Using Incongruous Genres to Explore Music Making with AI Generated Content

NICK BRYAN-KINNS, University of the Arts London, UK
ASHLEY NOEL-HIRST, Queen Mary University of London, UK
COREY FORD, Queen Mary University of London, UK

Deep learning generative AI models trained on huge datasets are capable of producing complex and high quality music. However, there are few studies of how AI Generated Content (AIGC) is actually used or appropriated in creative practice. We present two first-person accounts by musician-researchers of explorations of an interactive generative AI system trained on Irish Folk music. The AI is intentionally used by musicians from incongruous genres of Punk and Glitch to explore questions of how the model is appropriated into creative practice and how it changes creative practice when used outside of its intended genre. Reflections on the first-person accounts highlight issues of control, ambiguity, trust, and filtering AIGC. The accounts also highlight the role of AI as an audience and critic and how the musicians' practice changed in response to the AIGC. We suggest that our incongruous approach may help to foreground the creative work and frictions in human-AI creative practice.

CCS Concepts: • **Applied computing** → **Sound and music computing**; • **Human-centered computing** → *HCI design and evaluation methods*; • **Computing methodologies** → **Artificial intelligence**.

Additional Key Words and Phrases: generative AI, first-person accounts, music generation, datasets, Folk, Glitch, Punk

ACM Reference Format:
Nick Bryan-Kinns, Ashley Noel-Hirst, and Corey Ford. 2024. Using Incongruous Genres to Explore Music Making with AI Generated Content. In *Creativity and Cognition (C&C '24), June 23–26, 2024, Chicago, IL, USA.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3635636.3656198

1 INTRODUCTION

Generating music with computers and Artificial Intelligence (AI) has a history dating back to the dawn of modern computing in the 1950s. Similarly there are decades of research on interactivity of computer arts [19], computer based music making [29, 58], and Human-Computer Interaction (HCI) research on making music with computers [31, 51, 64]. With the recent dramatic growth in generative AI there has likewise been increasing HCI interest in how AI Generated Content (AIGC) could be used in the creative practice of music making e.g. [8, 30, 37, 61]. However, generative music models have currently been mostly evaluated in terms of their generative performance using musical metrics e.g. [4, 66] or controlled listening tests e.g [28, 54], with few human-centred studies on how these models are actually used by musicians.

In this paper we take an AI plugin trained on an Irish Folk music dataset [60] and explore how it is used and appropriated in the music making practices of two sets of musicians from different musical genres. We purposefully use incongruous genres of music as a form of playful and speculative HCI design thinking [33, 56]. Reflecting on first-person accounts of these explorations we offer insights into the use and limitations of the AI plugin and how it was appropriated by musicians into their creative practices. Specifically, we explore the following question in this paper: *How is AI Generated Content used in musicians' music making practices of a different genre?* We offer the following contributions:

- Reflective accounts of how a generative AI tool is used and appropriated outside its intended musical genre;
- Reflections on how music making practice changes in response to the introduction of an AI model;
- A playful approach to exploring the use and appropriation of incongruous AI Generated Content.

In this paper we first give a background on computer music making, specifically generative music, and how it has been studied in HCI. We then introduce our methodology and AI plugin and discuss why we take a first-person study approach. This is followed by two case studies of musicians' use of the generative music plugin. We reflect on the appropriation of the AI models and the impact on music making practices in the discussion.

2 BACKGROUND

Making music with real-time interactive systems has been envisioned since the early days of electronic computation [22]. We now have a wide range of interactive systems for music making, e.g. [51]. Digital tools are now part of complex music composition workflows, where musical ideas are developed and shared across numerous individuals, software, and environments [41]. Given this richness of music interaction, Human-Computer Interaction (HCI) with digital tools for music making has been explored in various ways from task-based evaluations of the usability of digital musical interfaces [64] to user experience and engagement [9, 52]. Recent developments in deep learning systems have led to increasingly convincing and high quality musical outputs [17], such as FolkRNN [60] which generates convincing Folk music, or DeepBach [28] for Choral music, or MIDIMe [20] which allows users to train a generative AI on their own music. Several generative AI music systems have also contributed to aspects of the music composition process. For example, by connecting two musical phrases (e.g. [50]), turning a melody into chords (e.g. [37]), or making it easier to combine generative outputs in a Digital Audio Workstation (DAW)[1] [8, 30]. Other systems allow users to navigate within an AI model to generate musical output. For example, Sonified Body [43, 44] in which a dancer's movement is used to explore the latent space of a generative music model.

2.1 Studies of the Use of Generative Music Systems in Musical Practice

Whilst there is an ever increasing number of generative music systems there is a lack of research on how these systems are used in musical practice and what effect these systems might have on music making. As Jourdan and Caramiaux [32] note, there has been a lack of user-centred evaluation of Machine Learning systems for music over the past ten years. Of the few works reflecting on the use of generative AI in music making practice (e.g. [37, 40]), Xambó [65] review the features of nine virtual agent systems used in Live Coding to generate music. Their review highlighted questions of machine musicianship, agency, and autonomy, and offered an analysis of the features of the generative music systems used – connecting with broader guidelines for Human-AI interaction [1]. Pachet et al. [47] reflect on the use of a generative music system designed to produce real-time accompaniment to musicians, allowing users to respond the AI in real-time. In this case the researcher-musician reflects on their use of their generative music system from a personal perspective, much as Murray-Browne and Tigas [43] reflect on their design and use of Sonified Body. In contrast, Ben-Tal et al. [6] examine how the outputs of their *FolkRNN* generative folk music system are used by themselves as performers and composers. Importantly, they also explore how FolkRNN is used by folk musicians either directly as part of their research or serendipitously through the FolkRNN webpage. Loth et al. [36] curate and edit

[1]A Digital Audio Workstation (DAW) is a type software widely used for music composition, editing, and production using audio or symbolic presentations.

generative music outputs in the progressive metal genre to create original compositions. They noted challenges with using AI in this way such as the generation of musical ideas not reasonably playable by a human or producing outputs which needed tweaking to sound more natural. Armitage and Magnusson [2] explored musicians' engagement with musical scores formed as real-time agents, for example encounters between two guitarists and these 'agential' scores. They highlight questions around the sense of agency with agent generated scores and how the forms of interaction between musicians and artificial agents might be studied and better understood.

3 METHODOLOGY

Given the exploratory nature of our enquiries into how AIGC might be used in music making practice we employ a subjective first-person research methodology rather than a more objective lab-based approach and take inspiration from practice-led research across HCI e.g. [24, 34–36, 59]. In doing so we aim to capture the nuanced minutiae of creative practice, considering the artist as both a researcher and participant [38, 45]. To explore the appropriation of AI in music making practice we purposefully engage musicians with an AI model trained on a different musical genre to their own. In doing so we take inspiration from playful approaches to speculative musical interface [33] and HCI design [56] with the "purpose of challenging [the musicians] to use an unfamiliar creative framework" (ibid.) as a driver to explore interaction and appropriation of an AI model. An overview of decades of traditions and approaches to exploring incongruous, speculative, surreal, and often absurd approaches to music making and technologies is given in [33]. Given the early stage of this research area and our exploratory approach we reflect post-hoc on similarities and differences to the congruous use folkRNN by folk musicians [6].

3.1 Musicians

We provided three practising musicians with an AI plugin (Section 3.2) to use over extended periods of time. Musicians were asked to use the AI plugin in their usual music making practice to create a new piece of music over a period of time commensurate with their usual music making practice and to document and reflect on their music making process as they used the AI plugin. These accounts form the basis of the cases studies in Sections 4 and 5. Our first case study (Punk) reports on two musicians' collaborative use of the plugin to compose Punk music for their band across four structured writing sessions. The second case study (Glitch) reports on one musician's use of the plugin to compose generative Glitch music over 6 months. Both genres are intentionally musically different to Irish Folk music e.g. in terms of musical structure and technique, and offer a glimpse into the incongruous use of AI from one genre in another. The explorations span different time frames to mirror the different musical practices of the case studies. Our first-person accounts offer subjective perspectives on the use of AI in music making practice.

Two of the musicians are co-authors of this paper as detailed in the case studies. Two of the musicians worked with their supervisor at Queen Mary University of London (QMUL), UK and both had undertaken advanced level study of AI model creation and deployment, though only the Glitch musician had previously used AI in their music making. The study was fully exempt from ethics review at QMUL.

