Tuneable Machine Learning in Musical Instruments: A Duoethnography

Chris Kiefer
c.kiefer@sussex.ac.uk
Experimental Music Technologies Lab, University of Sussex
UK

Adam Staff
adamjstaff@proton.me
Independent Researcher
UK

Andrea Martelloni
a.martelloni@sussex.ac.uk
Experimental Music Technologies Lab, University of Sussex
UK

Figure 1: The MEMLNaut Channel

Abstract

Results are presented from a duoethnographic study of musical practice with instruments using embedded, tuneable, machine learning (ML); tuneable because the musicians can optimise machine learning models within the instrument, and embedded because these processes are built into the interface and processes of the instrument. Two researchers investigated their systems through their individual musical practices, and juxtaposed their experiences in a duoethnographic report. Both instruments are based on the MEMLNaut hardware and NISPS ML system, which offer embedded reinforcement learning for the design of mappings. The first system is an instrument for creative studio production, and the second is a signal processor embedded within a string feedback instrument. The design of these systems is explained, followed by interviews between the two researchers. Analysis of these interviews highlights key issues in musical practice with tunable ML, concerning the creative opportunities presented by new workflows, challenges in engaging with reinforcement learning processes, the design of biases in ML systems, and strategies that embrace or mitigate the approximate nature of creative ML.

Keywords

Machine Learning, Mapping, Design, Ethnography

1 Introduction

Machine Learning (ML) can offer musicians novel tools and novel modes of creativity, with transformative potential in how we make music [18]. With well-considered design and workflow, ML can offer musicians modes of engagement with computer music with the potential to open up design and customisation of digital and hybrid musical instruments to a broader audience. For example, a musician can optimise a signal processing algorithm towards their creative goals by recording a data set and training a model, rather than directly specifying the algorithm, moving towards a more musical, performative approach. In this paper, we address the compelling challenges in designing creative ML interfaces, embedded in musical instruments.

This project focuses on the design of ITMLs: Instruments with Tuneable Machine Learning. ITMLs use ML as

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

a fundamental part of their working, in producing sound and/or in the design of interaction with players. ML is *tuneable* in that the whole processes of ML (training, inference, control) are embedded, standalone, within the instrument and controllable by the player. This exploration of ITMLs responds to challenges with our current creative machine learning systems. Many of these offer inference only, they do not offer the musician the opportunity to train or tune the models. This can disenfranchise musicians, who are unable to customise models for their own creative needs, and forces reliance on potentially opaque data biases [2] or provenance. When systems do have trainable models, there is a potential reliance on expertise in ML theory and well-practiced technical skills in order to experiment with customisation. There are also issues in creative workflow, with disembodied splits between performative data collection and the technical processes of data curation and training. Fortunately, tools such as Wekinator [6] and libraries for computer music environments such as FluCoMa [10] provide bridges between data science style approaches and the tools computer musicians might be more familiar with. However, what if the processes of both training and inference were embedded within musical instruments? What if the instrument itself provided the interface to ML processes? What are the novel creative opportunities that arise? How might musicians interact with these processes? What does it mean for a musician to train their musical instrument during practice or even in live performance? These are the broad questions that the Musically Embodied Machine Learning (MEML) project has been engaging with.

1.1 Exploring ITMLs

Tunable ML, when embedded in musical instruments, has the potential to offer new creative opportunities, by enabling complex processes to exist within an instrument that might be customisable through simple or intuitive embodied interactions. There are relatively few instruments that have embedded, tunable machine learning. Examples would be Elia and Overholt's use of training as part of performance with the digitally controlled no-input analog mixer [5], or Visi's Sophtar [19] which offers embedded Reinforcement Learning (RL).

MEML has been exploring the design of ML workflows integrated into ITMLs over the past two years. The project began by exploring technology platforms for embedded machine learning and signal processing, and evolving this platform with feedback from workshops. We decided on the constraint of working entirely with low-power compute platforms that would, for example, fit into the modular synthesis ecosystem or run on battery packs within a musical instrument, portability and convenience being a key concern for developing instruments that would be usable by a wide range of musicians. Through the process of workshops, we built a software library for trainable machine learning, optimised for the Raspberry Pi RP23xx microcontroller. We adopted this particular platform primarily for accessibility through the Arduino ecosystem, and because of dual core processing, allowing ML and signal processing or synthesis to run with a high degree of independence on the same device. All firmware and hardware is available open-source[1].

We are currently running three longer term co-research collaborations, co-designing ITMLs with practicing musicians, using this technology as a foundation to prototype a variety of different instruments. It is one of these collaborations that we report on here.

Before focusing on the collaboration, it's necessary to set the context and motivation of the technology that this collaboration explores, and our path towards this point.

