computing*, pages 129–138, Cham, Germany, 2013. Springer.
- [7] Amitai Etzioni and Oren Etzioni. Incorporating ethics into artificial intelligence. *The Journal of Ethics*, 21(4):403–418, 2017.

- [8] Andres Ferraro. Music cold-start and long-tail recommendation: Bias in deep representations. In *Proceedings of the 13th ACM Conference on Recommender Systems*, RecSys'19, pages 586–590, New York, NY, USA, 2019. ACM.
- [9] Bruce Ferwerda, Mark P. Graus, Andreu Vall, Marko Tkalcić, and Markus Schedl. How item discovery enabled by diversity leads to increased recommendation list attractiveness. In *Proceedings of the Symposium on Applied Computing*, SAC'17, pages 1693–1696, New York, NY, USA, 2017. ACM.
- [10] Bruce Ferwerda, Andreu Vall, Marko Tkalcić, and Markus Schedl. Exploring music diversity needs across countries. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*, UMAP'16, pages 287–288, New York, NY, USA, 2016. ACM.
- [11] George Hurlburt. How much to trust artificial intelligence? *IT Professional*, 19(4):7–11, 2017.
- [12] Dietmar Jannach, Lukas Lerche, Fatih Gedikli, and Geoffroy Bonnin. What recommenders recommend: An analysis of accuracy, popularity, and sales diversity effects. In *21st International Conference on User Modeling, Adaptation, and Personalization*, UMAP'13, pages 25–37, Berlin Heidelberg, Germany, 2013. Springer.
- [13] Yucheng Jin, Nava Tintarev, and Katrien Verbert. Effects of individual traits on diversity-aware music recommender user interfaces. In *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*, UMAP'18, pages 291–299, New York, NY, USA, 2018. ACM.
- [14] Shah Khusro, Zafar Ali, and Irfan Ullah. Recommender systems: Issues, challenges, and research opportunities. In *Information Science and Applications*, ICISA'16, pages 1179–1189, Singapore, 2016. Springer.
- [15] Matevž Kunaver and Tomaž Požrl. Diversity in recommender systems: A survey. *Knowledge-Based Systems*, 123:154–162, 2017.
- [16] Naomi Larsson. Live music acts are mostly male-only. what's holding women back? *The Guardian*, October 2017. <https://www.theguardian.com/inequality/2017/oct/12/tonights-live-music-acts-will-mostly-be-male-only-whats-holding-women-back>.
- [17] Bruno Lepri, Nuria Oliver, Emmanuel Letouzé, Alex Pentland, and Patrick Vinck. Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4):611–627, 2018.
- [18] Eli Pariser. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group, London, United Kingdom, 2011.
- [19] Vaughn Schmutz and Alison Faupel. Gender and cultural consecration in popular music. *Social Forces*, 89(2):685–707, 2010.
- [20] Stacy L. Smith, Marc Choueiti, and Katherine Pieper. Inclusion in the recording studio?: Gender and race/ethnicity of artists, songwriters & producers across 600 popular songs from 2012–2017. Report, Annenberg Inclusion Initiative, 2018. <http://assets.uscannenberg.org/docs/inclusion-in-the-recording-studio.pdf>.
- [21] Nava Tintarev, Matt Dennis, and Judith Masthoff. Adapting recommendation diversity to openness to experience: A study of human behaviour. In Sandra Carberry, Stephan Weibelzahl, Alessandro Micarelli, and Giovanni Semeraro, editors, *User Modeling, Adaptation, and Personalization*, UMAP'13, pages 190–202, Berlin Heidelberg, Germany, 2013. Springer.
- [22] Nicol Turner Lee. Detecting racial bias in algorithms and machine learning. *Journal of Information, Communication and Ethics in Society*, 16(3):252–260, 2018.
- [23] Saúl Vargas and Pablo Castells. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the 5th ACM Conference on Recommender Systems*, RecSys'11, pages 109–116, New York, NY, USA, 2011. ACM.
- [24] Yixue Wang and Emőke-Ágnes Horváth. Gender differences in the global music industry: Evidence from musicbrainz and the echo nest. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13 of *ICWSM'19*, pages 517–526. AAAI, 2019.
- [25] Christine T. Wolf, Haiyi Zhu, Julia Bullard, Min Kyung Lee, and Jed R. Brubaker. The changing contours of “participation” in data-driven, algorithmic ecosystems: Challenges, tactics, and an agenda. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 377–384, New York, NY, USA, 2018. ACM.
- [26] Wen Wu, Li Chen, and Yu Zhao. Personalizing recommendation diversity based on user personality. *User Modeling and User-Adapted Interaction*, 28(3):237–276, August 2018.
- [27] Patrizia Zandoni, Maddy Janssens, Yvonne Benschop, and Stella Nkomo. Guest editorial: Unpacking diversity, grasping inequality: Rethinking difference through critical perspectives. *Organization*, 17(1):9–29, 2010.