OTIAC – Co-Improvising With a Musical Agent in a Feedback-Based Guitar Performance

Claudio Panariello
claudio.panariello@univ-lille.fr
Univ. Lille, CNRS, Centrale Lille,
UMR 9189 CRIStAL, F-59000, Lille, France

Ken Déguernel
ken.deguernel@cnrs.fr
Univ. Lille, CNRS, Centrale Lille,
UMR 9189 CRIStAL, F-59000, Lille, France

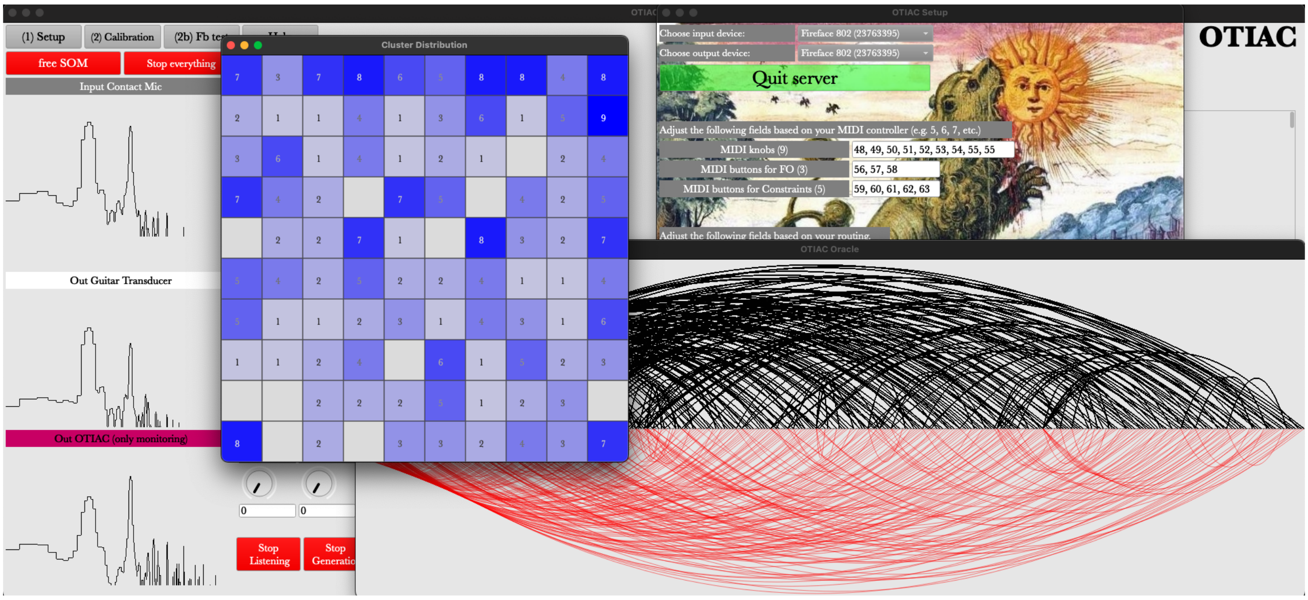

7	3	7	8	6	5	8	8	4	8
2	1	1	4	1	3	6	1	5	9
3	6	1	4	1	2	1		2	4
7	4	2		7	5		4	2	5
	2	2	7	1		8	3	2	7
5	4	2	5	2	2	4	1	1	4
5	1	1	2	3	1	4	3	1	6
1	1	2	4		6	1	5	2	3
		2	2	2	5	1	2	3	
8		2		3	3	2	4	3	7

Choose input device: Fireface 802 (23763395)
Choose output device: Fireface 802 (23763395)
Adjust the following fields based on your MIDI controller (e.g. 5, 6, 7, etc.)
MIDI knobs (9) 48, 49, 50, 51, 52, 53, 54, 55, 55
MIDI buttons for FO (3) 56, 57, 58
MIDI buttons for Constraints (5) 59, 60, 61, 62, 63
Adjust the following fields based on your routing.

Figure 1: Main user interface for OTIAC with additional subwindows showing setup, SOM clusters, and Factor Oracle.

Abstract

OTIAC —“O totaro int’a chitarra” (the octopus in the guitar), from Neapolitan folklore— explores co-improvisation between a guitarist playing on a feedback-augmented guitar, an electronics performer, and an artificial musical agent. The system combines online machine learning with algorithmic sequence generation: a Self-Organizing Map clusters live feedback-based audio in real-time, creating a learned symbolic vocabulary that feeds a Factor Oracle automaton. The oracle then generates and re-injects sonic material into the guitar’s body via transducer, creating a self-referential feedback ecosystem where all the agents shape the emerging performance.