3.2 Generative AI Plugin

We provided our researcher-musicians with a generative AI music plugin [3] as part of a larger research programme exploring the design and evaluation of explainable AI tools for the Arts [3, 10, 11, 13]. Whilst the plugin is not as finessed as commercial AI tools such as Google Magenta's plugin suite [55], it does offer some similar functionalities for melody generation.

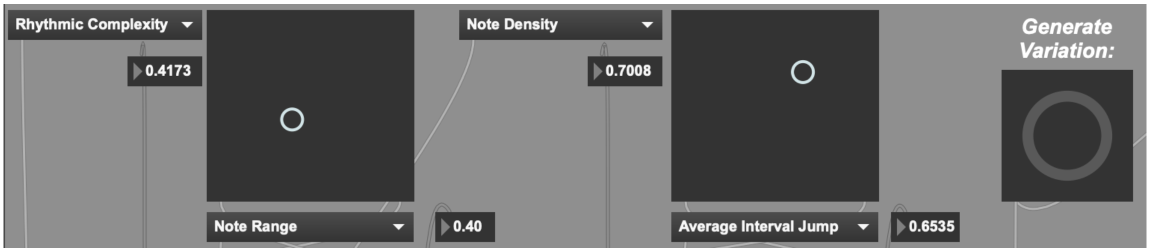

Fig. 1. Generative AI plugin user interface [3]

The user interface of the plugin is built as a Max4Live device[2] which can be used in Ableton Live[3], a popular music making and performance software, and is open-source, allowing us to contribute to a more open AI research community. The plugin uses a Machine Learning (ML) model to create variations of human-composed themes as MIDI[4] files. To influence the music generation users are able to move points on two two-dimensional input spaces in the user interface (UI), shown as two large black squares in Figure 1 and referred to as *pads*. The axes of each of the pads correspond to four metrics commonly used in AI music research [4, 66]: the left square offers *rhythmic complexity* (the amount of variety in note duration and syncopation) and *note range* (distance between highest and lowest notes); the right square offers *note density* (number of notes in a musical measure), and *average interval jump* (how large jumps are between the notes in a measure). When the "Generate Variation" button is clicked the ML architecture of the plugin generates a new MIDI file from the points in the 2D pads along with the MIDI input provided by the user.

The underlying ML architecture of the plugin uses a Variational Auto-Encoder (VAE) consisting of encoder and decoder blocks, which in this case are neural networks. The encoder compresses datasets into a smaller latent vector, which is then used an input to be reconstructed by the decoder. The latent vector can then be explored to produce novel outputs through the decoder. This plugin uses the MeasureVAE architecture [50] to construct measures of music and as with many other generative AI tools was trained on 20,000 monophonic Irish folk melodies [60]. Regularisation terms were imposed during the training of MeasureVAE so that changes in the first four dimensions of the latent vector are mapped to changes in the four musical metrics described above, following Pati and Lerch [48, 49].

4 CASE STUDY: PUNK

This section introduces the first-person account of how the AI plugin was used by two musicians in a Punk band. The writing in this section is by two musicians who used the plugin together. Writing is in the first-person, using "we" to refer to the two musicians.

4.1 Artist Background & Process

We are the guitarist and bass player of a Punk band which started making music and gigging in August 2022. The guitarist is an artist-researcher and author of this paper, is completing a HCI PhD thesis on AI music making, and has formal music composition training from their undergraduate degree in music technology. The bass player has an

[2]Note that we are using 'plugin' in the general sense to indicate a small piece of software used within a larger piece of software rather than a DAW specific 'plugin' format such as VST.
[3]https://www.ableton.com/en/
[4]MIDI is widely used a file format containing representations of musical notes.

undergraduate degree in Film, Theatre and Television, is a school teacher (ages 5 through 11 years old), and has actively played music across 11 bands from Punk to Folk to Jazz.

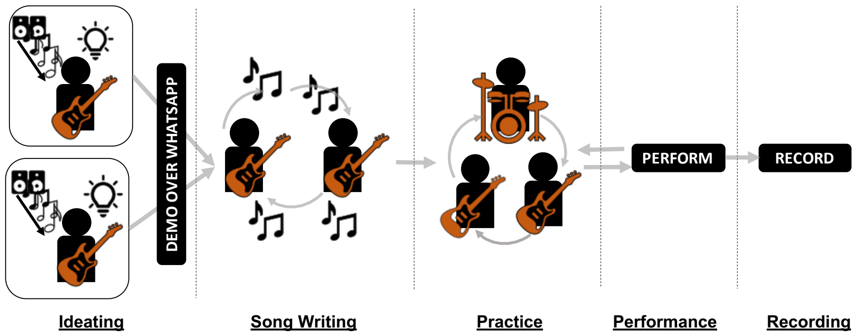

Fig. 2. The creative process for the Punk band case study.

Our typical composition process is shown in Figure 2. In the ideating phase, we usually write ideas at home, and share demos over WhatsApp[5], taking inspiration from existing Punk songs. We then meet in the songwriting phase to build upon our ideas until we have the full structure and lyrics for a song. Only then do we meet with the drummer to develop the songs more fully in the practice phase. Once the songs have been played to audiences in the performance phase, we then record our tracks. Once a song is brought to the practice phase, we rarely go back to the songwriting phase.

4.2 Documentation

The AI plugin and Ableton were setup on a laptop we typically use in our music making process and we received short initial instruction on how to use it. We then met across 4 sessions to use the plugin as part of our music making. These played out as follows, although we did not have set goals in advance:

- **Session 1 (33m):** We met to initially test and understand the plugin.
- **Session 2 (1h 33m):** We treated the AI as a way to generate ideas for the bass guitar. Choosing the bass here was an arbitrary choice – we felt we *had* to sacrifice one of our instruments to the AI system to initiate collaboration. After the session, the bassist took the ideas and created a demo at home.
- **Session 3 (1h 28m):** We used the AI to elaborate on the piece, writing a variation for our second verse by putting together generations of the AI, taking the riff from the previous session as the input seed.
- **Session 4 (2h 54m):** We set out to record our piece.

We decided to video record ourselves and the AI plugin to help us reflect on our music making later. We also noted moments which we felt were significant in our creative process by jotting down the current time whilst composing music. All data was then gathered in Miro[6] to create a timeline of how our songwriting evolved. This process was inspired by video-cued recall [14, 15] - the method is inconspicuous and does not disrupt the creative flow of our writing sessions, yet supports post-hoc data analysis. After each session we revisited the extracted video clips, writing

[5]https://web.whatsapp.com/
[6]https://miro.com/

reflections onto post-it notes in Miro. Inspired by the method in Lewis et al. [35], we then reviewed our post-it notes to generate a set of topics describing our experience with the AI plugin and how it informed our practice.

4.3 Findings

We identified four topics from our collected data which arereported here as written narrative, supported by "thick descriptions" [27] and visual artefacts. Audio examples can be found in the Appendix.

T1.1: Punk & Identity. Throughout the project we considered the AI outputs in relation to the band's musical identity. For example, we input riffs written by us into the plugin so we could have confidence that the AI outputs might resemble our style. That said, we were still hesitant to feed the AI our best ideas, as we knew the AI would take over or change our initial thoughts, and didn't have full *trust* in the AI to make good use of them. We also found ourselves cross-referencing Punk songs and other influences that we both liked with the AI's outputs. On reflection, this is something that we might do in our usual practice, but perhaps we had to do this more so with the AI as its output lacked an understanding of Punk and its conventions.

Consequently, when generating outputs, we first would try to steer the AI such that it created music closer to Punk conventions. For example, in the second session we initially set all the points of the pads to 50%. Next, we set the note range low to sit more like a bass (which would normally play the bottom notes of chords to emphasise the guitar), and then set the note density to high (the bass player in our band likes to play lots of notes in a phrase). The output we then found too jaunty for Punk (likely due to the folk data set), so we decided to decrease rhythmic complexity and note range to almost zero, appropriating the model to make the rhythm more simplistic like Punk. This worked as a fast single note phrase was created which we considered as a very Punk style phrase – see Figure 3. From this, we suggest that the rhythmic complexity values weren't very useful to us because, for the Punk genre, complex rhythms are often sparse or, when used, have the whole band playing the same rhythm together.

