The Living Looper: Rethinking the Musical Loop as a Machine Action-Perception Loop

Victor Shepardson
Intelligent Instruments Lab
Iceland University of Arts
Thverholt 11
Reykjavík, Iceland
victor@lhi.is

Thor Magnusson
Intelligent Instruments Lab
Iceland University of Arts
Thverholt 11
Reykjavík, Iceland
thor.magnusson@lhi.is

ABSTRACT

We describe the Living Looper, a real-time software system for prediction and continuation of audio signals in the format of a looping pedal. Each of several channels is activated by a footswitch and repeats or continues incoming audio using neural synthesis. The live looping pedal format is familiar to electric guitarists and electronic musicians, which helps the instrument to serve as a boundary object for musicians and technologists of different backgrounds to study the impact of machine learning on musical performance. Each Living Loop channel learns in the context of what the other channels are doing, including those which are momentarily controlled by human players. This leads to shifting networks of agency and control between players and Living Loops. In this paper we present the ongoing design of the Living Looper as well as preliminary encounters with musicians in a workshop and concert setting.

Author Keywords

Looper, neural synthesis, prediction, agency, machine learning

CCS Concepts

• **Applied computing** → **Sound and music computing;** Performing arts; • **Computing methodologies** → *Machine learning;*

1. INTRODUCTION

In this work, we consider that musical instruments are not merely instrumental. As Rodger et al. [39] point out, musical instruments cannot be understood simply as devices for achieving specific goals. Rather, instruments are ecological in that they offer different affordances to different players in different contexts. Waters uses the term "performance ecosystems" [46] to describe musical configurations which complicate the instrumentality of the instrument, the alterity of the environment and the agency of the performer.

NIME'23, 31 May2 June, 2023, Mexico City, Mexico.

We also note that tools have a kind of agency when they mediate social relations in the sense of actor-network theory [21]. This might be a straightforward delegation of human agency to material forces, the relational agency of culture and belief around the tool, or the unintended behavior of the tool as a natural object.

Biotechnologist Michael Levin argues for a cognitive gradualism implied by his practical work, where machines simulate life, life is engineered to build machines, and natural-artificial hybrids blur the line between the two. Levin suggests a less anthropocentric idea of agency based on degree of *persuadability*, and the *scale* of an agent's goals.

To persuade a rock, only blunt physical coercion will do, while a person may be persuaded by a subtle nuance of language. An acoustic musical instrument can be persuaded to sound only by mechanical actuation: blowing, scraping, plucking. A more agential instrument might be persuaded by the musical content of acoustic signals. The musician has agency on a human scale: he or she moves in spaces of musical signification, human social life, and somatic experience. An agential musical instrument might act in a space of acoustic signals, able only to bring about certain sonics. In this work, we are interested in agential instruments as a kind of microscope for machinic agency. Though far from the whole story, machinic agency cannot be neglected: it, too, is an actor in the network.

This paper is about a new musical interface called the Living Looper, the purpose of which is to amplify machine agency from the micro-world of signals and bits to the meso-world of sounds, where a musician can reach down from the macro-world of music to meet it. Since all that is horribly abstract, we have made our instrument superficially similar to an ordinary looping pedal, so that musicians can easily pick it up and experience its very different behavior. This allows our instrument to act as a boundary object [27] at which researchers and musicians of different backgrounds can meet.

A typical musical looping device is descended from vinyl records and magnetic tape, now often appearing as a digital foot pedal device, but inheriting the idea of storing some fixed-length input in a buffer as the underlying material of the loop. Now, machine learning algorithms move beyond memorization to generative *models* of data which extrapolate to new situations. What if a loop was not a buffer of recorded audio, but a model of that audio? The Living Looper asks, what would a live looping pedal for the age of machine learning be like?

Using a methodology of encounters [4], we place the instrument with musicians in various configurations, with attention to the differences between them. We report on the outcomes of some initial encounters, which generated feedback on the design of the Living Looper, insight into its role in performance ecologies, and descriptions of what it is like to play an agential instrument.

2. BACKGROUND

The Living Looper combines live looping pedals, cybernetic- and ecologically-inspired improvisation systems, and generative machine learning models. We briefly survey the histories and interconnections of these below.

