Exploring Musical Creativity in the Era of Generative Artificial Intelligence: A Qualitative Inquiry

Matthew Peachey
Graphics and Experiential Media Lab
Dalhousie University
Halifax, Canada
peacheym@dal.ca

Joseph Malloch
Graphics and Experiential Media Lab
Dalhousie University
Halifax, Canada
jmalloch@dal.ca

Advances in Artificial Intelligence (AI) are reshaping creative practices, particularly through generative systems capable of producing highly realistic outputs at a rapid pace. While these tools offer many new opportunities, they also present several challenges for musicians, whose workflows may be fundamentally changed by the emergence of such technologies. This paper reports on a qualitative user study ($n = 10$) exploring how musicians engage with generative AI in their creative practice or why they choose not to adopt these tools. Semi-structured interviews were analyzed following the Reflexive Thematic Analysis methodology, resulting in key themes such as a preference for co-creation and perspectives on creative ownership and agency. These findings both highlight concerns regarding the adoption of generative AI in musical practice as well as inform design implications for integrating AI tools more effectively into musical workflows.

1 INTRODUCTION

Generative models are a recent development in the rapidly advancing field of artificial intelligence (AI) and machine learning (ML). They are a class of ML architectures that learn the underlying structure of a dataset to generate new instances that plausibly resemble the original data (21). In recent years, they have produced high-quality results across modalities such as text, images, and video, leading to claims of improved productivity, streamlined workflows, and new approaches to digital content creation.

However, as Generative AI (GenAI) technologies advance rapidly, concerns have emerged regarding their implications for creative practitioners including musicians. While enthusiasm for AI continues to grow, so too do anxieties that end-to-end generation techniques may reduce the creative agency traditionally exercised by humans. This motivates the exploration of alternative interaction paradigms, particularly as prior work in creative domains suggests that discrete text-prompt interfaces do not adequately support exploratory creative processes (14; 15; 41; 1). To this end, this paper reports findings from a qualitative inquiry ($n = 10$) examining how experienced musicians perceive GenAI from a Human–Computer Interaction (HCI) perspective, including its current use and the challenges of integrating such systems into established workflows.

2 RESEARCH QUESTIONS

While this qualitative inquiry was designed to be highly exploratory with a focus on better understanding musicians and their wide range of expectations for GenAI, we did pose two research questions to serve as preliminary context when designing our study. Our research questions are as follows.

- **RQ1:** How do musicians think about or integrate generative technologies into their creative workflows?
- **RQ2:** What are the impacts on users when implementing continuous interactions for generative systems in the context of their creative workflows?

These research questions guided our choice of inquiry questions to ask of participants, but as will be discussed in a later section, thematic analysis revealed additional insights that go beyond simply answering these guiding research questions.

3 RELATED WORK

This research explores how musicians perceive the integration of GenAI into their creative workflows. To contextualize this work, we first outline key areas of related research. We begin with an overview of GenAI and its distinction from traditional machine learning methods, followed by a discussion of interaction techniques, including both continuous and discrete paradigms. We then examine how interaction techniques shape the design and use of GenAI systems, before reviewing applications of GenAI for creativity, with a focus on musical contexts and their associated challenges and opportunities. This related work provides a foundation for interpreting our qualitative findings within the broader intersection of HCI, GenAI, and musical creativity.

3.1 Generative Artificial Intelligence

Generative systems are typically powered by a distinct class of machine learning models that, unlike supervised learning (which maps input data to target outputs) or unsupervised learning (which identifies underlying patterns in a dataset), can produce highly realistic data (21). Foundational generative models, such as Variational Autoencoders (VAEs) (28) and Generative Adversarial Networks (GANs) (24), were among the first to demonstrate the capacity to produce realistic samples by leveraging latent representations and adversarial training paradigms. In contrast, modern systems such as ChatGPT are built on top of transformer-based architectures, reflecting a shift towards large-scale models trained on an immense amount of data. Generative models, irrespective of the modeling technique used, have seen a rapid up-tick in applications and use-cases across domains, yet are often still viewed as a black-box system, making them an important area of continued research throughout academia and industry alike, especially in the context of creativity.

3.2 Interaction Techniques

Interaction is a cornerstone of HCI and is defined by Hornbæk and Oulasvirta as the dynamic process of two entities influencing one another's behaviour over time (26). To guide the discussion presented in this paper we position user interactions as being categorized into two types: discrete and continuous. Discrete interactions involve synchronous exchanges of information between a user and a system (e.g. buttons or text-prompts). Conversely, continuous interactions occur when a user remains constantly engaged with a system over an extended period of time (e.g. sliders or motion-capture systems) (17).