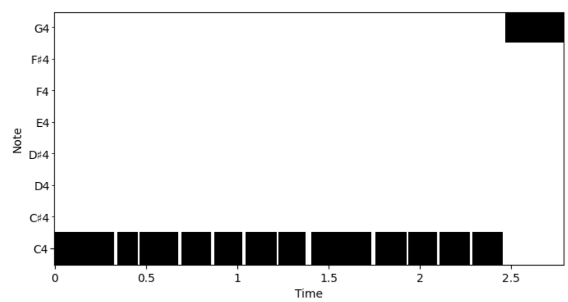

Fig. 3. A Punk style musical phrase generated in the Punk case study with the note range and rhythmic complexity set low, note density and average interval jump set high. The rectangles represent musical notes (pitches on vertical axis) arranged over time (horizontal axis; seconds).

The second verse of our song involves the guitar and bass taking turns, overlapping bar by bar different variations of the main riff, which we devised by appropriating MIDI outputs produced by the AI plugin, as shown in Figure 4. In Figure 4, the green phrase is human-generated, whereas the orange and blue phrases are AI-generated from the green phrase. We decided to have the guitarist start playing the first phrase (green - human generated), then when they played the second phrase (orange - AI-generated) the bass player starts playing the previous phrase over the top, then

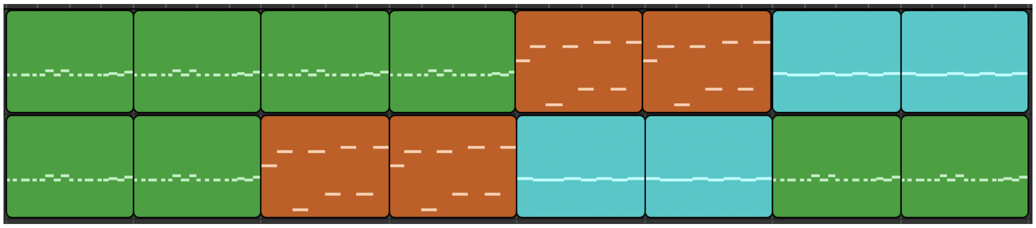

Fig. 4. Second verse of the song made in the Punk case study. The top row is the bass part. The bottom row is the the lead guitar part. The phrases are colour coded to show how the material overlaps with one another, described in Section 4.3.

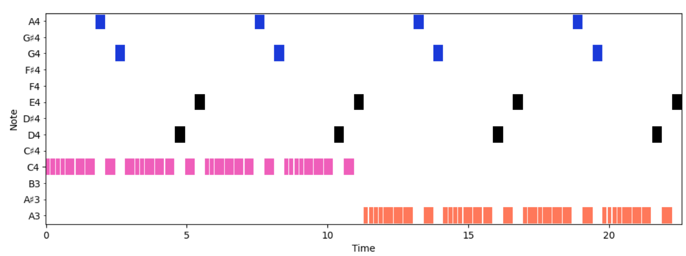

Fig. 5. The main riff for the song in the Punk case study. Notes A4 and G4 were taken from an AI output which sparked the inspiration for rest of the riff. Notes are colour coded for clarity.

the third phrase (blue - AI-generated), and so on. In this way, we ended up following a less typical structure to our usual songwriting, which would likely have repeated the first verse verbatim without these extra AI generations.

When playing the final demo to our drummer, he suggested that he didn't like this new AI section, although couldn't quite articulate why, other than that it sounded "wrong". It is possible that this is because the AI material is away from the genre or the style of our band, sounding out of place. We found this an exciting new direction, although perhaps it drifts too far away from our style to become a future feature of our music.

T1.2: Timbre. The main verses of the song we produced use a motif partly inspired by an AI-output shown in Figure 5. Specifically, we took inspiration from two notes of an AI-generated output we identified in the second writing session (in blue), and move between a C major chord (in pink) and an A minor chord (in orange). During the recording of our track, we decided to use a distortion effect on the guitar with the A minor chord, creating an alternation between (non-distorted) quiet and (distorted) heavy. It's plausible that we incorporated a distortion effect here to be able to add variety to the repetitive nature of the motif, which might be a side-effect of the plugin only generating short measures.

We found also that the AI-inspired melody needed a more powerful timbre to sound closer to the aggressive style of Punk music. A Punk band would likely use power chords (where notes are played with other notes at the same time, a fifth higher in pitch) to make the guitar more aggressive and heavier. In our case, trying to play the AI-generated outputs such as Figure 6a as power chords on the guitar, however, would've involved moving around the guitar neck too much, so we opted to use a digital octave effect (which adds a pitch an octave below the current note to add power to the overall sound, similarly to power chords). It is worth noting that using effects in addition to the guitar and amplifiers is something that, as a band, we have actively resisted, favouring a simple setup which is characteristic of the Punk genre and easier to carry to performance venues.

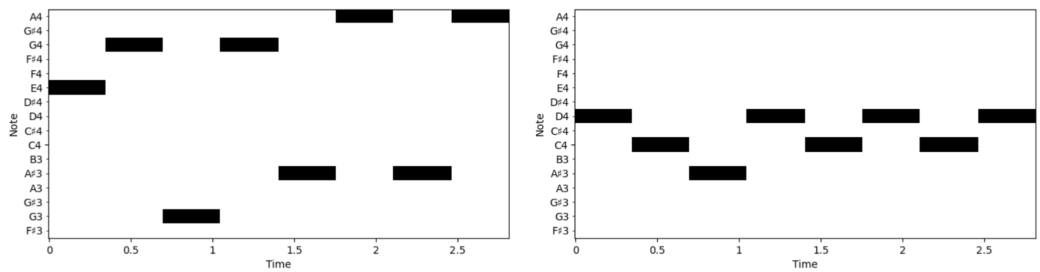

(a) (b)

Fig. 6. Figure (a) and (b) show variations on the main riff shown in Figure 5 which was input into the plugin. These correspond also to the orange (Figure (a)) and blue (Figure (b)) outputs in Figure 4.

T1.3: Looping. We found that the most effective way of selecting AI outputs we liked was to loop outputs generated by the AI and listening through a set of speakers whilst we played alongside with our musical instruments. This was closer to how we would write music normally. This looping was also easier in later sessions once we had become accustomed to controlling the AI, and were more able to “jam” with it, leaving loops running in the background whilst tweaking parameters and adding them to the live soundscape. For example, by the third session we were comfortable with dragging the MIDI outputs created by the AI into the session “live” and switching between them, without breaking up the flow of music coming from the speakers.

Through looping, we used the AI as an idea machine. We would play around looping ideas until we latched onto an idea. However, we largely decided to follow AI-generated ideas that we found “least worst” and then tinkered with the phrases as we saw fit, as opposed to hoping to find generated music which we fell in love with. For example, the main riff started off with the two blue notes in Figure 5, which we extended to include the two red notes when experimenting with the riff on guitar – finding a phrasing that was idiomatic for guitar (see Topic 4) and more in-keeping with Punk conventions.

T1.4: Playability. We found the MIDI outputs created by the AI plugin difficult to play on our instruments. The AI plugin was unaware of the fingering used on the bass guitar when generating its outputs, suggesting ideas that required large stretches across the neck of the instrument. As the guitarist phrased it in our recordings “a bass just wouldn’t do that”. We would also manually edit phrases to make them better fit our musical style. For example, we manually dropped phrases down an octave to be more bass-like. We also found that because the parameters on the plugin’s UI were entangled, changes in one of the metrics on the pads might not have any effect on the AI outputs depending on the values of the other metrics. For example, a small note density and high note range setting for the AI model might only generated MIDI files with a single note as the small note density limits the number of possible notes that could be generated. Despite this confusion, we found that the vagueness of the pads encouraged us to go back to the plugin and try to find new outputs, when we had a rough gist of what we were trying to produce. For example, although the settings on the plugin were not obvious in the first-session, we felt that play was encouraged by this because, even if we had trouble with the outputs, we could simply move the dots around on the pads and build some intuitive understanding of how the outputs might change.

5 CASE STUDY: GLITCH

This section introduces the first-person account of how a generative Glitch musician used the AI plugin in their creative practice. The writing in this section is written by the researcher-musician using first-person language.

