Leveraging small datasets for ethical and responsible AI music making

Nick Bryan-Kinns
University of the Arts London
London, UK
n.bryankinns@arts.ac.uk

Anna Wszeborowska
University of the Arts London
London, UK
a.wszeborowska@arts.ac.uk

Olga Sutskova
University of the Arts London
London, UK
o.sutskova@arts.ac.uk

Elizabeth Wilson
University of the Arts London
London, UK
e.j.wilson@arts.ac.uk

Phoenix Perry
University of the Arts London
London, UK
phoenix.perry@arts.ac.uk

Rebecca Fiebrink
University of the Arts London
London, UK
r.fiebrink@arts.ac.uk

Gabriel Vigliensoni
Concordia University
Montréal, Canada
gabriel.vigliensoni@concordia.ca

Rikard Lindell
Dalarna University
Falun, Sweden
rli@du.se

Andrei Coronel
Ateneo de Manila University
Quezon City, Philippines
acoronel@ateneo.edu

Nuno N. Correia
Tallinn University
Tallinn, Estonia
ITI/LARSyS
Lisbon, Portugal
nuno.correia@tlu.ee

Abstract

The impact of Artificial Intelligence is felt on every stage of contemporary musicking and is shaping our interaction with sound. Deep learning Generative AI (GenAI) systems for high-quality music generation rely on extremely large musical datasets for training. As a result, AI models tend to be trained on dominant mainstream musical genres, such as Western classical music, where large datasets are more readily available. In addition, the reliance on extremely powerful computing resources for deep learning creates barriers to use and negatively impacts our environment. This paper reports on contemporary concerns and interests of musicians, researchers, and music industry stakeholders in the responsible use of GenAI models for music and audio. Through analysis of focus group discussions and exemplar case studies of the use of GenAI in music making at a hybrid workshop of 148 participants, we offer insights into current discourses about the use of GenAI beyond dominant musical styles and suggest ways forward to increase creative agency in music making beyond the mainstream. Our findings highlight the value of small datasets of music for GenAI, the suitability of AI models for working with small datasets of music, and pose questions around what constitutes a 'small' dataset of music.

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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
AM.ICAD '25, June 30–July 04, 2025, Coimbra, Portugal
© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Author pre-print

CCS Concepts

• **Applied computing → Sound and music computing**; • **Computing methodologies → Artificial intelligence**; • **Human-centered computing**;

Keywords

Music, Artificial Intelligence, Responsible AI, Ethics, Generative AI, Small datasets, Low-resource AI models

ACM Reference Format:
Nick Bryan-Kinns, Anna Wszeborowska, Olga Sutskova, Elizabeth Wilson, Phoenix Perry, Rebecca Fiebrink, Gabriel Vigliensoni, Rikard Lindell, Andrei Coronel, and Nuno N. Correia. 2025. Leveraging small datasets for ethical and responsible AI music making. In *Proceedings of AudioMostly 2025 (AM.ICAD '25).* ACM, New York, NY, USA, 12 pages.

1 Introduction

In contemporary music making the impact of Artificial Intelligence (AI) is felt across creative practice, from composition [17, 29], interpretation [10, 52], improvisation [21, 40, 43], to accompaniment [34], and across the Music Industries from creation and production [26, 33], protection [12, 24], distribution [1], to consumption [39]. The breadth of reach of AI systems in music making raises pressing questions about how AI, especially Generative AI (GenAI) systems, impacts our interaction with sound and music, how we play together, and how GenAI might foster or hinder creativity. Music generation systems, available both commercially and in research labs [2], rely heavily on deep learning. For example, the Magenta suite[1] can be used to generate both symbolic music and audio files, while Udio[2] and Suno[3] can be used to generate full-length songs from text prompts directly in audio form. These systems depend on extremely large datasets. This reliance introduces several issues: bias in musical output due to training data availability and selection [5], limited accessibility and reduced creative agency when using such tools [47], significant environmental impact [14], and the growing centralisation and control of music-making infrastructure under a small number of large corporations.

To gain insights into the opportunities, concerns, and current practices around music making and sonic interaction with GenAI, we organised an open one-day hybrid participatory stakeholder workshop in London, UK, in July 2024, followed by a closed online data analysis workshop in October 2024 with invited expert researchers. These workshops invited the contemporary discourses of musicians, researchers, and music industry stakeholders around the use of GenAI models in music, audio, and sonic interaction. Based on the analysis of discussions and data collected from these workshops, this paper reflects on these discourses and offers insights into how GenAI could be deployed responsibly and ethically beyond dominant musical styles and large AI models, and suggests ways forward to increase creative agency in music making beyond the mainstream using small datasets and low-resource AI models.

2 Background

State-of-the-art GenAI systems for music, such as Magenta, Suno, or Udio, rely on extremely large datasets to train the models, whether they generate symbolic outputs such as MIDI or audio output. These datasets primarily represent dominant, mainstream musical genres where substantial training material is available [9], such as Western[4] Classical music, with datasets like MuseData[5] offering over 4,500 compositions and nearly a million notes just for J. S. Bach, or the Lakh MIDI Dataset [38][6] comprising over 170,000 unique files mostly of Western pop music. Research estimates that less than 6% of available music datasets used for AI model training represent non-Western music [31]. Whilst there have been initiatives to build music collections of smaller and less mainstream musical forms, such as Dunya [37] which brought together music corpora of five music traditions[7] these are not the norm and do not offer the scale of data needed to train commercial deep learning AI models. As a result, AI systems become biased toward genres with abundant data and struggle to generalise to other musical traditions [3, 13]. For example, it is not possible to use contemporary deep learning AI models to generate music such as Qin genre in China, nor genres of contemporary subcultures such as glitch or algorithmic music. In many cases, it is also simply not possible to train these GenAI models on genres beyond those that the AI models and architectures have been optimized to. For example, it may be difficult to train deep learning models on musical genres beyond the mainstream due to differences in musical tunings and timings [5], timbral qualities, and notation (or lack of it). Moreover, the time and effort required to construct large datasets of music can in itself be a barrier to using GenAI for music making, despite the potential for data augmentation to increase the size of small datasets [16].

2.1 Marginalisation and Bias

As GenAI increasingly becomes an integral part of music making [48], especially in tools such as digital audio workstations, the biased nature of GenAI models will further marginalise and exclude genres of music outside the dominant training sets - either the tools will not be usable with these genres, or the tools will implicitly push features of less mainstream genres to conform more to mainstream stylist features such as a 4/4 time signature or reliance on 12-note equal temperament musical scales and notations. Moreover, GenAI models developed for mainstream music generation are unlikely to be useful for systems for interacting with sound such as auditory display, new musical instruments, and so on as they restrict output to dominant musical forms. This severely constrains the sonic interaction design space when using these deep learning GenAI systems, for example constraining audio interface design to adhering to 4/4 time signatures.

In addition, the reliance on extremely large AI models and datasets needed for deep learning GenAI [44] requires substantial computational power for both training and inference, consuming large amounts of electricity. This creates financial and technological barriers to access and creates damaging environmental impact for deep learning GenAI music systems [14]. In response to this, researchers call for more thoughtful and considered use of deep learning GenAI systems to reduce environmental impact, contribute to climate solutions, and develop more responsible approaches to GenAI development and use [35].

And, finally, it is worth noting that the most advanced GenAI tools for music are owned and trained by large multinational companies that monopolise music generation. Most of these companies rely on huge training sets scraped from the internet with limited or no attribution about who created the original music [15]. These potential infringements of musicians' creative rights have substantial negative impacts for musicians "making it harder for human composers to gain recognition and earn a fair income" [18]. The UK Government's recent consultation on Copyright and AI [23] had three main objectives, one of which was "Supporting right holders' control of their content and ability to be remunerated for its use" reflecting the importance of protecting creative rights and livelihoods in the face of growing (mis)use of creative content to train large AI models. Whilst the consultation highlights the importance of fair use of creative materials, it is worth noting that experts in Responsible AI argue that the UK Government's proposed approach to copyright and AI would not meet this objective [41].

