

# The Future of Musical Knowledge in the Age of Machine Learning

Eamonn Bell<sup>1</sup>, Nikita Braguinski<sup>2</sup>, Luca Casini<sup>3</sup>, Sarah-Indriyati Hardjowirogo<sup>4</sup>,  
Rujing Stacy Huang<sup>5</sup>, Christopher White<sup>6</sup>, Miriam Akkermann<sup>7</sup>

<sup>1</sup>Durham University; <sup>2</sup>KHK Cultures of Research, RWTH Aachen; <sup>3</sup>KTH Stockholm; <sup>4</sup>Oldenburg University;  
<sup>5</sup>The University of Hong Kong; <sup>6</sup>University of Massachusetts Amherst; <sup>7</sup>FU Berlin

Machine learning and methods of artificial intelligence (AI) have become increasingly present for music practitioners, music researchers and music listeners. Going beyond the more widely-known cases of audio deepfakes and automatic music generation for commercial use, the advent of music AI also has a vast influence on music research in all its various contexts. This paper aims to outline some of the current developments in music research, focusing explicitly on intradisciplinary interactions between music history, music analysis, music theory, music technology, music production, and music informatics. It is based on the work results and discussions of the interdisciplinary and international Visiting Research Group “The Future of Musical Knowledge in the Age of Machine Learning,” hosted at ZiF Bielefeld for the period April-May 2023. The goal is to show the multifaceted nature of this field, highlight cross-cutting links and potential for mutual developments, as well as to take a position on what the growing entanglement of AI with music research might mean for the future.

## Introduction

The implementation of machine learning techniques in music production and music analysis has taken place in a variety of ways and across a wide range of fields, opening up both new perspectives on music research as well as providing new roles and ideas for research within an inter- and intra-disciplinary group of self-identified musicologists, artists, scientists, and engineers. Our paper discusses ways in which advances in artificial intelligence and the growing collection of data about musical behavior might influence the future course of development in musical thinking, focusing especially on discourses of music professionals working under the umbrella of music research, such as musicologists, music theorists, music informatics specialists, and musicians.

This paper will outline our arguments and positions concerning three central points: (1) the challenges, hopes, and fears related to music AI and research, (2) the disciplinary differences among music and AI researchers, and (3) the institutional context of this research between musicology, arts, science and engineering. Our goal is to provide a basis for further discussions on this emerging topic while providing a critical evaluation of recent developments in the field of music research against the background of machine learning.

## Background and motivation

### Human-centric approaches vs. technology

Insomuch as they remove the human agent as the direct creator of content, artificial intelligence tools undercut the traditional subject-oriented model of the musical audience (Briot et al., 2020; Cambouropoulos & Kaliakatsos-Papakostas, 2021; Christensen, 2002;

Miranda, 2021; Nierhaus, 2009; Sterne & Razlogova, 2019) and this change has the potential to foundationally reorient how scholars, listeners, and composers understand musical media. In recent years, discussions of these dynamics have been dominated by scholarly interest in the cultural and economic impact of creative applications (automatic generation of music) and recommender systems (used by the likes of Spotify and YouTube to steer music discovery) (Avdeeff, 2019; Born, 2021; Bown, 2021; Eriksson, Fleischer, Johansson, Snickars, & Vonderau, 2019; Gioti, 2020; Hesmondhalgh, Campos Valverde, Kaye, & Li, 2023; Lubart, 2005; Tatar & Pasquier, 2019). What remains too little understood and explored are future changes in disciplinary knowledge about music, represented in musicology and music theory. Our interests lie in changes to these more humanistic approaches to music, and it is exactly these approaches that have been understudied.

### The continuously changing nature of musical knowledge

Music theory in particular has long adapted in response to several technological upheavals, such as when sound recording technologies appeared at the end of the 19th century, or when new techniques of visualization allowed microtonal and microtemporal analysis. Then, in the mid-twentieth century, with the rise of scientism, cybernetics and information theory, and digital computation, formal knowledge about music became academically embedded in new institutional contexts (Akkermann, 2017; Bell, 2021; Haworth, 2021; Iverson, 2018; Miller, 2021).

In other words, despite its frequent focus on centuries-old traditions, musical theory never stands still. It is in a process of constant change, often following shifting musical and intellectual behaviors, and even sometimes guiding new musical practices. Importantly, these shifts and changes are often linked to the emergence of new musical technologies (Assayag et al., 2002; Braguinski, 2022).

More recently still, new tendencies in musicology and music theory, along with the belated recognition of the contribution of indigenous, subaltern, and marginalized musical traditions to institutional knowledge about music, have benefited from the explosion of the Web and the reduced costs of digital sound recording, distribution, and virtual pedagogy. In all these cases, music theory and musicology have been able to continue asserting their tools for the analysis and composition of music by adapting to these new circumstances.

Similarly, humanistic music research will almost certainly change—and is currently changing—to adapt, adopt, and react to contemporary AI technologies, which here includes the data-intensive use of computers to process culture and behavior, including applications of machine learning (ML) and deep learning to music. This process is, however, neither linear nor deterministic, and many different actors influence it. Analyzing these changes and making informed predictions necessitates an interdisciplinary approach. Given the technical nature of AI, these disciplines include not only music theory and musicology, but computer science, ethnomusicology, and even industrial perspectives (Hanemaayer, 2022; Rehding, 2021; Roberge & Castelle, 2021).