2 Machine Learning Approaches

The technology that supports this project has seen a transition from an initial focus on Interactive Machine Learning (IML) towards a new system NISPS (Neural Interactive Shaping of Parameter Spaces), inspired by Reinforcement Learning. Responding to the constraints of our hardware, our ML processes necessarily work with small datasets [3, 18], and relatively small deep neural networks, in order to effectively function as part of realtime interactive instruments within low memory and low power systems. As an example, our platform runs a five layer neural network with 6 inputs, 30 outputs and internal layers with 16 nodes, at around 100-150Hz. This scope enables the system to work with control rate data and processing, with an ML architecture that can work effectively. In terms of creative application, we have been exploring systems for deterministic (non-generative) multiparametric mapping. The systems take input from the musician (e.g. sensors, machine listening analysis) and project them into a new, usually higher dimensional space, mapping to sound control parameters. In this way, the neural network can become a core influence on the behaviour of an instrument; how it generates or processes sound, and how it responds to the musician.

In earlier stages of development, we explored workflows around IML [6], inspired by Fiebrink's *Wekinator*. The IML approach is characterised by iterative cycles of data collection or curation, optimisation of a model with this data, and musical evaluation of the model as it evolves. We built this workflow into an evolving hardware system, the MEMLNaut , which we used in co-research projects, and in workshops, to build a variety of experimental prototype instruments. These instruments enabled the musician to collect data points, train models and perform basic curation tasks (reseting memory, switching between datasets) using minimalist interfaces (switches, a joystick, potentiometers, and a small touch screen). Feedback from musicians showed that while this process was effective to some extent in designing mappings, there were also some challenges; when working with small interfaces, the data-centric approach of IML can become difficult to manage. Data curation can become challenging when the musician needs to keep track of the data they have recorded without the support of a more complex interface. The cyclic switches between data collection and model evaluation can also create challenges in fluidity of workflow. Therefore, we began to explore RL style approaches, asking if they have the potential to simplify how the musician interacts with ML, and if they

[1] https://github.com/MusicallyEmbodiedML

may be better suited to the minimal, embedded interfaces we are building in this project.

2.1 Reinforcement Learning

RL describes a family of methods for optimising the behaviour of an agent, through application of rewards and penalties based on actions taken within an environment [17]. There's a growing body of RL research in instrument design: Scorto et al's *Co-Explorer* [15, 16],Visi and Tanaka's *AIML* [20], and Smith and Freeman's *GestAlt* offer varied formulations for interactive design of mappings. Visi's *Sophtar*[19] and Pulz-Melbye's *Spectral Parrot* [12] demonstrate RL techniques embedded into physical instruments.

3 Design

NISPS evolved through testing in workshops, through co-research with musicians, and with reflections on creative practice from the researchers. The aim of this approach has been to progress the research by triangulating feedback from a variety of sources; fast experimental prototyping, in-depth collaborations, and reflections on everyday musical use.

As with the RL projects above, NISPS attempts to address the unusual formulation of RL when used for realtime design of mappings. Figure 2a shows a typical RL setup, where an agent produces actions in an environment, and receives feedback via an objective function. Through repeated feedback, the behaviour of the agent converges towards the desired function. In ITMLs, we have a different formulation (figure 2b), where the musician is providing feedback; the musician provides both the environment and feedback. They also share exploration with the agent. This tight connection in the RL system reflects how musicians have a strong embodied coupling with their instrument [8].

Figure 3 shows the NISPS workflow, and tables 1, 2, 3 and 4 describe the elements in the diagram.

When using NISPS, the musician designs mappings through ongoing exploration, and by giving rewards and penalties depending on preference for the current mappings. When the musician indicates a preference, this preference is for the current combination of input and output (e.g joystick position and resultant sound). In the background, the system optimises a neural network which takes input from the controls of the instrument, and whose outputs are mapped to signal processing or synthesis parameters. Further to this, NISPS offers creative low-level control of aspects of ML process, for example modulating the rate of learning, randomising the neural network or adding exploration noise to the output of the neural network. These additional controls are intended to allow the musician to explore different areas of the parameter space to find new mappings, and to facilitate varied strategies for optimising the neural network.