It is possible to fine-tune many deep learning models to produce different genres of music than they were originally trained on [8], to use style transfer to change generated music from one genre to another [32], and to use timbre transfer to apply the timbre of one sound to another [11]. Such approaches might offer ways

[1] https://magenta.tensorflow.org
[2] https://www.udio.com
[3] https://suno.com/home
[4] We acknowledge that the term *Western* is contentious here and use it as a well understood convention for music largely produced by European and North American musicians, whilst acknowledging that many of the forms of Western music originated beyond these geographical limitations and were appropriated by what we now refer to Western musical styles.
[5] MuseData https://musedata.org
[6] https://colinraffel.com/projects/lmd/
[7] https://dunya.compmusic.upf.edu

to generate music in non-dominant musical styles. However, the results of such fine-tuning and style transfer rapidly deteriorate in quality the further away the intended genre is from the genre(s) that the AI model was originally trained on.

2.2 Low-Resource AI Models

Recent advances in AI research explore 'low-resource' approaches to music classification and generation [25, 36] which have the potential to make small datasets usable in GenAI models thereby reducing AI bias. Here, the term *low resource* refers to models that are specifically designed or adapted to operate with limited computational resources (e.g. CPUs or embedded devices rather than cloud-based GPUs) and modest amounts of training data, which can often be several orders of magnitude smaller than standard deep learning models [7]. Techniques such as model pruning (removing redundant parameters) [42, 51], quantisation (reducing numerical precision) [42], knowledge distillation (training smaller models using the outputs of larger ones) [22], and efficient architecture design (like transformers optimised for edge devices) [27] make it possible to reduce model size and complexity while retaining performance.

This framing aligns with the broader paradigm of *frugal computing* [45] that recognises that computing resources current emissions are almost 4% of the world total, and further that by 2040 emissions from computing alone will account for more than half of the emissions budget to keep global warming below 1.5°C. (ibid.). Frugal computing emphasises the need for treating computational resources as finite and precious, and that should be used effectively where possible. Low-resource modelling attempts to shift towards this ethos by prioritising sufficiency over scale—developing models that perform effectively under constrained conditions rather than pursuing ever-larger architectures. This approach reframes technical efficiency as a form of environmental responsibility, particularly relevant in the arts and humanities, where computational needs often intersect with cultural, economic, and geographic constraints.

In the context of musical AI, unlike conventional deep learning systems that rely on massive datasets (e.g. millions of data points) and energy-intensive training and inference processes, *low-resource models* can work effectively with datasets as small as a few thousand data points, running locally on personal laptops or embedded platforms. However, such approaches have not been explored with datasets of music marginalised by AI and are not readily available for use by musicians. At the same time, researchers and artists have explored the use of small datasets with GenAI models, questioning assumptions about the value of using huge datasets with big AI models [47]. These approaches suggest artistic directions in which smaller models trained on small datasets are valued and explored, and open up discourses about how such approaches could reduce the marginalisation of music forms outside the mainstream. This paper builds on these directions to explore current discourses around the use and potential of GenAI for music making outside dominant musical genres and AI models.

3 Methods

To better understand contemporary concerns and practices of musicians and music industry stakeholders who create music and are working with GenAI, we convened an international network of partners in late 2023 as part of the 12-month MusicRAI project[8]. The MusicRAI project built an international community to address Responsible AI (RAI) challenges of bias in AI music generation and analysis. MusicRAI project partners included researchers, musicians, and music industry stakeholders from Canada, China, Estonia, Germany, Philippines, Sweden, UK, and the USA. We then organised a one-day hybrid workshop at the University of the Arts London, UK, in July 2024. We invited participants to the workshop by advertising in our own communities, on email lists, and through social media. For the workshop, we focused on exploring contemporary uses of AI with musical styles outside dominant musical forms as a way to probe questions of the challenges and concerns of working with major AI models. 148 people took part in the workshop, with 70 people participating in-person and 78 registered for online participation. Figure 1a illustrates the range of types of sectors that participants identified as being in as part of the workshop registration process - participants could select multiple sectors. The most common sectors of activity of our participants were 'Music' (94), 'Technology' (89) and 'Art and Design' (69), followed by 'Education' (53) and 'Media' (24) with a small number selecting 'Public sector' (3), 'Charity' (2), and 'Healthcare' (2), and 20 selecting 'Other'. The workshop was hosted by 4 facilitators. Figure 1b illustrates the overlap in sectors of the most popular sector types - for example, 53 people selected 'Music' and 'Technology', and 9 people selected 'Music', 'Technology', and 'Art and Design'. Data collection was undertaken in the workshop through:

- Two panel discussions about the responsible use of GenAI and music illustrated in figure 2a. One on "Challenges and Opportunities for Music Creation" and the second on "The Future of Music Creation". These panels brought together experts, including music industry leaders, musicians, and academics, to provoke audience discussion and inform topics in the focus groups.
- Eleven case studies of contemporary uses of GenAI beyond mainstream music genres, presented by participants to help provoke discussion and provide concrete examples in later focus groups, illustrated in figure 2b.
- Two focus group activities around topics of responsible use of GenAI and music.

Data collected from the case studies and focus groups was analysed (Section 3.3) to identify recurring topics and themes in contemporary discourses around GenAI and music (Section 4).

3.1 Case Studies

Participants who registered for the workshop were offered the opportunity to present a short case study of their use of AI models with forms of music beyond mainstream styles. The case study topics were reviewed by the workshop organisers and all case studies were invited to be presented at the workshop. We also invited our project partners to present case studies of their artistic practice and academic research if they wished to resulting in 11 case studies - 5 from participants and 6 from partners. These were intended to allow us to capture both current artistic practice and academic research in the area, and to offer talking points to provoke discussion in the focus groups.

[8]http://musicrai.org

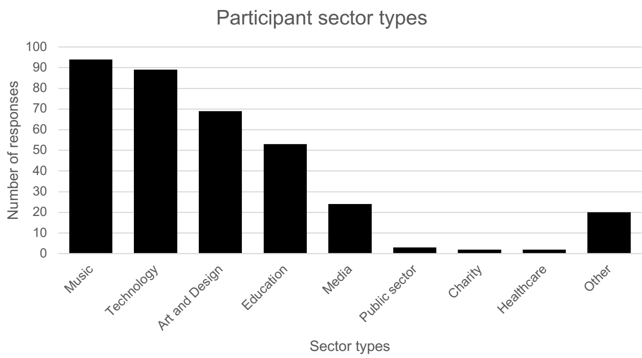

(a) Participant sector types at the workshop (participants could select multiple sectors)

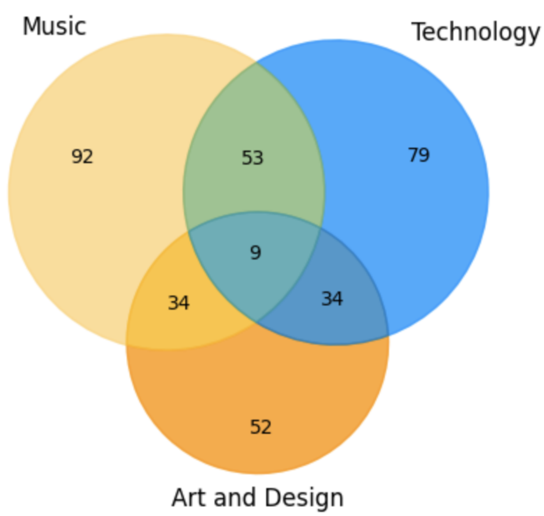

(b) Overlapping sector types for Music, Technology, and Arts and Design categories (numbers indicate number of responses)

Figure 1: Participant background and demographic data collected.