3.2.1 Interaction Techniques for AI Systems

The choice of interaction technique is crucial when designing AI systems. Höök outlines that control, predictability and transparency must be considered when designing intelligent interfaces (25). Although each are necessary for an effective AI system, the principle of control is influenced most by the choice of interaction technique. In AI systems, a discrete interaction may involve running inference on a single data point and synchronously receiving the result, termed Intermittent Human–AI Interaction by Van Berkel et al. (43). In contrast, Continuous Human–AI Interaction describes systems where users can iteratively explore an input space and observe how outputs change in response (43)[1]. Gammelgaard-Larsen et al. study how interaction methods shape users' perceived control, which Norman argues is central in interactions with intelligent systems (33). Continuous interactions may require users to relinquish some control, but precise feedback mechanisms can improve system understanding (23). Therefore, the choice of interaction technique is a key design consideration in AI-powered systems.

3.3 Generative AI and Creativity

As generative AI systems become embedded in creative practice across disciplines, researchers have explored how these technologies reshape longstanding assumptions about authorship (44), agency (4; 9), and responsible use of technology (10) in creative work.

Furthermore, there are several key areas for generative technology in music, including waveform-level audio synthesis and music generation (18; 16; 3; 40; 46; 19; 47), synthesizer patch generation (29; 35; 36), and physical interfaces such as *Stacco* (38; 39) or motion capture techniques (32) being used to explore a model's latent space. The breadth of musical GenAI applications already available implies a wide variation in how musicians perceive and adopt these tools in their workflows, underscoring the need for this type of study.

3.3.1 Limitations of Text Prompts for Creative Work

Current GenAI models heavily rely on text-prompts, which while excellent for democratizing access to the technology, do not necessarily support divergent thinking and therefore may limit what a user is able to create (15). Additionally, many artists will choose not to share the specific prompts used to generate their artwork, making it impossible to know the details of the process that went into the creation of a certain piece[2] (34). Researchers have proposed several different techniques for crafting prompts (5; 30; 45; 37; 20; 42; 13). However, other studies that have highlighted several limitations of using text prompts for AI driven creative practices. These include the largely trial and error process of writing prompts (14) and concerns that small changes to prompts may lead to large changes in generated output (15). Furthermore, text-prompts typically do not support visually driven workflows (1) and therefore users may find it difficult to accurately describe what they want a model to generate using text (41). Finally, it is unfair to expect a user to effectively describe output possibilities that they do not already know to exist (15), . Therefore, while text prompts are effective for basic interactions with generative models, other interaction techniques should be explored in order to determine the best method for supporting creativity.

4 METHODS

We conducted a qualitative inquiry using semi-structured interviews with experienced musicians in order to explore how participants currently integrate GenAI into their musical practices as well as explore any reasons as to why some musicians choose not to adopt the technology.

4.1 Participant Recruitment

We recruited skilled musicians with diverse musical backgrounds, influences, and genre preferences, including producers, instrumentalists, and music-technology scholars, to capture a wide range of perspectives. Prior studies (27; 22; 2) informed an ideal sample size of eight to twelve participants. Participants were recruited via prior study cohorts, mailing lists, and word of mouth. A brief screening survey, which also served to document informed consent, was used for eligibility. As shown in Table 1, ten participants were recruited, eight of whom had over five years of experience. While this recruitment strategy limited any traditional qualitative saturation or the generalization of findings, the many different backgrounds of participants led to the surfacing of key concerns and opportunities associated with GenAI in musical practices, providing an initial context for discussion and a basis for further exploration.

4.2 Semi-structured Interviews

Upon receipt of completed consent forms, the lead researcher contacted participants to schedule interviews. Sessions were conducted primarily via Microsoft Teams, with a small number held in person, and lasted 30–60 minutes. Each interview began with an overview of the session and question structure. The interview questions, organized into four subsections (Appendix A), first explored the creative workflows and musical influences of the participants. The second examined how musicians interact with generative tools, particularly with regard to interaction styles. The third focused on preferences for personalization in GenAI tools, and the fourth on openness to change and factors influencing it.

[1]It is important to make a clear distinction between continuous interactions and continuous input when discussing ML powered systems. Continuous inputs are themselves a sequence of continuous data, such as a sentence of text or a brief section of music. However, using a continuous input does not directly imply the presence of continuous interaction, as a continuous input can still be used to run a one-off, synchronously processed inference activity of a model, which is more indicative of a discrete interaction.

[2](i.e. did the artist get lucky during the initial prompting experiments or was a methodical creative process followed in order to achieve their results.)

PID	Type of Music	Years of Experience	Studies Music or Music Technology?
1	Electroacoustic, Bass Guitar, Heavy Metal Drumming	5+	Yes
2	Classical Piano, Indie Electronic, Ambient	5+	Yes
3	Ambient, Electronic	5+	Yes
4	Electroacoustic, Ambient, "New Music", Experimental	5+	Yes
5	Piano	2-5	Yes
6	RNB, Electronic, Synth Pop, Ambient, Indie, Folk	2-5	No
7	Rock (Math Rock, Alternative Rock, Post Rock), Pop	5+	No
8	Ambient, Alternative, Indie, Guitar, Piano	5+	No
9	Folk, Jazz, Ambient, Electroacoustic, Classical	5+	Yes
10	Indie-Rock, Folk, Experimental Hip-Hop	5+	Yes

Table 1: Overview of the musical experience of the participants that were interviewed for this qualitative inquiry.

