

# Opening the Design Space: Two Years of Performance with Intelligent Musical Instruments

Charles Patrick Martin  
charles.martin@anu.edu.au  
The Australian National University  
Canberra, ACT, Australia

## Abstract


Machine generation of symbolic music and digital audio are hot topics but there have been relatively few digital musical instruments that integrate generative AI. Present musical AI tools are not artist centred and do not support experimentation or integrating into musical instruments or practices. This work introduces an inexpensive generative AI instrument platform based on a single board computer that connects via MIDI to other musical devices. The platform uses artist-collected datasets with models trained on a regular computer. This paper asks what the design space of intelligent musical instruments might look like when accessible and portable AI systems are available for artistic exploration. I contribute five examples of instruments created and tested through a two-year first-person artistic research process. These show that (re)mapping can replace retraining for discovering AI interaction, that fast input interleaving is a new co-creative strategy, that small-data AI models can be a transportable design resource, and that cheap hardware can lower barriers to inclusion. This work could enable artists to explore new interaction and performance schemes with intelligent musical instruments.

## Keywords

generative AI, small data, human-AI interaction, intelligent musical instruments

## 1 Introduction

This paper explores how generative AI may be embedded within interactive musical systems for live performance through a series of case study instruments and performances where existing hardware and software instruments have been coupled with a Linux-based single board computer pre-loaded with generative AI software. A broad range of music technology research has applied AI techniques, and in particular deep learning, for tasks such as mapping, audio or sensor analysis and musical data generation [16]. While the analysis of gestures [50] remains a popular application, deep learning models now enable generation of audio [6], symbolic music [15], and interaction data [25]. Despite this wide range of research, relatively few intelligent musical instruments, that is, digital musical instruments (DMIs) integrating generative AI, are available to be used in ordinary musical practice. The goal of this research is to expand the space of intelligent musical instrument design by prototyping and performing with multiple new instruments, examining them through artistic research, leading to actionable design insights. The approach is to introduce a platform (IMPSY) [23] for creating intelligent instruments that is cheap, small and battery-powered, thus enabling



**Figure 1: The generative AI interactive music platform with a hardware synthesiser. This platform connects to other devices via MIDI, receiving signals from a human performer and send signals to control sound on the synthesiser. The software is pre-installed on a Raspberry Pi operating system image and runs on even the cheapest Raspberry Pi Zero 2 W (15USD).**

integration into a wider variety of musical setups and practices across a varied community of artists and experimenters. The new platform includes AI software for generating expressive musical signals from small artist-centred datasets, runs on inexpensive Raspberry Pi computers, and is configurable via a web interface. As an initial step towards identifying a wider variety of intelligent musical instrument designs, this paper reports a series of five instruments and performance experiences from a first-person perspective, illustrating different ways that this system can be incorporated into an artistic practice.

AI techniques have been applied in music throughout the history of computing [39]. Real-time interactive systems became more prominent from the 1990s such as Continuator [34], Voyager [20], GenJam [4] as well as tools such as Wekinator [12] for embedding machine learning models in new musical systems. New and more capable deep learning techniques have led to an expansion of interest in AI music making at CHI workshops [31], music industry white-papers [13], and the general media [30]. Generative AI is now applied in tasks such as sound design [18], symbolic music generation [15], and music production [10]. Its use in live musical performances is still only beginning to be explored in NIME and related communities through practice-based methodologies [e.g., 37, 43]. With many now framing generative AI as a threat to musicians' livelihoods [13] and fraught with ethical pitfalls [1], it is appropriate to introduce and study artist-centred approaches to creating music with AI.

This work introduces a small-data [48] approach to embedding AI within musical instruments where artists collect, curate, train, and deploy their own intelligent musical instruments. I define an intelligent musical instrument as an instrument where an AI system generates actions independently of a musician's actions. This definition would encompass Continuator [34] and Voyager [20] but exclude systems such as Wekinator [12] where AI is used for



This work is licensed under a Creative Commons Attribution 4.0 International License.

NIME '26, June 23–26, 2026, London, UK

© 2026 Copyright held by the owner/author(s).

mapping musicians' actions or interpreting sensor data but not generating independent actions. The platform introduced in this paper allows a generative AI algorithm to be incorporated within existing electronic music devices through standard MIDI communication. The AI algorithm is able to control not just notes but the significant variety of timbral parameters that are available in electronic instruments. Similarly, the AI platform can respond to the musician's actions on standard interfaces such as piano-style keyboards as well as other controls such as knobs, sliders, levers, or movement sensors that are typically part of electronic music devices (e.g., see Figure 1). The AI platform's interaction mappings are configured through a web interface which is also used to retrieve recorded data for re-training the system's machine learning model.

With this platform, intelligent musical instruments can be quickly prototyped and tested through musical performance. Incorporating generative AI within electronic instruments leverages musicians' existing significant skills and knowledge which could enable deeper co-creative designs. In this work, I describe the new AI music platform, showing that it is affordable and achievable for non-experts in AI to use. I document five intelligent musical instruments (the Intelligent Volca, MicroFreak, S-1, DAW, and Setup) and reflections on two years of performance practice with this system. The results contribute an expansion of the design space for intelligent musical instruments, showing that (re)mapping can replace retraining for discovering AI interaction, that fast interleaving between human and AI input is a new co-creative strategy, that small-data AI models can be a transportable design resource, and that cheap hardware can lower barriers to inclusion. More broadly, it has been argued that explorations with technology within artistic practices can contribute to HCI [14, 19, 44, 45]. This work uses artistic methodologies to contribute to the broader conversation on design principles and patterns for generative AI in HCI [51] by examining design possibilities that might encourage more widespread prototyping, hacking, and musicking with AI in musical performance.