Topics of the case studies included the use of small datasets of music, the use of open AI models both large and small, and musical tasks from composition to production, performance and installation. For example, case study topics included: supporting community-led language reclamation, musically embodied machine learning, machine learning with original Pilipino music, research and development at Steinberg[9] on using small datasets of music, steerability and embodiment of latent audio models through interactive machine learning, the open-source AI music generation system Stable Audio Open[10] by Stability AI which can be used to produce short form music, personalised AI with small datasets for artists, neural audio generation with minimal data, neural audio synthesis and small datasets, real-time site specific acoustic experiential qualities, user-centric intelligent and assistive multitrack music mixing, and the opportunities and challenges of using large foundation models with small datasets. Readers are directed to the workshop website[11] for details of all the case studies.

(a) Workshop panel and audience

(b) Case study of real-time site-specific acoustic experiential qualities

Figure 2: Participants at the workshop delivering case study presentations and panel discussion.

To illustrate the diversity of case studies, three examples are briefly summarised here. First, the *Clastic Music* project focused on modelling drum patterns and rhythms of non-mainstream genres that exhibit tempo octaves—where tempo can be perceived and embodied at its original speed, half speed, or double speed—and often use compound metre. To address this, the project created several small-scale datasets of these complex rhythmic patterns and developed a custom AI network architecture, resulting in a full-length music album presented internationally [46].

As a second example case study, researchers found that artists still often rely on physical medium storage and have developed their own search and composition strategies based on file management

[9]https://www.steinberg.net
[10]https://huggingface.co/stabilityai/stable-audio-open-1.0
[11]http://musicrai.org

tools [28]. In this case study the researchers explore what designs enable artists to use and reuse artistic material from their personal repositories, and what designs support reflection on artistic decisions and how this process could be reinforced by AI. For example, the researchers found that using Kohonen networks for features of audio files such as centroid, spread, and kurtosis, allowed users to create music from their dataset with serendipitous results [19].

Finally, the third example case study focused on research on Original Pilipino Music (OPM) - a genre of commercial music in the Philippines. OPM is inconsistently defined with one view being that OPM is commercial music (e.g. pop, rock, dance, hip-hop) that is original and composed or produced in the Philippines. In contrast, a more purist perspective considers OPM as contemporary popular music that incorporates strong elements of traditional Filipino music, particularly styles such as kundiman and harana. Classifying OPM subgenres via machine learning is a foundational step towards the automated generation of music that aligns with the OPM genre. However, this application of GenAI faces several challenges, particularly in the lack of availability and low quality of metadata associated with OPM music tracks. The challenge becomes even more pronounced when focusing on indigenous music (not OPM), as machine learning-ready datasets for indigenous Philippine music are still under development. Although government and academic agencies maintain databases of indigenous music sounds, these collections typically do not include complete music tracks.

3.2 Focus Groups Sessions

Two one-hour focus group sessions were organised in the style of a community workshop to capture current discourses around the use of GenAI for music. GenAI relies on i) Data to train the AI models and ii) AI models themselves, and so we oriented our two focus groups to capturing discourses and ethical concerns about: 1) Music datasets, and 2) AI models.

In each focus group session, the workshop participants were split into equal sized self-selecting groups and asked to explore topics around AI and music (outlined below) in 40 minutes documenting their discussions using post-it notes and large sheets of paper (Figure 3a). The participants were then asked to report back the main topics of their discussions to the whole workshop in a five-minute summary per group (Figure 3b). The reporting back to the group was video recorded and the post-it notes were photographed for later analysis.

3.2.1 Focus group #1 (5 groups): Understanding current discourse about music genres and datasets that are marginalised by deep learning approaches. Participants were first asked to discuss as a group the music styles and datasets that are marginalised by deep learning AI models focussing particularly on current approaches, opportunities, discourses, and challenges. Participants were then asked to write on post-it notes responses to the following prompts:

- The names and styles or genres of datasets used by participants or known to participants.
- The special features of these datasets of music e.g. features of the music itself, stylistic features, musical notations (or lack of), cultural context and meaning, history, instrumentation, access, etc.
- How these datasets are collected, processed, and shared.

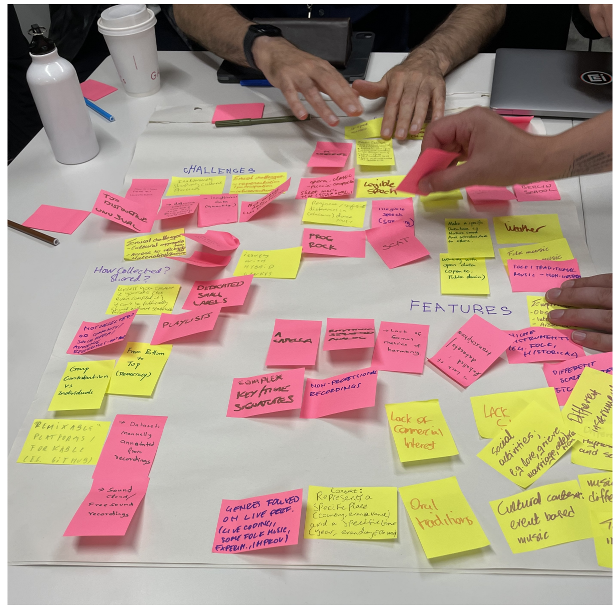

(a) Focus group brainstorming with post-it notes

(b) Reporting back on focus group discussions

Figure 3: Participants at the workshop documenting and presenting their discussions.

- The challenges with using these datasets with large AI models such as Udio or Suno, and why these challenges happen.

Participants were then asked to group their post-it notes together to capture the key features of these datasets and how might these be shared, searched, and used in AI models. One person from the group then reported back to the whole workshop a summary of their discussions and groupings.

3.2.2 Focus group #2 (4 groups): Understanding current discourse around AI models and architectures for generating music. In this activity participants were asked to discuss in their groups the opportunities and challenges of using AI models to work with datasets and styles of music identified in focus group #1 and to write on post-it notes responses to the following prompts:

- What might be suitable AI models and architectures to work with the datasets identified in focus group #1.
- Why these approaches might reduce bias in AI music.
- How these AI models might be used for other datasets.
- The special features of these AI models.

As in focus group #1 participants were then asked to group their post-it notes into key opportunities for using these AI models to work with marginalized datasets and how they might be made more available. One person from the group then reported back to the whole workshop a summary of their discussions and groupings.

3.3 Topic and Theme Identification

Immediately after the workshop, the facilitators summarised their own observations of the topics and themes that had emerged during the community workshop. Recordings of the focus group reporting sessions were then transcribed and an open coding approach [49] was used by two researchers to independently identify key topics that emerged in the discussions and the post-it notes of the focus groups. In the open coding the researchers sifted through the text to organise similar words and phrases into broad topics (ibid.). The topics were then added to a Miro[12] board by the two researchers as post-it notes and together they clustered the topics into themes as both a way of creating a set of themes and coordinating their independent analysis. The Miro board was pre-structured to reflect the workshop structure: Focus group #1, Focus group #2, Key takeaways, Opportunities, and Challenges. We added additional spaces for AI Model taxonomies, and AI Repository requirements in response to our overall research project themes.

A two hour online data analysis workshop was then held in October 2024 with 8 participants who had participated in the stakeholder workshop (project partners, 6 male, 2 female) and 4 facilitators (3 female, 1 male) to collaboratively refine the themes and topics captured in the Miro board. Participants were all senior researchers with extensive experience in AI and music (1 commercial researcher, 7 academic researchers, from institutions in Canada, Estonia, Germany, the Philippines, Sweden, and the UK). We used this expert participant approach to build on the skills and expertise of participants when analysing the grouping the data from the stakeholder workshop. Participants were given access to the following data to inform the collaborative data analysis:

- Video recordings of summaries of each focus group session in the July workshop. These are less than 5 minutes each and there were 9 summary videos.
- Photos of the flipcharts from each focus group session.
- Transcripts of each video.
- The Miro board with virtual post-it notes of topics identified by the researchers. Importantly, these topics were not clustered into themes.
- Participants were also encouraged to look back at their own notes from the workshop.

