

Report on the ISMIR 2020 special session: How do we help artists?

Christine Bauer
Utrecht University
c.bauer@uu.nl

Abstract

The 21st edition of the Annual Conference of the International Society for Music Information Retrieval (ISMIR) introduced so-called “special sessions,” giving room for discussion on various topics related to music information retrieval (MIR). I report on the activities related to this special session No. 7 with the title “How do we—in MIR research—help artists? Do we?”, and summarize the insights gained from the discussion. Being well-aware that the MIR community has come up with contributions that had a positive impact on artists, we particularly discussed the challenges that hinder MIR contributions from being adopted in the field and the problems when MIR innovations with presumably good intentions turn out to have a bad impact. Besides the many challenges, we also came up with ideas how we could move forward and better address artists’ needs.

1 Introduction

The 21st edition of the Annual Conference of the International Society for Music Information Retrieval (ISMIR) was the first one being held online. With the transfer to the virtual world, ISMIR became vISMIR. So-called “special sessions” were introduced giving room for discussion on various topics related to music information retrieval (MIR). I had the honor to host a special session (No. 7) that I entitled “How do we—in MIR research—help artists? Do we?” that was scheduled for the last day of the conference (15 October 2020). We used one of the conference’s Slack channels dedicated to discuss the topic with the community throughout the conference week. The copestone of the discussion was the special session that took place as a video conference via Zoom. While originally scheduled for one hour, we spontaneously extended our lively discussion to 1.5 hours.

In this event report, I summarize the topics discussed in both the Slack channel and in the video conference.

2 Starting Point

The special session (No. 7) with the title “How do we—in MIR research—help artists? Do we?” came with the following announcement:

In MIR research, we have a wide repertoire of methods to learn about music and about how it is used, etc. Yet, how can MIR be used to help current artists to move forward—in creativity or in making a career? Are we helping artists already? Could we do better?

And I started off the Slack discussion with the following comment:

My introductory question is very cynical—between the lines. I think the artists are rarely in the loop when we are doing our research in the MIR field. What is your stance on this? Can you convince me of the opposite? Do you have ideas to move forward? Maybe it's not even desired—in the MIR field—to help artists...?!'

3 MIR Contributions That Had an Impact on Artists

The community was clear that the MIR field came up with contributions that had a (positive) impact on artists.

Optical music recognition (OMR) (for recent reviews, see [Calvo-Zaragoza et al. \[2020\]](#) and [Shatri and Fazekas \[2020\]](#)) is a prominent example that enabled a variety of outcomes that support real-world tasks, including music search and processing (e.g., transposing), score-following, automatic page-turning. It is worthwhile to note that also other technologies and methods support these and similar tasks such as, for instance, score-following from raw sheet-music images (see e.g., [Henkel et al. \[2020\]](#) and [Shan and Tsai \[2020\]](#)).

Furthermore, participants indicated that some solo artists use automatic (expressive) accompaniment in their live performances. For a recent critical review of research in the field, see [Cancino-Chacón et al. \[2018\]](#).

We observe that a flood of music training apps has entered the market in the recent years, where those apps typically rely on a variety of MIR contributions with respect to beat tracking, pitch detection, etc., in real-time. While we noticed that most of those apps seem to target self-training, integrating such apps in other educational settings (e.g., at music schools) would basically be possible.

4 Challenges of MIR Contributions Being Adopted

The discussion also pointed to the challenges that hinder MIR contributions being adopted by artists. One problem is that research contributions typically do not go beyond the prototype level; yet, further development is needed so that artists can adopt innovations in practice. Moreover, we discussed that the apps provided are often very complex and lack usability, which hinders adoption.

Furthermore, for many tasks, MIR technology—although advanced over time—still performs relatively poorly compared to a trained human carrying out the same task.

A challenge also lies in clearly communicating what an app can do (and what not) and still conveying the message of innovativeness for marketing the app. For example, users may expect perfection (e.g., in a transcription task) instead of considering the app as a tool to support a human in performing the transcription task.