This paper presents an open-source SuperCollider implementation integrating real-time feature extraction, unsupervised clustering (SOM), and context-aware sequence generation (Factor Oracle) within a unified performance system, and reports on preliminary observations from practice-led research with performers. These include the emergence of asymmetric awareness between performers with and without visual access to system state, early performance strategies for balancing manual intervention with autonomous behavior, and empirical insights into the performative and collaborative dimensions of human–AI musical interaction.

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

Keywords

Augmented instruments, feedback, musical agents, human–AI musical interaction

1 Introduction

Music Generation Systems are able to compose music using computational means [11]. They can be constituted of one or several Musical Agents, *i.e.* artificial agents capable of creating music, individually or in collaboration with other agents [51]. The development of musical agents has motivated decades of research, from early seminal works by Cope and Lewis to recent approaches using machine learning [22] and corpus-based concatenative synthesis [40, 47]. However, such systems typically rely on pre-existing corpora, necessitating the construction of a priori knowledge about the musical material they are expected to listen to and co-improvise with. Consequently, they are not inherently suited to operate with unpredictable sonic material such as Larsen tones or feedback-based music.

In this paper, we introduce OTIAC, a musical agent designed for a feedback-augmented classical guitar. OTIAC’s architecture combines Self-Organizing Maps with Factor Oracle automaton. Unlike corpus-based systems that work offline on pre-recorded material, OTIAC operates exclusively online, learning and generating from live feedback-based guitar performance. The research pursues a dual objective: to present the system’s implementation, and to document the development and integration of feedback augmentation within the instrument itself.

Therefore our implementation is distinctive in three ways: (1) it provides a novel open-source Factor Oracle implementation for

SuperCollider,[1] enabling tight integration with real-time audio processing; (2) it does not rely on any pre-existing corpora, constructing all knowledge from the current feedback session; and (3) it is specifically tailored for feedback-augmented instruments, where continuous, spectrally-rich timbres replace discrete note events.

The musical agent is then evaluated through practice-led research with professional performers in a real-world performance scenario: OTIAC —"O totaro int'a chitarra" (the octopus in the guitar)— a feedback-based guitar performance conceived for the H[t] Duo.

The system's source code and documentation are available online.[2] A recording of a rehearsal session with the performers is available as Supplementary Material.

2 Related work

2.1 Music Generation Systems

Several frameworks have been proposed for categorizing music generation systems and musical agents. Herremans et al. [27] provide a taxonomy distinguishing between systems based on their input requirements, generation methodology, and interaction paradigm. Tatar and Pasquier [51] offer a typology of musical agents oriented toward Musical Metacreation, distinguishing agents by their degree of autonomy, learning capabilities, and interaction modes. Jung [28] proposes a design framework for Intelligent Music Performance Systems emphasizing the balance between system intelligence and performer agency.

In this work we use a musical agent based on Self-Organizing Maps (SOMs) and Factor Oracle (FO) automaton.

SOMs are unsupervised neural networks that cluster and visualize high-dimensional data into a 2D map while preserving topological relationships as faithfully as possible (*i.e.*, similar data points stay close on the map) [29]. SOMs are computationally efficient and are built as a grid of nodes: each node representing a possible category in the input space. It has been proposed that SOMs can be effective in implementing the memory properties of creative processes [7, 25]. In the music domain, however, only few cases have been found of musical agents employing SOMs: Smith and Deal [48] adopted SOMs to model the short term memory of a music agent; Martins and Miranda [34] employed a SARDNET (a variation of SOMs with an addition of temporality) for rhythm generation; Tatar and Pasquier proposed MASOM [50], a machine improvisation architecture for live performance which combines SOMs with Variable Markov Models; and Thelle and Pasquier presented Spire Muse [52], which builds upon MASOM architecture.

A Factor Oracle (FO) is a finite state automaton introduced by Allauzen et al. [1] and later applied to music generation [5]. A key advantage of FO is its lightweight nature and linear-time and linear-space learning complexity, allowing for adaptive applications such as live musical improvisation. FO has become a well-known technique extensively adopted in designing musical agents, such as OMax [3], Anticipatory Model [13], Audio Oracle [17], PyOracle,[3] Variable Markov Oracles [2, 57], and Probabilistic Factor Oracle [18]; and including also the extensions of the OMax family, like ImproteK [37, 38] (which introduces the scenario/memory generation model), Dicy2 [39], and Somax2 [4, 8].

