



MuRS 2024: 2nd Music Recommender Systems Workshop

Andres Ferraro
andres.ferraro@siriusxm.com
SiriusXM
Oakland, CA, USA

Peter Knees
peter.knees@tuwien.ac.at
TU Wien
Vienna, Austria

Lorenzo Porcaro
lorenzo.porcaro@proton.me
Joint Research Centre, European Commission
Ispra, Italy

Christine Bauer
christine.bauer@plus.ac.at
Paris Lodron University Salzburg
Salzburg, Austria

ABSTRACT

Music recommendation has been relevant to the Recommender Systems (RecSys) community since the early days. With the growth of music streaming platforms, algorithmic recommendations have become critical in the music industry. However, many challenges are still wide open in the area of music recommender systems. Such challenges are currently being addressed in several research communities, including and beyond the RecSys and the Music Information Retrieval (MIR) communities. The RecSys conference has traditionally not focused very much on music content understanding. In contrast, while music content understanding is central to the MIR community, research on recommender systems is not prominent in MIR research. The Music Recommender Systems Workshop (MuRS) aims at bridging the existing gap between the diverse research communities focused on the specific challenges of music recommender systems. The workshop provides a space for researchers and practitioners from multiple disciplines to jointly discuss and exchange perspectives and solutions, and to promote discussion from both academia and industry upon future research directions in the area of music recommender systems.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Applied computing** → **Sound and music computing**.

KEYWORDS

music recommender systems, music information retrieval, music streaming

ACM Reference Format:

Andres Ferraro, Lorenzo Porcaro, Peter Knees, and Christine Bauer. 2024. MuRS 2024: 2nd Music Recommender Systems Workshop. In *18th ACM Conference on Recommender Systems (RecSys '24)*, October 14–18, 2024, Bari, Italy. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3640457.3687097>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

RecSys '24, October 14–18, 2024, Bari, Italy

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0505-2/24/10

<https://doi.org/10.1145/3640457.3687097>

1 WORKSHOP DESCRIPTION AND RATIONALE

Music recommendation has been a prominent use case in the RecSys community since the early days [4, 20]. With the growth of music streaming platforms in the last twenty years, algorithmic recommendations have become critical in the music industry. With tenths of millions of music items being readily available for consumers, recommender systems are essential in helping to reduce choice overload. Further, beyond assisting the listener in their music discovery, recommender systems have expanded to many aspects of the musical experience. A virtuous influential circle between the music industry and technological research drove improvements both in music listening experiences and in general scientific knowledge in recommender systems. Many fundamental topics in RecSys have matured together with their applications in music streaming (e.g., collaborative filtering, user modeling, etc.), while some distinctive aspects of the music medium (i.e., often consumed in a sequence, passively, re-recommendation possible, etc. [17]) drove their own specific topics, such as playlist generation [5] or next-song recommendation [22].

Today, music recommendation is a vibrant research area, prolific with respect to new topics [8, 17] that led to novel contributions to the RecSys community. For example, in 2022 one of the best paper awards focused on understanding ways that recommendation can help users by nudging their music exploration preference [14], and in 2021 among the winners of the Women in RecSys Journal Paper of the Year there is an overview on the role that diversity by design have in music recommender systems [15]. However, many challenges are wide open, e.g., how to make fair recommendations considering the interests of multiple stakeholders [7, 11], how to accurately identify and represent users' characteristics and their context [13], how diversity in recommendation may impact users' listening experience [2, 16, 23], or what roles will large language models play in music recommender systems [12]. Additionally, there are open questions regarding the potential biases of novel algorithmic approaches considered state-of-the-art [9, 10], leading to an undesired impact on what people listen to, negatively affecting both listeners and music artists. And, more recently, an important question was raised on how the systems can be aligned with normative principles in democratic societies¹ [9, 10].

Such challenges are being addressed in diverse research communities beyond RecSys, and notably in the Music Information

¹<https://musicairesearch.wordpress.com/projects/interdisciplinary-interventions/>

Retrieval (MIR) community. However, today there is **no forum where all these challenges are discussed jointly**. For instance, in the RecSys conference, music-specific systems are less studied compared to other domains, and music content understanding is not a focus. At the same time, in the MIR community, research on recommender systems is less present compared to other topics. This leaves a research gap between the two communities, e.g., for exploring methods that take full advantage of collaborative information and music content (such as audio, lyrics, scores, or biographies) for generating and evaluating the recommendations.