2.1 Looping

Repetition is a fundamental aspect of music. Against a ground of manual repetition, repetitive audio technologies [22] have emerged many times. They have served many purposes and had periods of repetition across many timescales, from medieval hurdy-gurdys to player pianos to optical film sound. Pierre Schaeffer used the locked groove of phonographic records to facilitate reduced listening [20], later moving to loops of magnetic tape. Composers including Riley and Reich explored loops extensively in the 1960s, both in notation and tape. By the 1970s, Fripp and Eno were making ambient records with Oliveros and Riley's live tape looping system, Laurie Anderson was duetting with her 'self-playing violin' on street corners [3], and pioneers of turntablism and dub were looping tape in the studio and vinyl onstage. Digital sampling and delay effects became practical in the 1980s, and technology for live-sequenced reproduction of samples soon exceeded the basic loop, forming the basis of new electronic music genres. Streamlined live samplers or 'loopers' became available in the 1990s and were mass-produced by the early 2000s.

A live looping pedal can capture live input, play it back, and overdub further material. It can be controlled by the musician whose hands may be busy via a minimal footswitch interface, allowing close integration with e.g. electric guitar. Today 'live looping' is a popular genre with its own folk histories [17]. Contemporary live looping devices often include multiple channels of loops, quantization and synchronization of timing, variable speed and other effects.

2.1.1 Loopers as NIMEs

Loopers have appeared in NIME; Frisson et al. [16] survey a number of them. Many loop control signals or MIDI rather than audio, placing the looper in a 'mapping' layer between gesture and sound. Other looping NIMEs somehow re-imagine the control interface to a looper, which is unrelated to our approach; we leave the interface relatively intact while changing the function.

The Concentric Sampler [44] is a looper and granulator which explores the materiality of obsolete floppy disk technology. We explore machine learning algorithms in a similar way. In both cases, the outcome is not to extend or improve the looper, but to use looping as an interface for making an opaque technology sensible.

Intelligent Loopers such as Pachet's Reflexive Looper [34][28] or Wallace and Martin's PSCA system [45] replace loops with models at the meso-level of notes, chords and phrases. As its name suggests, the Reflexive Looper reflects back the style of the player. In contrast, our Living Looper is more concerned with the micro-level of texture and microsound [38]. Instead of detecting quantities like pitch or harmony, we use a learned process to encode audio to a continuous feature space in which micro-sonic details can be modeled. As its name suggests, the Living Looper makes sounds 'come alive' and wander away from the player.

2.2 Cybernetic music & predictive models

Frippertronics is often cited as the first live looping machine, but it was Riley and Oliveros' tape system that Eno and Fripp got started with. Pauline Oliveros's interest in long tape delays evolved to facilitate improvisation and deep listening [33]. In her music, loops move sound through time and space, creating new environments for improvisation and new relations between improvisers.

If the live looping pedal descends from the tape delay on the organological tree of life, feedback and cybernetic instruments are its cousins. There is no feedback without a loop: feedback instruments [41][13] could be seen as live looping on a micro-sonic time scale.

Audio feedback, whether acoustic, analog or digital, is just one sort of feedback loop. Many composers and musicians with an interest in cybernetics have noticed this, from Roland Kayn's networks of glacially shifting sound processes [36], to the "audible ecosystemics" of Agostino Di Scipio [11], and many others.

Invoking cybernetic concepts[40], Levin [23] places 'feedback' behaviors along a gradualist scale of cognition, being more cognitive than behaviors which don't depend on the past. More cognitive still are *predictive* behaviors, a special kind of feedback using the past to form expectations about the future and act on them. For Levin, cognition is everywhere in greater and lesser concentrations, from a particle, to a thermostat, to a cell, to a machine learning algorithm, to a person.

In computer music, prediction has been proposed as a unifying framework for *mapping* and *modeling* [29]. Notably, for machine learning, 'prediction' can mean inference of any unknown quantity, not only forecasting future conditions. From this perspective, 'interactive machine learning' tools such as Wekinator [14] apply prediction to mapping: 'what output goes with this input?' Meanwhile data-driven algorithmic composition systems are examples of modeling: 'what sound comes next?'

2.2.1 Neural Audio Synthesis

Deep learning-based generative models of control signals have been explored in recent NIMEs [18][30], but because of the high dimensionality of audio waveforms and the extreme subtlety of human audition, generative models of audio have been rare in the past. Sampling audio in real time has been possible only for simplistic generative models. However, new methods and infrastructure make it increasingly feasible to predict raw audio. This opens new opportunities for NIMEs: the distinction between mapping and modeling blurs as the richness of sonic interfaces [11] meets the machinic agency of learning algorithms.

With neural audio synthesis (NAS), probabilistic models built on deep neural networks [32] are used to generate audio. Compared to concatenative methods, NAS may produce more natural sounding results since there are no 'seams' between audio grains when sampling a generative model of the audio waveform. Compared to physical models, less knowledge of the acoustic phenomena is needed, and the inverse problem of inference from audio to parameters can be tackled by latent variable models which encode audio to a compressed or disentangled feature space.