2.2 Feedback-Based Music

Feedback has a long and deep presence in the music domain. A rich body of literature has documented the emergence of various approaches throughout history, from the experiments in their respective fields by Roland Kayn and Jimi Hendrix in the 1960s, to the more recent advances in Complex Adaptive Systems (CASes) [45, 46, 55]. Works by Alvin Lucier (I Am Sitting in a Room, 1969), Eliane Radigue (Opus 17, 1970), David Tudor (Microphone, 1973), Nicolas Collins (Pea Soup, 1974) [12], and more recently Di Scipio (Audible Ecosystemics series, 2002-2005) [15], Sanfilippo (Order from Noise, 2016-2019) [43], and utrumque duo (Elblaus and Eckel; Rundgång, 2020; St. Elisabeth Song Cycle, 2021) [19],[4] just to name a few, are key examples of how feedback-based music systems of different complexity have posed compelling technological challenges, while exploring questions of musical aesthetics and musical forms.

Feedback-based music systems have informed diverse applications of feedback in musical performance. One such application can be found in feedback-augmented instruments, which are a sub-category of augmented instruments that uses audio feedback to enrich the acoustic instrument's timbre. These instruments generally employ an actuator and a microphone (or any means that can transduce acoustic vibrations to an electrical signal), which are coupled through the resonant body of the instrument and connected in a feedback loop that includes analogue and/or digital signal processing [21]. In recent years, feedback-augmented string instruments have attracted considerable attention [33], as exemplified by Halldór Úlfarsson's halldorophone [54] and related family of feedback string instruments (see *e.g.* [14, 20, 31, 35, 41], or the work by Johannes Burström[5]).

Despite the proliferation of feedback string instruments, the classical guitar has received comparatively less attention in academic literature on feedback augmentation. Most documented guitar augmentation focuses on steel-string instruments (such as Harriman's Feedback Lap Steel [26]), on the electric guitar (like [32] or, outside of academic contexts, Christian Blandhoel's Feedbackers[6] [6]), or on customized tabletop string instruments (like Visi's Sophtar [56]). On the other hand, there are examples of classical guitar augmentation without taking into account the role of acoustic feedback as sound source (see e.g. [30, 36]).

As of yet, to the knowledge of the authors, there are no examples of feedback-augmented classical guitars in the academic literature.

3 OTIAC: System Architecture

We now present the ideas behind the hardware design, the design of the generation system, and the design of the interface presented to the performers.

3.1 Hardware design

The hardware design was conceived with the aim of achieving a deliberately minimal augmentation, requiring readily available materials and enabling easy replication by other performers. The augmentation is, in fact, stripped-down: a contact microphone is placed on the guitar's headstock, and a transducer is mounted on the top plate.

The classical guitar presents an interesting case for feedback augmentation due to its distinctive acoustic properties. The lower

---

[1] https://supercollider.github.io/

[2] https://github.com/claudiopanariello/otiac or https://www.algomus.fr/code/

[3] https://pypi.org/project/PyOracle/

[4] https://www.utrumque.com/

[5] https://johannesburstrom.se/works/solo/

[6] https://www.ijin.no/feedbacker.htm

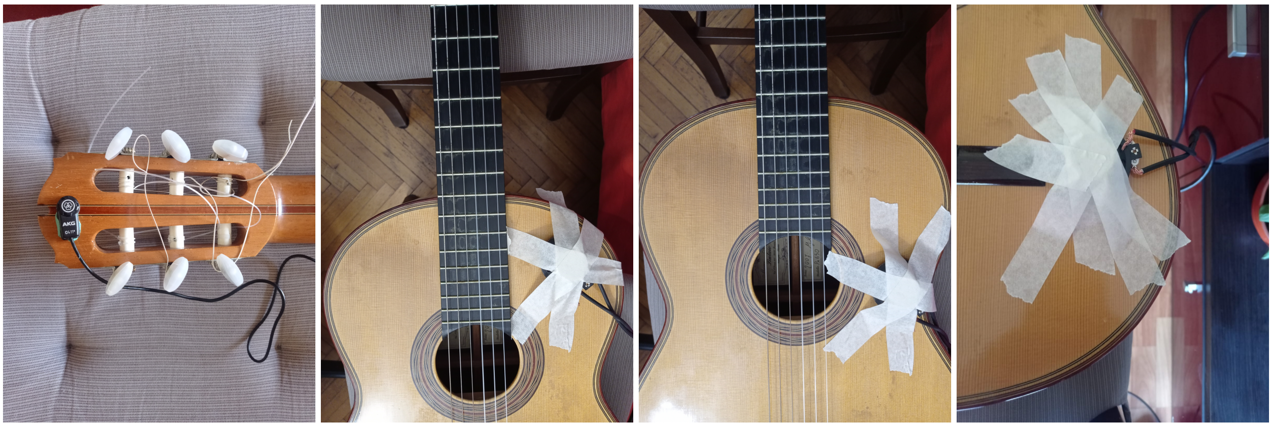

(a) (b) (c) (d)