5.1 Artist Background & Process

I am a Sonic Artist, improviser and composer. My background is in Philosophy and Computer Music. I am currently undertaking a PhD in HCI and Explainable AI, with a focus on musical systems. I compose interactive and generative music systems for both live and installation environments. My compositions employ sampling (manipulating prerecorded audio), Live Coding ((re)forming music software in real-time) and Glitch (exploring and leveraging "errors" in technological systems).

My compositional process tends to start with an abstract idea which I try to represent with a number of simple examples. These examples might come from collaborative songwriting sessions, algorithmic experimentation, or humming to myself on the bus home. Then, I analyse these sets of ideas, keeping a diary of the themes which cut across them. I use Max/MSP[7] to explore these ideas as musical algorithms created as patches[8].

5.2 Documentation

In preparation for the study, I installed the plugin and received a short instruction on how to use it. Between February-July 2023 (6 months), I kept track of my music-making process with the AI plugin, using a diary, audio recordings, screenshots, project files and sketches rather than video recordings used in Section 4. These choices in data collection were due to the nature of my compositional practice and the longitudinal nature of my music-making which typically extends over several months making video documentation and analysis cumbersome and prohibitively time consuming.

All data was collected into a folder on my laptop, where it was labelled and organised chronologically. I reviewed and analysed my reflections in between each session, adding memo notes to the folder in order to capture shifts in my compositional approach. This first-person approach was inspired by Benford et al. [7]. In line with the approach taken in the Punk case study (section 4) and inspired by Ellis et al. [23], Lewis et al. [35], I then reviewed my data, highlighting a set of topics which best described the range of experiences I had working with the AI plugin.

5.3 Findings

In my analysis, I identified three topics presented below as written narrative supported with "thick descriptions" [27] and visual artefacts. Audio can be found in the Appendix.

Before introducing the topics I would like to note for context that I was the first person from my music making community to use the plugin, and I wanted to use it in as novel a way as possible. In my initial explorations of the plugin I spent considerable time and effort editing the plugin's patch to try to create generative Glitch. Whilst I had very few musical outputs during this period, on reflection, I was finding the edges in the musical space of the AI through my explorations. The later stages of my process where I used the plugin with other patches and musical algorithms proved to me a more productive approach.

T2.1: Dissonance Between Glitch and AI Musical Dimensions

[7]Max/MSP is a visual programming language typically used for audio and multimedia in which data is passed between objects which process data or support user interaction. https://cycling74.com/products/max
[8]An end-user Max/MSP program is referred to as a patch.

Many AI models for real-time music generation allow for navigation around unlabelled dimensions of the AI's latent space, necessitating naive exploration [43]. In contrast, the plugin outlined in Section 3.2 offers four musically labelled semantic axes for navigating its latent space. Rather than getting inspiration from a generative AI's unpredictability, the plugin's labelled axes offered me some feeling of control. Interestingly this reduced my sense of play and exploration since I felt that I needed to have an idea of what I wanted to get out of the plugin *before* use. For this to be effective, I needed to create a mental model linking musical inputs and the labelled axes to the plugin's musical output. At first I struggled, but this was made easier when I thought deeper about the "rhythmic complexity" axis – a perceptually motivated descriptor introduced by Thul and Toussaint in [62]. This metric is closely linked to "euclidean rhythms"[63], which are a common compositional device in the Live Coding community that I am part of. Recognising this link inspired me to explore this approach in my own practice. I developed an algorithmic process to generate MIDI clips which were then "fed" to the plugin. First I used euclidean rhythms assigned to notes from a chosen chord and then used monophonic processing to turn these layers into a melody. I then used this melody as musical input to the plugin's VAE encoder. Once encoded, I would then steer the melody generation using the semantic axes to control the decoder. In this way both the melody of my algorithmic system and the latent space of the VAE could be controlled by me as the user. This allowed me to explore the limitations of the plugin's latent space. Once I had some mental model of the latent space, changing the semantic axes of the VAE decoder made for useful exploration of neighboring ideas. However, the generative system I developed generated fast, syncopated rhythms. This greatly contrasted the dataset of Folk songs on which the model was trained. As I explored further, outputs of euclidean rhythm generation could not be represented by the VAE tool. Here my compositional process was shaped in the negative – by the things the model struggled to generate.

This form of glitch was more than a technical glitch, it felt to me like the sonification of a dissonance between my glitch music making practice and the intended meaning of the labelled axes of the plugin's latent space. Reflecting on this further, my aim became to explore dissonance as far as possible – bringing to the fore the complexity and opaqueness of the plugin's VAE. In the glitches of the systems, I explored a dialogue between my individual artistic practice and the often uncanny, hegemonic applications of commercial generative AI [26, 53].

T2.2: Recursive Errors Inform Understanding There were a number of patches I developed through my creative process in which I explored the dissonant identity of the VAE through feeding an encoding process back into itself. I explored recursive generative systems which included the plugin in successive iterations. For example, Table 1 shows examples of the evolution of drum patterns through successive iteration. Initially, a seed would be given to the plugin (top left in the table), and then encoded and decoded by the plugin as shown in the Iteration 0 column (with 4 variations of the seed made by applying maximum Note Range, Rhythmic Complexity, Note Density, and Average Interval Jump in the plugin). These would then be transformed before being fed back into the plugin. A simple transformation might be transposing all notes up by one MIDI note. With other, more complex transformations used as I iterated further, small gaps between the seed and the VAE output would widen, shown in Iteration 1 and 2 columns of Table 1. Notes were rapidly introduced from outside the set of notes which I had originally written the phrases with, and complex rhythms would collapse into simple eighth notes. The musical output at certain locations in the VAE's latent space became chaotic when pushed beyond the conventions of the Folk training set – as patterns became more complex through recursions, the musical output from the plugin would become a whole new pattern which stopped resembling anything which came before it. As I grew sensitive to these points that I thought of as musical collapse during the composition, I gained a more intuitive understanding of the plugin's latent space and its limits.

Table 1. The effect of successive iterations on a seed pattern in the Glitch case study.

	Iteration 0	Iteration 1	Iteration 2
Seed	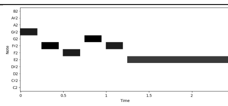	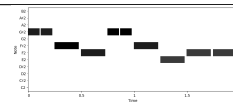	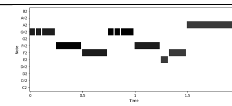
Seed + Max Note Range	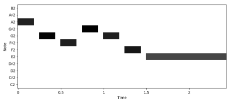	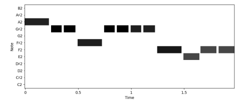	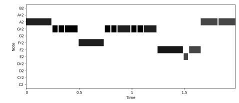
Seed + Max Rhythmic Complexity	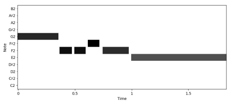	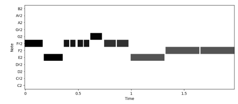	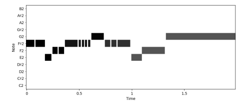
Seed + Max Note Density	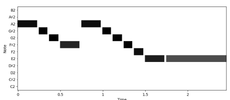	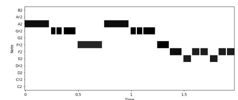	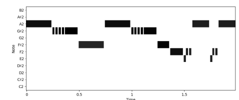
Seed + Max Average Interval Jump	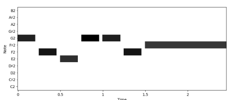	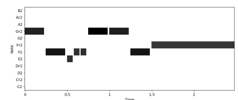	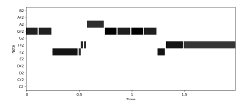

T2.3: Beyond AI Composition Whenever I tired of using the AI plugin, I began transcribing and learning songs on the guitar again and helping to write a jingle for a local radio station. I found freedom in my distance from the AI's "creativity support" though the plugin sat in the corner of my screen haunting me. Later I decided that these ideas would become offerings to the plugin - I needed a collaborator with some distance from the piece, or at least some temporal distance from my own ideas.