## 2 Related Work

Jourdan and Caramiaux identified gaps in interactivity and practice [16] in AI-supported musical expression, particularly where deep learning is applied. They argued that as deep learning systems have become more complex there have been fewer examples of application in long-term musical practice and artists have less interaction with data collection and training phases. This contrasts with early AI music systems such as Voyager [20] that were developed over multiple decades and parameter-mapping approaches where long term practice has been studied [11]. Tahiröglü's AI-terity instrument [47], a custom physical interface for exploring AI sound generation while also acting autonomously to engage the performer, has undergone artist-centred evaluation. The experience of performing with AI-terity was later analysed to understand how the system prompted "unfamiliar musical expectations" [46]. This role mirrors some of Parkinson and Dunning's findings for roles that musical machines, including intelligent musical instruments, might have [35]. They argue that such instruments can "evolve", "collaborate", and "learn" (among other roles). Generative AI may be well placed to help enhance these roles within new musical instruments. Recent practice-based approaches to intelligent instrument study have articulated how small datasets can be curated and applied with care and sensitivity [43]. Generative AI algorithms have also been embedded

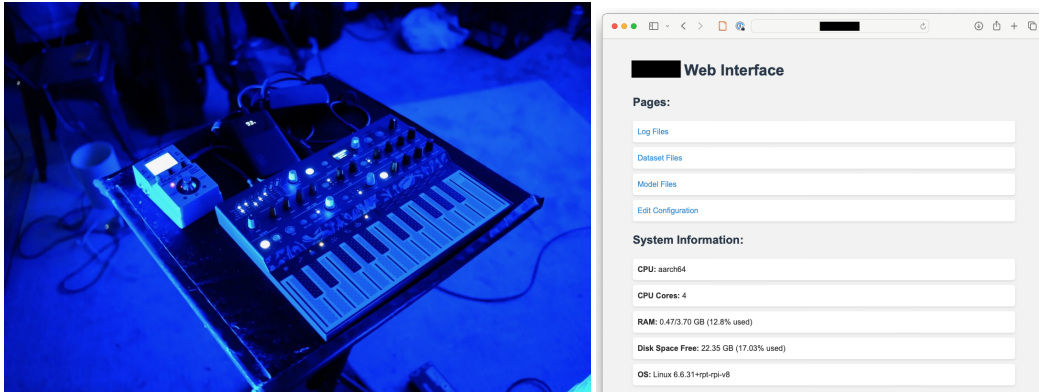
within software-based instruments. For example, Magenta Studio [40] seeks to embed generative AI algorithms within the Ableton Live digital audio workstation (DAW) software as a plugin that can generate and manipulate MIDI recordings. While these have been examined from an artist perspective [9], the focus is music production or composition not live performance.

Incorporating computer music systems into embedded hardware systems can allow better integration into artistic practices with many recent examples using Raspberry Pi [3] and Beaglebone [28] single-board computers. Research applying generative AI on such single board computers promises to allow these technologies to be embedded within self-contained musical instruments. This has been demonstrated for Raspberry Pi [24, 32] and Bela [36] platforms. Previous work has acknowledged the difficulties faced by instruments builders wishing to embed AI within new instruments in terms of data collection and training [25] and cross-compiling code [36]. These efforts are mirrored in other creative design fields [7]. Pelinski et al. [36] demonstrated a pipeline for collecting data and training machine learning models for musical instruments built using the Bela embedded platform [28]. Martin et al. [24] examined an embedded AI-based musical instrument and how different embodied representations of musical actions were perceived by improvisers. In the Sophtar, a custom string instrument with configurable interface components and embedded computer, Visi focuses on the evolving design process of a physical instrument with AI sound modelling and autonomous interactions [49]. In contrast with very large and costly industrial-scale AI, these projects have emphasised a small data [48] mindset where artists should collect, curate, train, and deploy their own AI systems.

This research explores generative AI instruments from a first-person or autoethnographic perspective [17] where instruments are created and tested through autobiographical design [33]. This is related to the practice-based or -led approach [42] and artistic research methods [e.g., 5] often applied within creative arts research where knowledge is distilled from an artistic process. Artistic processes have been previously applied as sources of HCI knowledge [45], for understanding intelligent instrument practice [37, 43], and for expanding instrument design spaces [5]. As the goal of this research is to expose design possibilities, it is appropriate to apply a first-person artistic perspective as an initial step with the ultimate goal of incorporating a wider range of perspectives through artist-centred methods.

## 3 System Design

The generative AI interactive music platform consists of Python software, a Raspberry Pi single board computer, and a custom operating system image that communicates to external electronic music devices over a USB, MIDI, or network connection. This setup is not intended to produce sound but only to control other electronic music devices and so the single board computer does not need to have an audio output. The platform is designed so that a musician can install the operating system image, which runs the AI software on boot, onto an SD card and then configure the platform from their computer using a web browser. After an initial setup the platform can be incorporated into an electronic music setup for performances and experiments without further configuration. If desired, data recorded to the SD card can be retrieved and used to train new AI models for the system. A diagram of the platform in use and the web interface is shown



**Figure 2: The stage setup for an intelligent musical instrument using the generative AI platform running on a Raspberry Pi 4 (left) and the web interface for configuration (right). The complete setup of synthesiser, Raspberry Pi, and effects pedal is battery powered. The software reacts to the musicians' actions through knobs and keys on the synthesiser front panel and controls notes and timbral parameters. Video examples are available here: <https://doi.org/10.5281/zenodo.19550146>**

in Figure 2 and a system diagram for a typical setup is shown in Figure 1.

### 3.1 Hardware

The platform uses the Raspberry Pi family of single board computers which are widely available and familiar for many electronic instrument designers. The software works on all Raspberry Pi models that support the 64-bit Raspberry Pi OS which ensures good compatibility with machine learning libraries for Python. This includes the inexpensive Raspberry Pi Zero 2 W (15USD as of 2026) that features a quad-core 64-bit ARM Cortex-A53 at 1GHz with 512MB of RAM and which is small enough (65mm x 30mm) to be incorporated inside new instrument designs.

Several methods of communicating with other devices from the Raspberry Pi are possible. MIDI is used as the primary communication format and the software can send MIDI messages via either a USB-connected MIDI interface, a direct USB MIDI connection to a hardware synthesiser, or the Raspberry Pi's serial (UART) output. The software can also communicate over a network using OSC (open sound control) or WebSockets messages. The smallest Raspberry Pi Zero models only support one USB host connection while larger models (e.g., 4 Model B or 5) have multiple USB host ports and an Ethernet interface so could support multiple communications channels simultaneously. For MIDI-over-serial connections, the MIDI connector can be soldered to the GPIO pads corresponding to the Pi's serial (UART) output. Providing a MIDI output from the Pi's 3.3V signals is straightforward requiring only two resistors [29]:  $10\Omega$  between the UART TX pin and pin 5 on the MIDI connector and  $33\Omega$  between the 3.3V output and MIDI pin 4. MIDI pin 2 is connected to ground. In the setup shown in Figure 4, the resistors are concealed within the MIDI connector.