4.3 Reflexive Thematic Analysis

We followed Braun and Clarke's Reflexive Thematic Analysis (RTA) (7; 8) to analyze the qualitative data collected through semi-structured interviews. Initial codes were first generated by each researcher based on the first two participant's data in order to get a sense of common codes between researchers. While Braun and Clarke insist that within RTA there is no "correct" or "accurate" way of coding qualitative data, and therefore that two researchers should not necessarily expect to come up with the same set of codes for a dataset (Braun and Clarke; 11), we reflected on our initial codes collaboratively to better understand how the data was being interpreted as suggested by Bryne (11). Figure 1 depicts how a set of nine initial candidate themes were created by grouping codes, and how after additional analysis a set of five refined themes represented our qualitative findings.

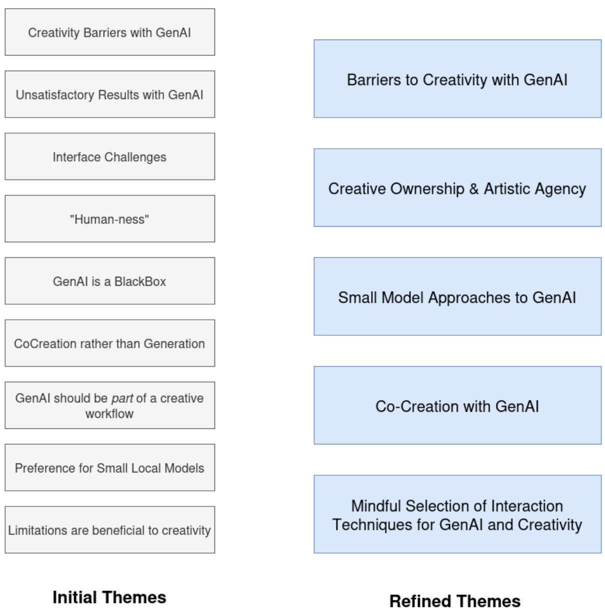

Figure 1: Nine initial themes were consolidated into five refined themes through iterative analysis and cross-theme code comparison.

5 RESULTS

The following sections detail the five key themes derived from a reflexive thematic analysis of the study dataset.

5.1 Theme 1: Barriers to Creativity with GenAI

The first theme emerging from the data concerns barriers to creativity that musician's observed when considering integrating GenAI into their workflows.

5.1.1 Generic Results

A common barrier that was repeatedly discussed by participants was that they felt that GenAI tools for music produce overly generic results. For example, the first participant talks about how they feel that there is a lack of options for developing their own music style when working with generative tools and how that results in generic sounding results.

> "*Just sort of generating the same thing over and over, then you're not really ever pushing some kind of musical envelope or developing your own style too, you get really pigeon holed in*" (P1).

Another participant observed that art, by its very nature, is not a field that can be optimized, and that no single correct solution can be identified. This participant's comment again relates to how artists feel about GenAI outputs being overly generic, as attempting to optimize musical outputs will often lead to generic results.

> "*Art is not one of those things you can't optimize… There are optimal solutions (plural) but there isn't one specific optimal solution. I hope these tools get better at giving you a variety of options. But I think pretty much everyone is aware that it's detrimental to make it do one thing forever.*" (P4).

A third participant noted that 'AI hype' has contributed to generic outputs, as some GenAI users appear to be striving for something that merely resembles popular music.

> "*I think there is a wave of people who aren't artists who are using tools to generate very generic songs where they don't have much vision in it, they just want something that resembles the idea of music. I think that's going to fizzle out like it's definitely mean that we'll have so much slop that it becomes insignificant*" (P7).

Thus, when GenAI produces outputs that are generic or redundant, it does not foster a sense of creativity or excitement that musicians can rally behind.

5.1.2 Interface Challenges

Another barrier to creativity that was repeatedly discussed was a lack of sufficient interfaces for engaging with GenAI. For example, one participant highlighted their dissatisfaction with a reliance on text-prompting for user input, suggesting that interface modality is not particularly useful in the context of musical creativity.

> "*The other thing is the interface. I don't think text is a great interface for anything to do with music*" (P2).

Furthermore, another participant talks about how GenAI tools don't support fine grained control, which we've already identified as a key pillar of a a user's perception of control over a system (33).

> "*The obvious one [drawback] for all generative tools is a lack of fine grained control. It's like that across the board*" (P5).

Each of these comments are similar to findings from previous literature that discuss the challenges of using text-prompting and discrete interactions for creative workflows, highlighting a similar concern in the domain of musical creativity.

5.2 Theme 2: Creative Ownership and Artistic Agency

The second theme focuses on ownership and agency as it relates to creative outputs from GenAI tools. Specifically, this theme reflects participants' views that current generative tools often fail to support their desires for creative ownership or artistic agency, largely due to the black-box nature of these systems.