With months of experience using the plugin, I was aware that my compositional ideas contained abstractions which the labelled dimensions of the plugin could not encode. I worried that the plugin's latent space encodings – optimised for musical elements which weren't as relevant to my ideas – might distort these abstractions, such as "modal interchange" or "rhythmic modulation" present in my composition. In my recordings of my guitar playing, I had also included explanations of my thought process for each idea, e.g. "the core idea here is outlining this Lydian sharp 5, and I'm starting every phrase on the offbeat". Aware that the plugin would discard such concepts, I clung to these larger structures so that I could reinstate them at a later point. The phrases in Figure 7a are grouped based on these concepts (represented by the different colours). Since they are generated by the plugin and do not retain the more guitar-specific aspects of the composition, I used sketches to communicate more abstract elements such as melodic contour, guitar-specific techniques and harmonic contexts. Figure 7b shows an example sketch – the arrows depict movement across the guitar neck more related to how notes on the guitar would be played intuitively on the guitar than could be captured by the MIDI generated by the plugin illustrated in Figure 7a. For each idea, I would send it through the plugin, before curating and transforming it by hand to match the initial description, label or drawing. This gave me the illusion that I had a collaborator who was interpreting my work, albeit without the influence of the original song. It also pro-actively gave me an audience, forcing me to consider the perspectives and processes of the system as well as my own. All of this collaboration meant that ideas didn't have to be "finished". I merely had to specify the parts of the melody which I felt were essential, then enforce these on any output from the plugin.

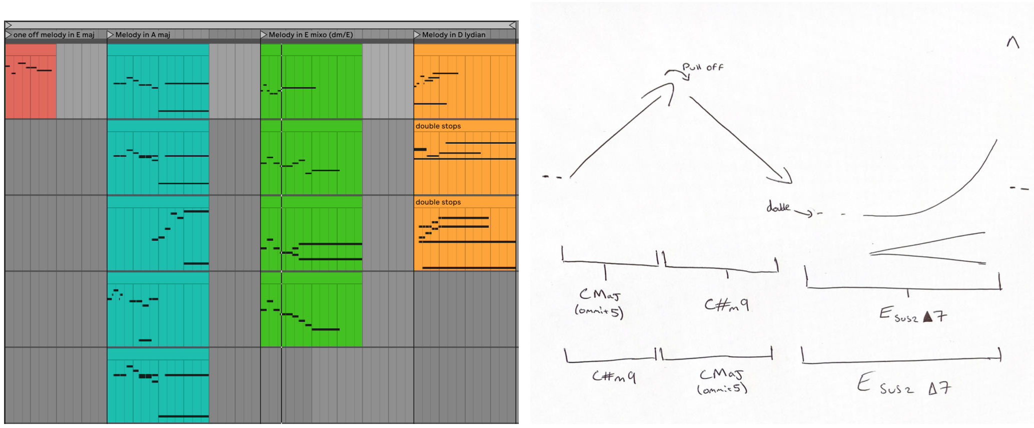

(a) (b)

Fig. 7. (a) A screenshot showing 13 musical ideas that were retained in the Glitch case study; (b) An excerpt from my notebook visually describing a musical gesture with guitar-specific techniques and a harmonic context

6 DISCUSSION

To recap, this paper details first-person accounts of musicians integrating a generative AI plugin trained on Folk music into their Punk and Glitch music making practices. In this section we reflect on how our musicians appropriated AI into their music making practices, then introduce reflections on the AI tool itself, propose ways in which our approach could be used by other researchers, and then discuss limitations and suggest directions for future work.

6.1 AI and Music Making Practices

Broadly, our musicians took different approaches to appropriate the AI to their practice. The Glitch artist strove to incorporate complex musical techniques such as algorithmic composition to appropriate the AI model, explore it, and push it to its limits, perhaps into a different genre. Their practice retained the iterative exploratory nature of their individual music making, but changed as they embraced the AI plugin as a critic or even audience offering a sounding board for their music making (T2.3). In contrast, the Punk musicians followed their usual music making approach but the AI tool pushed them to use musical effects that they would normally actively resist, and the final music composition was even considered “wrong” by the drummer (T1.1). The Punk musicians also manipulated different parameters to the Glitch musician – for example, the Punk musicians found rhythmic complexity was not a useful parameter (T1.1) given that Punk music typically follows strict rhythms, whilst rhythmic variations were central to Glitch music and resonated with the Glitch musician’s compositional inspiration from Toussaint [63] (T2.1). Similarly, the Punk musicians were less interested in exploring the limits of the AI’s latent space than the Glitch musician, and instead were mostly searching for outputs which fitted with their Punk aesthetic (T1.1). Indeed, the Glitch musician’s creative aim became over time to explore the “dissonance” between computer music making conventions and the explanations of the AI plugin (T2.1).

We note similar observations of the use of generative AI FolkRNN by Folk musicians [6] - a congruous setting in contrast to our incongruous approach. Firstly, there are similarities between how the Punk musicians chose outputs from our plugin to work into a composition to the way that folk musicians repeatedly listened to FolkRNN outputs to select a musical phrase to use in their composition. There were also similarities in how the Glitch musician pushed our

AI plugin beyond its musical limits, much as folk musicians could push FolkRNN outside of the Folk genre, by increasing the temperature of FolkRNN [6]. And, there are similarities in how the Glitch musician used our plugin's latent space to Ben-Tal et al. [6]'s example of musicians generating several outputs to identify usable regions of FolkRNN's generative space. Finally, there are similarities in how one folk musician developed a compositional process which initialised the generation process with "combinations that are uncommon in the original data" - much as the Glitch musician did with their algorithmic approach. The main differences we noted in how the folk musicians worked with their AI model, FolkRNN, were that folk musicians were noted to have "corrected some 'mistakes'" in the generated output - in this case mistakes in folk style [6]. This kind of activity was not seen as prominently in our incongruous approach.

6.2 Reflections on our AI Plugin

Here we note reflections on how our musicians used our AI plugin in their music making outlined in the preceding sections. We also suggest **AI design implications** based on these reflections which AIGC designers and researchers could use to inform future system design and studies.

Ambiguous AI. In keeping with research on creative uses of AI we found that the ambiguity of AI control positively contributed to the creative process [16, 25, 37] - in this case the ambiguity of the control pads. We suggest that for the Punk musicians the vagueness of the pads encouraged them to return again and again to the plugin and try to find new outputs (T1.1), yet they reported that they still felt that they had some control over what they were trying to produce. In this way some balance was found between control (the AI generation was not completely random) and serendipity (in finding surprising outputs to then adopt into their practice) [10, 16, 42]. Whilst the AI plugin had been designed with explainable AI design features in mind by imposing musically meaningful metrics on the manipulable dimensions of the latent space [12], there were many points in which both sets of musicians were unclear about the meaning of the controllable dimensions. The Punk musicians found certain outputs confusing, even if the sliders were acting "as intended" (T1.1). This led to the Punk musicians shying away from more complicated exploration of the AI plugin and instead appropriating the AI outputs mostly through "jamming" along using their instruments. In contrast, this confusion became more central to the Glitch musicians practice, developing feedback techniques to take advantage of these complex entangled dimensions. Indeed, as the Glitch musician developed their mental model of the latent space they reported that the semantic dimensions allowed for "useful exploration of neighbouring ideas" (T2.1). In both cases, the musicians spent substantial time listening to the AI-generated outputs on loop, as a way to decide whether to include the output into their music composition. **AI design implications:** Designing ambiguity into AI control can contribute to the creative process. By this we do not mean designing unlabelled controls, but rather suggest designing controls which give a gist of their effect.

Torrential AI. Both our musicians found that they had to narrow down the wealth of material that the AI plugin could create, similar to approaches such as Ben-Tal et al. [6], Loth et al. [36], Sturm [59], in generating AI material, cherry-picking the best outputs, and organising these into a music composition. We suggest that the effort to select the most suitable outputs from a huge set of outputs is exaggerated in our approach - both the Punk and Glitch musicians had to work hard to find suitable musical outputs that suited their respective genres whereas a folk musician would likely find many more suitable melodies in the outputs, e.g. Sturm [59]. **AI design implications:** Given the exaggerated effort needed to select suitable outputs we suggest that researchers designing human-AI interfaces could use our approach to stress test [5] user interfaces for navigating AI outputs.