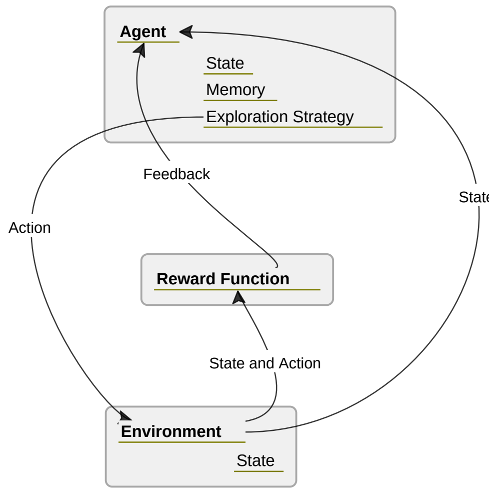

(a) A Typical RL Configuration

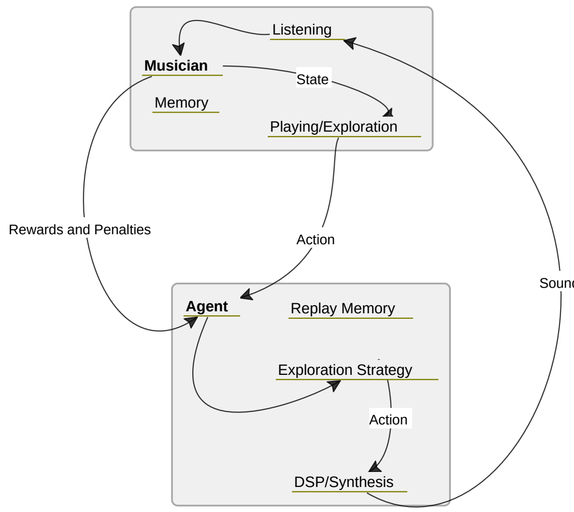

(b) RL Elements in NISPS

Figure 2: RL Elements

4 Study: A Duoethnography of NISPS In Practice

This study is an open-ended practice-based exploration of instruments using NISPS . We took a qualitative approach, that embraces the complex experiences and processes entangled in the use of musical instruments [13]. Two researchers worked with two different instruments, based on the same core machine learning technology, in order to build a more nuanced understanding of this technology in everyday musical contexts .

4.1 Duoethnography Method

Duoethnography explores a phenomonon through juxtaposition of the experiences of two individuals [14]. It could be seen as a layering and constrasting of autoethnographic accounts, that improves on autoethnography by reducing subjective bias and insularity [21]. Researchers work *"in tandem to untangle and disrupt meanings"* [4]. Meaning can

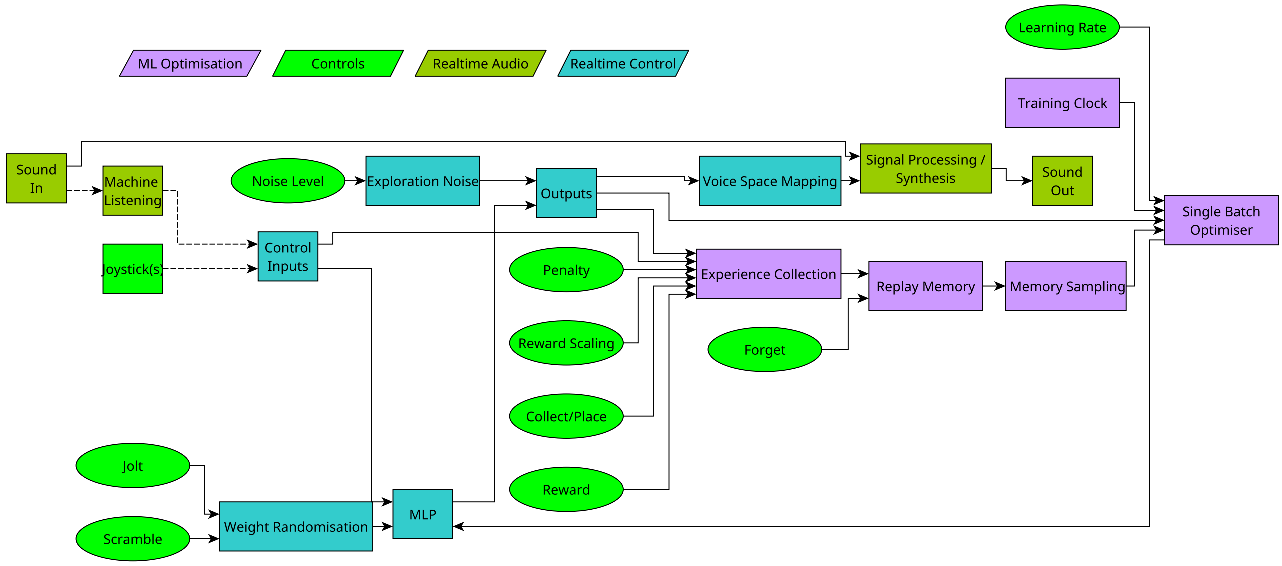

Figure 3: NISPS workflow

Table 1: Description of the RL system

Element	Function
Replay Memory	A collection of the most recent 64 input-output pairs, together with the reward value (from -1 to 1)
Training Clock	Triggers single batch optimisation at 100Hz
Memory Sampler	Collects a random batch of training items from replay memory
Single Batch Optimser	Optimises the neural network with a single batch of data

Table 2: Description of mapping system

Element	Function
MLP	A deep neural network
Control Inputs	Realtime input collected from joysticks or from machine listening analysis of audio input
Voice Space Mapping	Biases the mappings towards subsets on the possible mapping space, selected on the touch screen

be *"extended, deepened or transformed because participants build on each other's contribution"* [21].