## Snapshots of interdisciplinary music and AI research

Recent years have seen the emergence of a number of prominent interdisciplinary research programs focused on music and AI. These relatively large-scale projects and labs demonstrate a number of different approaches to the challenge of interdisciplinary collaboration. Alongside prominent academic research projects funded, e.g., by Horizon Europe, various commercial and semi-commercial projects also point towards new use cases and artistic implementations for AI tools.

### Within the academy

The project *Raising co-creativity in cyber-human musicianship*, led by PI Gerard Assayag, aims to study, model and develop musical co-creativity between humans and machines through the lens of improvisation, while *Music and Artificial Intelligence: Building Critical Interdisciplinary Studies*, headed by PI Georgina Born (UCL), is an ambitious attempt to develop a new methodology for critical studies of art and technology, with participating researchers drawn from a variety of disciplines in addition to artists and composers (Born, 2007). Notably, some of the work from this group is explicitly historical, focusing on the origins of music AI, and the relationship between aesthetics and political economy, while other researchers have addressed the global perspective on streaming platforms outside the well-studied global North. The team at *Music at the Frontiers of Artificial Creativity and Criticism*, led by Bob Sturm (KTH), investigates, among other things, how artificial intelligence can be used creatively to construct and interrogate the notion of a musical tradition itself. The lab has developed generative algorithms to produce “machine folk”, drawing on a large dataset of notated traditional music.

### Within industry

Most industry actors to this point have focused on projects aimed at supporting commercial products. Among the most popular applications are systems that can generate finished musical material (audio or symbolic) at the push of a button and systems that tackle music understanding tasks for recommendation purposes. In 2020, OpenAI released Jukebox, a system that generated music conditioned on style and lyrics provided by the user (Dhariwal et al., 2020). OpenAI has since ceased work on further music projects, but other companies have been able to produce similar results on a larger scale thanks to improvements in computational power. For example, Suno (now a partner of Microsoft) released a text-to-sound open-source model called BARK in 2023. This model in turn powers their music generation system, creating music conditioned on style and lyrics prompts. The system is commercial in nature and there is no research output supporting it.

Meta’s AI research and development division has developed three open-source music-related projects under the Audiocraft umbrella. Encodec is an efficient codec for music that is learned by a neural network. The discrete tokens output by the systems in turn power audiogen and musicgen, two systems for text-to-audio and text-to-music respectively (Défossez et al., 2022; Copet et al., 2024; Kreuk et al., 2023). Less product-focused than other industry actors, Google has also been active in the space with its Magenta lab, which has developed experimental research projects and applications for timbre transfer and

music generation with various neural network architectures, and was responsible for a 2019 Google Doodle that could harmonize material provided by the user in the “style” of J.S. Bach, an important milestone for public-facing human-AI musical creativity (C.-Z. A. Huang et al., 2019).

Recently, Google has also released test-conditioned models called AudioLM and MusicLM (available to the public as MusicFX). The latter leverages a joint music-language embedding called MuLan, also developed by Google, with the goal of supporting music understanding and creation tasks (Agostinelli et al., 2023; Q. Huang et al., 2022).

## Inter- and Intra-disciplinary Perspectives

### Eliciting challenges, hopes, and fears related to music AI and research<sup>1</sup>

To consider ways that humanistic music research might develop in response to innovations surrounding musical AI requires first developing a mutual understanding of the specific problems this technology poses in the current moment. In what follows we begin by outlining several such problems, and use these issues to consider broader concerns and desires surrounding AI technology.

One major concern is the extent to which contemporary AI models are *interpretable*, especially given that contemporary architectures do not always easily lend themselves to open interpretation, and open access to code and other such straightforward strategies for “transparency” fall short of the mark (Ananny & Crawford, 2018). Because of their opacity, contemporary AI tends to be less studied within the musical academy than less complex algorithms. Because their inner workings are more easily to describe and their interpreting their behaviors are more straightforward, music researchers gravitate toward older statistical and rule-based computational methods to study harmonic progressions and other basic music-theoretical objects and concepts, even at times finding grounds to re-evaluate those objects on a methodological level.

There are also several fears about the fallout of AI technology. Detecting AI-supported plagiarism is one particularly potent concern, especially given the imperfect use of algorithms to detect and mitigate the reuse of music (Lester & Pachamano, 2017). Additionally, in a landscape in which AI- and human-generated content is increasingly indistinguishable, it’s unclear how audiences might change their relationship and expectations around musical expression and live performance, and the extent to which AI could render certain musical occupations and roles obsolete. Other fears—such as the fear of having one’s personal information being exploited, perhaps as a consequence of a music AI data collection exercise—also arise around potential abuses of this technology.