### 3.2 AI model


Following Martin and Torresen [25], the new system uses a mixture density recurrent neural network (MDRNN) and the KerasMDN-Layer library [22]. This machine learning paradigm can generate a stream of musical data in free rhythm that models embodied musical gestures meeting the intention of controlling the parameters of electronic musical instruments in improvised

musical performances. The system uses an MDRNN that generates tuples of data that are interpreted as number of musical values and a time delta for when that value should occur in the future. The number of musical values is configurable within the software but is at a minimum one, i.e., the system models one musical parameter over time. The neural network is autoregressive, so takes the preceding value as input to generate the next, but it is also recurrent, so it stores a lossy history of generated values using LSTM units. In this research, small MDRNN models (e.g., 2 layers of 64 LSTM units) are generally used to model 1–8 parameters. Although these are very small compared to typical generative AI and large language models, they are sufficient for generating musically useful data in real-time on small hardware. The time for training a new model depends on the size of the training data and speed of the training computer, but effective models can be trained in under 30 minutes on a normal laptop.

### 3.3 Software

The system's software is a Python program running on the Raspberry Pi that generates musical data and sends it to external electronic musical instruments. This program listens to MIDI signals from the instrument and responds with AI-generated continuations of the performance state. MIDI note-on as well as control change messages are supported so the system is able to play notes as well as change the timbre of a connected instrument. As mentioned above, the AI model is capable of listening and responding to a number of channels of MIDI data simultaneously. The software records all data that is received from the MIDI interface and stores this as timestamped logs. This means that performing with this system builds up new datasets allowing new AI models to be trained.

An initial configuration is required to choose the MIDI interface and the specific MIDI signals to listen for and send out. Configuring interaction mappings is a design task that fundamentally changes how a created intelligent musical instrument will function. For example, it may be interesting to only assign timbral changes to the generative AI system while a performer controls notes (or vice-versa). Alternatively, a performer's notes can be cross-mapped to influence the timbral parameters changed by the generative AI platform, or the input and output of the AI model mapped to different devices. A simple web interface is provided for adjusting the configuration file, downloading logged



**Figure 3: Inference time for differently sized AI models on Raspberry Pis and an Apple MacBook Air (M1, 16GB). All Pis can run AI predictions in < 5ms.**

**Table 1: Boot times for different Raspberry Pi models measured from power-on to first MIDI output from the AI model.**

Raspberry Pi Model and RAM	Power-to-Sound (s)
Zero 2 W (512MB)	114
4 B (2GB)	78
5 (4GB)	38


data of previous performances and uploading new AI models. The configuration capability for the system evolved over the two-year time span considered in this paper, as described in Section 4.

The software for the system can be installed using Poetry or Python’s pip tool but this is a laborious process for new users. To accelerate setup with a new instrument, software is pre-installed on a Raspberry Pi OS image that can be flashed to an SD card. The generative AI interaction system and web server programs are started on boot and Ethernet-over-USB is enabled so that the Raspberry Pi can be directly connected to a computer for configuration. The software is open source [23] and available online<sup>1</sup>.

### 3.4 Benchmarks and speed testing

While quantitative testing is not a main focus of this paper, MDRNN inference speed testing previously reported in [25] was replicated to compare saved model formats and Raspberry Pi versions. The goal was to establish expectations for different system configurations in creative application compared with a normal laptop computer. Model formats compared were the Keras library’s native format (keras) and the optimised Tensorflow Lite format (tflite). The tflite format was faster on all platforms and even the cheapest Raspberry Pi Zero 2 W runs AI predictions in < 5ms (Figure 3), below the previous benchmark of 10ms [25]. For smaller models, the Raspberry Pi 5 runs AI predictions in less than 0.5ms. These tests revealed that while the Pi Zero 2 W had useful performance on smaller neural networks (considering number of LSTM units as a measure of size), it ran out of memory and did not complete tests on the larger 512-unit model. The boot-up time for different Raspberry Pis is another important aspect for integration in musical setups. The Zero 2 W was slow to boot at 114s while the Raspberry Pi 5 starts playing in 38s (Table 1).

<sup>1</sup>Source: <https://github.com/cmpmpercussion/impsy>, Raspberry Pi images: <https://github.com/cmpmpercussion/impsy-pi>



**Figure 4: Battery-powered human-AI interaction on a Korg Volca FM. The generative AI system running on the Raspberry Pi Zero controls pitch and rhythm via the MIDI connector while the musician edits the synth patch.**

## 4 Performance Experiences

This section describes the development of five new intelligent musical instruments using the generative AI platform and the experience of using them in performance over a period of approximately two years (2024-2026). These instruments were the Intelligent Volca, MicroFreak, S-1, DAW, and Setup. In this autobiographical artistic research, these instruments were developed for my own performance practice and tested in live gigs and recordings to expand knowledge about the design space for intelligent musical instruments. The performance style was free improvised music in solo, duo, and group situations. These performance experiences include solo recordings and concerts as well as applying intelligent instruments within ensembles of acoustic and electronic instruments. I introduce the instruments in roughly chronological order, reflecting on how they were used in performance and how the interaction mappings and creative approaches evolved over time. These reflections cover the 15 performances and recordings listed in Table 2. A demo video with each system is available as a supplementary media file with this paper and online<sup>2</sup>.

### 4.1 The Intelligent Volca

The Intelligent Volca was a proof-of-concept to establish whether real-time generative-AI MIDI signals from the Raspberry Pi Zero was feasible and interesting. MIDI output was obtained from the Raspberry Pi’s GPIO header UART pins and connected to the battery-powered Volca FM. The AI model was trained on a corpus of around one hour of musical human interaction with a single continuous controller. The AI model output (0.0–1.0) was converted to a MIDI pitch value from 0–127 and a corresponding MIDI note-on message scheduled to be sent at the time-delta generated, so the model could control pitch and rhythm on the Volca.

This initial experiment emphasised portability with a completely self-contained and battery-powered system; the Volca FM even has an internal speaker. The generative AI system played notes while I could adjust the synthesiser’s timbre or play overlapping notes on the keyboard; however, the AI model couldn’t respond to me as the Volca FM has no MIDI output. This meant that although this experiment started to explore different roles

<sup>2</sup>Video examples can be found here: <https://doi.org/10.5281/zenodo.19550146>

for a human and AI system in controlling the synth, the interaction was one way. The AI model, trained on expressive interaction with a continuous controller, tended to perform glissandi of various speeds, starting points and lengths on the Volca. These sounded somewhat unusual as the synthesiser re-triggered its envelope for each note. After recording a demo video with the system, I felt that the AI model might be more applicable controlling synthesis parameters where a musician would like a particular value to change smoothly in an expressive way.