As machine learning algorithms, software, and hardware evolve, it is becoming feasible to use high-fidelity NAS in real-time, with relatively low latency, on a laptop computer. NAS has begun to appear in the NIME literature [43][37] and real-time NAS is becoming a practical tool [8][25][1][10] in active use by musicians [9][19].

As deep learning algorithms proliferate, the quirks and failure modes of NAS also become more relevant. We suggest that incorporating novel technologies in a non-instrumental, idiosyncratic manner can contest solutionist narratives [31] around AI technology. Our intelligent looper doesn't solve any 'problem' which a dumb looper struggles with, but offers a different sort of experience to players.

3. THE LIVING LOOPER

The Living Looper resembles a live looping pedal in its basic form and function. In contrast to a standard looper, the Living Looper does not replay the recorded waveform. Instead, it attempts to fit a generative model which could produce the recorded sound, then uses it to continue the sound.

The Living Looper is a software instrument which transforms incoming audio and can be controlled via MIDI or OSC. For all prototypes reported on here, we paired it with electric guitars. Like a standard looper, players control our instrument by deciding when to record into it. In this work, it is configured with five parallel channels, or 'living loops'. The loops are controlled using five momentary switches, so that pressing a switch begins a recording on the corresponding channel, and releasing the switch completes the recording, at which point the loop begins to sound.

When a recording is completed, the Living Looper makes a model of it. Specifically, it maps the recent history of the sound, and also the sound of the other loops, to a predicted next few milliseconds. This mapping is chosen to minimize error with respect to what was sounding during the recording interval, with input from the player 'inhabiting' the targeted loop. Then the living loop continues the sound, using its new predictive model. Put differently, each living loop is an action-perception 'loop' which works to reproduce the musical 'loop'. By constantly reifying its predictions, it 'tries to keep things how they were' while the switch was pressed.

For a concrete example (Figure 1), consider just one living loop. I press the switch, play an arpeggio on the guitar, then release the switch. Now, the living loop constructs its model of what I played: a function is optimized to map each sequence of N microsounds (Section 3.1) in my recording to the very next microsound. Then the loop begins to use that function to produce the next instant of sound. It continues to predict from its past outputs, "photocopying a photocopy" as one player described it. Were it to perfectly fit the recording it would faithfully reproduce the sound, like a looper. But usually, like the photocopy machine it introduces some changes, and the loop mutates over time. This process tends toward a fixed point where the loop's function *does* perfectly match the sound it is producing – a sort of homeostasis.

When there are multiple loops in a Living Looper, they incorporate each other's recent history into their models (Section 3.3). A change in one loop can affect the others. As each loop falls toward a zero-error attractor, it is constantly perturbed by the other loops. Collectively, the living loops may settle into a low-surprise soundscape from which they are unlikely to deviate – until the player interferes. When the player records, their input replaces the targeted loop from the perspective of the other loops. So if I record into Loop 1, then Loop 2, 2 depends 1. If I begin to record again into Loop 1, Loop 2 is immediately affected. To 'erase' a loop, I tap its switch while muting the input (i.e., guitar strings), fitting its model to silence. Playing subtractively this way can have surprising effects on other loops. For example, if two loops have a similar structure, the second loop should rely on the first to make its predictions; when the first loop is silenced, the second is simplified.

3.1 Encoder-Decoder Model

Each living loop fits its predictive model in the instant after its switch is released. In such a short interval, it would be difficult to fit a detailed model for raw audio.

Instead, we used a pre-trained RAVE autoencoder [7] to map all inputs into a compressed latent space. The RAVE encoder converts audio input to streams of 'microsounds': feature vectors of $R = 16$ or so elements for each block of 2048 samples. Living loops fit their models to these streams, and the RAVE decoder takes their predictions back to audio. RAVE models can be run in real-time streaming mode on a laptop CPU [8].

3.2 Training Data

RAVE autoencoders must be pre-trained on a particular dataset, which determines the types of sounds a living loop can produce. Since the Living Looper does not rely on a pre-trained RAVE 'prior', only the encoder-decoder, our dataset focuses on timbre but not musical structure. We used a handmade dataset of electric guitar sounds including conventional playing techniques, but also various noises and preparations with objects like coins, paper, and knitting needles. About six hours of dry guitar sounds were recorded through noiseless magnetic pickups. When training, data augmentation was used to introduce some robustness to equalization, pickup position, gain, and saturation. Our guitar recordings are available online. [1]

A similar process could be followed to build a model for e.g. saxophone or no-input mixing board. Additionally, the Living Looper is compatible with pre-existing RAVE models (if trained with causal convolutions for low latency), which can be used creatively for timbre-transfer effects.

