MuRS: 3rd Music Recommender Systems Workshop

Andres Ferraro
Pandora/SiriusXM
Barcelona, Spain
andresferraro@acm.org

Lorenzo Porcaro
Sapienza University of Rome
Rome, Italy
lorenzo.porcaro@uniroma1.it

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

Abstract

Music recommendation has been a core area of interest within the recommender systems community since its early days. With the rise of music streaming platforms, algorithmic recommendations have become central to the music industry. However, critical challenges remain concerning the design, evaluation, and societal impact of music recommender systems. This third edition of the Music Recommender Systems Workshop (MuRS) centers on the growing influence of generative content on music recommendation. The rapid influx of AI-generated music transforms the streaming ecosystem, raising critical questions about discoverability, authenticity, and the curational role of recommender systems. The challenges and opportunities associated with AI-generated content extend beyond music and recommender systems, demanding transparent, fair, and accountable recommendation frameworks that reflect the interests of the diverse set of stakeholders.

CCS Concepts

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

Keywords

music recommender systems, music information retrieval, music streaming, generated content

1 Introduction

Background and Motivation. Music recommendation has been a core domain in the RecSys community since its inception [1, 15]. Over the last three decades, the rise of music streaming services has made algorithmic recommendation essential to the music industry. Research on recommender systems has evolved alongside these platforms, with new challenges continuously emerging.

The unique characteristics of music consumption—such as sequential listening, passive engagement, and repeated consumption—have inspired research into specialized topics like playlist generation [2], next-track prediction [17], and personalization [10].

Current trends include refining sequential recommendations [16], addressing algorithmic biases [4], and improving reproducibility through dataset comparisons [11].

Despite such advancements, many unresolved issues persist. These include aligning recommendations with multiple stakeholders' interests in a fair manner [3, 9] and assessing the impact of recommendation diversity on user experience [13, 14].

Pressing Challenges. A pressing new issue is the rise of AI-generated music and its implications for recommender systems.

Furthermore, concerns have been raised about the biases embedded in state-of-the-art recommendation algorithms [5]. These biases can shape what people listen to, potentially disadvantaging both listeners and artists. Researchers have recently begun questioning how recommender systems can align with broader societal and democratic values [6, 12].

Fragmentation Between Research Communities. Addressing such challenges requires interdisciplinary collaboration across multiple research domains and communities, particularly between Recommender Systems (RecSys) and Music Information Retrieval (MIR). However, these communities often operate in silos: RecSys has historically placed less emphasis on content understanding, while MIR focuses on audio and content analysis with limited engagement in recommendation. Bridging this divide is essential for advancing the field.

The MuRS Workshop Series. The MuRS workshop was launched to address this gap. The first two editions, held at RecSys 2023 [7] and 2024 [8] attracted balanced submissions from academia and industry. Topics covered included: (i) Biases in recommender systems and user behavior analysis; (ii) Algorithmic improvements (e.g., sequential recommendations, personality-aware models); (iii) Recommendations for artists in the process of creating music and related content; (iv) Music discoverability; (v) Multimodal music embeddings; (vi) Emotion prediction from music; (vii) Effects of negative feedback on diversity.

Goals for MuRS 2025. This third edition of MuRS focuses on the **growing impact of generative content on music recommendation**. The rapid influx of AI-generated music reshapes the streaming ecosystem, raising critical questions about discoverability, authenticity, and the curational role of recommender systems. While these challenges extend beyond music, these are a particularly pressing in the music domain. A dedicated forum is crucial to explore solutions, share insights, and develop frameworks that ensure recommender systems remain transparent, fair, and aligned

with user needs. This workshop provides a much-needed venue for addressing these pressing concerns, fostering interdisciplinary collaboration at a pivotal moment for the field.

2 Workshop Format

This half-day workshop follows a **hybrid format** (i.e., both online and in-person) and is held as part of the 19th ACM Conference on Recommender Systems (RecSys 2025) in Prague, Czech Republic.

The workshop resembles a mini-conference with a keynote presentation and a series of short oral presentations of accepted papers, each followed by a brief discussion. The keynote address is given by Cheng-Zhi Anna Huang, Massachusetts Institute of Technology (MIT), with the title "In Search of Human-AI Resonance".