## 4.2 The Intelligent MicroFreak and S-1

Hardware synths with MIDI outputs and inputs promised to support two-way human-AI interaction. The capacity for MIDI-over-USB connections was developed so the system could connect directly to synthesisers with built-in USB-MIDI interfaces where notes from their keyboard and control changes from their knobs and other interface elements can be sent and received over MIDI. In many cases, simultaneous control of the synthesiser from the hardware interface and from MIDI signals is possible allowing shared human-AI interaction. The interaction loop for the AI system responded in a call-and-response manner so that when the performer adjusted the synth controls and played notes, the AI system tracked these signals. When the human stopped for a certain amount of time the AI system took over the controls. My intention was to explore what music could be made by sharing control of the synths with this AI platform.

I first used an Arturia MicroFreak synthesiser with the generative AI system controlling eight parameters: note-on messages, and seven timbral parameters available to the performer via knobs on the synthesiser face. This instrument was used in individual and group improvisations with performers on a variety of other instruments. The instrument was portable (see Figure 5) with the tiny Raspberry Pi Zero and an effect pedal accompanying the small synthesiser. This AI model was trained on data sourced from my improvisation on eight continuous controllers. In performance, the AI model produced musically satisfying material. Unlike a human, the AI system is able to adjust many parameters simultaneously, resulting in an inhuman but exciting exploration of the synth parameter space. In non-AI performances, I would primarily play notes on the synthesiser with occasional timbral changes; however, the AI interactions update the timbral parameters in between phrases and in some cases in between notes, leading to varying and unique sounds. To enhance this effect, I tried setting the call-and-response switch-over time to a (very short) 0.1 seconds.

This performance concept was further explored with a Roland S-1 synthesiser where a similar mix of note and parameters signals were controlled by the AI system. The very small size of this synth compromises the playability of the keyboard but using the AI system led to less reliance on the keys. In performance, I focussed more on timbral parameter setting, allowing the AI system to improvise notes in between human-selected parameter changes which were also under the control of the AI model leading to surprising timbral outcomes.

## 4.3 The Intelligent Digital Audio Workstation

Much contemporary electronic music is performed on software synthesisers and audio files loaded into digital audio workstation (DAW) software [8, 38]. DAW software supports MIDI input and output to control parameters with internal routing and mapping possibilities. I experimented with connecting the system to a

DAW running on a computer or tablet via an inexpensive USB MIDI interface. This represents a convenient way to design new intelligent musical instruments and interactions as both the DAW and generative AI system are highly configurable.

An example is illustrated in Figure 6 where a DAW and software plugin host, Kymatica AUM (Audio Mixer), ran on an iPad with eight different sound generators available in channel strips with associated effects. The sound generators were a mix of software synthesisers and audio file players. The generative AI system was configured to respond to eight different MIDI control change channels and generate four different channels of MIDI notes and four channels of control change data. This input and output data is routed internally in AUM and assigned to control different parameters and notes within the sound generators. I could influence the AI generated data on the iPad touchscreen via a software MIDI controller or with a hardware MIDI controller connected to AUM with Bluetooth. This setup is very flexible in that different sound generators can be loaded or unloaded and the MIDI routing can be adjusted for different performance requirements without necessarily reconfiguring the generative AI software on the Raspberry Pi.

Mapping generative AI signals over a simple MIDI connection is an advantage for performing with tablets or laptops where installing a Python AI application is either impossible or inconvenient. I used this specific setup in solo and ensemble electronic music improvisations. The ability to re-map MIDI signals from the generative AI system to any parameter in the software synths was a way to evolve the intelligent instrument without retraining or modifying the generative AI system. I could map AI signals to areas where they were musically useful and swap software synths out with others or adjust them to sound significantly different. Having multiple synths loaded in a DAW allowed the output sound to be mixed and adjusted leading to quite different musical characteristics in between performances. While all parameters could be adjusted within the touchscreen interface, adding the hardware MIDI controller with physical knobs for influencing the generative AI system let me keep track of the autonomous components of the instrument.

## 4.4 Intelligent Setups

The above experiments and experiences focussed on a simple use-case where one synthesiser or computer was used for sound output as well as interface input; however, a natural extension is to support multiple interfaces simultaneously, each potentially capable of control input and sound output. In this kind of *intelligent setup*, the AI model mediates events in between the different interfaces. The intelligent DAW idea explained above could enable this kind of extension with MIDI routings within the DAW itself, but a better solution for a DAW-less setup was to extend the MIDI mapping capability of the AI system to enable multiple input and output routings.

This idea was originally motivated to support an additional controller for the Roland S-1 synth which has very small control knobs and no visible feedback for synth parameters. A first experiment with the intelligent setup combined the S-1 with a Behringer X-Touch Mini controller with 8 (larger) knobs with LED rings to provide feedback. The new mapping enabled the 8 knobs to control the AI model with human control events mapped directly through the S-1. Similarly, control on the S-1 was still available. MIDI outputs to control change could be restricted to a specific output range enabling control over the LED rings. This




Figure 5: Stage setups with an Arturia MicroFreak synthesiser (left) and Roland S-1 (right). The software ran on a Raspberry Pi Zero 2 W (left) or Raspberry Pi 4 (right) and in both cases a Zoom effects pedal provided additional audio effects. For both setups, the generative AI platform generated note data (tracking the keyboard) and seven timbral parameters (tracking knobs on the synthesiser interface).



Figure 6: Software synthesisers used in an intelligent instrument setup. Digital audio workstation (DAW) software ran on the iPad which is connected via MIDI to the generative AI system on a Raspberry Pi 4. A performer can interact with the instrument using the iPad touchscreen or a hardware controller interface.



Figure 7: An intelligent setup of the Roland S-1 and Keith McMillen QuNeo interfaces with a Raspberry Pi 5 as used in festival performances in January 2026

first setup worked well in performance, but the new flexibility of the MIDI mapping encouraged experimentation with interfaces beyond knobs. Substituting the X-Touch for a Keith McMillen QuNeo pad controller enabled control of and feedback from the 8 AI-linked parameters through touch sliders and additional note control with the drum pads.