3.3 Living Loop Models

Each living loop has a model of the sonic 'environment' within the instrument, which is made up of the player and the other loops. A loop acts on its own part of that environment to minimize surprise. The environment is not passive; whatever action a loop takes influences its next perception. Thus the setting is a closed action action-perception loop [24]. However, in this work, living loops passively learn from each recording, after which point they are at the mercy of the environment. In future work, we hope to explore active learning for living loops. For the prototypes described in this paper, we use a simple linear autoregressive model for each one.

To fit a loop, targets are gathered by forming the last N input vectors into an $N \times C$ matrix Y, where R is the dimensionality of RAVE features. Features are gathered by concatenating recent inputs with the the other loops along the RAVE feature (R) dimension, windowing into contexts of length C, and then flattening the feature and window dimensions to produce an $N \times LRC$ matrix X, where L is the number of loops. The mean across the time N dimension is also subtracted from X to zero-center the features.

The $LRC \times C$ parameter matrix W and bias vector b are computed to be the least-squares solution to $Y = XW + b$. The affine transformation $xW + b$ then maps windows of features x onto single feature vectors, i.e. predicts the next micro-sound from the N previous micro-sounds across all channels. In RAVE terminology, each living loop is es-

[1] `https://github.com/Intelligent-Instruments-Lab/IILGuitarTimbre`

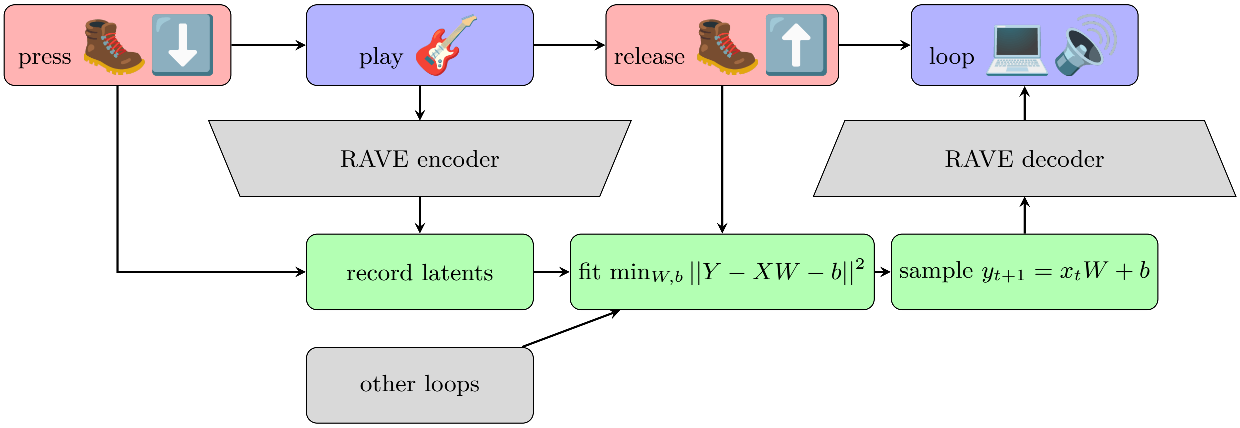

Figure 1: Creating a living loop (Section 3.3). Time flows left to right. Controls are red, sounds blue, living loop algorithm green.

sentially a (shallow, conditional) prior model fit to a tiny dataset made up of the looper state during the period the switch was pressed.

When iterated, the model described above does not tend to fall into a low-surprise attractor trajectory. Instead, the audio features tend to explode in directions presumably related to eigenvectors of the parameter matrix. To prevent this, we found it effective to post-process predictions with a sigmoid function (while pre-processing targets with the inverse to preserve the task):

$$S_p(x) = \begin{cases} 1 - 2\sqrt{-z} & x < -1 \\ 2\sqrt{z} - 1 & x > 1 \\ x & \text{Otherwise} \end{cases}$$

We also pre-processed features with $S_f(x) = \tanh(x/2)$.

3.4 Software

The core Living Looper is implemented using PyTorch [35]. We wrap the RAVE model and implement loop fitting as well as much of the control logic using TorchScript, a subset of Python which compiles to an intermediate representation interpreted by the PyTorch runtime. The PyTorch C++ API then interfaces the Living Looper with musical audio software. This allows core Living Looper features to be iterated upon rapidly in Python, while C++ wrappers need less attention.

We focused on a SuperCollider [2] UGen interface for connecting the Living Looper to control signals and further audio processing. The UGen, TorchScript model and our fork of RAVE can be found in our open source repository [2].

3.5 Hardware

We used a Keith McMillen Softstep 2 controller (Figure 2) to provide control switches. Two rows of five pads provide switches for each of five Living Loops. The software ran on an M1-based MacBook Pro.