Optimistic hopes complement these concerns. Some researchers have begun to find various ways of using neural networks and other deep-learning methods to garner musicological

---

<sup>1</sup> We thank members of the corresponding working group for their contributions to this section: Gérard Assayag, Nikita Braguinski, Mark Gotham, Brian Miller, Christopher Haworth.

insights (Cosme-Clifford et al., 2023; Cuthbert & Ariza, 2010; White & Quinn, 2018). Similarly, some researchers have begun to study how recommendation algorithms as a potential partner in listeners' music consumption, rather than a one-way, passive relation between the algorithm and the recipient of the recommendation (Seaver, 2022).

Importantly, each of these hopes and fears do not arise from the technology itself, but in the way that humans interact with that technology. These sorts of concerns, then, are best approached from a multidisciplinary perspective to not only understand AI's engineering, capacities, and performance, but how those capacities are situated within a human society. Each of these concerns is determined by the role in which an AI is used, and the goals and outcomes attributed to its use. One listener enjoying a new track by their favorite artists might be furious and embarrassed—perhaps feeling “tricked”—should they learn that the music was fabricated by an AI to emulate that artist, while another listener hearing the same track as an ignorable background soundtrack in a grocery store might have no qualms about the tune's provenance. Indeed, the shopper might even feel relieved to not be listening to a fresh underscoring to their grocery-store experience! The music's function and the different goals that both listeners' bring to both of these encounters respectively result in different outcomes.

These changing priorities could certainly affect content creators as well, with a goal in one constituency becoming a fear of another. For instance, film music studio managers may seek to save time, and ultimately labor costs, by delegating the relatively routine work of orchestration to a machine-learning algorithm that is trained to emulate a “house style.” Conversely, junior film composers and producers, who have traditionally begun their induction into the industry performing this work may fear their obsolescence as a consequence of automation. This is not an abstract scenario: automatic music arrangement continues to be the subject of active innovation and application within industry and without (Gotham et al., 2022; Pachet, 2016). Relatedly, though with appreciably different economic stakes, professional music theory educators may fear the introduction of tools that can convincingly harmonize four-part chorale melodies, whilst their students may rejoice.

## Reading across disciplines reveals different approaches to the musical uses of AI<sup>2</sup>

Disciplinary differences between music and AI researchers cause problems for the mutual intelligibility of their research outputs, which is a key impediment to understanding the future of both fields. This can manifest in several ways in reading and writing the relevant academic literature. Researchers may read literature for a variety of reasons: to determine the novelty of an approach, to learn new techniques, to understand alternate perspectives, and so on. These different goals engender different reading and writing strategies. Sometimes, scholars read pragmatically for an understanding of the intricacies of specific

---

<sup>2</sup> We thank members of the corresponding working group for their contributions to this section: Eamonn Bell, Luca Casini, Robert Lieck, Marcin Pietruszewski, Alexander Rehding.

implementation details as reported in, for example, a technical paper's methodology section.

The disconnects between these different individual and disciplinary goals can lead to frustration or dismissal when the writer's and reader's goals do not align. A humanities reader may be tempted to characterize STEM research as less concerned with sustained argument, while a researcher from a natural science department might find humanities writing frustratingly unconcerned with reproducible findings. The disconnect can even pervade misunderstandings between the very vocabulary being used. Humanities readers might look at a paragraph entitled "Discussion" and expect the prose to review the key problematics, history, and contexts of the research output, while a STEM researcher might find such details in a "Discussion" extraneous to a clear and concise presentation of a paper's methods and findings.

While seemingly superficial, these vocabulary differences can act as crucial nexuses for identifying the distances between disciplines, and as such can act as important bridges between these points of distance. Many keywords are shared between disciplines studying AI, for example, "noise", "model", "form", and "parameter", and recognizing the different meanings ascribed by different research communities can act as a point of discussion between these communities and serve as a way to elicit points of shared understanding and miscommunication. In other words, by drawing out competing and complementary meanings of carefully selected terms, the semantic gap between interlocutors could be made explicit, translations can be undertaken, and deeper interdisciplinary conversations be encouraged (Novak & Sakakeeny, 2015; Williams, 1985).

Similarly, one of the greatest challenges to building a shared understanding of music and AI research is the role of tacit knowledge. Disciplinary background, country of study, prior training, and professional setting all play crucial roles in coming to know the conventions of a particular body of the literature or discourse. Readers unfamiliar with disciplinary conventions struggle to identify when the omission or elision of an argumentative "move" in a paper is intentional or unintentional. In turn, this makes it hard to evaluate the quality of research in an unfamiliar domain. Again, these misunderstandings are not so much a challenge as an opportunity for an intervention, stimulating the development of new ways to design research outputs to facilitate greater access across the disciplinary divide.

### **(Re)envisioning music AI institutions<sup>3</sup>**

As it studies both how music is objectively constructed and how it is subjectively interpreted, the musical humanities sit at a disciplinary overlap in musical AI studies. As the focus is on a medium that is at once highly technical and deeply expressive, music studies is already a deeply multidisciplinary space. Humanistic music analysis focuses both on how music is built—its forms and structures—and musical expression—music cognition, narrative, embodiment, compositional intention, and the like. This balance between the music's objective materials and the ways humans use those materials makes it

---

<sup>3</sup> We thank members of the corresponding working group for their contributions to this section: Miriam Akkermann, Eric Drott, Sarah Hardjowirgo, Rujing Stacy Huang, Christopher W. White



amenable to posing and answering questions related to technologies like LLMs that produce objective structures with the goal of being useful to human users.