We solicit the following types of paper contributions:

- *Regular papers* (max. 8 pages), which may include work in progress and preliminary results. Already (pre-)published results will also be considered.
- *Position papers* (max. 4 pages) describing original ideas, perspectives, research visions, and open challenges.

Submissions undergo single-anonymized peer review by at least two Program Committee (PC) members and are selected based on relevance to the workshop, quality, novelty, clarity, and discussion potential. Once the workshop is concluded, we will include accepted papers in the proceedings, which shall be submitted to ceur-ws.org[1] for online publication.

3 Workshop Topics

The workshop is open to research from various academic disciplines and industry. Topics of interest include, but are not limited to, the following music recommendation topics that address:

- Fundamentals
  - Sequential music recommendation
  - Bandits and reinforcement learning for recommendation
  - Large language models for recommendation
  - Multi-stakeholder and multi-objective music recommendation
  - Music representation learning and music similarity metric learning
  - Music content understanding and automatic tagging
  - Methods that mitigate cold-start and popularity bias
  - Content-based/Hybrid methods that leverage multi-modal information
  - Listener taste modeling
  - Listener intent modeling (session-level or long-term) and context understanding
  - Fairness, transparency, interpretability, & explainability at scale,
  - Algorithmic biases and fairness
  - Online and offline evaluation of music recommender systems
  - Engineering aspects of music recommendation at very large scale
  - User studies on music consumption
- Applications
  - Playlist generation and continuation
  - Algorithmic radio programming
  - Visual recommendations and homepage personalization
  - Music discovery
  - Music search and browsing
  - Conversational interaction with systems
  - Virtual reality and listening experiences
  - Music recommendation in social media
  - Recommender systems in the live music industry
  - Recommender systems for record labels
  - Recommender systems for music creation and generation
- Societal aspects
  - Cross-cultural music recommendation
  - Local music recommendation
  - Socially-aware music recommender systems
  - Studies of the societal impact of algorithmic music recommendation
  - Ethics of music recommender system

4 Program Committee

For the Program Committee (PC), we have invited recognized researchers working in fields related to the workshop's topic. Furthermore, the PC covers experience in academia, industry (e.g., streaming platforms), and public service (e.g., policy-oriented research) and embraces different seniority levels.

The Program Committee (PC) for MuRS 2025 includes:

- Darius Afchar (Deezer Research, France)
- Filippo Betello (Sapienza University of Rome, Italy)
- Clara Borrelli (Apple, USA)
- Fabien Gouyon (SiriusXM, USA)
- Jaehun Kim (SiriusXM, USA)
- Dominik Kowald (Know-Center and Graz University of Technology, Austria)
- M. Jeffrey Mei (SiriusXM, USA)
- Marta Moscati (Johannes Kepler University Linz, Austria)
- Bruno Massoni Sguerra (Deezer Research, France)
- Viet Anh Tran (Deezer Research, France)
- Doug Turnbull (Ithaca College, USA)
- João Vinagre (Joint Research Centre, European Commission)
- Eva Zangerle (University of Innsbruck, Austria)

5 Organizers

The organizers bring complementary experiences and perspectives. Andrés Ferraro was the initiator of the first edition of MuRS. For the second edition of MuRS, he took Christine Bauer and Lorenzo Porcaro aboard, and in this third edition, the three continue as an organization team.

Andrés Ferraro[2] is a Senior 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 was also an organizer of DEFirst[4], a reading group at Mila, focused on discussions around fairness and recommender systems research.

Lorenzo Porcaro[5] is Marie Skłodowska-Curie Postdoctoral Fellow at the Department of Computer, Control and Management Engineering (DIAG) of the Sapienza University of Rome (Italy), where is PI of the project titled Algorithmic Auditing for Music Discoverability (AA4MD). His work focuses on assessing the impact of recommender systems on their users. He holds a PhD in Information and Communication Technologies 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 Workshop on Human-Centric Music Information Research (HCMIR), co-located with ISMIR. He also served as a lecturer at the ACM Europe Summer School on Recommender Systems 2024.