We also experimented with a two-player setup (Figure 3). Players sat across from each other, one in front of the controller as normal and one behind it. Each player used one row of switches; pressing a switch on their own side would select their audio input, as well as activate the corresponding loop.

Guitar signals were kept as dry as possible to match the training data, and a DI box was used to run Living Loops

[2]https://github.com/victor-shepardson/living-looper

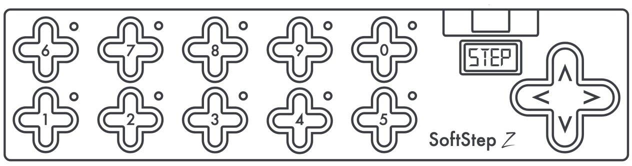

Figure 2: The MIDI foot controller used in our prototype (manufacturer's illustration). The ten pads were configured to send a message on press and release. The four-way pad and continuous control features were not used.

Figure 3: Two players using the Living Looper in duo mode during a concert.

into the same guitar amplifiers as the dry signal. When two loudspeakers or two guitar amplifiers were available, we spread the five Living Loops across the stereo field to help players separate them.

4. ENCOUNTERS

Because musical instruments are ecological in character, use by diverse players in various contexts will elicit different design specificities [39].

We use a methodology of *encounters* [4], a heterogeneous collection of events for players to meet the instrument with different goals in different environments. This opens the door to a *diffractive* [5][15][42] reading of the encounters, where the differences between configurations are attended to.

Through these encounters, we ask: does the Living Looper 'work well enough' that players remain engaged and imagine using the instrument later on? What specific improvements or design changes do the encounters suggest? Further, what unexpected new affordances appear when the instrument moves from a design setting to an musical one, and where do players perceive agency in the performance?

4.1 Workshop & concert

We hosted a two-day workshop for several local guitarists (P1-P6) to encounter the Living Looper, then perform with it in a concert. One additional musician, P7, used a live-coding environment instead of a guitar. Two other guitarists attended the workshop but did not perform in the concert. The primary researcher 'R' organized the workshop and also performed.

P3 and P7 were musicians and research colleagues in the NIME field. P1 was a designer from a different research group. P2 was a lifelong amateur musician, P5 was a musician and student, and P4 and P6 were professional guitarists. P4 and P5 feature live looping prominently in their musical practices. All except P6 had some contact with our research group before the event.

Musical styles ranged from improvisation to punk, metal, rock, live looping, and pop. Educational backgrounds were mixtures of jazz, classical and DIY. Participants described themselves variously as professionals, amateurs, composers, and improvisers.

On the first day, the workshop met in our research lab. R demonstrated the Living Looper prototype,then handed it over to workshop participants to encounter for the first time. Each participant took a turn while the others watched and commented. Later, P3 and P4 tried playing together as a duo sharing one Living Looper.

On the second day, the workshop met in a music school practice room. R explained a refined duo mode based on the previous day's experiment, and participants took turns playing in duos with all discussing.

P4 and P6 experimented with additional guitar effects placed before the looper. P4 and P5 also experimented with silencing the dry guitar signal.

A concert took place the next day at a local venue for experimental music.Two guitar amplifiers were used for sound.Compared to the two workshop days, the sound level was high, the lighting was low, and there was an audience. The extra effects and dry mix control from the workshop were made available to performers.

R first performed solo with the looper, then P3 joined in a duo. From there, the concert took an 'exquisite corpse' form, with one player at a time being replaced so a duo was usually playing. At one point, both performers left the Living Looper playing alone onstage.

4.2 Survey & Interviews

A few days later, participants were invited to do a survey on their experience and/or an in-person interview. P1, P5, P6 and P7 answered the survey, and P1 and P2 gave interviews.

Interviews were semi-structured following the survey, giving participants a chance to clarify or elaborate. Then the interviewee had a chance to play the Living Looper again while talking aloud, to elicit anything the questions may not have.

Interviews and talk-aloud sessions were audio-recorded and later transcribed. We roughly thematised excerpts from the survey and transcripts around our research questions.

4.2.1 *Usability*

All players reported that the instrument was fun to play, they were immersed while playing it, and did not lose interest.

Players felt the instrument was easy to get started with. P1 described it as fast to reach a playful interaction, but found a demonstration helpful: "after watching someone fill the channels you immediately get it." P7 wrote, "it very quickly gave interesting results no matter what it was thrown."

Some agreed that 'I needed more time to learn to play it well', while others were neutral, perhaps feeling that they had already reached the limits of control. P1 and P2 speculated whether mastery would be possible with practice.

Participants were generally enthusiastic about the instrument's unpredictability. P5 "liked not having complete control and being surprised". P2 even "wished I had less control, [...] wished it would morph more, do more crazy stuff."