Figure 2: Preliminary prototyping phase showing the contact mic and different placements of the transducer, which is temporarily fixed to the guitar using paper tape. (a) The contact mic is placed on the back of the guitar's headstock. (b) The transducer is placed at the top part. (c) The transducer is placed besides the sound hole. (d) The transducer is placed besides the bridge. The pictures were taken by P1 on his guitar.

string tensions typical of nylon strings (50–80 N compared to 100–180 N for steel strings [23]) allow the classical guitar top plate to be built thinner and more flexibly braced (fan-braced compared to the crossed bracing of a steel-string guitar), contributing to the instrument's particularly rich resonant response to external excitation via transducer. Furthermore, the absence of any electromagnetic pickup chain means that feedback in this context is purely acoustic-mechanical in nature, resulting in Larsen tones whose timbral character is shaped entirely by the instrument's physical resonances rather than by electronic transduction chains. These properties make the classical guitar body a particularly rich and sensitive medium for transducer-driven feedback excitation.

Since the goal was to generate feedback within the instrument itself, hardware placement proved critical and required multiple iterative prototyping phases. Ultimately, a condenser contact microphone was positioned on the headstock (see Fig. 2a) to capture the neck's tonal character and enable direct feedback control (*e.g.*, damping the neck with the hand interrupts feedback). The transducer (a 25 mm exciter with 20 W power and 4 Ω impedance) was mounted on the top plate, and its placement required extensive testing across multiple positions. An early insight was that feedback could also be modulated by interacting directly with the top plate, such as pressing nodal points that emerge during resonant excitation of the guitar body. Eventually, the final transducer position was the one shown in Fig. 2d, on the bottom part and close to the bridge. This location proved optimal for feedback generation while also corresponding to the guitar's top plate principal resonant modes [23], which are exploited musically throughout the work.

Fig. 2 shows a preliminary prototyping stage, in which the transducer is temporarily fixed to the guitar using paper tape. The rudimentary setup was intentionally designed for rapid iteration and non-invasive experimentation, allowing flexible repositioning of the transducer without risk of damage to the instrument. Future iterations will explore more robust yet reversible mounting strategies, such as removable adhesive materials or a custom lightweight support structure.

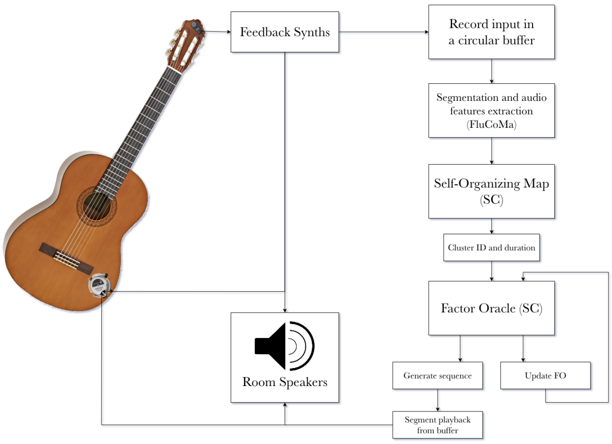

Figure 3: Overview of the overall signal flow in OTIAC.

3.2 System Design

OTIAC's processing pipeline follows this sequence: audio input from a contact microphone on the guitar's headstock is continuously analyzed to extract spectral and timbral features; these features are fed to an SOM that clusters similar sonic events while learning their distribution online; detected onsets trigger event classification, storing each event with its cluster assignment and temporal metadata; these classified events build an FO modeling sequential relationships in the performed material. These stages go under the *Listening Mode*. Finally, the *Generation Mode* can be activated, where the oracle generates new sequences that are matched to recorded audio segments and re-injected into the guitar via transducer, completing the feedback loop. The overall signal flow is shown in Fig. 3.