For some apps, it may be too risky to use those in live situations, not because the app’s performance itself but because of the environmental settings that it should be used in. For example, let us consider the use of a digital sheet music system for an entire orchestra or big band at a big live event. If there is a problem in the environment—which may be as simple as one cable being accidentally disconnected—the entire orchestra will not longer be able to access any sheet music at all [Bauer, 2012]. Overall, there is typically little flexibility for last-minute changes, while last-minute changes are rather the norm than the exception at live events.

5 Impact of MIR on Artists

In the discussion, we addressed many MIR innovations that have a direct or indirect impact on artists—yet, not necessarily always a good one.

For instance, full automation may eliminate jobs; jobs that were carried out by humans before. When considering the implications, we need to distinguish between the automation of tasks and the automation of full jobs. Essentially, the purpose of many current apps are meant to support humans in performing their tasks (e.g., better, quicker). We agreed that the main focus should lie on the automation of tasks rather than jobs, to support artists in their tasks and activities.

However, we also found that innovations with good intentions turned out to have a bad impact. An example for such negative implications is that copyright bots repeatedly block non-copyright-protected music [Brodeur, 2020]. Similarly, also content creators suffer from the platform’s semi-automated strike and appeal policies; for instance, chatbots frequently lead to a takedown of educational content that should fall under “fair use” (i.e., falling under the exemptions of most copyright laws that allow the use of copyrighted material for specific educational purposes and settings) [Alexander, 2019]. Thus, while this technology may have been developed with the intention to protect artists (or rather, their intellectual property), it turns out to work against them. The major underlying problem for such complications is that in many research communities—and this includes the MIR community—, we do not reflect deeply on what could go wrong (the risks) with our technologies when implemented in real-world scenarios and on the impact it may have on the various stakeholders (e.g., artists, music consumers, copyright management organizations). The question remains whether we need to reflect more on our contributions or earlier in the process—or (possibly) both.

Further concerns relate to the implications of music recommender systems for artists. The integration of recommenders into the majority of today’s music platforms may support artists in reaching a larger audience or achieving more intense listeners. Still, we were divided whether only a small number of artists is supported by current systems [Bauer, 2019; Im et al., 2019]. We agreed, though, that the MIR community has ever-since engaged in research to promote long-tail content (“digging into the music landscape”) [Celma, 2010; Craw et al., 2015; Turnbull et al., 2008]. A critical path could be to find a balance in satisfying users and providing artists visibility [Mehrotra et al., 2018].

Another concern relates to the way that music platforms currently employ MIR contributions: The platforms seem to detach the user experience from the artists. The currently employed access-based business model (compared to an ownership-based model) incentivizes consumption [Arditi, 2018] and is less focused on the artists than other business models based on, e.g., (physical) record sales. A participant pointed out,

Many people do not even care which artists they listen to under such a system.

There was a vivid discussion about finding “fairer” ways of compensating artists. For example, user-centric payouts seem to be fairer for a set of genres; yet, for some genres only [Dredge, 2020]. Another idea proposed by one participant was to integrate a “donate/support” button, providing an easy-to-use opportunity to support a specific artist. We also discussed portals that offer “fairer” compensation schemes for artists than the current major players, with particular pointers to Bandcamp’s model as well as Resonate.io’s ‘stream-to-own model.’ However, such portals are often niche portals; yet, users seem to prioritize choice over fairness to artists. A participant pointed out that the major hurdle for niche players arises from the licensing scheme and the costs involved; they would want to offer all kinds of content but cannot. A possible solution (if the licensing scheme is not changed) could be to find a way to integrate items into a (niche) platform that are sold elsewhere. In the discussion of the many structural reasons why many creators do not get paid their fair share, a participant pointed out that one of the major problems for fair royalty payments is the lack of comprehensive metadata [Deahl, 2019]. The MIR community could greatly contribute here. Yet, a participant commented to observe a lack of interest in (or underestimates the relevance of) the topic in the MIR community.

In the context of compensation, fairness, copyright, an unsolved issue is how to deal with false copyright claims and the situation that there are no, or only few, consequences for fraud.

6 Moving Forward

Besides the many challenges, we also came up with ideas how we could move forward and better address artists’ needs.