A duoethnographic approach has useful potential in the study of experiences of musicians and instruments, particularly in the case, common in our field, of collaborations between researchers and non-academic musicians. It can dissolve the hierarchical labels of *researcher* and *research participant* [21], acknowledging a balanced contribution of individual domain-expertise into the research. Duoethnography can be viewed as a form of Dialogic Design [22], which accepts the researcher's influence, and also rejects tokenistic forms of participation.

Table 3: Description of button functions

Button	Function
Reward	Store the current input-output pair in the replay memory with a positive reward
Penalty	As above, with a negative reward
Collect/Place	Positively reward the current output, but with a different input position
Forget	Clear the replay memory
Scramble	Randomise the weights of the neural network (and therefore the mapping)
Jolt	Add a small amount of noise to the neural network weights

Table 4: Description of continuous controls

Control	Function
Learning Rate	Modulate the learning rate of the RL process
Reward Scale	Set the magnitude of future rewards and penalties
Exploration Noise	Progressively add Orstein-Uhlenbeck noise [11] to the outputs
Joystick(s)	Dual 2D Joysticks, optionally mapped to neural network inputs

Here, we document our collaboration in using and developing the MEMLNaut and NISPS in musical practice, by reflecting on our experiences. Our roles in the collaboration are as follows: Staff is a producer and studio engineer. Kiefer is a researcher-musician and co-designer of the MEMLNaut with Martelloni. Our collaboration began in

February 2025, when we began to discuss how the MEMLNaut could be adapted as a creative studio production system, an idea first raised in an early workshop. Through a series of firmware adaptations and hardware modifications, we built the MEMLNaut Channel, an instrument with similar DSP process to a studio mixer channel, with mapping design and control built with ML. The design explores creative studio practice and the 'studio as an instrument' philosophy [1]. The MEMLNaut Channel offers DSP blocks for pre and post gain, four band parametric equalisation and compression, with voice spaces that bias the mappings towards the configuration of various different consoles. For fine control, each DSP block can be bypassed. The control input is from two two-dimensional joysticks, which map through the neural network to 24 parameters in the DSP chain.

In parallel, Kiefer built a customised version of the MEMLNaut with audio-reactive effects for a feedback-string instrument, the Xiasri. The system was continually adapted throughout 2025, and played at three public performances, both solo and as part of ensembles. Both instruments use the same core NISPS technology, which has been evolving throughout this collaboration. Kiefer 's system is specialised to work as part of the Xiasri feedback loop; its control parameters are the outputs of machine listening analysis of the incoming audio signal, and the signal processing is a series of delays, filters and pitch shifting. It is controlled by a MIDI footpedal so that the player can interact with NISPS while bowing the instrument.

Staff and Kiefer can be viewed as both co-researchers and co-designers, within a wider network of stakeholders in this project. Our collaboration has been continually documented through a series of discussions and shared media on an online server.

This report is a reflection on our activities up to this point. In choosing how to bring these reflections together, we look to Fiebrink and Sonami's 2020 paper [7], which examined their collaboration through the format of written interviews. They emphasised the importance of the first-hand voice of participants through the interview responses.

Inspired by their method, our reflection process took place as follows: In January 2026, we met in a recording studio, with the intention of an in-depth discussion of our experiences. Both of our instruments were set up for support of the discussion through practical exploration or demonstration. Our discussion was recorded and transcribed. Following this, we independently read the transcription, and then co-edited a set of interview questions, with the guideline of prompting a response of around 1000 words. We independently wrote responses to these questions, and then met again to discuss the key themes in the responses, which led to the discussion and analysis below. The interviews and discussion now follow.

5 Interviews

The interviews are unedited, to foreground the voice of the respondents.

5.1 Staff

In one sentence, describe the MEMLNaut Channel . The MEMLNaut Channel is a maximum-velocity channel strip for rapid exploration of real-time sound objects.

How have you been using MEMLNaut ? From the idea of a MEMLNaut firmware for the recording studio we specified a DSP channel strip consisting of saturation, EQ, filtering, and compression. The MEMLNaut maps the many parameters of this DSP onto four dimensions of control on two XY joysticks. Reinforcement learning shapes this mapping based on the preference of the operator. We created voice spaces based on historic console designs, which delivered reasonable sounds at all but the most extreme settings. I used this system in production sessions and mixing sessions, processing takes while recording or printing recorded material through the device. Sometimes a fixed processing was decided on for an entire performance, sometimes adjustments to the processing were performed while recording. During this process, the output sound world was familiar to me and collaborators, but the journey from sound conception to its realisation was non-linear and explorative. My experience was of floating over a relief map, moving through biomes, weather patterns, times of day. Qualia like fidelity, brilliance, or slap become conurbations to swoop over, testing accent, cuisine, night life, until I find the perfect place to pitch a tent.