Christine Bauer[6] is a Professor of Interactive Intelligent Systems at the Department of Artificial Intelligence and Human Interfaces (AIHI) at the University of Salzburg, Austria, and a Co-Lead of the focus area "InterMediation. Music—Effect—Analysis" at the inter-university organization Arts & Knowledges. Within her research field, she has focused on context-aware music recommender systems. Core interests are fairness and multi-method evaluations of recommender systems. She is a very active member of 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 Summer Schools on Recommender Systems in 2024, 2023 and 2019.

Acknowledgments

References

[1] 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. doi:10.1007/978-3-642-13287-2

[2] 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. doi:10.1145/3240323.3240342

[3] Karlijn Dinnissen and Christine Bauer. 2022. Fairness in music recommender systems: A stakeholder-centered mini review. *Frontiers in Big Data* 5, Article 913608 (2022), 9 pages. doi:10.3389/fdata.2022.913608

[4] Andres Ferraro, Michael D. Ekstrand, and Christine Bauer. 2024. It's Not You, It's Me: The Impact of Choice Models and Ranking Strategies on Gender Imbalance in Music Recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (*RecSys '24*). Association for Computing Machinery, New York, NY, USA, 884–889. doi:10.1145/3640457.3688163

[5] 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. doi:10.1145/3523227.3551476

[6] 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, Article 10 (2024), 32 pages. doi:10.1145/3643138

[7] Andres Ferraro, Peter Knees, Massimo Quadrana, Tao Ye, and Fabien Gouyon. 2023. MuRS: Music Recommender Systems Workshop. In *Proceedings of the 17th ACM Conference on Recommender Systems* (Singapore) (*RecSys '23*). Association for Computing Machinery, New York, NY, USA, 1227–1230. doi:10.1145/3604915.3608750

[8] Andres Ferraro, Lorenzo Porcaro, Peter Knees, and Christine Bauer. 2024. MuRS 2024: 2nd Music Recommender Systems Workshop. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (*RecSys '24*). Association for Computing Machinery, New York, NY, USA, 1202–1205. doi:10.1145/3640457.3687097

[9] 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. doi:10.1007/978-3-030-85616-8_33

[10] Yu Liang and Martijn C. Willemsen. 2023. Promoting Music Exploration through Personalized Nudging in a Genre Exploration Recommender. *International Journal of Human–Computer Interaction* 39, 7 (2023), 1495–1518. doi:10.1080/10447318.2022.2108060

[11] Kristina Matrosova, Lilian Marey, Guillaume Salha-Galvan, Thomas Louail, Olivier Bodini, and Manuel Moussallam. 2024. Do Recommender Systems Promote Local Music? A Reproducibility Study Using Music Streaming Data. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (*RecSys '24*). Association for Computing Machinery, New York, NY, USA, 148–157. doi:10.1145/3640457.3688065

[12] Lorenzo Porcaro, Carlos Castillo, and Emilia Gómez. 2021. Diversity by Design in Music Recommender Systems. *Transactions of the International Society for Music Information Retrieval* 4, 1 (Nov 2021), 114–126. doi:10.5334/tismir.106

[13] Lorenzo Porcaro, Emilia Gomez, and Carlos Castillo. 2022. Diversity in the Music Listening Experience: Insights from Focus Group Interviews. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) (*CHIIR '22*). Association for Computing Machinery, New York, NY, USA, 272–276. doi:10.1145/3498366.3505778

[14] 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, Article 3 (March 2024), 47 pages. doi:10.1145/3608487

[15] Upendra Shardanand and Pattie Maes. 1995. Social information filtering: algorithms for automating "word of mouth". In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, CO, USA) (*CHI '95*). ACM Press/Addison-Wesley Publishing Co., USA, 210–217. doi:10.1145/223904.223931

[16] Viet-Anh Tran, Guillaume Salha-Galvan, Bruno Sguerra, and Romain Hennequin. 2024. Transformers Meet ACT-R: Repeat-Aware and Sequential Listening Session Recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (*RecSys '24*). Association for Computing Machinery, New York, NY, USA, 486–496. doi:10.1145/3640457.3688139

[17] 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. doi:10.1007/s13735-019-00169-8

---

[3]https://www.meetup.com/grupo-em-sao-paulo-de-audio-signal-processing
[4]https://www.youtube.com/channel/UCI9jOq79MHUD3vnh3jUUirg
[5]Web: https://lorenzoporcaro.me
[6]Web: https://christinebauer.eu