In contrast to descriptions of cherry-picking the best outputs [6, 36, 59], our musicians described their AI output selection process and criteria more in terms of negative affect. For example, the Glitch musician focused on trying to find areas of the latent space where the music was chaotic (T2.2), using this as a driving motivation for their work. Similarly, the Punk musicians selected outputs which they felt were "least worst". **AI design implications:** This use of negative affect to describe selection criteria suggests scope for alternatives to AI user interfaces navigation to the "best output" - for example, supporting navigation to awkward, unusual, or ill-fitting outputs around the edges of intended AI output might provide interesting and stimulating outputs for artists, much as artists search for AI glitches [13, 18].

Hungry AI. Both sets of musicians described their interaction with the AI in a way that conjures ideas of a hungry or demanding AI. For example, the Glitch musician made "offerings to the plugin" (T2.3) and "fed" the plugin (T2.1), and the Punk musicians "had to sacrifice one of our instruments" to the AI. Possibly the difficulty of using the Folk AI in Glitch and Punk music making caused the musicians to perceive the AI as being demanding. Or, maybe the imposition of the AI model into existing creative practices led to this feeling. Alternatively, it may suggest some theistic conceptions of AI by the musicians cf. [57]. **AI design implications:** Whatever the reasons for perceptions of a hungry AI are, we suggest that an incongruous approach in AI design and study could be used to prompt and provoke discussions around perceptions of AI power and demands.

(Un)Trustworthy AI. The musicians all reported some difficulties in trusting the AI model. The Punk musicians were most explicit about this, stating that they were hesitant to use the AI with their best ideas (T1.1), suggesting that they distrusted the AI and were interested in keeping some creative control [1, 46]. The incongruity of the AI generated music might exaggerate this lack of trust and hamper the development of trust in the AI if it continually makes "bad" use of creative effort. The Glitch musician also expressed caution and the need to spend time building trust of the AI by creating a "mental model linking inputs and the axes to their output" (T2.1). **AI design implications:** We suggest that by using our incongruous approach we exaggerate AI trust issues. This could be used to allow researchers more time and space to investigate how trust develops (or does not) between a human and an AI in creative practice. In essence, an incongruous approach could provide a context for in-depth exploration of trust in human-AI partnerships.

Shallow AI. The Glitch musician felt the need to add extensive sketches and notes to their compositions to allow them to recreate musical features not captured or generated by the AI (T2.3). For the Punk musicians they felt the need to include a digital octave effect to compensate for the shortcomings of the AI generated content despite the use of effects being something that they had "actively resisted" in the past (T1.2). Perhaps, they felt that they *needed* extra equipment to work with the repetitive, looping, nature of the plugin's short measures. We suggest that the incongruous nature of the activity highlighted the use of musical features and artefacts which might not have been observed with congruous AI Generated Content. **AI design implications:** We suggest that an incongruous approach might highlight representations or elements that are lacking in AI user interfaces and thereby suggest user interface or AI features to be designed.

Critical AI. The Glitch musician noted their perception of the AI model as offering a "collaborator" and an "audience" for their compositions (T2.3). This is despite the AI model being training on a different genre of music. Potentially there are similarities here to the value of interdisciplinarity in creative work cf. [39]. The Glitch musician also found that they could use the plugin as a critic, with the AI interpreting the musician's musical contributions through a Folk style and reflecting their ideas back through a different musical lens to prompt their creativity. **AI design implications:** We suggest that our incongruous approach offers opportunities for researchers to study the effect of critical or unexpected forms of AI feedback. It may even be useful in informing the design of AIGC systems tailored to prompting and provoking creativity by, for example, generating feedback that is unexpected or outside the norms and conventions of the genre being worked on.

6.3 Limitations and Future Work

We suggest that the incongruity of the genres in our approach foregrounds the creative work of the musicians and offers opportunities for researchers to study the development of trust and intuition between artists and AI in creative practice. To be clear - we are not arguing that by using incongruous genres we would expect musicians to invent some new genre(s) such as Punk-Folk or Glitch-Folk. We are also not arguing that our incongruous genres approach is a tool for early stage speculative design cf. [33, 56]. We are suggesting that the approach offers opportunities to explore how AI Generated Content is used in creative practice and how creative practice responds to AIGC. From our perspective our incongruous approach exaggerates the frictions that occur between people and AIGC offering researchers opportunities to study these interactions in depth.

The different timescales, approaches, and forms of documentation in our case studies reflected the different ways in which these musicians created music in their genres. Our approach was also shaped by the early and exploratory nature user studies of AIGC at this time - in many ways our study is about generating research questions about how musicians interact with AI over an extended period of time rather than developing generalisable insights about music making with AI. These points clearly limit the generalisability of our results, and whilst we do not see the two case studies as directly comparable, they do illuminate different perspectives on how the AI plugin was used and appropriated. For instance, we suggest that the Glitch musician might have been able to identify opportunities to use the plugin as part of their existing computer-based algorithmic music making practice. This contrasts the Punk musicians who avoided technological complexity where possible, instead jamming on their instruments with the AI. Future work needs to undertake a more consistent approach to the data collection and mix of incongruous genres to allow comparisons between the musicians and genres. Future work also needs to undertake some baseline comparisons to allow direct comparison between congruous and incongruous approaches. Our reflection on Ben-Tal et al. [6]' studies of congruous use of generative music AI in Section 6.1 suggest some similarities and differences, but more controlled studies are required to explore these better.

Whilst we suggest that our incongruous approach foregrounds the frictions in creative practice with AI it does not necessarily foreground the day-to-day reality of how AI is used in creative practice. To do this we as researchers needs to balance more rigorous autoethnographic and ethnographic approaches with our first-person incongruous approach.

We also wish to note that, although it might seem obvious that there are barriers to using AIGC in one genre to construct music for another, there are a number of considerations for why reusing generative models from one genre in another might have advantages beyond provoking creative practice and speculative design thinking. First, models such as MeasureVAE [50] used by the AI plugin we tested [3] require huge training datasets which are only available for a small number of well researched genres such as Folk [60]. This could potentially lead to the marginalization of minority cultures and subgenres of music where large datasets are not available. Future work needs to explore how the AI models can be tailored to under-represented genres of music such as Glitch or minority cultures. Secondly, training such deep learning models consumes large amounts of resources [21] which could be avoided by reusing and re-appropriating existing models. Again, future work needs to explore AI models which can build on existing trained models and datasets to reduce the environmental impact of AI models.

7 CONCLUSION

This paper presented two first-person accounts of how a generative AI plugin trained on Folk music was appropriated by musicians from incongruous genres of Glitch and Folk. We found unexpected ways that our musicians used the plugin to search for musical outputs which made sense in their genre, and how the use of the AI plugin changed their music making practice. From these reflections we suggested design implications for future AIGC systems. Our incongruous approach highlighted the musicians' lack of trust in AI, their approaches to cherry-picking AI output, the role of ambiguity in their music making practices, and how musicians perceived AI as a hungry consumer of their creativity. Whilst we do not argue that our incongruous approach will spark the next genre-breaking moment in music, we do suggest that using incongruous genres when exploring the use of AIGC can foreground the creative work and frictions between people and AI.

With the rapid development of generative AI we find ourselves at an exciting and potentially seismic point in the age old evolution of music making. It is our responsibility as musicians and researchers to explore the possibilities, challenges, and unexpected opportunities that AI might bring. In light of this we hope that our practice-led approach contributes in some small way to the growth of more human-centred Artificial Intelligence for the arts.

CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Nick Bryan-Kinns: Conceptualisation, Methodology, Validation, Resources, Writing - Original Draft, Writing - Review & Editing, Supervision, Funding acquisition; **Ashley Noel-Hirst**: Methodology, Investigation, Data Curation, Writing - Original Draft; **Corey Ford**: Methodology, Validation, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Project administration.

ACKNOWLEDGEMENTS

Corey Ford and Ashley Noel-Hirst are research students at the UKRI Centre for Doctoral Training in Artificial Intelligence and Music, supported jointly by UK Research and Innovation [grant number EP/S022694/1] and Queen Mary University of London. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. Special thanks to Callum Young for their bass playing and deep thinking. Many thanks to Alan Chamberlain and Jeba Rezwana for giving feedback on early drafts of this work, and to Berker Banar for his support with the AI plugin development and use.