Participants were asked as a group to review the data from focus group #1 and then focus group #2 and to cluster the post-it notes of topics into themes. Based on the participants' review of the data and their personal notes from the workshops, participants were encouraged to add new post-it notes of topics that they thought had been overlooked in the analysis. Once the clustering was complete, participants compared the clusters to the themes created by the researchers in the initial data analysis. A refined clustering of topics into themes was consolidated and agreed upon by participants by the end of the online workshop. In this way, we aimed to combine the researchers' in-depth analysis of the post-it notes generated in the data collection workshop with individual recollections and perceptions on the topics by the partners who are all experts in the field.

4 Analysis

Figure 4 shows the clusters of post-it notes developed in the online data analysis workshop. The final analysis included 329 post-it notes grouped into 40 topics (Figure 4). The rectangular shapes are the pre-structuring of the Miro board to reflect the structure of the data collection workshop. Eight broad themes emerged from the online data analysis workshop - these are summarised in the sections below. Topics are highlighted in bold in the descriptions and where quotes are given they are from post-it notes, from the data collection workshop or from transcriptions of the focus group reporting. Figure 5 summarises the positive values identified when using small datasets with low resource AI models for music making. The overlapping sections in the diagram bring increased value and we suggest that the intersection of all four circles offers the greatest positive value to music making.

Figure 4: Themes and topics identified in collaborative data analysis

[12]https://miro.com

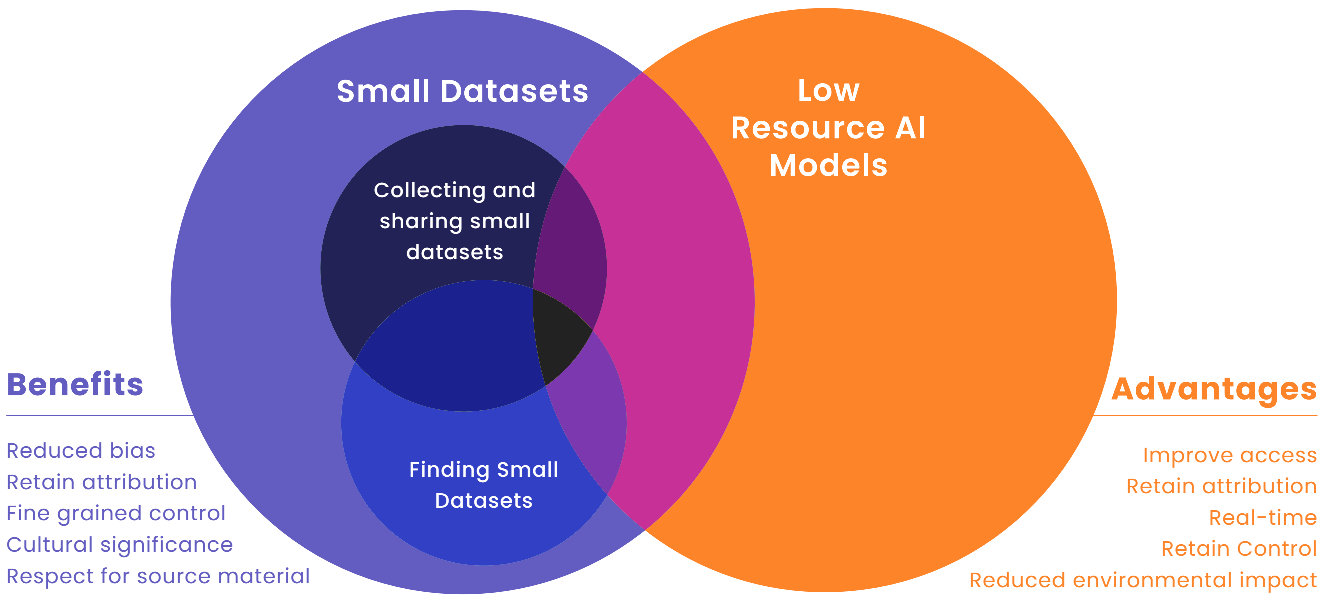

Figure 5: Positive values of using small datasets and low-resource AI models for music making

4.1 Theme: What is a 'Small' Dataset of Music

The concept of 'small' music datasets emerged as a key theme in discussions, reflecting broader questions about its role in musical creativity, AI applications, and the preservation of diverse traditions. This theme explored what defines a small dataset and how they differ from large-scale musical corpora.

Participants raised questions about **what constitutes a 'small' dataset** of music. Whilst it may be "music which doesn't fit a particular genre" there were concerns that focussing on genre was not useful in understanding what a small dataset is - **a small dataset is not defined by its genre**. There are, for example, a wide variety of genres and styles that might be considered small datasets including: traditional music, music performed by traditional or self-made instruments, non-western scales and microtonality, rap/ hip-hop/ grime, percussive guitar, sub-genres of electronic music (e.g. ambient, experimental), non-rhythmical or non-conventional structures of music (e.g. soundscapes, textures), non-diatonic scales, microtonality, forms of harmony which are not traditional western style, field recordings, vocal recordings including legible and illegible speech, Turkish Makam, and more generally niche styles referring to either traditional or sub-genres. Instead it was suggested that small datasets could be usefully thought of in terms of more **granular features of the dataset** such as: musical context (e.g. recording location, time, scenario, mood, or person); use of melody; pitches used; chord progressions; arrangements; musical structures; variability in time signatures; dynamics; the role of gestures; and tone. From this point of view, a small dataset is a set of musical examples which are consistent within themselves, rather than trying to generalise across different styles, genres, rhythms, sound types, and so on.

The use of a range of **representations** in small datasets was also considered a defining feature that differentiated small datasets from large datasets of music, which typically contain solely audio or symbolic representations. For example, whether the representation included symbolic scores and notations, audio recordings, multimodal information, or other metadata about the music, its cultural context, musical instruments used, musical performance techniques and instrument articulation, and so on. The **purpose of small datasets** was also explored as a distinguishing feature. For example, whether the datasets are collected for archives, live performance, or exhibitions.

4.2 Theme: Collection and Sharing of Small datasets

Discussions highlighted ethical concerns around the collection and use of small music datasets, particularly regarding issues of agency, consent, cultural appropriation, and data representation. This theme explored how these datasets could be managed in ways that respect the originating musical communities while balancing accessibility and preservation.

Participants raised concerns about the potential for collecting small datasets from communities becoming 'helicopter' research where the datasets are used and exploited by other academic, artistic, and commercial groups without any benefit or control being offered to the source musical community - reflecting established concerns about where the agency and value lays in cross-cultural research endeavours [4]. Similarly there were concerns about cultural appropriation of music and the misuse of datasets. To address these issues participants emphasised the need for small datasets to be **collected and used within musical communities** themselves. At their smallest these become **personally curated datasets** which might be collected over extended periods of times, such as collection over a 20 year timespan. However, it was noted that whilst these forms of datasets are often well documented and annotated, they are typically idiosyncratically annotated and would need improved annotations to increase their use and potential for sharing. Where datasets are collected by groups outside the originating musical community the importance of **informed consent** for use was emphasised, along with the need for **retractable consent** and detailed attribution of the sources of music in the datasets as could be achieved using structured datasheets [20]. However, it was noted that when small datasets are shared, their **categorisation is reductive** and does not preserve or attempt to represent their cultural meaning or context. Indeed, it was noted that all data representations are reductive and inherently leave things out. This reductive nature is potentially more problematic for under-represented forms of music where context may be important to support culturally sensitive use. Finally, licencing, copyright, or government policies which were established to protect **intellectual property may unintentionally become barriers** to access and sharing of small datasets by the people who created them.