5.2.1 Lack of Ownership or Agency reduces feelings of creativeness

Participants discussed how they felt that if they did not feel a sense of ownership or agency over the output of a GenAI tool, that they would not consider themselves to be *doing art* nor would they consider themselves to be a *creator*. For example, one participant spoke specifically about need for an artistic vision.

> "*What makes art art is that there is someone in the driver's seat and there is like a visionary as an expression of something*" (P7).

Furthermore, much like many other creative domains, music is a highly personal and individual field. Several participants explicitly reported that an artist's lived experience, or "the human soul", is essential to producing compelling art.

> "*I feel as if the key to music is the soul that the artist puts into it, and with AI you don't have that. So I don't like to incorporate AI creatively into my music because I feel like that takes away from the soul of it.*" (P8).

However, as is the case for musical instruments such as modular synthesizers, even when a system appears to be opaque, it may still afford creative opportunities to musicians, so long that they are involved in the process of determining the outputs of that system. In response to Q1.7, a participant articulated the limits of this approach.

> "*Oh yeah, I mean I definitely would, I mean this is a thing where it's still there are some things that are still blurry for me that I just like haven't come to a happy answer with even talking with academic folks of like 'what is the huge difference between an agent model I don't understand and a Eurorack module that I don't understand? [...] But I'm also you know not just in like 'hit generate' and not touch it and not mix it kind of thing.*" (P4).

In a similar vein to the previous quote, another participant discusses how when using AI, they feel as though they are more of an evaluator or co-creator, again suggesting a lacking sense of creative ownership when using generative tools.

> "*I don't really feel like a creator. If I'm using AI I would say I would fall between the range of evaluator and co-creator at most [...] personally, I wouldn't consider it art if you're just wholesale using something AI generated.*" (P3).

And finally this participant highlighted that personalization of generative models is one potential method to move beyond generic results towards more artistic outputs.

> "*So in terms of personalization, I think that is one way to get out of the AI slop kind of thing where everything sounds generic and it sounds like AI.*" (P9).

5.2.2 Frustrations with the Black-box nature of Generative AI

Our analysis further reveals that musicians' perceptions of a lack of ownership and agency when working with GenAI tools often directly stems from the black-box nature of AI systems. While P4 alluded to the fuzziness of working with opaque systems, several other participants expressed explicit concerns. For example, one participant described feeling that black-box systems produced lackluster results.

> "*I think the black box kind of just like, please make me a song that is like, this is kind of fruitless*" (P7).

Another participant expressed concern throughout their interview that an over reliance on GenAI tools may limit musicians' ability to fully understand how their instruments work. For instance, in the context of using GenAI to generate digital synthesizer patches a reliance on GenAI may limit how much is learned about the

> "*[A generative tool may] obfuscate a lot of the synth things, such that you wouldn't actually lean how a synth works*" (P6).

And that participant later responded with the following quote, discussing how that too much so called 'vibe-making-music' may force a user to become to reliant on their tools such that they'll never be able to move beyond a GenAI facilitated workflow, even if they were interested in doing so.

> "*[...] another challenge, more long term, would be people not wanting to take the training wheels off, similar to how vibe-coding we might have vibe-making-music. Then if you're [generating] drums,*

guitar, bass or whatever, when it comes to a performance they just can't perform it. Maybe they're relying on it [GenAI] too much." (P6).

Therefore, while GenAI has the potential to serve as a powerful tool within creative workflows, it is also important that continued effort be directed toward reducing its perceived black-box nature in order to better support musical adoption.

5.3 Theme 3: Small Model Approaches to GenAI

Very large models have become commonplace in discussions of AI, particularly in the context of generative AI and Large Language Models (LLMs). However, several of our participants discussed their preference for smaller, more focused models, when considering how they'd integrate GenAI into their creative workflow.

> "*I don't love the huge model approach [for creativity].*" (P5).

Furthermore, another participant reflected on their preference for smaller models that are trained on 'local data', implying a desire to shift away from the large corporate models in favour of models trained on data that they curate as part of their larger creative process.

> "*I'd definitely push for a more locally grown AI*" (P1).

As discussed in the following sections, there are also number of practical reasons that musicians in particular are interested in small model approaches to GenAI.

5.3.1 Creativity Benefits from Limitations

Creativity is often built on the basis of artistic limitation, including in the digital era (31). Limitations also allow artists to work within more manageable creative bounds, especially when using digital creativity support tools (12). One participant explicitly stated how limitations are essential to art.

> "*So limitation is more useful than everything to the average person*" (P9).

In the context of GenAI tools, imposing limitations on the range of what a model can generate may be one solution for increasing the perception of agency and ownership. For example, a participant suggested the following for a preferred GenAI system that limits its output based on a user's input.

> "*you play at the system and then it understands that's what you're going for, rather than sort of going into the whole range of its potential*" (P1).