Additionally, with its focus on the pedagogy of these music materials, many fields within the musical humanities have a front-row seat to issues surrounding AI's role in education and teaching, with AI making it both easier to engage users in making music, while supplanting the human understanding, fluency, and skills needed to make that music. With users being able to make new music at the click of a button, AI tools can spark interest in music making and provide access to users with little musical background. This access can foster learning aesthetic concepts in early music education and give musical novices the interest and encouragement to engage in music studies. But these tools can also provide shortcuts that might sap the motivation behind some music studies, especially skills-based tasks important to the development of fluency within technical domains. If the analytical and compositional tasks that are the focus of teaching can now be completed by music AI tools—even by general-purpose pretrained large language models (LLMs) that have not been specialized to music tasks—this could demotivate learners and reduce the opportunities for students to practice the craft associated with completing routine tasks. Such a dialectic tracks the fast-moving discourse around the AI tools in teaching core computer science competencies, such as programming, and within education more broadly (Becker et al., 2023; Kasneci et al., 2023).

On the other hand, different domains of musical activity make different demands on learners and creators. An activity like underscoring social media videos, for example, involves a more utilitarian application of musical skills than an activity like live, in-person improvisation. The former has modes of composition, performance, and listening that are quite distinct from the latter, which is more associated with human creative expression and “art”. Consequently, AI tools may accordingly see more widespread adoption in settings with less bespoke or expressive goals. Here, with its focus on music's role in the human experience, the musical humanities are ideally situated to investigate beliefs about the intrinsic value of human creative labor and expression, and how those interact with style and the cultural norms.

While the musical humanities might be an excellent location to begin building cross disciplinary dialogue relating to music AI, there exist multiple layers of tension to the politics of interdisciplinarity, multidisciplinary, and transdisciplinarity. There are ontological, epistemological, and methodological differences between disciplines (e.g., between MIR and empirical musicology). There are different approaches to the same field across continents: the relative value assigned to qualitative as against quantitative approaches to musicology is different in Germany than in the UK and the US, for example. Furthermore, university funding structures are not always conducive to interdisciplinary collaboration, and competition between opposing nation-states can even pose barriers to international scholarly collaborations. There are yet further potential conflicts of interest between industry-based and academia-based researchers; intersectoral traffic, with doctoral trainees and faculty members specializing in music AI moving freely between industry and academia, is especially notable (Born, 2020).

Any uptick in interest in music AI will surely pose challenges to academic departments, especially concerning academic infrastructure and disciplinary boundaries. We can expect new questions about musical fluency, musicianship, skills, and about the kind of musical knowledge that should be the focus of academic music programs. Will knowledge of musical composition go the way of cursive and Latin, relegated to highly specialized curriculums? Since LLMs rely on an *already-made* dataset and because automated generation disconnects musical production from its original human cultural context, we may risk ossifying musical content and values. Will an infinite repository of high-quality AI music diminish our value of live performance? Without humans creating the music, will listeners dissociate musical styles—like hip-hop, rap, or jazz—from their original cultural contexts?

AI may also uniquely impact the employment and enrolment structure of humanities programs. While the current interest in AI may well generate more student interest in music studies, it could also give greater credence to beliefs about a widening power imbalance between STEM disciplines and the humanities. At the same time, the difficulties that most institutions, which are deeply entwined with their material conditions, face when tasked to systematically transform themselves in response to cultural changes writ large, are beyond question.

## Outlook

The overall observations combine into a heterogeneous picture of opportunities, hopes, and risks. However, we have several large-scale outlooks on these landscapes.

For one, the rise of interest in AI will introduce new methods in various research cultures, making a new, stable, and interdisciplinary vocabulary surrounding AI a moving target. However, this fluctuation is a feature rather than a bug. As our own disciplines grow to meet this new interest, music and AI educators should strive to familiarize themselves and introduce academic conventions from each other's disciplines early in the research training journey. This could manifest in cross-department journal clubs or extracurricular events, if only to denaturalize disciplinary ways of reading and writing encountered in typical research training settings.

That said, several stable aspects of musical AI must also be recognized. For one, AI is a computational technology. Building greater AI literacy means that music teachers and students should equip themselves with basic programming and data skills especially if they expect to encounter such tools in their research or creative work. However, music is also fundamentally an expressive and human endeavor. STEM programs focusing on AI should therefore include studies on musical interpretation, cognition, and creativity.

To achieve this balance, researchers and universities should, therefore, avoid treating knowledge hierarchically and work to promote an equal flow of knowledge, especially between STEM fields and the humanities. This flow, however, would require a reorientation of many dynamics in the current landscape of musical AI, because, at present, the balance tips far more toward industrial and technical concerns. Within the academy, for



instance, the AI research within computer science often exerts more influence on humanities AI research than music scholarship has on AI work in computer science. At the broader social level, governments tend to put more pressure on the academy than on the tech industry, while the government itself receives lobbying and support from the tech industry. Industry also continuously exploits the intellectual resources of academia, often with minimal reciprocal investment, either financially or in terms of their insights from research and development.