APPENDICES

Audio files can be found at: https://on.soundcloud.com/CqMXWNwmiNc6wyxx9.

REFERENCES

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3290605.3300233

[2] Jack Armitage and Thor Magnusson. 2023. Agential Scores: Artificial Life for Emergent, Self-Organising and Entangled Music Notation. In *Proceedings of the International Conference on Technologies for Music Notation and Representation – TENOR'2023*, Anthony Paul De Ritis, Victor Zappi, Jeremy Van Buskirk, and John Mallia (Eds.). Northeastern University, Boston, Massachusetts, USA, 51 – 61.

[3] Berker Banar, Nick Bryan-Kinns, and Simon Colton. 2023. A Tool for Generating Controllable Variations of Musical Themes Using Variational Autoencoders with Latent Space Regularisation. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 13 (June 2023), 16401–16403. https://doi.org/10.1609/aaai.v37i13.27059

[4] Berker Banar and Simon Colton. 2022. A Systematic Evaluation of GPT-2-Based Music Generation. In *Artificial Intelligence in Music, Sound, Art and Design*, Tiago Martins, Nereida Rodríguez-Fernández, and Sérgio M. Rebelo (Eds.). Springer International Publishing, Cham, 19–35.

[5] Boris Beizer. 1990. *Software Testing Techniques* (2nd ed.). Van Nostrand Reinhold.

[6] Oded Ben-Tal, Matthew Harris, and Bob Sturm. 2020. How Music AI Is Useful: Engagements with Composers, Performers and Audiences. *Leonardo* 54 (09 2020), 1–13. https://doi.org/10.1162/leon_a_01959

[7] Steve Benford, Chris Greenhalgh, Andy Crabtree, Martin Flintham, Brendan Walker, Joe Marshall, Boriana Koleva, Stefan Rennick Egglestone, Gabriella Giannachi, Matt Adams, Nick Tandavanitj, and Ju Row Farr. 2013. Performance-Led Research in the Wild. *ACM Trans. Comput.-Hum. Interact.* 20, 3, Article 14 (jul 2013), 22 pages. https://doi.org/10.1145/2491500.2491502

[8] Renaud Bougueng Tchemeube, Jeffrey John Ens, and Philippe Pasquier. 2022. Calliope: A Co-Creative Interface for Multi-Track Music Generation. In *Creativity and Cognition* (Venice, Italy) *(C&C '22)*. Association for Computing Machinery, New York, NY, USA, 608–611. https://doi.org/10.1145/3527927.3535200

[9] Dom Brown, Chris Nash, and Tom Mitchell. 2017. A User Experience Review of Music Interaction Evaluations. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Aalborg University Copenhagen, Copenhagen, Denmark, 370–375. https://doi.org/10.5281/zenodo.1176286

[10] Nick Bryan-Kinns. 2024. Reflections on Explainable AI for the Arts (XAIxArts). *Interactions* 31, 1 (jan 2024), 43–47. https://doi.org/10.1145/3636457

[11] Nick Bryan-Kinns, Berker Banar, Corey Ford, Courtney N. Reed, Yixiao Zhang, Simon Colton, and Jack Armitage. 2021. Exploring XAI for the Arts: Explaining Latent Space in Generative Music. In *Proceedings of the 1st Workshop on eXplainable AI Approaches for Debugging and Diagnosis (XAI4Debugging@NeurIPS2021)*. https://xai4debugging.github.io/files/papers/exploring_xai_for_the_arts_exp.pdf

[12] Nick Bryan-Kinns, Berker Banar, Corey Ford, Courtney N. Reed, Yixiao Zhang, Simon Colton, and Jack Armitage. 2021. Exploring XAI for the Arts: Explaining Latent Space in Generative Music. In *1st Workshop on eXplainable AI approaches for debugging and diagnosis (XAI4Debugging@NeurIPS2021)*. https://xai4debugging.github.io/files/papers/exploring_xai_for_the_arts_exp.pdf

[13] Nick Bryan-Kinns, Corey Ford, Alan Chamberlain, Steven David Benford, Helen Kennedy, Zijin Li, Wu Qiong, Gus G. Xia, and Jeba Rezwana. 2023. Explainable AI for the Arts: XAIxArts. In *Proceedings of the 15th Conference on Creativity and Cognition* (Virtual Event, USA) *(C&C '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3591196.3593517

[14] Linda Candy. 2006. Practice based research: A guide. *CCS report* 1, 2 (2006), 1–19.

[15] L. Candy, S. Amitani, and Z. Bilda. 2006. Practice-led strategies for interactive art research. *CoDesign* 2, 4 (2006), 209–223. https://doi.org/10.1080/15710880601007994

[16] Baptiste Caramiaux and Sarah Fdili Alaoui. 2022. "Explorers of Unknown Planets": Practices and Politics of Artificial Intelligence in Visual Arts. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 477 (nov 2022), 24 pages. https://doi.org/10.1145/3555578

[17] Filippo Carnovalini and Antonio Rodà. 2020. Computational Creativity and Music Generation Systems: An Introduction to the State of the Art. *Frontiers in Artificial Intelligence* 3 (2020), 14. https://doi.org/10.3389/frai.2020.00014

[18] Minsuk Chang, Stefania Druga, Alexander J. Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie J Cai, and Michael Terry. 2023. The Prompt Artists. In *Proceedings of the 15th Conference on Creativity and Cognition* (Virtual Event, USA) *(C&C '23)*. Association for Computing Machinery, New York, NY, USA, 75–87. https://doi.org/10.1145/3591196.3593515

[19] Stroud Cornock and Ernest Edmonds. 1973. The Creative Process Where the Artist Is Amplified or Superseded by the Computer. *Leonardo* 6, 1 (1973), 11–16. http://www.jstor.org/stable/1572419

[20] Monica Dinculescu, Jesse Engel, and Adam Roberts (Eds.). 2019. *MidiMe: Personalizing a MusicVAE model with user data.*

[21] Constance Douwes, Philippe Esling, and Jean-Pierre Briot. 2021. Energy Consumption of Deep Generative Audio Models. arXiv:2107.02621 [cs.LG]

[22] J. P. Eckert. 1946. Continious variable input and output devices. In *The Moore School Lectures*, M. Campbell-Kelly and M. Williams (Eds.). The MIT Press, Cambridge, Massachusetts, 393–423.

[23] Carolyn Ellis, Tony E. Adams, and Arthur P. Bochner. 2011. Autoethnography: An Overview. *Historical Social Research / Historische Sozialforschung* 36, 4 (138) (2011), 273–290. http://www.jstor.org/stable/23032294

[24] Sarah Fdili Alaoui. 2023. *Dance-Led Research.* Ph. D. Dissertation. Université Paris Saclay (COMUE).

[25] Corey Ford and Nick Bryan-Kinns. 2022. Speculating on Reflection and People's Music Co-Creation with AI. In *Generative AI and HCI Workshop at CHI 2022* (Online). 6.

[26] Matthew Garvin, Ron Eglash, Kwame Porter Robinson, Lionel Robert, Mark Guzdial, and Audrey Bennett. [n. d.]. Counter-Hegemonic AI: The Role of Artisanal Identity in the Design of Automation for a Liberated Economy. In *AI and the Future of Creative Work*. Routledge.

[27] Clifford Geertz. 2008. "Thick Description: Toward an Interpretive Theory of Culture". In *The Cultural Geography Reader*. Routledge. Num Pages: 11.

[28] Gaëtan Hadjeres, François Pachet, and Frank Nielsen. 2017. DeepBach: a Steerable Model for Bach Chorales Generation. arXiv:1612.01010 [cs.AI]

[29] James Harley. 2002. The Electroacoustic Music of Iannis Xenakis. *Computer Music Journal* 26, 1 (2002), 33–57. http://www.jstor.org/stable/3681399

[30] Samuel J Hunt, Tom Mitchell, and Chris Nash. 2020. Composing computer generated music, an observational study using IGME: the Interactive Generative Music Environment. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 61–66. https://doi.org/10.5281/zenodo.4813222 ISSN: 2220-4806.

[31] Alexander Refsum Jensenius and Michael J. Lyons (Eds.). 2017. *A NIME Reader*. Current Research in Systematic Musicology, Vol. 3. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-47214-0

[32] Théo Jourdan and Baptiste Caramiaux. 2023. Machine Learning for Musical Expression: A Systematic Literature Review. In *New Interfaces for Musical Expression*.