4.3 Theme: Finding Small Datasets

Small datasets are **scattered around the Internet** which participants noted makes them difficult to find. Moreover, small datasets of music are typically organised using simple taxonomies, such as broad genre categories presented as text file names or brief text descriptions that reduce the heritage, background and cultural history on which the datasets were created often to a single filename. Participants noted that it could be interesting to explore how unsupervised machine learning and classification techniques could be used to classify or match small datasets in order to find some coherence and **decolonise taxonomies** away from genre based classifications and their inherent hierarchical and colonial era biases. Some participants expressed a more radical view that maybe datasets should not be collected and shared, but rather used solely by individuals in their music making process.

4.4 Theme: Using Small Datasets

Data preparation issues were a key challenge in using small datasets highlighted by participants. For example, audio quality issues around noise in recordings, the use of non-permanent media to collect audio, data pollution, lack of segmentation of audio recordings, and recordings which include large amounts of noise or silence (for example field recordings). In addition, the difficulty of using small datasets with AI models creates barriers to participation and representation and may further marginalize creative communities who are beyond the mainstream. The nature of small datasets means that they **might not be usable for all music making tasks**. For example, smaller datasets might not be suitable for AI composition tasks but might instead be useful for other purposes such as timbre transfer which could then be applied to an AI model that can compose music. Moreover, some datasets might be too distinctive or unusual to be useful in music making at all. Participants also noted a distinction between **what is captured in small datasets and what is generated** using these datasets - these are not necessarily the same, meaning that a small dataset of music might be used to generate music with different musical features. Finally, concerns were raised about the financial costs of storing and using datasets - who pays for the hosting and managing of datasets, combined with a **lack of commercial interest** in the use of small datasets, and scarce funding for developing AI models tailored to working with small data. This lack of interest limits the range of available AI tools that are able to work with small datasets. Indeed, it was noted that current narratives on AI and music mostly focus on large datasets and the wider music community, rather than thinking more granularly. For example, exploring how small datasets and low-resource AI model approaches could foster new creative tools rather than imposing big data tools on specialised music communities.

4.5 Theme: Value of Small Datasets

This theme captures the advantages that small music datasets offer in both AI music making and creative practice. For example, supporting deeper engagement with cultural context, providing fine-grained creative control, and resisting the homogenising effects of large-scale deep learning models.

Participants noted that the value of working with small datasets includes taking into account under-represented and **marginalized musical communities**, and **reducing bias** of AI music making more broadly. Smaller datasets are often more structured making them easier for an AI model to learn, and are **more nuanced** than large datasets which typically smooth out nuances in training data. Small datasets are also easier for musicians and artists to **access** as they can be used locally on a personal computer rather than access being mediated through cloud services. Communities and archives value small datasets of music for their **unique properties and cultural significance** and are often dedicated and enthusiastic about the music found in small datasets. Finally, working creatively with small datasets has advantages over working with deep learning models trained on large datasets including: greater **understanding and respect** of the source material and its origins, increased **fine-grained control** of AI in creative practice, and **production of forms of music beyond the homogenising trends** of deep learning models. Indeed, music generated from small datasets could be much more interesting to creatively work with and produce more interesting musical results than from a large homogenous pop music datasets.

4.6 Theme: AI Models for Small Datasets

Whilst music AI systems are dominated by closed commercial tools such as Suno or Udio, approaches that are more suited to use as **tools for working with small datasets** of music include architectures such as RNN, transformers, variational autoencoders, diffusion models, LSTM, and models such as RAVE [6] and SampleRNN [30]. These are typically **low-resource architectures** and models requiring significantly less computing power and resources to train and generate music than commercial deep learning approaches. It was also noted that many models from the past decades of AI research were effective and efficient for specific musical tasks or music forms and that better use could be made of these existing tried-and-tested AI models and architectures - do not be afraid to use **'old' architectures**.

4.7 Theme: Selecting AI Models

Participants noted that the **suitability of tools is not determined solely by the AI model**, but rather by the support framework around the tools. For example, participants noted that the following were important factors when selecting AI tools and models for working with small datasets of music: workflow integration; data representation(s) supported; user interfaces; ownership and federation of data; policies around data security; and whether the tool is proprietary. **Types of data** in this context include: audio; symbolic data e.g. MIDI; gestural data e.g. for musical instruments; and multimodal data to capture the context of music.

4.8 Theme: Value of Low-Resource AI Models

In addition to a substantially **reduced environmental impact** due to lower energy requirements for training and generation, most low-resource models are open-source, offering greater transparency of use and reduced financial cost. The use of open-source low-resource models and architectures offers opportunities to **improve access** to GenAI. The use of low-resource models with small datasets offers greater opportunities for sandboxing their use and **retaining control** of access to the datasets which do not need to be transferred to large corporations for AI model training. Similarly, **attribution of authorship** of content in the training datasets, generated music, and the models themselves is easier to support and trace with open-source low-resource models and small datasets than with closed deep learning models training on unknown enormous datasets. Finally, low-resource models are often able to produce music in **real-time** on personal computers, making them ideal for performance and improvisation. However, many open-source models rely on high levels of technical expertise such as Python scripting and command line tools to run them, meaning that there is a real lack of off-the-shelf tools for non-developers.

5 Discussion

Academic discourse around datasets and their use in AI training promotes mechanisms for greater documenting [20] and access to these datasets partly as a way to reduce bias in AI models. Intriguingly, no workshop participant proactively argued for greater access to other people's datasets. Instead, participants focussed on the value of small datasets in music making practice and especially the **value of personal, curated datasets** which might have been collected over decades. Given that our workshop participants were mostly musicians and technologists, this suggests that **sharing of datasets is not a high priority** for musicians as might be expected from academic literature. Instead, these small datasets form an integral part of an individual's music making practice rather than being seen as a separate resource to be found and drawn on. This more individual or personal view of small datasets offers an opportunity for re-framing sounds made using small datasets as a compositional space for exploration and creativity. This in turn offers opportunities for the creative use of low-resource AI models which focus on much smaller aspects of music making. Indeed, historically, there has been an affordance of machine learning approaches to making music, for example experimental music, that does not fit with making music in dominant genres.

5.1 Challenges and Opportunities

Participants noted that there is a lack of focus on small datasets when designing new AI architectures which nowadays rely on extremely large training datasets and can produce high quality audio outputs. This results in less advanced development of low-resource AI models compounding the comparatively lower quality of low-resource AI model audio generation. As noted in our analysis, low-resource models are typically not suitable for all aspects of music making, but often perform well or better than deep learning models on specific musical tasks with small datasets of music. And, as noted by participants, the features of low-resource audio generation become part of a creative aesthetic of music made this way. The challenge here is to be able to **identify which low-resource model would work best** with the specific features of a small dataset. The opportunity, on the other hand, is to be able to **create music outside mainstream musical styles** and norms with a transparently and ethically collected dataset. This also offers artists the opportunity to explore and underscore **alternative applications of AI** for music beyond commercially focussed end-to-end audio generation of popular musical genres. There is also an opportunity to study small datasets using low-resource models in ways that would not be possible using deep learning models, offering **fresh insights into under-represented forms of music** which are simply not captured in large datasets for deep learning approaches.

Creating personal, curated small datasets of music offers creative opportunities, but the data and metadata created and used can often be of poor quality. Whilst some research has undertaken to make available carefully curated datasets, or corpora of music (e.g. [37]), these are few and far between. Data representations are necessarily reductive and the challenge is how to select **which aspects of the music to capture** in the datasets. There are opportunities to **improve the quality of datasets** and reduce data pollution by developing and offering preprocessing workflows and standardizations of data and metadata formats. Similarly, there are opportunities to **define AI model training metrics that are perceptually relevant** to musicians rather than being focussed on established Music Information Retrieval tasks and competitions such as the Music Information Retrieval Evaluation eXchange (MIREX) competition tasks[13]. However, a challenge that cuts across these opportunities is the often **unusual and distinctive nature of music** in small datasets making them unamendable to one-size-fits-all approaches.