In practice, the intelligent setup idea allowed more freedom in choosing how I could control an intelligent instrument and what a performance might look like. With the QuNeo and S-1, I could spend bigger chunks of performances playing (by, e.g., triggering lots of notes from the drum pads), effectively stopping AI control of the S-1. The QuNeo control strips made it easier to select reasonably specific values of multiple parameters quickly while the S-1 was more useful for slow sweeps through parameters. With this setup, availability of different control gestures led to different ways of interacting with and managing the unfolding AI process and more ways of understanding the AI model within a performance. Over three different performances, this particular setup started to feel like a more comfortable and expressive instrument with potential to sit within and drive different improvisational scenarios. It's worth noting that at this point, I was more likely to perform with a Raspberry Pi 4 or 5 rather than the Zero 2 W used in early prototypes. Both the boot and inference times

were faster on the full-size Raspberry Pis, a practical and musical consideration for performance.

The S-1/QuNeo intelligent setup was used in the final performance experience listed in Table 2. This was, notably, an improvisation festival where I used the same instrument in three performances with different groups of musicians including acoustic instruments, acoustic instruments with electronic effects and augmentations, and electronic musical instruments. The challenge here was to manage the, sometimes unruly, intelligent musical instrument in an environment where the *vibe* of a particular performance might be unknown in advance and range from sensitive sparse textures to a full-on wall of noise. With the visual feedback and more flexible controllers in the intelligent setup, this finally felt achievable.

## 5 Discussion

This work has described a two-year practice of intelligent instrument building and exploration in performances and recordings. The goal of this process was to expand the design space of intelligent musical instruments by developing and evolving many prototypes. In this section, I seek to distill the development and performance experiences into design implications relating to

**Table 2: Configuration of the performances, recordings, and demonstrations considered as part of the first-person artistic research process in this work throughout 2024–2026. This list does not include rehearsal or practice sessions that also took place. A "\*" in the configuration column denotes where more than one performer used an intelligent instrument studied in this paper, other duo and group performances involved one intelligent instrument (performed by the author) with other performers using acoustic and conventional electronic instruments.**

Date	Type	Configuration	Intelligent Instrument
February 2024	Recording	Solo	Volca
May 2024	Recording	Solo	MicroFreak
May 2024	Performance	Group*	MicroFreak, DAW (AUM, Live)
June 2024	Performance	Duo	MicroFreak
June 2024	Performance	Group	MicroFreak
September 2024	Performance	Solo	MicroFreak, S-1
September 2024	Performance	Duo	MicroFreak
October 2024	Performance	Group	S-1
December 2024	Recording	Duo*	DAW (AUM, Live)
January 2025	Recording	Solo	S-1, DAW
August 2025	Performance	Group	S-1, DAW
November 2025	Performance	Duo	Setup (S-1/xTouch)
November 2025	Performance	Duo	Setup (S-1/QuNeo)
December 2025	Recording	Duo	Setup (S-1/QuNeo)
January 2026	Performance	Group	Setup (S-1/QuNeo)

how generative AI models may be used in a productive way for intelligent instrument design.

### 5.1 Mappings and Flexible Configuration

A main finding in this research was that expanding the mapping capability of the generative AI platform, that is enabling changes in how the AI model is connected to physical controls and instrument parameters, paid dividends in terms of different instrument designs. This reflects on calls to expand interactivity when operating with ML models [16]. Connecting the generative AI outputs to simultaneously control multiple timbral parameters led to new musical potential in the intelligent MicroFreak. Adapting the platform to a DAW exposed a convenient way to play with how AI inputs and outputs could affect live performance. Supporting simultaneous IO from multiple devices in intelligent setups expanded the gestural affordances for managing and tracking AI behaviours.

Adjusting the mapping of control parameters to and from a generative AI model allowed me to find where a model could have the best musical impact within an intelligent instrument. The earlier prototypes demonstrated that the AI actions were more usefully focussed on what a human could not do (e.g., turn five knobs at once) rather than focussing on note production. Mapping the AI outputs to a visual indicator turned out to support better real-time understanding.

This research suggests that re-mapping may be an alternative to retraining where training might be expensive or inconvenient. Early small-data ML models [11] can be retrained within seconds, while the models in this project take minutes to hours, and audio generation models may take days [43]. Supporting remapping before retraining led to a faster and environmentally aware [26] prototyping process.

*Design Implications:* Flexible (re)mapping enables discovery of productive AI interactions. AI actions should be focussed on what a musician cannot do and visual feedback can be critical.

### 5.2 AI Models as a Transportable Instrument Component

The generative AI model in this work was trained entirely on self-produced data using standard computer systems, a clear contrast with the industrial-scale machine learning approach and its many ethical problems. This connects well with Vigiensoni's small data mindset [48] and challenges the assumption that AI necessarily requires unethically sourced data and unsustainable computational resources. A finding of this research was that one trained generative AI model turned out to be useful across a range of intelligent musical instrument prototypes. While other small-data music models have been published [e.g., 43], this paper demonstrates the utility that this approach can deliver. The trained model itself became a design component similar to the many interchangeable devices, modules and interfaces that electronic musicians use to assemble new interfaces. The artistic process revealed new uses for the same model and the mapping functionality extended its utility into applications quite disconnected from the originally collected data. Through the process of performance and practice I got to know this model better and found better ways to apply it. This speaks to sustainability [26]. A useful generative model can, like a trusted effects pedal or a filter module, be a transportable component within an electronic music practice.

*Design Implications:* Small-data generative AI should be applied in individual creative practices. Models can become a transportable design resource for intelligent musical instruments.

### 5.3 Co-Creative Performance

Throughout this research, the interaction strategy between AI and human input has used a typical call-and-response mode not dissimilar from previous instruments such as Continuator [34] or EMPI [24]. The idea of very fast switchover between human and AI control of an instrument in this research appears to be novel, contrasting with recent approaches using footswitches to alternate AI listening and generation [41], and to produce a unique interactive experience. Rather than sharing a musical

process with a separate agent, the intelligent instrument feels like a continually changing device, like a free-running oscillator, feedback system, or sequencer, that a performer can guide and adjust, but not fully control. The ability to instantly take over control through natural performance gestures helped address the challenge of guiding or *rescuing* [43] intelligent instruments away from unwanted sounds. The idea of shared control has been previously studied, but present approaches in LLM-based AI tend towards complete control of creation with only high-level management through prompting. This research demonstrates that new kinds of co-creative experiences are still possible, musically productive, and fun.