Considering specific tasks, specific settings, specific situations. Different tasks, settings, or situations come with different requirements. Thus, if we want to support artists, we need to address those specific needs and requirements. One size does not fit all. For instance, requirements may depend on the profession, proficiency level, instrument, music styles, solo/band/ensemble/orchestra, etc.

Reaching out to artists—talk to them. If we want to support artists, we need to involve them. First, we need to reach out to artists for *eliciting their demands and requirements*. Artists are the only ones that genuinely understand what they are doing; they can tell what tasks they consider complex, repetitious, etc. Artists can tell how things work in the field—in the context of use. Hence, it is advisable to reach out to artists already before any development starts. Second, we need to involve artists in our *evaluation* efforts. More specifically, we need to evaluate our research and developments with artists that have the profiles that we target with our research or development. Integrating artists already early in the process, may provide early feedback and give direction. In all efforts, we should be aware that *one size does not fit all*. Different artists have different needs. For instance, we need to consider the different levels of proficiency and the different kinds of tasks. Often, there is even a wider variety of stakeholders involved (e.g., instrumentalists, singers, songwriters, educators, students). For example, for an educational app, learners and

educators most probably have different needs and may deliver different insights on what could support the learning process; also, parents may play a role in shaping these requirements.

Embracing methodology used in the human-computer interaction (HCI) field. The methodology embraced in the HCI community (e.g., interviews, co-design, interactive prototyping) may enrich what we are doing in MIR. We may learn and employ such methods ourselves or collaborate with experts in those fields.

Overall, I want to emphasize that music is embedded in a social and economic ecosystem. MIR research impacts all. To conclude, I want to quote Cynthia Liem’s words that very well express the timeliness and urgency to put effort in supporting artists:

[...] in locked-down times like these, when stages are closed and audiences are confined... having technology supporting artists in still being able to make and share music may be more important than it’s ever been...

Acknowledgements

I would like to acknowledge Anna Gatzoura, who volunteered in making the special session’s video conference technically happen, and who supported me in moderating the discussion. I am also grateful to Andrés Ferraro for taking notes during the video conference so that I could concentrate on the discussion. Further, I would like to thank the program chairs—Julie Cumming, Jin Ha Lee, Brian McFee, Markus Schedl—for giving me the opportunity to host a special session at vISMIR 2020. Most importantly, this work would not have been possible without the thoughtful and inspiring contributions of all the participants in the Slack channel and the video conference session. Together, we will move forward.

References

- Julia Alexander. Youtubers and record labels are fighting, and record labels keep winning: The battle over copyright continues. *The Verge*, May 2019. URL <https://www.theverge.com/2019/5/24/18635904/copyright-youtube-creators-dmca-takedown-fair-use-music-cover>.
- David Arditi. Digital subscriptions: The unending consumption of music in the digital era. *Popular Music and Society*, 41(3):302–318, 2018. doi: 10.1080/03007766.2016.1264101. URL <https://doi.org/10.1080/03007766.2016.1264101>.
- Christine Bauer. *Bands as Virtual Organisations: Improving the Processes of Band and Event Management with Information and Communication Technologies*. Electronic Business. Peter Lang, Frankfurt am Main, Berlin, Bern, Bruxelles, New York, Oxford, Wien, 2012. ISBN 978-3-631-63057-0. doi: 10.3726/978-3-653-01223-1. URL <https://doi.org/10.3726/978-3-653-01223-1>.