Figure 1 schematizes an idealized, balanced system of influences. Industry and academia partner in equal relationship, as does the government and academia. Computer science and music both engage in their own academic endeavors but share an overlapping interdisciplinary space. These spaces then affect teaching and research, which in turn produce students that perform within academia, as well as artistic and other non-academic venues. The media and public influence all these actors, but are also influenced by pedagogy, research, and artistic productions.

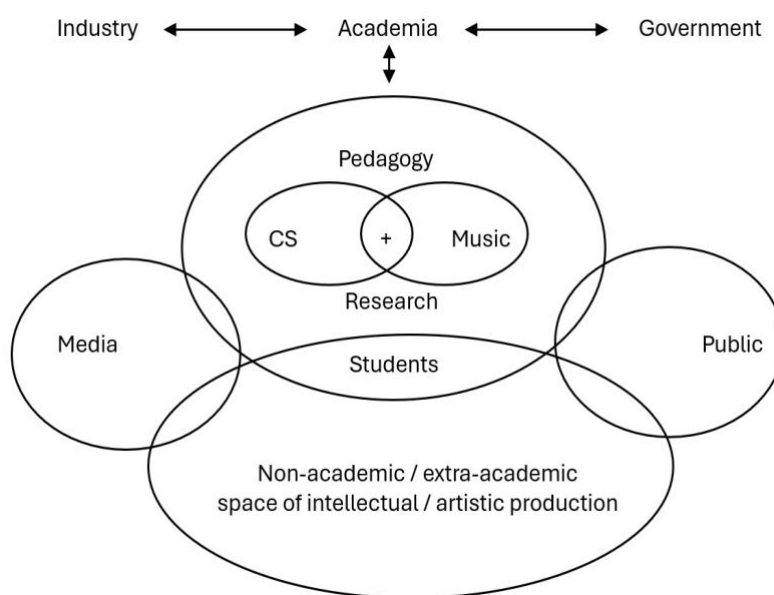


Fig 1. An ideal dynamic structure where there is an increasing amount of shared space and knowledge between institutions that currently relate to one another in lopsided power dynamics.

The group discussed which questions in the fields of music theory and musicology the research community can hope to ask of music AI and what developments in music AI do researchers, across disciplinary divides, require to support such work. The state of the art for a particular AI application is likely to change over the course of the longer timeline of a typical humanistic project. A music theorist or musicologist is faced with the challenge of keeping up with new developments without necessarily giving up the benefits of slower, more considered reflection characteristic of much high-quality humanistic work. Here the role of close-reading newly published research within multidisciplinary and interdisciplinary settings comes into view. Such exercises are valuable not only for generating answers to questions that non-experts may have about their peers' fields, but also (and perhaps more importantly) for generating questions that strike at the heart of

how and why the different disciplines contributing to a given field prosecute their research in a particular way and constitute their objects of study as such.

To date, musicologists and computer scientific researchers instantiate what might be called a “textualist” approach to the study of music, which emphasizes the role of audio and/or score-based formats in understanding music. Contemporary musicology has long put into question any sharp distinction between the textual aspect (for example, as an audio-visual recording posted on a social media platform) and the contextual aspect (for example, the plain-text comments that surround it) of music cultures, online or otherwise. Recent developments in AI research, such as the rise of multi-modal machine learning pretrained on large datasets, seem to give an operational basis for such a viewpoint.<sup>4</sup> This is because these tools can be said to no longer make a mathematical distinction between audio, text, and image and thus between musical text and musical context. Researchers can now use these AI-based tools to arrive at contextual, rather than purely textual, questions about music.<sup>5</sup>

A further additional challenge to the mutual comprehension of researchers representing different disciplines is that definitions of central terms such as “artificial intelligence”, “data”, “information”, or “knowledge” are still unstable with regard to the fast-moving field of music technology and are in need of further sharpening through collective discourse. As our discussions revealed, some terms are particularly evocative (such as “parameter” and “model”), as they bear quite different meanings to different constituencies. We advocate for future work on an interdisciplinary music and AI glossary, which can be used by researchers to both solidify their understanding of specific contested terms as well as to identify those concepts that have the greatest potential to stimulate discussion between humanistic and scientific music AI researchers. We also recognize a potentially broader appeal to such work, as terms with long, critical histories in the study of sound like “noise” (Kane, 2019; Thompson, 2017) and “vibe” (James, 2019) can even provide an entry point into critical AI studies more generally, via a set of metaphors that animate the analysis of music AI.

Finally, the group observed the existence of what can be called the naturalizing impulses of our given research disciplines. Together we ask the question of how the disciplines contributing to music AI research can be brought to defamiliarize their respective worldview, given the developments in technology and drawing productively on the discourses available in other disciplines. We think interdisciplinary institutions, like ZiF, effectively nourish new kinds of integrated spaces that support discussions that ultimately lead to programs that can train scholars at the intersections of fields. This will, of course,

---

<sup>4</sup> Google’s MuLan and Spotify’s LLARK models, discussed above, are some of many examples of research in this paradigm, as are many of the models used in popular text-to-image generators.