The workshop aims at bridging the existing gap between the different research communities focused on the specific challenges of music recommender systems. It provides a space for researchers and practitioners from multiple disciplines to jointly discuss and exchange perspectives and solutions, and to promote discussion from both academia and industry upon future research directions in the area of music recommender systems. With this, the workshop contributes to the continued development of the field of recommender systems. As the main program of RecSys rarely features papers from the music domain, MuRS is complementary to the topics of the main conference program. Further, with a panel, we particularly delve into the role and implications of music recommender systems as AI-generated music content proliferates.

Before MuRS, RecSys workshops on the topic proposed here date back to 2011–2012, and the latest effort at the RecSys conference to provide an exhaustive view on music recommendation dates to a 2017 tutorial [17]. Therefore, we believe that the second edition of MuRS offers a bridge between separate pockets of otherwise very related research.

2 WORKSHOP TOPICS

The workshop invited submissions on methods for music recommender systems that address but are not limited to:

- **Fundamentals**, such as bandits and reinforcement learning for music recommendation, sequential music recommendation, large language models for music recommendation, multi-stakeholder and multi-objective music recommendation, music representation learning and metric learning, automatic music tagging, methods that mitigate cold-start and popularity bias, content-based/hybrid methods that leverage multi-modal information, listener intent modeling (session-level, and long-term) and context-understanding, fairness, transparency, interpretability & explainability at scale, algorithmic bias, online and offline evaluation of music recommender systems, engineering aspects of music recommendation at very large scale, and user studies of music consumption
- **Applications**, such as playlist generation and continuation, algorithmic radio programming, visual recommendations and homepage personalization, music discovery, music search and browsing, conversational music interactions, virtual reality and music recommendation, music recommendation in social media, recommender systems in the live music industry, recommender systems for record labels, and recommender systems for music creation and generation

- **Societal aspects**, including cross-cultural and local music recommendation, socially-aware music recommender systems, studies of the societal impact of algorithmic music recommendation, and ethics of music recommender system development.

3 ORGANIZERS

The organizers bring complementary experience and perspectives, which makes them a particularly suitable consortium for the workshop topic.

Dr. Andrés Ferraro² (*corresponding organizer*, E-mail: andresferraro@acm.org) is a Research Scientist at SiriusXM. He completed his PhD at the Department of Information and Communication Technologies and Engineering of the Universitat Pompeu Fabra, Spain. His thesis uncovers multiple dimensions in which music recommender systems affect artists and proposes alternatives to mitigate such problems. After that, he was a Postdoctoral Fellow at McGill University and Mila (Quebec AI Institute), Canada, where he was part of an interdisciplinary project, rethinking music recommender systems by considering new and alternative conceptions from the social sciences and humanities, informed by non-profit systems and critical debates over bias and discrimination. He is co-organizer of LatAm Bish Bash,³ a series of meetings and networking events that connect engineers, researchers, and students working on music and audio signal processing. He is also an organizer of DEFIRST⁴, a reading group at Mila, focused on discussions around fairness and recommender systems research.

Dr. Lorenzo Porcaro⁵ is a Research Scientist at the European Commission's Joint Research Centre (JRC). His work focuses on assessing the impact that recommender systems have on their users, with a focus on human rights violations and discrimination. He holds a PhD in Information and Communication Technologies, and Master's degrees in Sound and Music Computing (M.Sc.) and Intelligent Interactive Systems (M.Sc.) from Universitat Pompeu Fabra (Spain). His research interests include recommender systems, social computing, human-computer interaction, and music information retrieval. He co-organized the tutorial *Trustworthy MIR: Creating MIR applications with values*, and the 2nd Workshop on Human-Centric Music Information Research (HCMIR23), co-located respectively at ISMIR 2022 and 2023. He is also the co-founder of *Italian Music Tech* meetup, a series of networking events for Italian-speaking music technology enthusiasts. Recently, he received a postdoctoral fellowship funded under the European Commission's Marie Skłodowska-Curie Actions, during which he will work on a project titled "Algorithmic Auditing for Music Discovery (AA4MD)" [6].