Another participant explicitly discussed a desire for parameterization of GenAI models, where a user-facing interface would allow them to interact with the models in a manner analogous to playing an acoustic instrument.

> "*Can I get a box that has limited things that I can change, like parameters on that? Can I like literally play and perform with it?*" (P9).

This theme highlights that many musicians are interested in moving away from the 'large models' approach to GenAI in favour of smaller models trained on their own curated data. Tools built on top of these types of models can enable artists to feel as though they are still in control of what the AI tool creates, as well as support both co-creative activities and continuous interactions, as will be discussed in future sections.

5.4 Theme 4: Co-Creation with GenAI

Another noteworthy aspect of participant responses was their preference to engage with GenAI in a co-creative manner, whereby human artists continued to create using tools powered by GenAI, rather than relying on the fully end-to-end generation techniques that are typical of many common GenAI applications. One participant commented explicitly on how a level of artistic input is necessary to produce meaningful results.

> "*I do think there is a level of artistic input that you need to be able to give the systems to get something interesting back out.*" (P1).

Another participant described how they view AI as either a tool which can support bringing their creative ideas to life or as a system that aims to create its own ideas.

> "*There is two very different sides [to AI]. There is the side that actually generates it [...] and then there is the part where you have an idea in your head, like a certain sound you're looking for and you say "AI make this" or "give me just a little bit" that part I think is fine, in fact it's a great tool and there is nothing like that right now.*" (P7).

Finally, participants expressed detailed opinions on how their ideal workflow when considering GenAI is to use these tools as a small part of their existing workflow, rather than seeing their creative practice being fully replaced by AI tools.

> "*I think it would be almost impossible to get me to go all the way and say, OK, let's just let's just generate a whole song. I would be open to, you know, small little tools here and there that could improve, you know, just my basic day-to-day or if I need a snare that's really specific, I need a certain sound. I think that would be great. And I'd definitely be open to using that.*" (P8).

Another participant also highlighted the fine line between 'grunt labour' and 'creative work' in the context of GenAI usage.

> "*I think my current stance on AI as it relates to the creative process is I would probably use tools that take away grunt labour, but anything that feels like it's doing creative labour (which is a fine line I think) I would probably vow to stay away from for the rest of my life and risk not being employed.*" (P10).

5.4.1 GenAI as a tool within a creative practice

An important aspect of co-creation with GenAI is determine the balance of how much AI to use within a workflow. Participants discussed how their preference is to use GenAI tools as a small portion of their workflow, either to support tedious or monotonous work or to be inspired when facing creative blocks, but not to fully take over creative tasks.

"*I think going back to that thing like computers and tools and machines can take away the grunt work. That's great. But if they're starting to get into the stuff that humans kind of differentiate themselves from machine and like trying to do creative work. That's where it makes me feel kind of grossed out and I don't. I don't want to participate in that kind of thing.*" (P10).

Another participant commented on how GenAI tools can support completing tasks based on a limited set of potential outcomes given their input scenario.

"*But in terms of idea generation, I think that you want like very very limited things [...] like 'oh I have these two chords like what comes next'.*" (P9).

And another participant quoted how GenAI tools, and the way in which they learn underlying representations of data, allow for a musician to think about their work in previously unexplored ways.

"*The co-creation aspect of it [re: GenAI]. It gives me a different way to kind of look at the source material I have and kind of rethink it in different ways.*" (P4).

This theme highlights that co-creation with GenAI is a preferred approach among musicians. By treating generative tools as a small component within a larger creative workflow, musicians can remain at the center of the creative process.

5.5 Theme 5: Mindful Selection of Interaction Techniques for GenAI and Creativity

Participants reflected on how interaction styles influenced their thoughts about GenAI and creativity. Participants discussed challenges of discrete interactions, in this case text-prompts, noting a lack of fine-grained control over the generated outputs.

"*You can keep clicking generate and a lot of times and it'll give you some garbled mess. It'll give you something that looks kind of right, but like, not quite right. And you just kind of have to keep generating. So now you sit there for like 20 minutes just hitting generate because because it's not giving me something I'm satisfied with. [...] I feel like is not very conducive to like, you know, that iterative, even that curatorial approach [to creativity]*" (P3).

Another participant discussed their distate for text-prompting in general.

"*You'd have less control over it [AI model], especially when you're starting out. Like I don't know much about prompt engineering, but the term itself makes me roll my eyes*" (P6).

Two participants responded simply about their preferences for continuous interactions, stating that the interaction style provided them with a method to be more creative.

"*I would stick with continuous because I think it would fosters a lot more creativity*" (P3).

"*With continuous interaction, you're still being creative, you're still giving input*" (P6).

And finally, one participant expressed how compelling musicianship may be tightly coupled with the notion of continuous interaction due to how a well-trained performer plays by feel, a method not supported by discrete, text-based interfaces.

"*Music is played a lot by feel, like you can definitely overthink things and creativity can kind of come out spontaneously. So I think a continuous system lets you kind of explore by intuition or know with by your gut. Rather than cerebrally typing in prompts.*" (P5).