<sup>5</sup> One of many indicators that suggest that the trend of multi-modality has arrived for computational music researchers is the recent announcement of a guest-edited special issue of *Transactions of the International Society for Music Information Retrieval* (TISMIR) on multi-modal MIR. See <https://transactions.ismir.net/announcements#tismir-special-collection---call-for-papers> for the Call for Papers.

require major changes to curriculum design, systems of evaluation, and existing discipline-specific publication models that pose barriers to collaborations, specially between STEM disciplines and the humanities. We also suggest the development of healthier modes of industry-academia partnerships, especially when those in the music academia are not traditionally used to being industry adjacent.

We cannot pretend to have had a comprehensive discussion of the many fruitful bases for closer interaction between music researchers and computer scientists, and certain key topics have not been discussed. A notable irony is that little music was either listened to or discussed during the visiting research group's activities, yet close reading of academic literature about music and music AI is self-evidently strengthened by close listening of relevant musical works and/or the systems that generate them. Most notably, however, a thorough consideration of the ethics of music AI is yet to be developed. This should comprise not only a careful and inclusive discussion about the ethical stakes of future AI innovations, but also a complete consideration of the ethical and political dimensions of the inputs, processes, and outputs of music AI research as it is actually done today. This includes, but is hardly limited to, the ongoing collection and consolidation of human cultural data, which can be both personally and contextually highly sensitive, the economic and ecological cost of storing and processing data for music AI, and the use of generative music AI and related techniques with the sole purpose of deskilling and marginalizing working musicians. Relatedly, more work on best practice around music datasets, their production, circulation, and organization, should be informed by close, interdisciplinary collaboration, rather than the exigencies of a particular application of music AI in a particular research context ([Lee, Cooper, & Grimmelmann, 2024](#)).

Though brief, the Bielefeld initiative was a valuable first step in convening a community of researchers with a shared interest in exploring the implications for music AI on knowledge about music. By actively encouraging critical exchanges between and within disciplines, the group retread much familiar ground in the conduct of interdisciplinary research. However, it was nevertheless felt that music AI offers much grist to the mill for both humanities and scientific researchers, and as such, clearly challenges researchers to question the value and very necessity of being merely one or the other. Members of the research group continue to meet on a regular basis, to advance the research agenda set out here, to sustain existing links between participants and to grow new ones. Future meetings on the topic are planned in the coming years, and researchers and practitioners who identify with our objectives and working practices are very welcome to join in the shared journey.

## Acknowledgements

We greatly acknowledge the financial and logistical support of *Zentrum für interdisziplinäre Forschung* (ZiF), which hosted the Visiting Research Group “The Future of Musical Knowledge in the Age of Machine Learning” during the period April-May 2023. We particularly thank Anika Haverig, Marina Hoffman, Rebecca Moltmann, and Jan Begrich. Without their support and expertise, the fruitful discussions which led to this paper would not have taken place in such a productive and welcoming atmosphere. Finally, we

especially thank for their input Brian Miller and Mark Gotham who contributed to developing an initial version of this paper.

## Ethics Statement

This paper has been written based on the work results of an interdisciplinary and international Visiting Research Group funded by ZiF at the University of Bielefeld, for the period April-May 2023. Preparatory as well as wrap-up and paper writing meetings were held online, for the on-site workshop, all travel, accommodation and allowance costs have been covered by ZiF. The research group's activities were organized along Bielefeld University standards including sustainability, accessibility and inclusivity guidelines.

Private data of participants have only been handled by ZiF in order to facilitate participating in the Research Group's actions. All involved researchers are dedicated to following their respective institutions' code of good research practice.

The convenors of the Visiting Research Group actively encouraged upcoming researchers from different countries and research fields to participate in the activities; they furthermore aimed to involve artists in the discussion and explicitly encouraged women to join the workshop in order to provide visibility against the background of a still rather low gender balance in the field of music and AI. There is no human or animal experimental research involved. The authors declare no conflicts of interest.

## References

- Agostinelli, A., Denk, T. I., Borsos, Z., Engel, J., Verzetti, M., Caillon, A., ... Frank, C. (2023s). MusicLM: Generating Music From Text. <https://doi.org/10.48550/arXiv.2301.11325>
- Akkermann, M. (2017). *Zwischen Improvisation und Algorithmus: David Wessel, Karlheinz Essl und Georg Hajdu*. Schliengen: Edition Argus.
- Ananny, M., & Crawford, K. (2018). Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
- Assayag, G., Feichtinger, H. G., & Rodrigues, J. F. (Eds.). (2002). *Mathematics and Music*. Berlin, Heidelberg: Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-662-04927-3>
- Avdeeff, M. (2019). Artificial intelligence & popular music: SKYGGE, flow machines, and the audio uncanny valley. *Arts*, 8(4), 130. MDPI. <https://doi.org/10.3390/arts8040130>.
- Becker, B. A., Denny, P., Finnie-Ansley, J., Luxton-Reilly, A., Prather, J., & Santos, E. A. (2023). Programming Is Hard - Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation. *Proceedings of the 54th ACM Technical Symposium on Computer*