Dr. Peter Knees⁶ is an Associate Professor of the Faculty of Informatics of TU Wien, UNESCO Chair on Digital Humanism, and coordinator of the SIG Digital Humanism of TU Wien's Center for Artificial Intelligence and Machine Learning (CAIML). Previously, he was a University Assistant at JKU Linz and Visiting Assistant Professor at Georgia Institute of Technology. For two decades, he

²Web: <https://scholar.google.com/citations?user=WQglBowAAAAJ>

³<https://www.meetup.com/grupo-em-sao-paulo-de-audio-signal-processing>

⁴<https://www.youtube.com/channel/UCI9jOq79MHUD3vnh3jUUIrg>

⁵Web: <https://lorenzoporcaro.me>

⁶<https://scholar.google.com/citations?hl=en&user=MtyaO2cAAAAJ>

has been an active member of the Music Information Retrieval research community, reaching out to the related fields of multimedia and text information retrieval, recommender systems, and the digital arts. His research activities center on music search engines and interfaces as well as music recommender systems, and on smart(er) tools for music creation.

Apart from serving on the program committees of major conferences in the field, he has organized several workshops on the topics of media retrieval: AdMIRE: Advances in Music Information Research series 2009-2012; Adaptive Multimedia Retrieval 2010; Social Media Retrieval and Analysis (SIGIR 2014, ICDM 2015); Collaborating with Intelligent Machines: Interfaces for Creative Sound (CHI 2015); SOAP: Surprise, Opposition and Obstruction in Personalized and Adaptive Systems (UMAP 2016, 2017); MILC: Intelligent Music Interfaces for Listening and Creation (IUI 2018, 2019, 2023); 1st Austrian Workshop on Music Information Retrieval (2018); international Vienna Workshops on Digital Humanism 2019, 2020 (online), 2022; the RecSys Challenge 2019 and accompanying RecSys Workshop, and the first edition of MuRS at RecSys 2023.

Dr. Christine Bauer⁷ is a Professor of Interactive Intelligent Systems at the Department of Artificial Intelligence and Human Interfaces (AIHI) at the Paris Lodron University Salzburg (PLUS), Austria. Within her research field, she has focused on context-aware music recommender systems. Core interests are fairness in algorithmic decision-making and multi-method evaluations of recommender systems. She is very active in the RecSys community. Besides serving on the Editorial Board of ACM Transactions on Recommender Systems (TORS), she was the co-chair for the Doctoral Symposia at RecSys 2023 and 2021, and a speaker at the ACM Summer School on Recommender Systems in 2023 and 2019. Furthermore, she has co-organized the workshop series Perspectives on the Evaluation of Recommender Systems (PERSPECTIVES) at RecSys 2021–2023, the Workshop on Intelligent User-Adapted Interfaces: Design and Multi-Modal Evaluation (IUadaptMe 2019) at UMAP 2019, and Dagstuhl seminars related to the evaluation of recommender systems.

4 WORKSHOP FORMAT AND ACTIVITIES

4.1 Targeted audience and setting

The workshop is organized as a **full-day** event as part of the RecSys 2024 conference, following a **hybrid format** (i.e., both online and in-person at RecSys).

In general, the workshop resembles a mini-conference with a series of 10–15 minutes oral presentations of accepted papers, a keynote (20–30 minutes), a panel, and proceedings of accepted papers. Between each presentation, there is time for Q&A and discussion. Additionally, during breaks there is time for discussion between participants and speakers.

Besides the authors presenting their accepted papers, the workshop is open to all RecSys participants who register for workshop attendance.

Furthermore, to foster a diverse attendance with respect to locations and time zones, we intend to record the talks (after explicit speaker consent) and make them available online for a limited time.

⁷Web: <https://christinebauer.eu>

4.2 Invited speaker and panel

The workshop's **keynote speaker** is **Nancy Baym**⁸. Nancy Baym is a Principal Researcher at Microsoft Research New England. Her work focuses on interpersonal relationships and new technologies. A pioneer in the field of internet research, Baym wrote some of the first articles about online communities in the early 1990s. With Jean Burgess, she is the author of *Twitter: A Biography* (2020, NYU). Other books include *Playing to the Crowd: Musicians, Audiences, and the Intimate Work of Connection* (2018, NYU), *Personal Connections in the Digital Age* (2010, Second Edition 2014, Polity), *Internet Inquiry: Conversations About Method* (co-edited with Annette Markham, 2010, Sage), and *Tune In, Log On: Soaps, Fandom and Online Community* (2000, Sage).