*Design Implication:* Very fast interleaving between AI and human inputs produces a collaborative, productive, and fun musical experience.

#### 5.4 Accessible Intelligent Instruments

An early motivation for this project was the discovery that the cheapest Raspberry Pi Zero 2 W (15USD) might be able to run small generative AI models, addressing cost issues as an accessibility barrier [21] for intelligent instrument design. The Intelligent Volca, MicroFreak, and S-1 confirmed that the Zero 2 W worked in different performance situations while revealing the limitations in terms of boot time when using this very low-end hardware. An advantage of the cheap Raspberry Pis was that I could set up multiple for use in group performances, loaning to collaborators, or dedicate to different instruments. The idea of dedicating single-board computers to different interfaces was established with Satellite CCRMA [2], but it's certainly easier to do this in practice with the cheapest hardware. Of course, I ended up using the more expensive Raspberry Pi 4 and 5 in my personal creative practice, prioritising speed.

While low costs help to make intelligent instruments more inclusive [16], a focus on the “new” raises questions of sustainability [26, 27]. Taking advantage of MIDI as the main communication standard in this project accelerated opportunities to prototype new designs but also sustainably reuse and retrofit existing instruments with a generative AI model. It was easier to prototype intelligent instruments by connecting existing hardware to an existing model in a new way, putting the focus on whether the interaction is musically satisfying. This approach built on previous use of OSC [25] but required more careful configuration approach to define mappings. It's notable that MIDI was particularly useful for connecting the system to a computer for the Intelligent DAW where existing MIDI mapping functionality in Live and AUM was available.

*Design implication:* Using cheap single board computers and standard MIDI lowers barriers to sustainable and inclusive prototyping.

#### 5.5 Limitations

This research has explored intelligent instrument performance from a first-person perspective. This approach is appropriate for understanding an exploratory design process over a long period of time (2 years). In future, it would be beneficial to extend this research to explore intelligent instrument designs from other creative practitioners. Reflecting on instruments from musicians of different backgrounds would likely lead to additional design implications. In this research, the potential to update the AI model over multiple practice and performance sessions was not

explored in a systematic way. Again, the evolution of a small-data AI model could be an additional aspect to the intelligent instrument design space.

## 6 Conclusions

This paper has presented a new generative AI platform for designing and experimenting with new intelligent musical instruments. I have argued that the design space for intelligent musical instruments using present AI techniques is still only beginning to be explored and that small-data approaches, such as this system, can lead to artist-centred musical AI. Experience using this platform in five prototype instruments over two years of live performances and recordings has shown that (re)mapping can replace retraining for discovering AI interaction, that fast input interleaving can be a co-creative strategy, that small-data AI models can be a transportable design resource, and that cheap hardware can lower barriers to inclusion. In the hardware synth examples, AI-control over notes and timbral parameters led to unique performances and sounds, making up for limitations in synth interface design. The intelligent DAW demonstrated the flexibility of the system when coupled with highly configurable DAW MIDI mappings, expanding the design possibilities for interactions with a DAW as an instrument. The intelligent setup allowed combinations of hardware interfaces to open up multiple co-creation styles with the AI system. This work contributes design possibilities explored in one creative practice with this platform through a first-person reflective process to the ongoing conversation within the computer music and broader HCI communities about how generative AI might be applied in interactive systems. Future work could examine the impact of updating models over time with artist-sourced data, and engage artists in co-design processes for creating new intelligent musical instruments.

## 7 Ethical Standards

This paper involved autobiographical design and study of new interfaces from the author's perspective. Where performances were in group settings, other group members were involved as creative collaborators rather than research participants and have been acknowledged below. The machine learning models used in this research were created with the author's own data.

## Acknowledgments

My thanks to collaborators and creative partners who made music with me as part of this project: Alec Hunter, Yichen Wang, Sandy Ma, Richard Johnson and the SoundOut Collective.

## References

- [1] Julia Barnett. 2023. The Ethical Implications of Generative Audio Models: A Systematic Literature Review. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montréal, QC, Canada) (AI/ES '23). Association for Computing Machinery, New York, NY, USA, 146–161. <https://doi.org/10.1145/3600211.3604686>
- [2] Edgar Berdahl and Wendy Ju. 2011. Satellite CCRMA: A Musical Interaction and Sound Synthesis Platform. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME '11)*, Alexander Refsum Jensenius, Anders Tveit, Rolf Inge Godøy, and Dan Overholt (Eds.). University of Oslo, Oslo, Norway, 173–178. [http://www.nime.org/proceedings/2011/nime2011\\_173.pdf](http://www.nime.org/proceedings/2011/nime2011_173.pdf)
- [3] Edgar Berdahl, Spencer Salazar, and Myles Borins. 2013. Embedded Networking and Hardware-Accelerated Graphics with Satellite CCRMA. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME '13)*, Woon Seung Yeo, Kyogu Lee, Alexander Sigman, Haru Ji, and Graham Wakefield (Eds.). KAIST, Daejeon, Republic of Korea, 325–330.
- [4] John A. Biles. 2007. Improvizing with Genetic Algorithms: GenJam. In *Evolutionary Computer Music*, Eduardo Reck Miranda and John A. Biles (Eds.).