Whilst sharing of datasets was not a high priority for participants, there was some interest in exploring other people's datasets. Not so much to train their own AI models but rather to explore how other musicians had used the datasets in their music making. One challenge to this is that whilst these personal datasets are often well annotated **they are typically idiosyncratically described**. An opportunity arises here to use machine learning **classification techniques to generate taxonomies** of existing small datasets based on musical features within the datasets rather than genre or style tags. This would also have the advantage of decolonizing metadata and reducing the bias imposed by genre labels. There are also opportunities here for **federated learning** to be employed across small datasets, i.e. rather than sharing small datasets, AI

[13]https://www.music-ir.org/mirex/wiki/MIREX_HOME

models are trained across individually held small datasets to allow **individual curation and ownership of datasets**, hand-in-hand with the scaling and quality advantages of training AI models across multiple datasets.

There are ethical concerns with the use of deep learning models that are trained on **enormous datasets of unattributed music** echoing Epple et al. [15], the **risk of cultural-appropriation** in AI music generation echoing concerns in other forms of cross-cultural creative practice [4], and the **barriers to the use of AI models caused by limited representation** of musical forms beyond the mainstream. Small datasets and low-resource models offer opportunities to address these concerns by offering more **ethical approaches to dataset collection** for example by individual musicians working together to collect datasets in person, robust **attribution** of data origins by carefully documenting in a variety of media the sources of the music, more transparent connections between audio generation and source material, and **informed and revocable consent** through bespoke rights management approaches. To strengthen these opportunities and reduce barriers to more fair and equitable music making with AI, open-source guides and best practices need to be developed and shared on ethical music making practices.

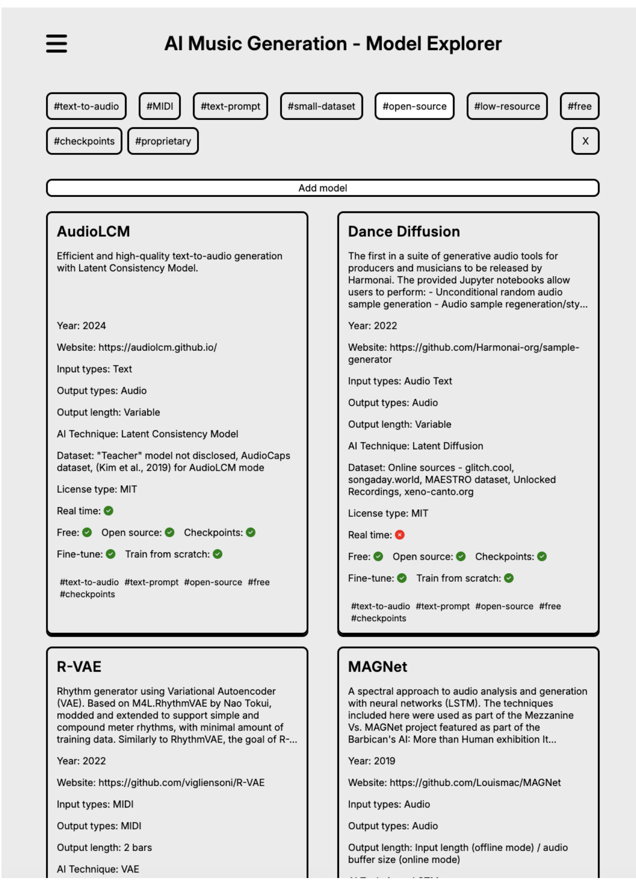

Figure 6: The MusicRAI open repository of AI models for music generation (cropped from screenshot)

5.2 An Open Repository of Generative AI Models for Music

To respond to the challenge of finding AI models suitable for different musical tasks and features we developed an open repository (repo) of generative AI models[14] for music - the MusicRAI open repository illustrated in figure 6. This repo supports open and editable lists of AI models, their training datasets, possibility of fine-tuning and control, licensing, and input and output types. As illustrated in the figure, the collection can be filtered by these properties, helping creative practitioners identify models which could be trained from scratch or fine-tuned on small datasets. It aims to promote low-resource AI models and link to open-access, non-profit archival repositories and AI models which are mostly excluded from other listings. We invite readers to add to the repository to increase the breadth of models capture and shared there.

6 Limitations and Future Work

Whilst the workshop attracted 148 people, generalising the results discussed in this paper to larger music and AI community would require confirmation through larger-scale studies. For example, a structured survey to explore the themes identified in this paper with a wider range of people both in terms of geographical location and music and AI expertise and interests. In addition, questions of the value of small-datasets and low-resource AI model approaches could be further explored through empirical studies or experiments to strengthen the evidence for the value of these approaches. For example, through empirical investigations comparing low-resource versus deep-learning AI models for music making for music generation tasks.

It is worth noting that work reported in this paper did not apply any musicological or algorithmic analyses to the datasets of music discussed in the workshops. This is a potential area for future workshop and would offer insights into how small datasets of music could be better categorised and shared beyond conventional genre labels.

7 Conclusions

Our position, informed by our participatory workshop and subsequent data analysis workshop, is that what is needed are more small dataset and open low-resource AI model approaches to music generation. We suggest from our analysis of current discourses in this paper that this combination offers a more responsible approach to AI music generation which is: **Less biased** with AI models offering access to GenAI for a more inclusive set of musical styles; **Less energy consuming**, reducing financial barriers to use, mitigating environmental impact, and creating more open access to GenAI; **More transparent** in the data collection methods used, recognizing the original creators of music in GenAI model training, and including creative attribution and content provenance across multiple levels of data, AI generation, and stages of music workflow; **More accessible and explainable** by using low-resource models which are simple, quick to train, and low-cost; **More democratic** in its dataset collection and attribution through personally collected,

[14] http://musicrai.org

curated, and documented datasets, including more individual, idiosyncratic, artist-led, grass-roots approaches to dataset collection and music generation. To achieve economies of scale in AI model training with small datasets a **federated learning** approach [50] might help manage AI training and access to personally held small datasets. The value of small datasets and low-resource AI models to sonic interaction design is ripe for further exploration. These approaches offer opportunities for more interactive and responsive audio generation, for example in new musical instruments and auditory displays where deep learning AI models are often too slow, too computationally intensive, and too tied to dominant musical forms to be useful.

Importantly, we believe that to realise more responsible use of AI for music generation there is a need to create educational tutorials (e.g. videos) for artists about data curation, model training, and ethical approaches to custom AI models. To be useful for musicians, this would also need to include guides on how to progress through the whole GenAI music workflow from the start with dataset preparation to the finish with AI model selection, training, fine-tuning, and music generation. For small datasets, this would need to be combined with guides about finding the best AI model and data representation for a particular use case e.g. a guide providing an overview of available open AI models and how they can be used to generate music along with guides to how much data would be needed to make the AI models functional for which purposes. Or, guides for what kind of dataset would be needed for a particular AI model, and what are the pros and cons of each AI models and what are privacy and regulatory concerns around their use. We have made a first step to making suitable low-resource AI models more findable by creating our open access repository of low-resource AI models at the project website[15].

Our position aligns with the AIforMusic "Principles for Music Creation with AI"[16] which in 2024 established guidelines for the responsible use of AI in music. Especially, the AIforMusic principles that argue for recognition of the human-created works, the importance of trustworthy and transparent AI, and valuing the perspectives of musicians, creatives, and songwriters in the development and use of AI in music making.

Ethical clearance

Ethical clearance for the workshop and data collection was granted by the workshop host University's ethics committee. All participants were informed of the purpose of the workshop and the forms of data collected in the workshop. Participants gave informed consent for their contributions and images to be recorded and used for research purposes and were able to withdraw their consent or leave the workshop at any point. Food and refreshments were provided. Participants did not receive financial incentives to attend the workshop.