Further, the program includes a **panel** discussion with experts from academia, industry, and policy, engaging in a dialogue about the implications of music recommender systems in regulatory terms, emphasizing the visibility and discoverability of music, as well as the ethical considerations surrounding AI-generated music content and its proliferation through recommender systems.⁹ Given the mounting scrutiny that such systems are undergoing from regulatory bodies and policy-makers within and beyond the realm of music, our objective is to facilitate a forum for discussion. In this forum, the panelists and attendees have the opportunity to delve into the technical intricacies of developing recommender systems and their potential societal impacts, both positive and negative. This paves the way for an exchange between practitioners, academics, industry professionals, and policy-makers.

5 SUMMARY OF THE FIRST EDITION OF MURS

The MuRS Workshop held at RecSys 2023 was the first edition of what we expect to be a prosperous series of workshops fully dedicated to music recommender systems. In the previous edition, we received a total of 15 submissions, of which 6 were accepted for inclusion in the proceedings (40%). Most of the authors were from the USA; other countries with several submissions are Austria, UK, and Japan. Submissions were fairly balanced from industry and academia. Beyond the authors of accepted papers and invited speakers, further RecSys participants (from both academia as well as industry) also joined the workshop and, with this, the workshop was well attended.

The main categories of works presented at the workshop were: (i) Investigating different biases from recommenders, and analyzing different behaviors in consumption and recommendations [1, 3, 21]; (ii) algorithmic focused research, discussing improvements for solving a technical problem such as sequential recommendation[18] or aligning recommendations with the user's personality traits[19]; and (iii) providing recommendations for the artists in the process of creation of the music and related content[24]

As the first edition of MuRS succeeded in bringing the sub-communities of RecSys and MIR together to jointly discuss the challenges of music recommender systems, MuRS stands on a solid basis. The organizers are committed to drawing a larger audience from those sub-communities.

⁸Web: <https://www.nancybaym.com/>

⁹https://www.europarl.europa.eu/doceo/document/A-9-2023-0388_EN.html

ACKNOWLEDGMENTS

This research was funded in whole, or in part, by the Austrian Science Fund (FWF) [P33526]. For the purpose of open access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission. This publication was supported by the Excellence in Digital Sciences and Interdisciplinary Technologies (EXDIGIT) project, funded by Land Salzburg under grant number 20204-WISS/263/6-6022. This work is partially supported by the HUMAINT programme (Human Behaviour and Machine Intelligence), Joint Research Centre, European Commission. The views expressed are purely those of the author and may not in any circumstances be regarded as stating an official position of the European Commission.