*Science Education V. 1*, 500–506. New York, NY, USA: Association for Computing Machinery.  
<https://doi.org/10.1145/3545945.3569759>

Bell, E. (2021). Cybernetics, Listening, and Sound-Studio Phenomenotechnique in Abraham Moles's *Théorie de l'information et perception esthétique* (1958). *Resonance*, 2(4), 523–558. <https://doi.org/10.1525/res.2021.2.4.523>

Birtchnell, T. (2018). Listening without Ears: Artificial Intelligence in Audio Mastering. *Big Data & Society*, 5(2), 2053951718808553. <https://doi.org/10.1177/2053951718808553>

Born, G. (2007). (Im)Materiality and Sociality: The Dynamics of Intellectual Property in a Computer Software Research Culture. *Social Anthropology*, 4(2), 101–116.  
<https://doi.org/10.1111/j.1469-8676.1996.tb00319.x>

Born, G. (2010). For a Relational Musicology: Music and Interdisciplinarity, Beyond the Practice Turn. *Journal of the Royal Musical Association*, 135(2), 205–243.  
<https://doi.org/10.1080/02690403.2010.506265>

Born, G. (2020). *Diversifying MIR: Knowledge and Real-World Challenges, and New Interdisciplinary Futures*. 3(1), 193–204. <https://doi.org/10.5334/tismir.58>

Born, G. (2021). Artificial intelligence, music recommendation, and the curation of culture: A white paper. *Schwartz Reisman Institute for Technology and Society, University of Toronto, CIFAR*.

Bown, O. (2021). *Beyond the Creative Species: Making Machines that Make Art and <usic*. MIT Press.

Braguinski, N. (2022). *Mathematical Music: From Antiquity to Music AI*. London New York: Routledge, Taylor & Francis Group.

Briot, J.-P., Hadjeres, G., & Pachet, F.-D. (2020). *Deep Learning Techniques for Music Generation*. Cham: Springer.

Cambouropoulos, E., & Kaliakatsos-Papakostas, M. (2021). Cognitive Musicology and Artificial Intelligence: Harmonic Analysis, Learning, and Generation. In E. R. Miranda (Ed.), *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity* (pp. 263–281). Cham: Springer. [https://doi.org/10.1007/978-3-030-72116-9\\_10](https://doi.org/10.1007/978-3-030-72116-9_10)

Christensen, T. (Ed.). (2002). *The Cambridge History of Western Music Theory*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CHOL9780521623711>

Copet, J., Kreuk, F., Gat, I., Remez, T., Kant, D., Synnaeve, G., Défossez, A. (2024). *Simple and Controllable Music Generation*. arXiv. <https://doi.org/10.48550/arXiv.2306.05284>

- Cosme-Clifford, N., Symons, J., Kapoor, K., & White, C. Wm. (2023). Musicological Interpretability in Generative Transformers. *2023 4th International Symposium on the Internet of Sounds*, 1–9. <https://doi.org/10.1109/IEEECONF59510.2023.10335202>
- Cuthbert, M. S., & Ariza, C. (2010). music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data. *Proceedings of the 11th International Society for Music Information Retrieval Conference*. <https://doi.org/https://doi.org/10.5281/zenodo.1416114>
- Défossez, A., Copet, J., Synnaeve, G., & Adi, Y. (2022). *High Fidelity Neural Audio Compression*. arXiv. <https://doi.org/10.48550/arXiv.2210.13438>
- Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., & Sutskever, I. (2020). *Jukebox: A Generative Model for Music*. arXiv. <https://doi.org/10.48550/arXiv.2005.00341>
- Eriksson, M., Fleischer, R., Johansson, A., Snickars, P., & Vonderau, P. (2019). *Spotify Teardown: Inside the Black Box of Streaming Music*. MIT Press.
- Gardner, J., Durand, S., Stoller, D., & Bittner, R. M. (2024). *LLark: A Multimodal Instruction-Following Language Model for Music*. arXiv. <https://doi.org/10.48550/arXiv.2310.07160>
- Gioti, A.-M. (2020). From artificial to extended intelligence in music composition. *Organised Sound*, 25(1), 25–32.
- Gotham, M. R. H., Song, K., Böhlefeld, N., & Elgammal, A. (2022). Beethoven X: Es Könnte Sein! (It Could Be!). *Proceedings of the 3rd Conference on AI Music Creativity*. AIMC 2022. <https://doi.org/10.5281/zenodo.7088335>
- Hanemaayer, A. (Ed.). (2022). *Artificial Intelligence and Its Discontents: Critiques from the Social Sciences and Humanities*. Cham: Springer International Publishing. <https://doi.org/10.1007/978-3-030-88615-8>
- Haworth, C. (2021). Music and Cybernetics in Historical Perspective: Introduction to the Special Issue Edited by Christopher Haworth and Eric Drott. *Resonance*, 2(4), 461–474. <https://doi.org/10.1525/res.2021.2.4.461>
- Hesmondhalgh, D., Campos Valverde, R., Kaye, D., & Li, Z. (2023). The impact of algorithmically driven recommendation systems on music consumption and production: A literature review. *UK Centre for Data Ethics and Innovation Reports*.
- Huang, C.-Z. A., Hawthorne, C., Roberts, A., Dinculescu, M., Wexler, J., Hong, L., & Howcroft, J. (2019). *The Bach Doodle: Approachable music composition with machine learning at scale*. arXiv. <https://doi.org/10.48550/arXiv.1907.06637>
- Huang, Q., Jansen, A., Lee, J., Ganti, R., Li, J. Y., & Ellis, D. P. W. (2022). MuLan: A Joint Embedding of Music Audio and Natural Language. <https://doi.org/10.48550/arXiv.2208.12415>