-
- Christine Bauer. Allowing for equal opportunities for artists in music recommendation. In *1st Workshop on Designing Human-Centric Music Information Research Systems (wsHCMIR 2019), satellite event to 20th Annual Conference of the International Society for Music Information Retrieval (ISMIR 2019)*, wsHCMIR 2019, pages 16–18, 2019. URL <https://arxiv.org/abs/1911.05395>.
- Michael Andor Brodeur. Copyright bots and classical musicians are fighting online. the bots are winning. *The Washinton Post*, May 2020. URL https://www.washingtonpost.com/entertainment/music/copyright-bots-and-classical-musicians-are-fighting-online-the-bots-are-winning/2020/05/20/a11e349c-98ae-11ea-89fd-28fb313d1886_story.html.
- Jorge Calvo-Zaragoza, Jan Hajič Jr, and Alexander Pacha. Understanding optical music recognition. *ACM Computing Surveys*, 53(4):1–35, July 2020. ISSN 0360-0300. doi: 10.1145/3397499. URL <https://doi.org/10.1145/3397499>.
- Carlos E. Cancino-Chacón, Maarten Grachten, Werner Goebel, and Gerhard Widmer. Computational models of expressive music performance: A comprehensive and critical review. *Frontiers in Digital Humanities*, 5:25, 2018. ISSN 2297-2668. doi: 10.3389/fdigh.2018.00025. URL <https://www.frontiersin.org/article/10.3389/fdigh.2018.00025>.
- Òscar Celma. *Music Recommendation and Discovery: The Long Tail, Long Fail, and Long Play in the Digital Music Space*. Springer, Berlin Heidelberg, 1st edition, 2010. ISBN 3642132863. doi: 10.1007/978-3-642-13287-2. URL <https://doi.org/10.1007/978-3-642-13287-2>.
- Susan Craw, Ben Horsburgh, and Stewart Massie. Music recommendation: Audio neighbourhoods to discover music in the long tail. In Eyke Hüllermeier and Mirjam Minor, editors, *Case-Based Reasoning Research and Development*, ICCBR 2015, pages 73–87, Cham, Germany, 2015. Springer International Publishing. ISBN 978-3-319-24586-7. doi: 10.1007/978-3-319-24586-7_6. URL https://doi.org/10.1007/978-3-319-24586-7_6.
- Dani Deahl. Metadata is the biggest little problem plaguing the music industry: It’s a crisis that has left, by some estimations, billions on the table unpaid to musicians. *The Verge*, May 2019. URL <https://www.theverge.com/2019/5/29/18531476/music-industry-song-royalties-metadata-credit-problems>.
- Stuart Dredge. Deezer still pushing for user-centric payouts: ‘we will continue fighting...’. *Musically*, October 2020. URL <https://musically.com/2020/10/01/deezer-still-pushing-for-user-centric-payouts-we-will-continue-fighting/>.
- Florian Henkel, Rainer Kelz, and Gerhard Widmer. Learning to read and follow music in complete score sheet images. In *Proceedings of the 21st International Society for Music Information Retrieval Conference*, ISMIR 2020, pages 780–787, 2020. URL <https://arxiv.org/abs/2007.10736>.
- Hyunsuk Im, Haeyeop Song, and Jaemin Jung. The effect of streaming services on the concentration of digital music consumption. *Information Technology & People*, 2019. doi: 10.1108/ITP-12-2017-0420. URL <https://doi.org/10.1108/ITP-12-2017-0420>.

-
- Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, pages 2243–2251, New York, NY, USA, 2018. ACM. ISBN 9781450360142. doi: 10.1145/3269206.3272027. URL <https://doi.org/10.1145/3269206.3272027>.
- Mengyi Shan and TJ Tsai. Improved handling of repeats and jumps in audio-sheet image synchronization. In *Proceedings of the 21st International Society for Music Information Retrieval Conference, ISMIR 2020*, pages 62–69, 2020. URL <https://arxiv.org/abs/2007.14580>.
- Elona Shatri and György Fazekas. Optical music recognition: State of the art and major challenges. In Rama Gottfried, Georg Hajdu, Jacob Sello, Alessandro Anatrini, and John MacCallum, editors, *Proceedings of the International Conference on Technologies for Music Notation and Representation, TENOR 2020*, pages 175–184, Hamburg, Germany, 2020. Hamburg University for Music and Theater. ISBN 978-3-00-066930-9. URL <https://arxiv.org/abs/2006.07885>.
- Douglas Turnbull, Luke Barrington, and Gert RG Lanckriet. Five approaches to collecting tags for music. In *Proceedings of the 9th International Conference of Music Information Retrieval*, volume 8 of *ISMIR 2008*, pages 225–230, 2008. URL http://ismir2008.ismir.net/papers/ISMIR2008_128.pdf.