Nevertheless, most reported feeling the instrument was sensitive to what they played and agreed that "its behavior made sense." P7 wrote, "though the overall pattern of behaviour was clear [...] I couldn't predict exactly how it would respond to the next input, which was satisfying and engaging." For P2, "it never did something really shocking [...] I didn't expect that, but it makes sense".

P1 and R noted that the evolving volume of loops could sometimes hide them behind louder loops. R concluded the stereo spread was not adequate for keeping track of multiple loops, and some visualisation would be helpful. P1 suggested that ambient lighting could indicate when the loop controls were in use without being distracting to the players or audience.

Participants also saw applications to their own practices. P2 suggested using the looper for buildups in death metal songs. P5 "would use it to create a soundworld that I could then work with (like layering things on top)." P7 wrote, "I can see it being extremely useful in texture-driven composition, ambient music and film scoring."

4.2.2 *Living Loop Behavior*

During the first workshop day, P4 described the Living Looper as making sense out of nonsense and vice-versa: it would develop noisy inputs into a drone or rhythm, with structured input eventually be destroyed by the same process. Several players pointed to the semi-predictable effects of adding and subtracting loops as the most interesting part of the instrument.

When asked their understanding of what the Living Looper did, P6 emphasized the unity of the loops: "I think it was taking whatever I put into it and morphing it together somehow."

In contrast, P5 thought of the loops as communicating individuals: "the different loops interact in a way and take inspiration from one another" and capable of novelty: "The looper suggests new musical material". P7 wrote that it was "commingling inputs, transforming the material in each slot as it went. The final loop often seemed to converge around the shortest, most prominent phrase/motif in the input." P2 noted a tendency for established loops to pull new loops into orbit.

P6 described it as "more like a morpher, morphing things together." P2 described also described the looper as "morphing" and the loops as "converging" or "melding", saying it was "feeding on what you're doing." For P1 it was "like an exchange with the machine. I really like how interaction is fluid."

When asked if the Living Looper was different when playing in a duo, most agreed. However, P1 found the duo and solo experiences similar in that "when you play the looper with someone, [...] it's interesting not knowing what's going to happen between you two, and I felt it was the same with the looper when we played alone with it." For P7, "the 'duo' actually felt more like a trio, two humans and one Living Looper." P5 wrote, "it reminds me of playing with another person." For P2, "you are definitely playing 'with it' " but "I think 'that's a machine'. I don't personify it really."

5. DISCUSSION & FUTURE WORK

The ergomimetic [26] design of our interface seemed to do its job, in the sense that players were able to start using it immediately, adapting quickly from looper to Living Looper. However the functionality of the Living Looper is drastically different, which users found surprising and exciting.

Players seemed to readily perceive the different agencies of the looper, individual loops, themselves, and other performers. Though each living loop is formally defined in software, that is only half the story; a loop's behavior is also constituted by the inputs to its predictive model, including streams of audio features produced by the other loops (Section 3.3). When the sound reaches an observer, those feature streams have been commingled into an auditory scene. Whether Living Loops go about their looping separately or fuse together into a single sound object is constantly being negotiated and never fully settled.

In Levin's [23] terms, they form a collective intelligence,[3] where multiple agents have reduced individuality and share a homeostatic goal: living loops can depend on others for their structure, and the population of loops seems to often settle into an shared attractor trajectory (Section 4.2.2). Meanwhile, the player is like one of Levin's larger agents, "warping the option space" of living loops to indirectly pursue their own musical ends.

In Karen Barad's [5] terms, living loops wouldn't interact but intra-act: they produce and change each other; they can't be laid down like ordinary loops and expected to stay put. One loop will subsume another, or do something unexpected which players have to react to. Players noticed this, and even identified it as the most compelling aspect of the instrument (Section 4.2.2). The gradual morphing-together of loops, and the unpredictable effects of replacing loops formed the basis for a performance practice.

Live looping is often a solo method, but our encounters saw the Living Looper placed into a performance ecology [46] of group improvisation, because the social environment of the workshop encouraged collaboration and the prototype couldn't yet be sent home with individuals. The Living Looper workshop brought together musicians who might not otherwise have played together, on the basis of their shared curiosity. By all accounts, this curiosity was sustained throughout the encounters, with participants eager to have more time exploring partial control over loop behavior. To make it accessible for longer-term engagement, the software should be more portable so musicians can use it in their environments of choice and have the opportunity to pursue the mastery some players speculated about (Section 4.2.1).

Though we focused on a Living Looper for electric guitar in this work, future versions of the Living Looper will include more instrument models, or perhaps 'universal' acoustic models building on neural audio codec research [12]. Using infra-instruments [6] could put more focus on the looper itself.