Describe the strongest similarities and differences between our two systems. Many differences came from contrasts in levels of chaos inherent in each system. With feedback instruments like the Xiasri, much of the play space involves discovering, expanding, manipulating and exploring modes of stability in an otherwise chaotic system. In contrast, my MEMLNaut was processing already stable sounds, working to shape them to the production without undermining their positive qualities. My challenge was to explore a large possibility space – none of which was broken or even especially bad – and find within it a relatively small ideal position. In both systems, the MEMLNaut extends existing tools, offering in its training and mapping a many-limbed assistant not quite smart enough to boss us around: Doctorow's ideal centaur. Both systems' processing is somewhat oxymoronic: Chris' MEMLNaut providing swooshy, smooshy creative effects in a system that otherwise tries to tame chaos; in mine, precise tools for correction are controlled chaotically, via the aforementioned mischievous assistant. Despite chaos, indirectness and assistance, the user is always the leading part of the feedback loop: rewarding the model, moving the controls, and always imposing an aesthetic judgement through listening, making decisions, and responding.

How does the using the MEMLNaut system differ to conventional studio practice? As a standalone, real-time device, the MEMLNaut superficially resembles a piece of studio outboard. My MEMLNaut is doing processing that would be familiar to users of mixing consoles, and by design is not producing manipulations that are uncomfortable or dangerous. However, through the reduction in the number of physical controls and the controls' mapping, precise and deliberate intervention, typical of an engineer's use of a mixing console, is rendered impractical or impossible. The

MEMLNaut tips the balance of power away from training in tools and toward a well-trained ear. To use with deliberacy, the operator needs a sharply focused conception of the ideal sound object, and the ability to hold this idea in mind while auditioning the sound world offered by the MEMLNaut. When the two shapes match, or a better alternative is offered, the operator's job is done. Crucially, unlike doing this work using conventional tools, the journey is not a straight line. Like a meteor, I spiralled toward my goal, occasionally thrown with velocity away, before coming back in again. This was the frustration and joy of using the MEMLNaut. Chris and I discussed a mode for the dual-joysticks where the one could provide fine-tuning – promising to offer the benefits of high velocity and high precision.

How have you used NISPS ? Have you settled on any particular strategy for training? Are there strategies that didn't work? In studio use, I found that the path towards the desired sound was indirect, and the size of the desirable sound area was small. I favoured strategies that made the mapping stable, still, quiet, slow. Trying to hit a moving target with a ricochet seemed too daunting a feat. I set a slow learning rate, and I used negative reinforcement 90% of the time, reducing the places that good sound could go, and the speed at which it could escape. When time was tight, or for supporting or background parts, I would forgo training altogether, randomising and exploring at high speed. Training - negative reinforcement again - could correct otherwise useful mappings that would top out in certain positions. Grabbing and repositioning sounds in the mapping space was especially useful for performative modes. Both me and clients found much utility in playing the MEMLNaut as an extension of the input instrument, and being able to choose the control position of areas of processing improved the gestural capacity of the instrument. Truly the studio as an instrument.

What's are the challenging aspects of working with the MEMLNaut ? In a producer/client setting, the nonlinear route to good sound was sometimes uncomfortable, as exploration through the MEMLNaut's sound space showed us many undesirable manipulations on our way to pleasurable ones. This was particularly acute for complex sources, before the initially high variation and velocity settled into detailed adjustments. My discomfort at being responsible for things sounding bad – even if we were tangentially moving towards good sound – surfaced many things. When I worked as a service provider, it challenged my pride in being able to control my tools. Being exposed to some of the MEMLNaut stack during prototyping, there were times when I thought to reach for lower-level solutions to my productions problem: an adjusted voice space, additions to or restructuring of the DSP chain. Such distractions should be resisted in the studio, where one can lose hours to a misjudged sidequest. Another puzzle was to find the right place for the MEMLNaut in production: I recall one early moment battling against the indirect control to set a precise corrective high-pass filter. Clearly, there are times that require a tool of specificity, and knowing when and when not to reach for the MEMLNaut will take as much experience as with any tool.

How has the system surprised you? Having initially seen the MEMLnaut operating in chaotic or open creative contexts, I expected the system to be a mismatch to, and a novelty within the controlled world of the studio. I instead discovered something with a great deal of utility, in no small part because of its contrast to existent studio tools. Such tools bring the unarguable benefits of multisensory feedback to the producer, but these are transformations of the sound that the majority of the audience can never experience. I found that rapid auditioning through gradients of randomisations, and the center-staging of listening showed where the production could improve, and where my aesthetic and aural skills could grow. The MEMLNaut nudged me always towards open listening: while I worked to bring its processing into focus, it resisted, loosening my strict idea of a correct processing, suggesting alternatives, offering a vibe in place of an idea.

In the future, how might the MEMLNaut fit into your studio workflow? With my current acuity on the instrument, high-velocity / low-precision tasks are ideal, even more so in low-tech environments. For example, smaller demo and writing sessions, smaller live performances, and proof-of-concept mixes. I have experienced no studio tool that comes close in the speed of control – just as long as errors in fine details are not so costly to the production. For some lofi aesthetics, this may never be a problem. There is also a pedagogical potential in foregrounding the ear training in music production, and could imagine developing workshops that used the device in contrast to low-velocity / high-precision tools.