REFERENCES

- [1] Archana Anandakrishnan and Aaron Carter-Enyi. 2023. Afropop Locally and Globally: Top Artists, Countries and Microgenres. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo. <https://doi.org/10.5281/zenodo.8372158>
- [2] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic Effects on the Diversity of Consumption on Spotify. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 2155–2165. <https://doi.org/10.1145/3366423.3380281>
- [3] Christine Bauer and Andres Ferraro. 2023. Strategies for Mitigating Artist Gender Bias in Music Recommendation: A Simulation Study. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo, 5 pages. <https://doi.org/10.5281/zenodo.8372477>
- [4] Oscar Celma. 2010. *Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space* (1st ed.). Springer Berlin Heidelberg, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-13287-2>
- [5] Ching-Wei Chen, Paul Lamere, Markus Schedl, and Hamed Zamani. 2018. Recsys Challenge 2018: Automatic Music Playlist Continuation. In *Proceedings of the 12th ACM Conference on Recommender Systems (Vancouver, BC, Canada) (RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 527–528. <https://doi.org/10.1145/3240323.3240342>
- [6] CORDIS. 2024. Algorithmic Auditing for Music Discoverability (AA4MD). <https://doi.org/10.3030/101148443>
- [7] Karlijn Dinissen and Christine Bauer. 2022. Fairness in music recommender systems: A stakeholder-centered mini review. *Frontiers in Big Data* 5, Article 913608 (2022), 9 pages. <https://doi.org/10.3389/fdata.2022.913608>
- [8] Andres Ferraro. 2019. Music cold-start and long-tail recommendation: bias in deep representations. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 586–590.
- [9] Andres Ferraro, Gustavo Ferreira, Fernando Diaz, and Georgina Born. 2022. Measuring Commonality in Recommendation of Cultural Content: Recommender Systems to Enhance Cultural Citizenship. In *Proceedings of the 16th ACM Conference on Recommender Systems (Seattle, WA, USA) (RecSys '22)*. Association for Computing Machinery, New York, NY, USA, 567–572. <https://doi.org/10.1145/3523227.3551476>
- [10] Andreas Ferraro, Gustavo Ferreira, Fernando Diaz, and Georgina Born. 2024. Measuring Commonality in Recommendation of Cultural Content: Recommender Systems to Enhance Cultural Citizenship. *ACM Transactions on Recommender Systems* 2, 1 (2024). <https://doi.org/10.1145/3643138>
- [11] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. What is fair? Exploring the artists' perspective on the fairness of music streaming platforms. In *IFIP Conference on Human-Computer Interaction*. Springer International Publishing, Cham, Germany, 562–584. https://doi.org/10.1007/978-3-030-85616-8_33
- [12] Deepesh V. Hada, Vijaikumar M., and Shirish K. Shevade. 2021. ReXPlug: Explainable Recommendation Using Plug-and-Play Language Model. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (Virtual Event, Canada) (SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 81–91. <https://doi.org/10.1145/3404835.3462939>
- [13] Peter Knees, Markus Schedl, Bruce Ferwerda, and Audrey Laplante. 2019. *User awareness in music recommender systems*. De Gruyter Oldenbourg, Berlin, Boston, Chapter 9, 223–252. <https://doi.org/10.1515/9783110552485-009>
- [14] Yu Liang and Martijn C. Willemsen. 2022. Exploring the Longitudinal Effects of Nudging on Users' Music Genre Exploration Behavior and Listening Preferences. In *Proceedings of the 16th ACM Conference on Recommender Systems (Seattle, WA, USA) (RecSys '22)*. Association for Computing Machinery, New York, NY, USA, 3–13. <https://doi.org/10.1145/3523227.3546772>
- [15] Lorenzo Porcaro, Carlos Castillo, and Emilia Gómez. 2021. Diversity by Design in Music Recommender Systems. *Transactions of the International Society for Music Information Retrieval* (2021). <https://doi.org/10.5334/tismir.106>
- [16] Lorenzo Porcaro, Emilia Gómez, and Carlos Castillo. 2024. Assessing the Impact of Music Recommendation Diversity on Listeners: A Longitudinal Study. *ACM Transactions on Recommender Systems* 2, 1 (2024). <https://doi.org/10.1145/3608487>
- [17] Markus Schedl, Peter Knees, and Fabien Gouyon. 2017. New Paths in Music Recommender Systems Research. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (Como, Italy) (RecSys '17)*. Association for Computing Machinery, New York, NY, USA, 392–393. <https://doi.org/10.1145/3109859.3109934>
- [18] Pavan Seshadri and Peter Knees. 2023. Leveraging Negative Signals with Self-Attention for Sequential Music Recommendation. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo. <https://doi.org/10.5281/zenodo.8372449>
- [19] Yuuki Tachioka. 2023. Conditioning of variational autoencoder by user traits for item recommendation. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo. <https://doi.org/10.5281/zenodo.8372443>
- [20] Marco Tiemann and Steffen Pauws. 2007. Towards Ensemble Learning for Hybrid Music Recommendation. In *Proceedings of the 2007 ACM Conference on Recommender Systems (Minneapolis, MN, USA) (RecSys '07)*. Association for Computing Machinery, New York, NY, USA, 177–178. <https://doi.org/10.1145/1297231.1297265>
- [21] April Trainor and Douglas Turnbull. 2023. Popularity Degradation Bias in Local Music Recommendation. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo. <https://doi.org/10.5281/zenodo.8372473>
- [22] Andreu Vall, Massimo Quadrana, Markus Schedl, and Gerhard Widmer. 2019. Order, context and popularity bias in next-song recommendations. *International Journal of Multimedia Information Retrieval* 8, 2 (2019), 101–113.
- [23] Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the Guides: Disentangling Human and Algorithmic Curation in Online Music Consumption. In *Proceedings of the 15th ACM Conference on Recommender Systems (Amsterdam, Netherlands) (RecSys '21)*. Association for Computing Machinery, New York, NY, USA, 380–389. <https://doi.org/10.1145/3460231.3474269>
- [24] Alexander Williams, Stefan Lattner, and Mathieu Barthet. 2023. Sound-and-Image-informed Music Artwork Generation Using Text-to-Image Models. In *Music Recommender Systems Workshop (MuRS 2023)*. Zenodo. <https://doi.org/10.5281/zenodo.8372471>