Iverson, J. (2018). *Electronic Inspirations: technologies of the Cold War musical avant-garde*. New York, NY: Oxford University Press.

James, R. (2019). *The Sonic Episteme: Acoustic Resonance, Neoliberalism, and Biopolitics*. Durham, NC: Duke University Press.

Kane, C. L. (2019). *High-Tech Trash: Glitch, Noise, and Aesthetic Failure*. Oakland, California: University of California Press.

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., ... Kasneci, G. (2023). ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education. *Learning and Individual Differences*, 103, 102274.  
<https://doi.org/10.1016/j.lindif.2023.102274>

Kreuk, F., Synnaeve, G., Polyak, A., Singer, U., Défossez, A., Copet, J., ... Adi, Y. (2023). *AudioGen: Textually Guided Audio Generation*. arXiv.  
<https://doi.org/10.48550/arXiv.2209.15352>

Lee, K., Cooper, A. F., & Grimmelmann, J. (2024). Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain (The Short Version). *Proceedings of the Symposium on Computer Science and Law*, 48–63. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3614407.3643696>

Lester, T., & Pachamanova, D. (2017). The Dilemma of False Positives: Making Content ID Algorithms more Conducive to Fostering Innovative Fair Use in Music Creation. *UCLA Entertainment Law Review*, 24(1). <https://doi.org/10.5070/LR8241035525>

Lubart, T. (2005). How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies*, 63(4-5), 365–369.

Miller, B. A. (2021). Leonard Meyer's Theory of Musical Style, from Pragmatism to Information Theory. *Resonance*, 2(4), 475–502. <https://doi.org/10.1525/res.2021.2.4.475>

Miranda, E. R. (2021). *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*. Cham: Springer.

Nierhaus, G. (2009). *Algorithmic Composition: Paradigms of Automated Music Generation*. Wien: Springer.

Novak, D., & Sakakeeny, M. (Eds.). (2015). *Keywords in Sound*. Durham, NC: Duke University Press.

Offert, F. (2023). Can We Read Neural Networks? Epistemic Implications of Two Historical Computer Science Papers. *American Literature*, 95(2), 423–428.  
<https://doi.org/10.1215/00029831-10575218>

- Pachet, F. (2016). A Joyful Ode to Automatic Orchestration. *ACM Transactions on Intelligent Systems and Technology*, 8(2), 18:1–18:13. <https://doi.org/10.1145/2897738>
- Rehding, A. (2021). Music Theory's Other Nature: Reflections on Gaia, Humans, and Music in the Anthropocene. *19th-Century Music*, 45(1), 7–22. <https://doi.org/10.1525/ncm.2021.45.1.7>
- Roberge, J., & Castelle, M. (Eds.). (2021). *The Cultural Life of Machine Learning: An Incursion into Critical AI Studies*. Cham, Switzerland: Palgrave Macmillan. <https://doi.org/10.1007/978-3-030-56286-1>
- Seaver, N. (2022). *Computing Taste: Algorithms and the Makers of Music Recommendation*. Chicago, IL: University of Chicago Press.
- Steinmetz, C. J., Bryan, N. J., & Reiss, J. D. (2022). Style Transfer of Audio Effects with Differentiable Signal Processing. *Journal of the Audio Engineering Society*, 70(9), 708–721. <https://doi.org/10.17743/jaes.2022.0025>
- Sterne, J., & Razlogova, E. (2019). Machine Learning in Context, or Learning from LANDR: Artificial Intelligence and the Platformization of Music Mastering. *Social Media + Society*, 5(2), 2056305119847525. <https://doi.org/10.1177/2056305119847525>
- Tatar, K., & Pasquier, P. (2019). Musical agents: A typology and state of the art towards musical metacreation. *Journal of New Music Research*, 48(1), 56–105.
- Thompson, M. (2017). *Beyond Unwanted Sound: Noise, Affect and Aesthetic Moralism*. New York: Bloomsbury Academic.
- White, C. W., & Quinn, I. (2018). Chord Context and Harmonic Function in Tonal Music. *Music Theory Spectrum*, 40(2), 314–335. <https://doi.org/10.1093/mts/mty021>
- Williams, R. (1985). *Keywords: A Vocabulary of Culture and Society*. Oxford University Press, USA.