In summary, this theme encapsulates how the choice of interaction technique directly influences musician's feelings about GenAI systems and creativity.

6 DISCUSSION

We have articulated five key themes that synthesize the findings of our qualitative inquiry. Across these themes, a central tension emerges between the potential of what GenAI offers musicians and the wariness of musicians to adopt that technology into their own workflow.

6.1 Limitations of Current AI Systems for Musical Creativity

Current discourse on GenAI often prioritizes efficiency and productivity as primary indicators of system success. However, our findings show that for musicians these are secondary to producing work that is personally meaningful and artistically impactful. This reframes the core design challenge: not to generate more output faster, but to develop AI systems that meaningfully augment existing creative workflows.

Our findings also reveal a key limitation of current generative systems: they obscure the musician's influence over the output, weakening perceptions of authorship and creative responsibility. This is partly driven by reliance on text-prompt interfaces and perceptions of output homogeneity. As a result, AI tools are frequently evaluated and marketed in terms of speed and volume, while neglecting their impact on creative agency and workflow integration.

6.2 Designing GenAI to Support Creative Workflows

Rather than prioritizing output speed or volume, future GenAI tools should be designed to align with the ways musicians naturally explore, refine, and develop ideas within their existing creative workflows. The following design implications aim to capture key principles for developing GenAI systems that more effectively support musicians' creative workflows.

6.2.1 Prioritize Creative Agency and Ownership over Full Automation

To support existing creative practices, GenAI systems should enable musicians to guide and shape outputs, consistent with Norman's notion of perceived control (33), ensuring they complement rather than replace human decision-making. Participants emphasized that music, like other art forms, cannot be optimized computationally, and that human subjectivity is central to its expressive value. Accordingly, AI tools for creative support should avoid automating away this subjectivity and instead promote ownership and agency by making the relationship between user input, decision-making, and model output more transparent.

6.2.2 Design for Specific Use Cases

Many participants emphasized the value of using GenAI as a small, targeted component within their creative workflows. Furthermore, prior work has shown that small-data approaches, which are best suited for specific use cases, can improve users' understanding of model behaviour (10). Participants also noted that the boundary between helpful tools that reduce repetitive tasks and overbearing systems that generate complete outputs with minimal input is often blurred, given the diversity of individual creative practices. While small model architectures trained on user-curated data may support specific musical use cases, designers must avoid overconstraining these systems, as repetitive or uninteresting outputs can discourage integration into creative workflows.

6.2.3 Select Interaction Techniques with Purpose

Supporting creative practice with GenAI requires careful selection of interaction techniques that align with musicians' expectations and existing workflows. Participants consistently reported that text prompts, despite being the de facto standard for GenAI, do not adequately support their creative processes, particularly due to limited fine-grained control over outputs. While we do not dismiss the role of discrete interactions in AI music tools, they were frequently described as insufficient for expressive creative work. In contrast, well-designed continuous interactions can enable musicians to perform AI systems as instruments, supporting more direct and expressive engagement. Accordingly, designers should prioritise interaction paradigms that better align with musicians' creative intentions and practices.

7 CONCLUSION & FUTURE WORK

This paper presents a qualitative inquiry into how experienced musicians perceive the integration of GenAI tools into their creative workflows. As an exploratory study, it opens several avenues for future research. First, as generative tools continue to evolve, it will be important to revisit these questions to capture shifting perspectives. Accordingly, the views reported here may change as technology develops and should be revisited over time. Second, there is a need to design and evaluate GenAI tools that build on the approaches proposed in this paper, particularly those leveraging smaller models and supporting continuous, performative interactions aligned with existing musical practices.

In conclusion, while musicians are generally open to experimenting with new technologies to complement their creative workflows, care must be taken to ensure that human artistry remains central to the creative process. In doing so, artists can leverage generative models without sacrificing the essence of creative practice in an increasingly digital era.

8 ETHICAL STANDARDS

This research was conducted with approval from Dalhousie University's Research Ethics Board (REB# 2025-7723).

REFERENCES

[1] Akverdi, C. and Baykal, G. E. (2024). Generative ai tools in design fields: Opportunities and challenges in the ideation process. In *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction*, NordiCHI '24 Adjunct, New York, NY, USA. Association for Computing Machinery.

[2] Bennett, C. L., Shelby, R., Rostamzadeh, N., and Kane, S. K. (2024). Painting with cameras and drawing with text: Ai use in accessible creativity. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–19.

[3] Bińkowski, M., Donahue, J., Dieleman, S., Clark, A., Elsen, E., Casagrande, N., Cobo, L. C., and Simonyan, K. (2019). High fidelity speech synthesis with adversarial networks. *arXiv preprint arXiv:1909.11646*.

[4] Bomba, F. and De Angeli, A. (2025). Agency and authorship in ai art: Transformational practices for epistemic troubles. *International Journal of Human-Computer Studies*, page 103652.