Acknowledgments

With many thanks to the participants who took part in our workshops. Many thanks to our project partners: Central Conservatory of Music (China), Tallinn University (Estonia), Georgia Tech (USA), University of British Columbia (Canada), Concordia University (Canada), Ateneo de Manila University (Philippines), Dalarna University (Sweden), Music Hackspace (UK), DAACI (UK), Steinberg (Germany), Bela (UK). Many thanks to our stakeholder workshop panellists: François Pachet (Spotify Creator Technology Research Lab), Rebecca Fiebrink (University of the Arts London), Nuno Correia (Tallinn University), Phoenix Perry (University of the Arts London), Paul McCabe (Roland and AI For Music), Hazel Savage (SoundCloud and Musiio), Daisy Rosenblum (University of British Columbia), and CJ Carr (Dadabots). Many thanks to our case study presenters: Gabriel Vigliensoni, Daisy Rosenblum, Benjamin Timms, Soumya Sai Vanka, Rikard Lindell, Andrea Martelloni, Julian Parker, Lizzie Wilson, Andrei Coronel, Mark Gotham, Yiwei Ding, Moisés Horta Valenzuela, and Louis McCallum. And, many thanks to the performers who performed live music with AI models and small datasets at the close of the workshop: digital selves, Portrait XO, Dadabots, and Gabriel Vigliensoni.

This project was funded by Responsible Artificial Intelligence (RAI) UK International Partnerships (UKRI EPSRC grant reference EP/Y009800/1).

[15]http://musicrai.org
[16]https://aiformusic.info

References

[1] Berklee. 2025. *LibGuides: Virtuosity with Artificial Intelligence (AI): Explorations and Challenges in Music, Performance and Higher Education: Music Business, Management, & Marketing.* https://guides.library.berklee.edu/artificial-intelligence/music-business

[2] Jean-Pierre Briot and François Pachet. 2020. Deep learning for music generation: challenges and directions. *Neural Computing and Applications* 32, 4 (Feb. 2020), 981–993. https://doi.org/10.1007/s00521-018-3813-6

[3] Nick Bryan-Kinns, Zijin Li, and Xiaohua Sun. 2020. On Digital Platforms and AI for Music in the UK and China. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK.

[4] Nick Bryan-Kinns, Wei Wang, and Tie Ji. 2022. Qi2He: A co-design framework inspired by eastern epistemology. *International Journal of Human-Computer Studies* 160 (April 2022), 102773. https://doi.org/10.1016/j.ijhcs.2022.102773

[5] Nick Bryan-Kinns, Bingyuan Zhang, Sngyan Zhao, and Berker Banar. 2023. Exploring Variational Auto-Encoder Architectures, Configurations, and Datasets for Generative Music Explainable AI. *Machine Intelligence Research* (2023).

[6] Antoine Caillon and Philippe Esling. 2021. RAVE: A Variational Autoencoder for Fast and High-Quality Neural Audio Synthesis. https://doi.org/10.48550/arXiv.2111.05011 arXiv:2111.05011 [cs, eess]

[7] Chunlei Chen, Peng Zhang, Huixiang Zhang, Jiangyan Dai, Yugen Yi, Huihui Zhang, and Yonghui Zhang. 2020. Deep learning on computational-resource-limited platforms: A survey. *Mobile Information Systems* 2020, 1 (2020), 8454327.

[8] Shuiyi Chi and Hao Chen. 2025. Music style transfer and creation method based on transfer learning algorithm. *Journal of Computational Methods in Sciences and Engineering* (Feb. 2025), 14727978251318819. https://doi.org/10.1177/14727978251318819 Publisher: SAGE Publications.

[9] Miguel Civit, Javier Civit-Masot, Francisco Cuadrado, and Maria J. Escalona. 2022. A systematic review of artificial intelligence-based music generation: Scope, applications, and future trends. *Expert Systems with Applications* 209 (Dec. 2022), 118190. https://doi.org/10.1016/j.eswa.2022.118190

[10] Palle Dahlstedt. 2021. Musicking with algorithms: Thoughts on artificial intelligence, creativity, and agency. *Handbook of artificial intelligence for music: Foundations, advanced approaches, and developments for creativity* (2021), 873–914.

[11] Nils Demerlé, Philippe Esling, Guillaume Doras, and David Genova. 2024. Combining audio control and style transfer using latent diffusion. https://doi.org/10.48550/arXiv.2408.00196 arXiv:2408.00196 [cs].

[12] Andrew DeWitt. 2024. Identifying Cover Songs, Live Performances, AI Clones and More. https://www.audiblemagic.com/2024/02/07/identifying-cover-songs-live-performances-ai-clones-and-more/

[13] Anahita Doosti and Matthew Guzdial. 2023. Transfer Learning for Underrepresented Music Generation. In *Proceedings of the 2023 International Conference on Computational Creativity.*

[14] Constance Douwes, Giovanni Bindi, Antoine Caillon, Philippe Esling, and Jean-Pierre Briot. 2023. Is Quality Enough? Integrating Energy Consumption in a Large-Scale Evaluation of Neural Audio Synthesis Models. In *ICASSP 2023 -*

2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)). IEEE, Ixia-Ialyssos (Rhodes), Greece. https://doi.org/10.1109/ICASSP49357.2023.10096975

[15] Pascal Epple, Igor Shilov, Bozhidar Stevanoski, and Yves-Alexandre de Montjoye. 2024. Watermarking Training Data of Music Generation Models. https://doi.org/10.48550/arXiv.2412.08549 arXiv:2412.08549 [cs] version: 2.

[16] Lucas Ferreira-Paiva, Elizabeth Alfaro-Espinoza, Vinicius M. Almeida, Leonardo B. Felix, and Rodolpho V. A. Neves. 2022. A Survey of Data Augmentation for Audio Classification. *Congresso Brasileiro de Automática - CBA* 3, 1 (Oct. 2022). https://doi.org/10.20906/CBA2022/3469 Number: 1.

[17] Corey Ford, Ashley Noel-Hirst, Sara Cardinale, Jackson Loth, Pedro Sarmento, Elizabeth Wilson, Lewis Wolstanholme, Kyle Worrall, and Nick Bryan-Kinns. 2024. Reflection Across AI-based Music Composition. In *Proceedings of the 16th Conference on Creativity & Cognition*. 398–412.

[18] Mark Fox, Ganesh Vaidyanathan, and Jennifer L. Breese. 2024. THE IMPACT OF ARTIFICIAL INTELLIGENCE ON MUSICIANS. *Issues In Information Systems* 25, 3 (2024), 267–276. https://doi.org/10.48009/3_iis_2024_121

[19] Austin A Franklin, Henrik Frisk, and Rikard Lindell. 2024. Sonic Serendipity: Embracing Discovery in File Finder-Based Improvisation. , Article 34 (September 2024), 6 pages. https://doi.org/10.5281/zenodo.13904836

[20] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. https://doi.org/10.48550/arXiv.1803.09010 arXiv:1803.09010 [cs].

[21] Toby Gifford, Shelly Knotts, Jon McCormack, Stefano Kalonaris, Matthew Yee-King, and Mark d'Inverno. 2018. Computational systems for music improvisation. *Digital Creativity* 29, 1 (2018), 19–36.

[22] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision* 129, 6 (2021), 1789–1819.

[23] UK Government. 2024. Copyright and AI: Consultation. https://assets.publishing.service.gov.uk/media/6762c95e3229e84d9bbde7a3/241212_AI_and_Copyright_Consultation_print.pdf Publisher: Intellectual Property Office.