5.2 Kiefer

How have you been using the MEMLNaut? I've been experimenting with the MEMLNaut as a part of a feedback string instrument, named the Xiasri. The MEMLNaut takes an audio signal as input, and performs machine listening analysis to create inputs for the neural network. The neural nework outputs are mapped to a set of effects that process the XIASRI signal, before feeding the processed sound to exciters in the instruments body. The instrument has a set of footswitches for controlling the RL processes, for *reward*, *penalty*, *scramble* and *forget*. This means that I can play the instrument to explore mappings, which also controlling NISPS to shape those mappings.

Could you describe the strongest similarities and differences between our two systems? Perhaps there's a difference in goals. In the Xiasri, the goal is very loosely defined; I have been openly exploring mappings with NISPS , without any particular idea of where the system may end up; although this idea seems to strengthen as the mappings are refined. With the studio system, you might have a stronger idea from the start as to what you need the MEMLNaut to do, but maybe the system opens up that space. With feedback, NISPS is making a complex process simpler, but perhaps in the studio we're complexifying aspects of a well-established traditional approach. Both of our approaches seem to value the mapping approach, the aesthetic value of engaging dynamically with a large DSP space, gesturally.

What was your experience of moving between modes of MEMLNaut design and use? It can be a real challenge mixing different modes or headspaces, sometimes getting lost in detailed technical work when I want to be playing, which can be frustrating. Sometimes technical work is necessary, but sometimes it's unnecessary feature creep, and I'm learning to accept the instrument 'as is' and not try to change it but to play it more instead. The more I follow this cycle, I find a deeper understanding of the idea of 'composing instruments', how musical and technical aspects of instrument development can be balanced and how they are intimately linked. Working with ITMLs creates an interesting zone between these modes of working, because the instrument is customisable with gesture as well as code and electronics.

How did playing alongside the MEMLNaut compare to previous iterations of Xiasri? It's been quite transformative, using dynamic effects that respond to the feedback loop in which they are embedded. This technique can create really lively responses which have all sorts of interesting areas and trajectories to play and explore. The challenge is how to tune such a system, and NISPS offers a useful solution because I'm able to (necessarily) play the instrument as part of the mapping design process; I can indicate preferences with a foot controller while bowing the strings. In older versions, I would go through a cycle of changing settings (via code or GUI) and then playtesting them. With NISPS , this process becomes a slowly evolving dialogue, it feels more musical compared to the old modify-evaluate approach.

How did the limitations of the MEMLNaut provide challenge and inspiration to your playing? A functional limitation of the MEMLNaut is the relatively small amount of DSP power. This means that the machine listening algorithms are quite simple, but they also very reactive and slightly quirky. The multi-effects are simple – for the Xiasri, it uses components of reverb algorithms – delay lines, allpass and comb filters, pitch shifting. These constraints have led be to focus more on the dynamic possibilities of these simple effects, rather than the design of the effects chain. In feedback instruments, you can get a lot 'for free' from simple processes as part of a more complex system, so this approach of 'under-engineering' the DSP chain fits well.

What is your experience of training? It feels like the training process has opened-up some possibilities and caused challenges in other areas. The process of moving from random places in the parameter space and shaping the mappings feels intuitive – especially with an instrument that has no 'right answer' to the mappings. It feels effective to slowly evolve and explore the mappings dynamically through playing the instrument and expressing preference, it's like co-exploring with the machine. It's quite novel to feel the mappings changing as you play them, and to explore them as you play, and slowly constrain them towards a setting that feels good to play. The challenge is fine-tuning; as the mapping becomes more refined, I start to develop a stronger idea of how I want the mapping to change, and this becomes difficult sometimes to guide the machine towards this because of the imprecise nature of the process. So either I need to accept this imprecise nature, or develop the training process so there are ways to fine tune the mapping. However, a fine-tuning approach would likely require more controls and new processes, which might counteract what I really enjoy about NISPS , which is the musical dialog with the machine and simplicity of shaping the mappings.

If we had the luxury to come back to you after 8 years of playing the Xiasri and the MEMLNaut, what outcomes do you expect or hope for? I would hope to have a more intuitive feeling for training the system, and to have also smoothed out or found a better balance in some of the tensions in training; approximation, precision, emergence and musical intention.

What does virtuosity look like on your instrument? I think potential for virtuosity lies in the depths of training process; as I play the MEMLNaut more, I'm continually finding new ways in how to engage with training. For example, the different strategies of expressing feedback, learning when to persist and pursue a particular path, or when to change and try something new; when to abandon and start again, or when to hang on to something. In this way, training is a process you can practice and improve at, and musical exploration and improvisation are part of that process, playing the instrument in a way which tests the edges of the mappings so you can choose where to optimise.