[33] Giacomo Lepri, Andrew Mcpherson, and John M. Bowers. 2020. Useless, not Worthless: Absurd Making as Critical Practice. *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (2020). https://api.semanticscholar.org/CorpusID:218663663

[34] Makayla Lewis. 2023. AIxArtist: A First-Person Tale of Interacting with Artificial Intelligence to Escape Creative Block. In *The 1st International Workshop on Explainable AI for the Arts at ACM Creativity and Cognition 2023*. https://xaixarts.github.io/accepted-2023/Lewis-XAIxArts-2023-

Paper.pdf

[35] Makayla Lewis, Miriam Sturdee, Mafalda Gamboa, and Denise Lengyel. 2023. Doodle Away: An Autoethnographic Exploration of Doodling as a Strategy for Self-Control Strength in Online Spaces. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 414, 13 pages. https://doi.org/10.1145/3544549.3582747

[36] Jackson Loth, Pedro Sarmento, CJ Carr, Zack Zukowski, and Mathieu Barthet. 2023. ProgGP: From GuitarPro Tablature Neural Generation To Progressive Metal Production. arXiv:2307.05328 [cs.SD]

[37] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J. Cai. 2020. Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376739

[38] Andrés Lucero, Audrey Desjardins, Carman Neustaedter, Kristina Höök, Marc Hassenzahl, and Marta E. Cecchinato. 2019. A Sample of One: First-Person Research Methods in HCI. In *Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion* (San Diego, CA, USA) *(DIS '19 Companion)*. Association for Computing Machinery, New York, NY, USA, 385–388. https://doi.org/10.1145/3301019.3319996

[39] Lena Mamykina, Linda Candy, and Ernest Edmonds. 2002. Collaborative creativity. *Commun. ACM* 45, 10 (oct 2002), 96–99. https://doi.org/10.1145/570907.570940

[40] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano Rodriguez, Matthew Yee-King, and Mark d'Inverno. 2019. *In a Silent Way: Communication Between AI and Improvising Musicians Beyond Sound*. Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10.1145/3290605.3300268

[41] Glenn McGarry, Peter Tolmie, Steve Benford, Chris Greenhalgh, and Alan Chamberlain. 2017. "They're All Going out to Something Weird": Workflow, Legacy and Metadata in the Music Production Process. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) *(CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 995–1008. https://doi.org/10.1145/2998181.2998325

[42] Herbert Molderings. 2010. *Duchamp and the aesthetics of chance: art as experiment.* Columbia University Press.

[43] Tim Murray-Browne and Panagiotis Tigas. 2021. Emergent Interfaces: Vague, Complex, Bespoke and Embodied Interaction between Humans and Computers. *Applied Sciences* 11, 18 (2021). https://doi.org/10.3390/app11188531

[44] Tim Murray-Browne and Panagiotis Tigas. 2021. Latent Mappings: Generating Open-Ended Expressive Mappings Using Variational Autoencoders. *NIME 2021*. https://doi.org/10.21428/92fbeb44.9d4bcd4b https://nime.pubpub.org/pub/latent-mappings.

[45] Maarit Mäkelä. 2007. Knowing Through Making: The Role of the Artefact in Practice-led Research. *Knowledge, Technology & Policy* 20, 3 (Oct. 2007), 157–163. https://doi.org/10.1007/s12130-007-9028-2

[46] Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3173574.3174223

[47] François Pachet, Pierre Roy, Julian Moreira, and Mark d'Inverno. 2013. Reflexive Loopers for Solo Musical Improvisation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 2205–2208. https://doi.org/10.1145/2470654.2481303

[48] Ashis Pati and Alexander Lerch. 2019. Latent space regularization for explicit control of musical attributes. In *ICML Machine Learning for Music Discovery Workshop (ML4MD), Extended Abstract, Long Beach, CA, USA*.

[49] Ashis Pati and Alexander Lerch. 2020. Attribute-based Regularization of Latent Spaces for Variational Auto-Encoders. arXiv:2004.05485 [cs.LG]

[50] Ashis Pati, Alexander Lerch, and Gaëtan Hadjeres. 2019. Learning to traverse latent spaces for musical score inpainting. *arXiv preprint arXiv:1907.01164* (2019).

[51] Ivan Poupyrev, Michael J. Lyons, Sidney Fels, and Tina Blaine (Bean). 2001. New Interfaces for Musical Expression. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, Washington) *(CHI EA '01)*. Association for Computing Machinery, New York, NY, USA, 491–492. https://doi.org/10.1145/634067.634348

[52] P. J. Charles Reimer and Marcelo M. Wanderley. 2021. Embracing Less Common Evaluation Strategies for Studying User Experience in NIME. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Shanghai, China, Article 12. https://doi.org/10.21428/92fbeb44.807a000f

[53] Paola Ricaurte. [n. d.]. Ethics for the Majority World: AI and the Question of Violence at Scale. 44, 4 ([n. d.]), 726–745. https://doi.org/10.1177/01634437221099612

[54] Adam Roberts, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. 2018. A Hierarchical Latent Vector Model for Learning Long-Term Structure in Music. (2018). https://doi.org/10.48550/ARXIV.1803.05428 Publisher: arXiv.

[55] Adam Roberts, Claire Kayacik, Curtis Hawthorne, Douglas Eck, Jesse Engel, Monica Dinculescu, and Signe Nørly. [n. d.]. Magenta Studio: Augmenting Creativity with Deep Learning in Ableton Live. In *Proceedings of the International Workshop on Musical Metacreation (MUME)*.

[56] Jennifer G. Sheridan. 2014. *Digital Arts Entrepreneurship: Evaluating Performative Interaction*. Springer International Publishing, Cham, 243–259. https://doi.org/10.1007/978-3-319-04510-8_16

[57] Beth Singler. 2020. "Blessed by the algorithm": Theistic conceptions of artificial intelligence in online discourse. *AI and Society* 35 (2020), 945–955. https://doi.org/10.1007/s00146-020-00968-2

[58] Paul Steinbeck. 2022. *5 George Lewis, Voyager*. University of Chicago Press, Chicago, 109–120. https://doi.org/10.7208/chicago/9780226820439-007

[59] Bob Sturm. 2022. Generative AI helps one express things for which they may not have expressions (yet). In *Workshop on Generative AI and HCI at the CHI Conference on Human Factors in Computing Systems 2022*.

[60] Bob L. Sturm, J. F. Santos, O. Ben-Tal, and I. Korshunova. 2016. Music transcription modelling and composition using deep learning. *ArXiv* abs/1604.08723 (2016).

[61] Renaud Bougueng Tchemeube, Jeffrey Ens, Cale Plut, Philippe Pasquier, Maryam Safi, Yvan Grabit, and Jean-Baptiste Rolland. 2023. Evaluating Human-AI Interaction via Usability, User Experience and Acceptance Measures for MMM-C: A Creative AI System for Music Composition. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. ijcai.org, 5769–5778. https://doi.org/10.24963/ijcai.2023/640

[62] Eric Thul and Godfried T. Toussaint. 2008. On the relation between rhythm complexity measures and human rhythmic performance. In *Proceedings of the 2008 C3S2E conference (C3S2E '08)*. Association for Computing Machinery, New York, NY, USA, 199–204. https://doi.org/10.1145/1370256.1370289

[63] Godfried Toussaint. 2005. The Euclidean Algorithm Generates Traditional Musical Rhythms. In *Proceedings of BRIDGES: Mathematical Connections in Art, Music and Science*.

[64] Marcelo Mortensen Wanderley and Nicola Orio. 2002. Evaluation of Input Devices for Musical Expression: Borrowing Tools from HCI. *Computer Music Journal* 26, 3 (2002), 62–76. http://www.jstor.org/stable/3681979

[65] Anna Xambó. 2022. Virtual Agents in Live Coding: A Review of past, present and future directions. *e-Contact!* 21, 1 (2022).

[66] Li-Chia Yang and Alexander Lerch. 2020. On the evaluation of generative models in music. *Neural Computing and Applications* 32, 9 (May 2020), 4773–4784. https://doi.org/10.1007/s00521-018-3849-7