- Springer London, London, 137–169. [https://doi.org/10.1007/978-1-84628-600-1\\_7](https://doi.org/10.1007/978-1-84628-600-1_7)
- [5] John M Bowers. 2023. A Hapless But Entertaining Roar: Developing a Room Feedback System through Artistic Research and Aesthetic Reflection. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Miguel Ortiz and Adnan Marquez-Borbon (Eds.). Mexico City, Mexico, Article 71, 10 pages. <https://doi.org/10.5281/zenodo.11189266>
  - [6] Antoine Caillon and Philippe Esling. 2021. RAVE: A variational autoencoder for fast and high-quality neural audio synthesis. *CoRR abs/2111.05011* (2021). arXiv:2111.05011
  - [7] Michelle Carney, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/3334480.3382839>
  - [8] Matthew Clauhs and Brian Dozoretz. 2023. The DAW Revolution. In *The Routledge Companion to Creativities in Music Education*, Clint Randles and Pamela Burnard (Eds.). Taylor & Francis Group. <https://doi.org/10.4324/9781003248194>
  - [9] Dylan Davis. 2022. Compute and Resonate: An Ongoing Experiment in Creating Acid Music Using Accessible Artificial Intelligence and Computer-Based Generative Tools. In *Designing Interactions for Music and Sound*, Michael Filimowicz (Ed.). Routledge.
  - [10] Emmanuel Deruty, Maarten Grachten, Stefan Lattner, Javier Nistal, and Cyran Aouameur. 2022. On the Development and Practice of AI Technology for Contemporary Popular Music Production. *Transactions of the International Society for Music Information Retrieval* 5, 1 (Feb. 2022), 35. <https://transactions.ismir.net/article/10.5334/tismir.100/>
  - [11] Rebecca Fiebrink and Laetitia Sonami. 2020. Reflections on Eight Years of Instrument Creation with Machine Learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 237–242. <https://doi.org/10.5281/zenodo.4813334>
  - [12] Rebecca Fiebrink, Dan Trueman, and Perry R Cook. 2009. A Meta-Instrument for Interactive, On-the-fly Machine Learning. In *Proc. NIME*. 280–285.
  - [13] Goldmedia. 2024. *AI and Music: Market Development of AI in the Music Sector and Impact on Music Creators in Australia and New Zealand*. Technical Report. <https://www.apraamcos.com.au/about/supporting-the-industry/research-papers/aiandmusic>
  - [14] Myoungsoon Jeon, Rebecca Fiebrink, Ernest A. Edmonds, and Damith Herath. 2019. From rituals to magic: Interactive art and HCI of the past, present, and future. *International Journal of Human-Computer Studies* 131 (2019), 108–119. <https://doi.org/10.1016/j.ijhcs.2019.06.005> 50 years of the International Journal of Human-Computer Studies. Reflections on the past, present and future of human-centred technologies.
  - [15] Shulei Ji, Xinyu Yang, and Jing Luo. 2024. A Survey on Deep Learning for Symbolic Music Generation: Representations, Algorithms, Evaluations, and Challenges. *Comput. Surveys* 56, 1 (2024), 1–39. <https://doi.org/10.1145/3597493>
  - [16] Théo Jourdan and Baptiste Caramiaux. 2023. Machine Learning for Musical Expression: A Systematic Literature Review. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 13 pages. [http://nime.org/proceedings/2023/nime2023\\_46.pdf](http://nime.org/proceedings/2023/nime2023_46.pdf)
  - [17] Annika Kaltenhauser, Evropi Stefanidi, and Johannes Schöning. 2024. Playing with Perspectives and Unveiling the Autoethnographic Kaleidoscope in HCI – A Literature Review of Autoethnographies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 819, 20 pages. <https://doi.org/10.1145/3613904.3642355>
  - [18] Purnima Kamath, Fabio Morreale, Priambudi Lintang Bagaskara, Yize Wei, and Suranga Nanayakkara. 2024. Sound Designer-Generative AI Interactions: Towards Designing Creative Support Tools for Professional Sound Designers. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 730, 17 pages. <https://doi.org/10.1145/3613904.3642040>
  - [19] Laewoo (Leo) Kang, Steven J. Jackson, and Phoebe Sengers. 2018. Intermodulation: Improvisation and Collaborative Art Practice for HCI. 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/10.1145/3173574.3173734>
  - [20] George E. Lewis. 2000. Too Many Notes: Computers, Complexity and Culture in “Voyager”. *Leonardo Music Journal* 10 (2000), 33–39. <https://doi.org/10.1162/096112100570585>
  - [21] Alex Lucas, Miguel Ortiz, and Franziska Schroeder. 2020. The Longevity of Bespoke, Accessible Music Technology: A Case for Community. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 243–248. <https://doi.org/10.5281/zenodo.4813338>
  - [22] Charles Martin. 2019. *cpmpercussion/keras-mdn-layer*. <https://doi.org/10.5281/zenodo.1482347>
  - [23] Charles Patrick Martin. 2026. *IMPSY: An Intelligent Musical Instrument Platform*. <https://doi.org/10.5281/zenodo.2580175>
  - [24] Charles Patrick Martin, Kyrre Glette, Tønnes Frostad Nygaard, and Jim Torresen. 2020. Understanding Musical Predictions With an Embodied Interface for Musical Machine Learning. *Frontiers in Artificial Intelligence* 3 (2020). <https://doi.org/10.3389/frai.2020.00006>
  - [25] Charles Patrick Martin and Jim Torresen. 2019. An Interactive Musical Prediction System with Mixture Density Recurrent Neural Networks. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 260–265. <https://doi.org/10.5281/zenodo.3672952>
  - [26] Raul Masu, Nicolò Merendino, Antonio Rodà, and Luca Turchet. 2024. Sustainable Internet of Musical Things: Strategies to Account for Environmental and Social Sustainability in Network-Based Interactive Music Systems. *IEEE Access* 12 (2024), 62818–62833. <https://doi.org/10.1109/ACCESS.2024.3393468>
  - [27] Raul Masu, Fabio Morreale, and Alexander Refsum Jensenius. 2023. The O in NIME: Reflecting on the Importance of Reusing and Repurposing Old Musical Instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Miguel Ortiz and Adnan Marquez-Borbon (Eds.). Mexico City, Mexico, Article 14, 10 pages. <https://doi.org/10.5281/zenodo.11189120>
  - [28] Andrew McPherson, Robert Jack, and Giulio Moro. 2016. Action-Sound Latency: Are Our Tools Fast Enough?. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Queensland Conservatorium Griffith University, Brisbane, Australia, 20–25. <https://doi.org/10.5281/zenodo.3964611>
  - [29] MIDI Manufacturers Association. 2014. MIDI 1.0 Electrical Specification Update. <https://www.midi.org/specifications-old/item/midi-din-electrical-specification>
  - [30] Blake Montgomery. 2023. Music publishers sue Amazon-backed AI company over song lyrics. *The Guardian* (oct 2023). <https://www.theguardian.com/technology/2023/oct/19/music-lawsuit-ai-song-lyrics-anthropic>
  - [31] Michael Müller, Lydia B Chilton, Anna Kantosalo, Charles Patrick Martin, and Greg Walsh. 2022. GenAICHI: Generative AI and HCI. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 110, 7 pages. <https://doi.org/10.1145/3491101.3503719>
  - [32] Torgrim Rudland Næss and Charles Patrick Martin. 2019. A Physical Intelligent Instrument using Recurrent Neural Networks. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 79–82. <https://doi.org/10.5281/zenodo.3672874>
  - [33] Carman Neustaedter and Phoebe Sengers. 2012. Autobiographical design in HCI research: designing and learning through use-it-yourself. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne, United Kingdom) (DIS '12). Association for Computing Machinery, New York, NY, USA, 514–523. <https://doi.org/10.1145/2317956.2318034>
  - [34] François Pachet. 2003. The Continuator: Musical Interaction With Style. *Journal of New Music Research* 32, 3 (2003), 333–341. <https://doi.org/10.1076/jnmr.32.3.333.16861>
  - [35] Adam Parkinson and Graham Dunning. 2025. Improvising with Machines: A taxonomy of musical interactions. *Organised Sound* (2025), 1–12. <https://doi.org/10.1017/S1355771824000268>
  - [36] Teresa Pelinski, Rodrigo Diaz, Adan L. Benito Temprano, and Andrew McPherson. 2023. Pipeline for recording datasets and running neural networks on the Bela embedded hardware platform. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 160–166. [http://nime.org/proceedings/2023/nime2023\\_22.pdf](http://nime.org/proceedings/2023/nime2023_22.pdf)
  - [37] Nicola Privato and Thor Magnusson. 2024. Querying the Ghost: AI Hauntography in NIME. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, S M Astrid Bin and Courtney N. Reed (Eds.). Utrecht, Netherlands, Article 63, 7 pages. <https://doi.org/10.5281/zenodo.13904901>
  - [38] Anders Reuter. 2022. Who let the DAWs Out? The Digital in a New Generation of the Digital Audio Workstation. *Popular Music and Society* 45, 2 (2022), 113–128. <https://doi.org/10.1080/03007766.2021.1972701>
  - [39] Curtis Roads. 1985. Research in music and artificial intelligence. *ACM Comput. Surv.* 17, 2 (June 1985), 163–190. <https://doi.org/10.1145/4468.4469>
  - [40] Adam Roberts, Jesse Engel, Yotam Mann, Jon Gillick, Claire Kayacik, Signe Nørly, Monica Dinulescu, Carey Radebaugh, Curtis Hawthorne, and Douglas Eck. 2019. Magenta Studio: Augmenting Creativity with Deep Learning in Ableton Live. In *Proceedings of the International Workshop on Musical Metacreation (MUME)*. [http://musicalmetacreation.org/buddydrive/file/mume\\_2019\\_paper\\_2/](http://musicalmetacreation.org/buddydrive/file/mume_2019_paper_2/)
  - [41] Victor Shepardson, Halla Steinunn Stefánsdóttir, and Thor Magnusson. 2025. Evolving the Living Looper: Artistic Research, Online Learning, and Tentacle Pendula. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Doga Cavdir and Florent Berthaut (Eds.). Canberra, Australia, Article 36, 6 pages. <https://doi.org/10.5281/zenodo.15698851>
  - [42] Hazel Smith and Roger T. Dean. 2009. Introduction: Practice-led Research, Research-led Practice – Towards the Iterative Cyclic Web. In *Practice-led Research, Research-led Practice in the Creative Arts*, Hazel Smith and Roger T. Dean (Eds.). Edinburgh University Press, 1–38.
  - [43] Halla Steinunn Stefánsdóttir and Thor Magnusson. 2025. Of altered instrumental relations: a practice-led inquiry into agency through musical performance with neural audio synthesis and violin. *Frontiers in Computer Science* Volume 7 - 2025 (2025). <https://doi.org/10.3389/fcomp.2025.1578595>
  - [44] Miriam Sturdee, Makayla Lewis, Mafalda Gamba, Thuong Hoang, John Miers, Ilja Šmorgun, Pranjal Jain, Angelika Strohmayer, Sarah Fdili Alaoui, and Christina R Wodtke. 2022. The State of the (CHI)Art. In *Extended Abstracts*