6 Discussion

After reading each others interview responses, we met to discuss them, and through discussion (we did not use formal analysis), we noted the following themes.

6.0.1 Speed vs Accuracy. The MEMLNaut system offers methods for exploring and shaping mappings in large parameter spaces at speeds that would be impossible compared to exploration with individual controls. However, the cost of this scope is in fine precision. This positions NISPS as a useful prototyping system, but improved strategies are needed to enable detailed tuning of mappings (*"from 75% to 99%"* (Staff)).

6.0.2 Novel Workflows. For both instruments, NISPS offers a significantly different workflow to previous practice. The MEMLNaut Channel offers a stark contrast to standard studio practice approaches, and dissolves habitual, trained practices. On the Xiasri, NISPS integrates mapping design into a continuous process instead of an edit-evaluate cycle. Both require a deeper focus on listening and performativity. This contrasts to more traditional approaches; in studio practice, this moves the value of mastery from skills in top-down, precision use of production tools towards skills in emergent listening and reacting. With the Xiasri, focus is moved towards playing and exploring the edges and depths of mappings, with the NISPS controls integrating into the wider dynamics of the instrument.

6.0.3 Approaches to NISPS . The authors developed differing strategies to interacting with RL. Staff tended to explore spaces quickly with the scramble function, and then start to pin down the mappings using negative feedback, also modulating the learning rate throughout this process. Kiefer used a mixture of scramling, along with positive and negative feedback, and would sometimes pause playing to slow down the learning rate when needing more detailed work on the mapping. These approaches demonstrate how NISPS has the flexibility to be adapted creatively in different contexts. It also shows how training strategy responded to varied goals; for Staff , the goal was to move swiftly towards an effective mapping; for Kiefer , training was an open-ended process, with the goal emerging through experimentation.

6.0.4 The 'Scenic Route'. Training workflow in NISPS is indirect, it might be more akin to a spiraling path towards a goal, sometimes moving in the perceived *wrong direction* before correcting course. Sometimes *good* settings pass by too quickly and would be difficult to return to. For both authors, this created challenges when required to make fast changes, when with a client in the studio or when adjusting mappings in a sound check. Staff discussed the need to *"to reach for lower-level solutions "* in addition to the instrument, and, reflecting on the interview, discussed how it can be challenging to fit NISPS into the 'service' role of the producer.

6.0.5 Mapping and Signal Processing Biases. The design of both systems involved considerable pre-planning in biasing the signal processing algorithms. Bias emerges not only from signal processing design, but from non-linear response of the neural network architecture, the way in which the network responds to weight randomisation, and the hardcoded mappings from neural network outputs to signal processing parameters. In the studio system, processing was biased to be similar to typical mixing desks, with voicespaces biasing further to emulate settings from certain consoles. The Xiasri design took inspiration from the building blocks of reverb units. Both designs constrained the effects towards *safe* spaces, which is important when NISPS might traverse many different areas of the possible space. The Xiasri was constrained to avoid unreactive saturating feedback and to encourage complexity. Staff reported that the MEMLNaut Channel always produced *safe* settings, but then the challenge was to move from *safe* to *good*. By defining these biases, we can raise questions about the balance of risk and safety in mapping design, and also acknowledge the significant role of offline pre-planning required to produce an effective realtime system.

6.0.6 Barriers to Entry. The NISPS style interactivity, with a focus moved from top-down precision towards emergent listening and reacting, may have the potential to break down barriers to entry. This was demonstrated when one of Staff 's clients was able to manipulate audio processing, without need for training with studio production tools.

7 Conclusions

The MEML project explores the creative opportunities arising from embedding ML within the design of musical instruments. We have presented the motivation and technology foundations of the product, and presented a duoethnographic report on musical practice with the MEMLNaut and NISPS. Analysis of this report contributes to wider conversations in a small but expanding area of research into embedded RL approaches in musical instrument design.

Before returning to the research questions, it should be noted that some of the responses in the report refer more to the nature of NISPS as a non-linear multiparametric mapping tool and some refer more towards the role and implications of ML in shaping these mappings. These two aspects are co-dependent, and ground the scope of the reported results here.

The two instruments probed quite different aspects of use, one which challenged traditional and convention in the setting of studio production, and the other which explored open-ended improvisatory musicianship with a feedback system. The juxtaposing reflections offered insights into the original research question about novel creative opportunities and interaction design for tunable ML. The results show how NISPS offers new workflows to musicians, offering fast but approximate methods for tuning mappings. It entangles mapping and performance within the instrument, which encourages the musician into an evolving dialogue with the RL system. A key challenge is the strong aesthetic of NISPS, which is unconventional and constrains the musician into a performative, reactive mode of working. There are however advantages to this mode, which has the potential to reduce dependence on specific skills training in order to work creatively with mapping design.

Focusing on questions of interaction design, the results emphasise the role of pre-biasing of mappings in machine learning systems, and the balance of risk and safety in exploration. They demonstrate a need to improved workflow in fine-tuning mappings compared to broader approximate shaping of parameter spaces, and show that better fine-tuning strategy could inspire more confidence from musicians.