Players also suggested some visual elements to help distinguish the contents of different loops, and the states of the controller (Section 4.2.1). A future version of the instrument might explore bespoke controller and loudspeaker designs to address the reported difficulty separating loops. Embodying living loops in space by physically distributing their speakers and controllers could improve on the current stereo approach. If players had to walk between loops to control them, it might also encourage a more deliberate and listening-oriented style.

[3]Levin uses 'agency' and 'intelligence' somewhat interchangeably

6. CONCLUSION

We described the Living Looper, a new human-machine interface for musical expression. It draws inspiration from live looping pedals, which also ground it in familiarity. If a looper makes music from technologies which are precise, stable, and mechanistic, the Living Looper does so with learning algorithms which are approximate, malleable, and agential.

We elicited design feedback from musicians and began to explore our central research question of what agency looks like in a musical instrument. Players were enthusiastic, suggesting practical uses for it and contributing design suggestions. We observed a shifting distribution of agency between players, the Living Looper, and individual loops.

Future work will refine the instrument to enable more real-world use and more diverse and longitudinal user studies.

7. ACKNOWLEDGMENTS

This research is supported by the European Research Council (ERC) as part of the Intelligent Instruments project (INTENT), under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

INTENT is also supported by an NVIDIA hardware grant of two A5000 GPUs.

The workshop and concert were supported by a grant from the Icelandic government via the Rannís Music Fund.

8. ETHICAL STANDARDS

The Living Looper workshop was free to attend, and quasi-public in the sense that drop-ins were welcomed. Participants in the survey and interviews received an information sheet and signed a consent form. Participation in the workshop or concert bore no obligation to participate in the study. Musicians who did perform at the concert were paid for their labor out of both ticket sales and a music fund grant from the Icelandic government.

9. REFERENCES

[1] Neutone by qosmo. https://neutone.space/. Accessed: 2023-01-31.

[2] *The SuperCollider book*. MIT Press, Cambridge, Mass, 2011.

[3] L. Anderson. Self-playing violin. https://www.moma.org/collection/works/109919, 1974.

[4] J. Armitage, T. Magnusson, V. Shepardson, and H. Ulfarsson. The proto-langspil: Launching an icelandic nime research lab with the help of a marginalised instrument. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, The University of Auckland, New Zealand, Jun 2022.

[5] K. Barad. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. duke university Press, 2007.

[6] J. Bowers and P. Archer. Not hyper, not meta, not cyber but infra-instruments. In *Proceedings of the 2005 conference on New interfaces for musical expression*, page 5–10. Citeseer, 2005.

[7] A. Caillon and P. Esling. Rave: A variational autoencoder for fast and high-quality neural audio synthesis. *arXiv:2111.05011 [cs, eess]*, Nov 2021. arXiv: 2111.05011.

[8] A. Caillon and P. Esling. Streamable neural audio synthesis with non-causal convolutions. *arXiv:2204.07064 [cs, eess, stat]*, Apr 2022. arXiv: 2204.07064.

[9] C. J. Carr and Z. Zukowski. Generating albums with samplernn to imitate metal, rock, and punk bands. *arXiv:1811.06633 [cs, eess]*, Nov 2018. arXiv: 1811.06633.

[10] J. Chowdhury. Rtneural: Fast neural inferencing for real-time systems. *arXiv:2106.03037 [eess]*, Jun 2021. arXiv: 2106.03037.

[11] A. Di Scipio. 'sound is the interface': from interactive to ecosystemic signal processing. *Organised Sound*, 8(3):269–277, 2003.

[12] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi. High fidelity neural audio compression. (arXiv:2210.13438), Oct 2022. arXiv:2210.13438 [cs, eess, stat].

[13] A. Eldridge, C. Kiefer, D. Overholt, and H. Ulfarsson. Self-resonating Vibrotactile Feedback Instruments ||: Making, Playing, Conceptualising :||. In *NIME 2021*, jun 9 2021. https://nime.pubpub.org/pub/6mhrjiqt.

[14] R. Fiebrink. Wekinator 2.0, 2015.

[15] C. Frauenberger. Entanglement hci the next wave? *ACM Transactions on Computer-Human Interaction*, 27(1):2:1–2:27, Nov 2019.

[16] C. Frisson, M. Bredholt, J. Malloch, and M. Wanderly. Maplooper: Live-looping of distributed gesture-to-sound mappings.

[17] M. Grob. Growth due to limitations. http://www.livelooping.org/history_concepts/theory/growth-along-the-limitations-of-the-tools/, 2009. Accessed: 2023-01-25.