[5] Brade, S., Wang, B., Sousa, M., Oore, S., and Grossman, T. (2023). Promptify: Text-to-image generation through interactive prompt exploration with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA. Association for Computing Machinery.

[Braun and Clarke] Braun, V. and Clarke, V.

[7] Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101.

[8] Braun, V. and Clarke, V. (2022). Conceptual and design thinking for thematic analysis. *Qualitative psychology*, 9(1):3.

[9] Bryan-Kinns, N., Perry, P., Wilson, E., and Wszeborowska, A. (2025a). Human-centered ai in arts. In *Handbook of Human-Centered Artificial Intelligence*, pages 1–32. Springer.

[10] Bryan-Kinns, N., Wszeborowska, A., Sutskova, O., Wilson, E., Perry, P., Fiebrink, R., Vigliensoni, G., Lindell, R., Coronel, A., and Correia, N. N. (2025b). Leveraging small datasets for ethical and responsible ai music making. In *Proceedings of the 20th International Audio Mostly Conference*, pages 70–81.

[11] Byrne, D. (2022). A worked example of braun and clarke's approach to reflexive thematic analysis. *Quality & quantity*, 56(3):1391–1412.

[12] Candy, L. (2007). Constraints and creativity in the digital arts. *Leonardo*, 40(4):366–367.

[13] Chang, M., Druga, S., Fiannaca, A. J., Vergani, P., Kulkarni, C., Cai, C. J., and Terry, M. (2023). The prompt artists. In *Proceedings of the 15th Conference on Creativity and Cognition*, CC '23, page 75–87, New York, NY, USA. Association for Computing Machinery.

[14] Dang, H., Mecke, L., Lehmann, F., Goller, S., and Buschek, D. (2022). How to prompt? opportunities and challenges of zero-and few-shot learning for human-ai interaction in creative applications of generative models. *arXiv preprint arXiv:2209.01390*.

[15] Davis, R. L., Wambsganss, T., Jiang, W., Kim, K. G., Käser, T., and Dillenbourg, P. (2024). Fashioning creative expertise with generative ai: Graphical interfaces for design space exploration better support ideation than text prompts. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.

[16] Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., and Sutskever, I. (2020). Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*.

[17] Doherty, G. and Massink, M. (1999). Continuous interaction and human control. In *Proceedings of the XVIII European annual conference on human decision making and manual control*, pages 80–96.

[18] Engel, J., Agrawal, K. K., Chen, S., Gulrajani, I., Donahue, C., and Roberts, A. (2019). Gansynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710*.

[19] Evans, N., Haki, B., and Jordà Puig, S. (2024). Groovetransformer: a generative drum sequencer eurorack module.

[20] Feng, Y., Wang, X., Wong, K. K., Wang, S., Lu, Y., Zhu, M., Wang, B., and Chen, W. (2023). Promptmagician: Interactive prompt engineering for text-to-image creation. *IEEE Transactions on Visualization and Computer Graphics*.

[21] Foster, D. (2022). *Generative deep learning*. "O'Reilly Media, Inc.".

[22] Fox, M., Vaidyanathan, G., and Breese, J. L. (2024). The impact of artificial intelligence on musicians. *Issues in Information Systems*, 25(3).

[23] Gammelgård-Larsen, A., van Berkel, N., Skov, M. B., and Kjeldskov, J. (2024). Designing for human-ai interaction: Comparing intermittent, continuous, and proactive interactions for a music application. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '24, New York, NY, USA. Association for Computing Machinery.

[24] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

[25] Höök, K. (2000). Steps to take before intelligent user interfaces become real. *Interacting with computers*, 12(4):409–426.

[26] Hornbæk, K. and Oulasvirta, A. (2017). What is interaction? In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 5040–5052.

[27] Jourdan, T. and Caramiaux, B. (2023). Culture and politics of machine learning in nime: A preliminary qualitative inquiry. In *New interfaces for musical expression (NIME)*.

[28] Kingma, D. P. and Welling, M. (2022). Auto-encoding variational bayes.

[29] Le Vaillant, G., Dutoit, T., and Dekeyser, S. (2021). Improving synthesizer programming from variational autoencoders latent space. In *2021 24th International Conference on Digital Audio Effects (DAFx)*, pages 276–283. IEEE.

[30] Liu, V. and Chilton, L. B. (2022). Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, pages 1–23.

[31] Mansoux, A., Howell, B., Barok, D., and Heikkilä, V.-M. (2023). Permacomputing Aesthetics: Potential and Limits of Constraints in Computational Art, Design and Culture. In *Ninth Computing within Limits 2023*. LIMITS. https://limits.pubpub.org/pub/6loh1eqi.

[32] Murray-Browne, T. and Tigas, P. (2021). Latent mappings: Generating open-ended expressive mappings using variational autoencoders.

[33] Norman, D. A. (1994). How might people interact with agents. *Communications of the ACM*, 37(7):68–71.