The results question broader notions of how we define mastery in musical instruments, resonating with parallel conversations in other fields and notions of *unmastery* [e.g. 9]. This leads to questions about how we can develop strategies and perspectives that address the challenges of the inherently appromixate nature of creative machine learning.

8 Ethical Standards

This project has received ethics approval from the University of Sussex Social Sciences and Arts Research Ethics Committee, reference ER/CK84/5.

Acknowledgments

This project is funded by the UK Arts and Humanities Research Council, reference AH/Y004051/1.

References

[1] Adam Patrick Bell. 2018. *Dawn of the DAW: The studio as musical instrument.* Oxford University Press.

[2] Nick Bryan-Kinns, Berker Banar, Corey Ford, Courtney N Reed, Yixiao Zhang, and Jack Armitage. 2024. Explainable AI and Music. In *Artificial Intelligence for Art Creation and Understanding.* CRC Press, 1–29.

[3] Nick Bryan-Kinns, Anna Wszeborowska, Olga Sutskova, Elizabeth Wilson, Phoenix Perry, Rebecca Fiebrink, Gabriel Vigliensoni, Rikard Lindell, Andrei Coronel, and Nuno N Correia. 2025. Leveraging small datasets for ethical and responsible AI music making. In *Proceedings of the 20th International Audio Mostly Conference.* 70–81.

[4] Dawn Burleigh and Sarah Burm. 2022. Doing duoethnography: Addressing essential methodological questions. *International Journal of Qualitative Methods* 21 (2022), 16094069221140876.

[5] Gianluca Elia and Dan Overholt. 2024. Musicking with dynamical systems: introducing a digitally-controlled analog no-input mixer. In *Proceedings of the 19th International Audio Mostly Conference: Explorations in Sonic Cultures.* 22–36.

[6] Rebecca Fiebrink and Perry R Cook. 2010. The Wekinator: a system for real-time, interactive machine learning in music. In *Proceedings of The Eleventh International Society for Music Information Retrieval Conference (ISMIR 2010)(Utrecht)*, Vol. 3. 2–1.

[7] 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.

[8] Thor Magnusson. 2009. Of epistemic tools: Musical instruments as cognitive extensions. *Organised Sound* 14, 2 (2009), 168–176.

[9] Adam Pultz Melbye and Simon Waters. 2023. *A Continually Receding Horizon*. Ph. D. Dissertation. Queen's University Belfast.

[10] Ted Moore, James Bradbury, Pierre Alexandre Tremblay, and Owen Green. 2021. Making Machine Learning Musical: Reflections on a Year of Teaching FluCoMa. *Journal SEAMUS* 32, 1-2 (2021), 24.

[11] Johannes Nauta, Yara Khaluf, and Pieter Simoens. 2019. Using the Ornstein-Uhlenbeck Process for Random Exploration.. In *COMPLEXIS*. 59–66.

[12] Adam Pulz Melbye. 2026. Spectral Parrot. https://www.adampultz.com/spectralparrot.

[13] Matthew Rodger, Paul Stapleton, Maarten Van Walstijn, Miguel Ortiz, and Laurel Pardue. 2020. What makes a good musical instrument? a matter of processes, ecologies and specificities. In *Proceedings of the international conference on New Interfaces for Musical Expression, NIME 2020*. 484–490.

[14] Richard D. Sawyer and Joe Norris. 2012. *Duoethnography*. Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199757404.001.0001

[15] Hugo Scurto. 2022. Musicking Deep Reinforcement Learning. In *7th International Conference on Technologies for Music Notation and Representation (TENOR)*.

[16] Hugo Scurto, Bavo Van Kerrebroeck, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Designing deep reinforcement learning for human parameter exploration. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 1 (2021), 1–35.

[17] Ashish Kumar Shakya, Gopinatha Pillai, and Sohom Chakrabarty. 2023. Reinforcement learning algorithms: A brief survey. *Expert Systems with Applications* 231 (2023), 120495.

[18] Gabriel Vigliensoni and Rebecca Fiebrink. 2024. Data-and interaction-driven approaches for sustained musical practices with machine learning. *Journal of New Music Research* 53, 1-2 (2024), 19–32.

[19] Federico Visi. 2024. The Sophtar: a networkable feedback string instrument with embedded machine learning. In *International Conference on New Interfaces for Musical Expression (NIME), Utrecht, Netherlands, September 4-6, 2024*. International Conference on New Interfaces for Musical Expression, 142–148.

[20] 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.

[21] Amelia Walker and Corinna Di Niro. 2019. Creative duoethnography: a collaborative methodology for arts research. *Text* 23, Special 57 (2019), 1–17.

[22] Eevee Zayas-Garin and Andrew McPherson. 2022. Dialogic design of accessible digital musical instruments: Investigating performer experience. (2022).