[24] Amjad Hanif. 2024. *New Tools to Protect Creators and Artists*. https://blog.youtube/news-and-events/responsible-ai-tools/

[25] Yun-Ning Hung, Chao-Han Huck Yang, Pin-Yu Chen, and Alexander Lerch. 2022. Low-Resource Music Genre Classification with Advanced Neural Model Reprogramming. https://doi.org/10.48550/arXiv.2211.01317 arXiv:2211.01317 [cs] version: 1.

[26] Sony Computer Science Laboratories Inc. 2022. *Flow Machines / Stories / Sony Design*. https://www.sony.com/en/SonyInfo/design/stories/flow-machines/

[27] Yangluxi Li, Huishu Chen, Peijun Yu, and Li Yang. 2025. A review of artificial intelligence in enhancing architectural design efficiency. *Applied Sciences* 15, 3 (2025).

[28] Rikard Lindell and Henrik Frisk. 2023. The Unfinder: Finding and reminding in electronic music. In *The 16th International Symposium on Computer Music Multidisciplinary Research*. Tokyo, Japan.

[29] Omar Lopez-Rincon, Oleg Starostenko, and Gerardo Ayala-San Martín. 2018. Algoritmic music composition based on artificial intelligence: A survey. In *2018 international conference on electronics, communications and computers (CONIELECOMP)*. IEEE, 187–193.

[30] Soroush Mehri, Kundan Kumar, Ishaan Gulrajani, Rithesh Kumar, Shubham Jain, Jose Sotelo, Aaron Courville, and Yoshua Bengio. 2017. SampleRNN: An Unconditional End-to-End Neural Audio Generation Model. https://doi.org/10.48550/arXiv.1612.07837 arXiv:1612.07837 [cs].

[31] Atharva Mehta, Shivam Chauhan, Amirbek Djanibekov, Atharva Kulkarni, Gus Xia, and Monojit Choudhury. 2025. Music for All: Exploring Multicultural Representations in Music Generation Models. https://doi.org/10.48550/arXiv.2502.07328 arXiv:2502.07328 [cs].

[32] Sreetama Mukherjee and Manjunath Mulimani. 2022. ComposeInStyle: Music composition with and without Style Transfer. *Expert Systems with Applications* 191 (April 2022), 116195. https://doi.org/10.1016/j.eswa.2021.116195

[33] Shahan Nercessian. 2018. iZotope and Assistive Audio Technology. https://www.izotope.com/en/learn/izotope-and-assistive-audio-technology.html

[34] Philippe Pasquier, Arne Eigenfeldt, Oliver Bown, and Shlomo Dubnov. 2017. An Introduction to Musical Metacreation. *Comput. Entertain.* 14, 2, Article 2 (jan 2017), 14 pages. https://doi.org/10.1145/2930672

[35] Alison Pease and Arnold Pease. 2023. Computational Creativity and the Climate Crisis: 14th International Conference on Computational Creativity. *Proceedings of the Fourteenth International Conference on Computational Creativity, ICCC 2023* (2023), 293–297. https://computationalcreativity.net/iccc23/papers/ICCC-2023_paper_114.pdf Publisher: Association for Computational Creativity.

[36] Teresa Pelinski, Victor Shepardson, Steve Symons, Franco Santiago Caspe, Adan L Benito Temprano, Jack Armitage, Chris Kiefer, Rebecca Fiebrink, Thor Magnusson, and Andrew McPherson. 2022. Embedded AI for NIME: Challenges and Opportunities. In *NIME 2022*. PubPub, The University of Auckland, New Zealand. https://doi.org/10.21428/92fbeb44.76beab02

[37] Alastair Porter, Mohamed Sordo, and Xavier Serra. 2013. Dunya: A System to Browse Audio Music Collections Exploiting Cultural Context.. In *Proceedings of the 14th International Society for Music Information Retrieval Conference*. ISMIR, 101–106. https://doi.org/10.5281/zenodo.1417355

[38] Colin Raffel. 2016. *Learning-Based Methods for Comparing Sequences, with Applications to Audio-to-MIDI Alignment and Matching*. PhD Thesis. Columbia University, USA.

[39] Praveen Ravichandran, Marco De Nadai, Divita Vohra, Sandeep Ghael, Manizeh Khan, Paul Bennett, Tony Jebara, and Mounia Lalmas-Roelleke. 2024. Contextualized Recommendations Through Personalized Narratives Using LLMs. https://research.atspotify.com/2024/12/contextualized-recommendations-through-personalized-narratives-using-llms/

[40] Jeba Rezwana and Mary Lou Maher. 2023. Designing creative AI partners with COFI: A framework for modeling interaction in human-AI co-creative systems. *ACM Transactions on Computer-Human Interaction* 30, 5 (2023), 1–28.

[41] Anna-Maria Sichani, Paula Westenberger, Nick Bryan-Kinns, Mercedes Bunz, Clementine Collett, Bahareh Heravi, Kate M. Miltner, Caterina Moruzzi, and Beverley Alice Townsend. 2025. BRAID researchers' response to UK Government copyright and AI consultation. (Feb. 2025). https://doi.org/10.5281/zenodo.14945987 Publisher: Zenodo.

[42] Suryabhan Singh, Kirti Sharma, Brijesh Kumar Karna, and Pethuru Raj. 2022. Pruning and quantization for deeper artificial intelligence (AI) model optimization. In *International Conference on Robotics, Control, Automation and Artificial Intelligence*. Springer, 933–945.

[43] Dimitrios Smailis, Areti Andreopoulou, and Anastasia Georgaki. 2021. Reflecting on the musicality of machine learning based music generators in real-time jazz improvisation: a case study of OMax-improteK-Djazz. In *Proceedings of the 2nd Conference on AI Music Creativity (MuMe+ CSMC)*.

[44] Vanessa Utz and Steve DiPaola. 2023. Climate Implications of Diffusion-based Generative Visual AI Systems and their Mass Adoption. In *Proceedings of the 14th International Conference on Computational Creativity (ICCC'23)*.

[45] Wim Vanderbauwhede. 2023. Frugal Computing–On the need for low-carbon and sustainable computing and the path towards zero-carbon computing. *arXiv preprint arXiv:2303.06642* (2023).

[46] Gabriel Vigliensoni, Louis McCallum, Esteban Maestre, and Rebecca Fiebrink. 2022. R-VAE: Live Latent Space Drum Rhythm Generation from Minimal-Size Datasets. *Journal Of Creative Music Systems* 1, 1 (2022). https://doi.org/10.5920/jcms.902

[47] Gabriel Vigliensoni, Phoenix Perry, and Rebecca Fiebrink. 2022. A Small-Data Mindset for Generative AI Creative Work. In *Generative AI and HCI - CHI 2022 Workshop*. ACM, Online.

[48] Megan Wei, Mateusz Modrzejewski, Aswin Sivaraman, and Dorien Herremans. 2024. Prevailing Research Areas for Music AI in the Era of Foundation Models. https://doi.org/10.48550/arXiv.2409.09378 arXiv:2409.09378 [cs].

[49] Michael Williams and Tami Moser. 2019. The Art of Coding and Thematic Exploration in Qualitative Research. *International Management Review* (2019). https://www.semanticscholar.org/paper/The-Art-of-Coding-and-Thematic-Exploration-in-Williams-Moser/c0a0c26ac41cb8beb337834e6c1e2f35b91d071d

[50] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated Machine Learning: Concept and Applications. https://doi.org/10.48550/arXiv.1902.04885 arXiv:1902.04885 [cs].

[51] Seul-Ki Yeom, Philipp Seegerer, Sebastian Lapuschkin, Alexander Binder, Simon Wiedemann, Klaus-Robert Müller, and Wojciech Samek. 2021. Pruning by explaining: A novel criterion for deep neural network pruning. *Pattern Recognition* 115 (2021), 107899.

[52] Haizi Yu, Heinrich Taube, James A Evans, and Lav R Varshney. 2020. Human evaluation of interpretability: The case of AI-generated music knowledge. *arXiv preprint arXiv:2004.06894* (2020).