- of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 104, 6 pages. <https://doi.org/10.1145/3491101.3503722>
- [45] Miriam Sturdee, Makayla Lewis, Angelika Strohmayer, Katta Spiel, Nantia Koulidou, Sarah Fdili Alaoui, and Josh Urban Davis. 2021. A Plurality of Practices: Artistic Narratives in HCI Research. In *Proceedings of the 13th Conference on Creativity and Cognition* (Virtual Event, Italy) (C&C '21). Association for Computing Machinery, New York, NY, USA, Article 35, 14 pages. <https://doi.org/10.1145/3450741.3466771>
- [46] Koray Tahiroğlu. 2022. Evolving musical expectations: Mutual correlation between a human musician and an AI musical instrument. In *Embodied gestures*. Technische Universität Wien, 30–36. [https://doi.org/10.34727/2022/isbn.978-3-85448-047-1\\_4](https://doi.org/10.34727/2022/isbn.978-3-85448-047-1_4)
- [47] Koray Tahiroğlu, Miranda Kastemaa, and Oskar Koli. 2021. AI-terity 2.0: An Autonomous NIME Featuring GANSpaceSynth Deep Learning Model. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Shanghai, China, Article 80. <https://doi.org/10.21428/92fbeb44.3d0e9e12>
- [48] Gabriel Vigliensoni, Phoenix Perry, and Rebecca Fiebrink. 2022. A small-data mindset for generative AI creative work. In *Proceedings of the CHI 2022 Workshop on Generative AI and HCI*. <https://doi.org/10.5281/zenodo.7086327>
- [49] Federico Visi. 2024. The Sophtar: a networkable feedback string instrument with embedded machine learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, S M Astrid Bin and Courtney N. Reed (Eds.). Utrecht, Netherlands, Article 22, 7 pages. <https://doi.org/10.5281/zenodo.13904810>
- [50] Federico Ghelli Visi and Atau Tanaka. 2021. Interactive Machine Learning of Musical Gesture. In *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*. Springer, 771–798. [https://doi.org/10.1007/978-3-030-72116-9\\_27](https://doi.org/10.1007/978-3-030-72116-9_27)
- [51] Justin D. Weisz, Jessica He, Michael Muller, Gabriela Hofer, Rachel Miles, and Werner Geyer. 2024. Design Principles for Generative AI Applications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 378, 22 pages. <https://doi.org/10.1145/3613904.3642466>