[18] L. Hantrakul. Gesturernn: A neural gesture system for the roli lightpad block. In T. M. Luke Dahl, Douglas Bowman, editor, *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 132–137, Blacksburg, Virginia, USA, June 2018. Virginia Tech.

[19] M. Horta. Okachihuali. http://moiseshorta.audio/obra/okachihuali/. Accessed: 2023-01-31.

[20] B. Kane. *Sound Unseen: Acousmatic Sound in Theory and Practice*. Oxford University Press, 2014.

[21] B. Latour. Technology is society made durable. *The sociological review*, 38(1 suppl):103–131, 1990.

[22] C. Levaux. The forgotten history of repetitive audio technologies. *Organised Sound*, 22(2):187–194, 2017. Publisher: Cambridge University Press.

[23] M. Levin. Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 2022.

[24] D. Y. Little and F. T. Sommer. Learning and exploration in action-perception loops. *Frontiers in Neural Circuits*, 7:37, 2013.

[25] Magenta. Realtime ddsp neural synthesizer and effect. https://github.com/magenta/ddsp-vst, Jan 2023. Accessed: 2023-01-31.

[26] T. Magnusson. Ergodynamics and a semiotics of instrumental composition. *Tempo*, 73(287):41–51, Jan 2019.

[27] T. Magnusson. *Sonic writing: technologies of material, symbolic, and signal inscriptions*. Bloomsbury Academic, 2019.

[28] M. Marchini, F. Pachet, and B. Carré. Rethinking reflexive looper for structured pop music. In *NIME 2017*, 2017.

[29] C. P. Martin, K. O. Ellefsen, and J. Torresen. Deep predictive models in interactive music. *arXiv preprint arXiv:1801.10492*, 2018.

[30] C. P. Martin and J. Torresen. An interactive musical prediction system with mixture density recurrent neural networks. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, page 260–265, Porto Alegre, Brazil, Jun 2019.

[31] F. Morreale, S. M. A. Bin, A. McPherson, P. Stapleton, and M. Wanderley. A nime of the times: Developing an outward-looking political agenda for this community. In R. Michon and F. Schroeder, editors, *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 160–165, Birmingham, UK, July 2020. Birmingham City University.

[32] K. P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023.

[33] P. Oliveros. The expanded instrument system (eis). *International Computer Music Conference Proceedings*, 1991, 1991.

[34] F. Pachet, P. Roy, J. Moreira, and M. d'Inverno. Reflexive loopers for solo musical improvisation. In *CHI 2013*, 2013.

[35] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[36] T. W. Patteson. The time of roland kayn's cybernetic music. *Springerin*, 18(2):62, 2012.

[37] T. Pelinski, V. Shepardson, S. Symons, F. S. Caspe, A. L. B. Temprano, J. Armitage, C. Kiefer, R. Fiebrink, T. Magnusson, and A. McPherson. Embedded ai for nime: Challenges and opportunities. In *International Conference on New Interfaces for Musical Expression*. PubPub, 2022.

[38] C. Roads. *Microsound*. The MIT Press, 2004.

[39] M. Rodger, P. Stapleton, M. van Walstijn, M. Ortiz, and L. S. Pardue. What makes a good musical instrument? a matter of processes, ecologies and specificities. In R. Michon and F. Schroeder, editors, *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 405–410, Birmingham, UK, July 2020. Birmingham City University.

[40] A. Rosenblueth, N. Wiener, and J. Bigelow. Behavior, purpose and teleology. *Philosophy of Science*, 10(1):18–24, Jan 1943.

[41] D. Sanfilippo and A. Valle. Feedback systems: An analytical framework. *Computer Music Journal*, 37(2):12–27, Jun 2013.

[42] H. Scurto, B. Caramiaux, and F. Bevilacqua. Prototyping machine learning through diffractive art practice. In *Designing Interactive Systems Conference 2021*, DIS '21, page 2013–2025, New York, NY, USA, Jun 2021. Association for Computing Machinery.

[43] K. Tahiroğlu, M. Kastemaa, and O. Koli. Al-terity: Non-rigid musical instrument with artificial intelligence applied to real-time audio synthesis. In R. Michon and F. Schroeder, editors, *Proceedings of the International Conference on New Interfaces for Musical Expression*, page 337–342, Birmingham, UK, Jul 2020. Birmingham City University.

[44] T. Tate. The concentric sampler: A musical instrument from a repurposed floppy disk drive. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, The University of Auckland, New Zealand, June 2022.

[45] B. Wallace and C. P. Martin. Comparing models for harmony prediction in an interactive audio looper. In *EvoMUSART*, 2019.

[46] S. Waters. Performance ecosystems: Ecological approaches to musical interaction. *EMS: Electroacoustic Music Studies Network*, page 1–20, 2007.