3.2.1 Audio Features. Audio features are computed using the FluCoMa toolkit's SuperCollider implementation [53] with fixed settings, using a sampling rate of 48 kHz, window size of 1024 samples (*i.e.* 21 ms), and hop size of 512 samples (11 ms). The extracted feature set comprises 20 dimensions: loudness, true

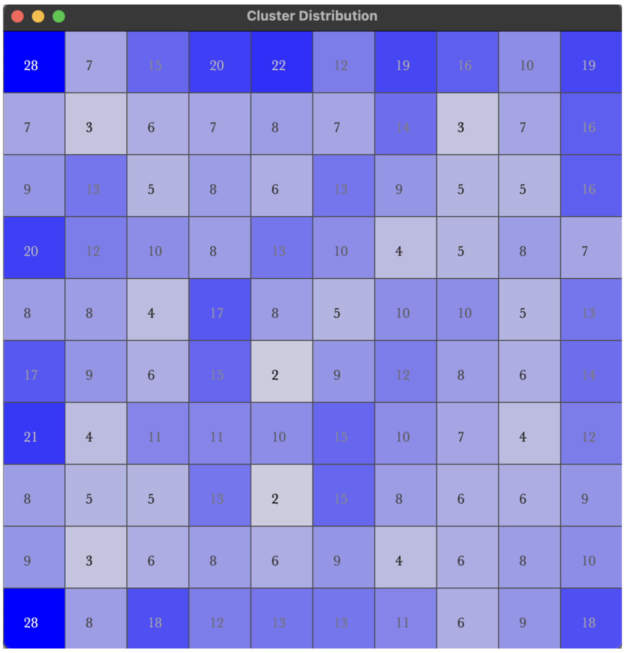

Cluster Distribution

28	7	15	20	22	12	19	16	10	19
7	3	6	7	8	7	14	3	7	16
9	13	5	8	6	13	9	5	5	16
20	12	10	8	13	10	4	5	8	7
8	8	4	17	8	5	10	10	5	13
17	9	6	15	2	9	12	8	6	14
21	4	11	11	10	15	10	7	4	12
8	5	5	13	2	15	8	6	6	9
9	3	6	8	6	9	4	6	8	10
28	8	18	12	13	13	11	6	9	18

Figure 4: Visualization of a 10×10 SOM grid. Each square represents a node colored according to how many events have been classified in that specific cluster ID.

peak, six spectral shape descriptors: spectral centroid, spectral spread, normalized skewness, normalized kurtosis, rolloff, flatness; and 12 Mel-frequency cepstral coefficients (MFCCs). Note that the 0th MFCC coefficient, which correlates with the overall energy in the spectrum, is excluded from the feature set since dedicated loudness descriptors already capture this information. All features are normalized to the range $[0, 1]$ before being passed to the SOM. Spectral features use exponential or linear scaling as appropriate to their perceptual characteristics.

The incoming signal is segmented employing FluCoMa's novelty detection algorithm, based on [24], which identifies significant changes in the spectral content. A Schmidt trigger with hysteresis (0.05 onset threshold, 0.1 offset threshold) debounces the raw novelty signal, preventing spurious re-triggering. Additionally, silence detection (with an amplitude threshold of 0.001) ensures that the system captures also meaningful pauses. The combination of novelty detection and silence detection provides a robust event segmentation that responded to both timbral changes and dynamic envelope characteristics typical of feedback.

The incoming audio signal is continuously written to a circular buffer (default: 10 minutes capacity) independent of event detection, ensuring that the system can retrieve any segment of recent performance history when generating playback.

3.2.2 *Self-Organizing Map (SOM).*

The system uses the SOM implementation in SuperCollider realized by Dan Stowell, [7] which provides two UGens: `SOMTrain` for online training and `SOMRd` for classification.

The training of `SOMTrain` occurs only when new events are detected (gated by the onset detector), ensuring that computational resources are allocated to genuine sonic events rather than continuous background processing. Each training iteration updates the best-matching unit and its neighbors according to standard SOM learning rules. On the other hand, `SOMRd` queries the network taking the current 20-dimensional feature vector and returning the grid coordinates (x, y) of the best-matching node (the node whose weight vector has minimum Euclidean distance to the input). These coordinates are converted to a single cluster ID, yielding values in the range $[0, 99]$.

OTIAC employs a 10×10 grid topology, yielding 100 nodes in a two-dimensional lattice. The SOM is initialized using random grid. Fig. 4 shows a SOM with 100 nodes trained through OTIAC. Each square represents a node colored according to how many events have been classified in that specific cluster ID.