[34] Oppenlaender, J. (2022). The creativity of text-to-image generation. In *Proceedings of the 25th International Academic Mindtrek Conference*, Academic Mindtrek '22, page 192–202, New York, NY, USA. Association for Computing Machinery.

[35] Peachey, M., Oore, S., and Malloch, J. (2023). Creating latent representations of synthesizer patches using variational autoencoders. In *2023 4th International Symposium on the Internet of Sounds*, pages 1–7. IEEE.

[36] Peachey, M., Oore, S., and Malloch, J. (2025). Evaluating low-dimensional latent representations as a creative interface for digital synthesizers. In *The Sixth Conference on AI Music Creativity*, pages 1–15.

[37] Peng, X., Koch, J., and Mackay, W. E. (2024). Designprompt: Using multimodal interaction for design exploration with generative ai. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, pages 804–818.

[38] Privato, N., Lepri, G., Magnusson, T., and Einarsson, E. T. (2024). Sketching magnetic interactions for neural synthesis. *TENOR*, page 89.

[39] Privato, N., Magnusson, T., and Einarsson, E. T. (2023). The magnetic score: Somatosensory inscriptions and relational design in the instrument-score. In *Proceedings of the International Conference on Technologies for Music Notation and Representation–TENOR*, pages 36–44.

[40] Roberts, A., Engel, J., Raffel, C., Hawthorne, C., and Eck, D. (2018). A hierarchical latent vector model for learning long-term structure in music. In *International conference on machine learning*, pages 4364–4373. PMLR.

[41] Sanchez, T. (2023). Examining the text-to-image community of practice: Why and how do people prompt generative ais? In *Proceedings of the 15th Conference on Creativity and Cognition*, pages 43–61.

[42] Torricelli, M., Martino, M., Baronchelli, A., and Aiello, L. M. (2024). The role of interface design on prompt-mediated creativity in generative ai. In *Proceedings of the 16th ACM Web Science Conference*, pages 235–240.

[43] Van Berkel, N., Skov, M. B., and Kjeldskov, J. (2021). Human-ai interaction: intermittent, continuous, and proactive. *Interactions*, 28(6):67–71.

[44] Wang, C. (2025). Beyond authorship: Human-ai collaboration and open creative processes in music. In *Journées d'Informatique Musicale*.

[45] Wang, Z., Huang, Y., Song, D., Ma, L., and Zhang, T. (2024). Promptcharm: Text-to-image generation through multi-modal prompting and refinement. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.

[46] Warren, N. and Çamcı, A. (2022). Latent drummer: A new abstraction for modular sequencers.

[47] Weber, A., Alegre, L. N., Tørresen, J., and Castro da Silva, B. (2019). Parameterized melody generation with autoencoders and temporally-consistent noise. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 174–179. Universidade Federal do Rio Grande do Sul.

A APPENDIX | SEMI-STRUCTURED INTERVIEW QUESTIONS

Section 1: Working with Generative AI in a Creative Context
1. Please share a bit about your creative process (i.e., what are your tools, techniques, inspirations, etc.)? 2. When using Generative AI as part of your workflow, do you feel like you're doing art? 3. What are the opportunities or advantages that Generative AI has offered (or can offer) your creative workflow? 4. What have been (or could be) the challenges of integrating Generative AI into your creative workflow? 5. What do you think is the importance of Human in the Loop when using Generative AI within a creative context? 6. When working with Generative AI, do you feel like you are a creator? 7. When using Generative AI as part of your workflow, do you feel like you did art?

Section 2: Continuous vs Discrete Interactions
1. Has your experience working with or thinking about AI tools for creativity been focused on the use of discrete or continuous interactions? 2. Which of these (discrete or continuous) interactions would fit better into your creative workflow? 3. In your opinion, what are some benefits or drawbacks of discrete interactions? 4. In your opinion, what are some benefits or drawbacks of continuous interactions? 5. If you could choose, would you prefer to have the ability to explore the entirety of what a generative model could output? Or would you rather be exposed to new possibilities in smaller increments? What might be the benefits of each approach?

Section 3: Need for Personalization of Tools
1. Is it important to you that Generative AI tools incorporate some kind of personalization? Can you provide some examples of when/how? 2. Do you feel as though current Generative AI tools are personalized in such a way that they support individual creative workflows or creative practitioners? 3. How would you expect your Generative AI powered tools to adapt to how you use them over time? What kind of personalization would you expect? 4. Do you have an example of a tool (Generative AI or otherwise) that you've used that you felt did a good job of personalization?

Section 4: Openness to Change
1. How likely do you believe yourself to be able to change your mind about using AI in a creative context? 2. Have you changed your mind about using or not using Generative AI in your creative workflow in the past year? 3. Do you believe that you would be open to changing your mind again? If so, what might influence that change?

Table 2: The set of questions that were used to guide the semi-structured interview. There were four subsets each focused on invoking a discussion around a central theme of the inquiry, however participants were free to answer each question as they felt most appropriate.