3.2.3 *Event Representation and Storage.*

Each detected onset triggers the creation of a composite event structure stored in a list:

`[durationCategory, clusterID, metadataDict],`

where **durationCategory** is a symbol categorizing the event's temporal extent based on the inter-onset interval:

- `\very_short` (< 0.5s);
- `\short` (0.5–1.0s);
- `\medium` (1.0–3.0s);
- `\long` (3.0–5.0s);
- `\very_long` (5.0–7.0s);
- `\sustained` (> 7.0s).

clusterID is the SOM node index (0–99) representing the event's timbral class; and **metadataDict** is a dictionary containing: the exact duration in seconds, buffer position (sample index in the circular audio buffer), onset timestamp, SOM grid coordinates, the full 20-dimensional feature vector, and current training iteration.

By storing both discrete symbols (duration category and cluster ID) and continuous metadata (exact timing, features, buffer positions), the system enables the Factor Oracle to model sequences symbolically while still retrieving the original audio segments with their full timbral characteristics.

3.2.4 *Factor Oracle implementation.*

OTIAC features a novel SuperCollider implementation of the FO model, developed specifically for this project. The `FactorOracle` class accepts sequences of events in the form `[durationCategory, clusterID]`. For example, an event might be represented as `[\sustained, 6]`, where 6 is the SOM node index and `\sustained` indicates a duration exceeding 7 seconds.

Each event's unique symbol—a composite key combining duration category and cluster ID (*e.g.*, `"sustained_6"`)—defines a state transition. The algorithm incrementally constructs the forward transitions and the suffix links. The FO employed in OTIAC automatically updates whenever the event count reaches a multiple of N events, incorporating newly performed material while preserving learned sequential relationships from earlier in the performance. With $N = 1$ the FO updates after each event, providing immediate responsiveness. However, the slower temporal evolution characteristic of feedback-based material, with sustained tones and gradual transformations, benefits from a larger accumulation window. Therefore setting a higher value for N allows the system to accumulate sufficient events to capture these slow changes. OTIAC's FO update is set to $N = 30$.

When in *Generation Mode*, OTIAC employs context-based sequence generation triggered by incoming events. The process unfolds as follows: after a new event is detected and classified by SOM, the returning cluster ID (combined with the event's duration category) defines the current context. The oracle searches for the state that best matches this context by finding the longest

---

[7] https://github.com/supercollider/sc3-plugins/tree/main/source/MCLDUGens/sc

matching suffix in the learned sequence. Starting from this state, the oracle generates a new sequence of events. Each generated event is matched to a recorded audio segment in the circular buffer and the various segments finally are played back with crossfading envelopes, routed to the guitar transducer and the room speakers.

3.3 Interface Design

`OTIAC_GUI` provides a comprehensive graphical interface designed to make the system accessible to performers who may not be familiar with SuperCollider's text-based programming environment. The interface manages audio routing, MIDI mapping, real-time monitoring, and performance parameters while balancing immediate accessibility with protection against accidental changes during performance.

Since OTIAC relies on a number of dependencies, namely `JSONlib`,[8] FluCoMa [53], and custom SuperCollider classes, the GUI constructor performs automatic dependency checking on launch.

3.4 Main Interface Layout

The main interface is organized into three functional zones: setup and initialization controls at the top, real-time monitoring in the center, and performance controls at the bottom. Fig. 5 shows the interface with the system in action.

Setup and Initialization Controls:

- **Setup:** Opens a configuration window (Fig. 6) for selecting audio interfaces, configuring MIDI controller mappings, setting channel routing, defining resonance frequencies, and managing configuration presets via JSON files.
- **Calibration:** Initiates a guided calibration process to establish optimal input/output levels for feedback generation. The system provides real-time feedback in the log window, instructing the performer to adjust the transducer level until achieving the target range where Larsen tones activate reliably.
- **Fb test:** Activates a basic feedback loop between the contact mic and the transducer to verify the audio routing and test fundamental feedback generation before engaging the full system.
- **Initialize SOM:** Instantiates the SOM and audio analysis synths, preparing the system to receive, analyze, and cluster incoming audio from the contact microphone. It also handles the FO build and the subsequent sequences generation.
- **Start synths and timer:** Launches all resonator synths, the adaptive filter, the mixer, and starts the timer.

Real-time Monitoring: The center section features five frequency analyzers displaying spectral content of different signal paths: contact microphone input, transducer output, left and right speaker outputs, and FO monitoring bus. Adjacent to these are utility buttons:

- **Draw clusters:** Shows the current distribution of events across the 10×10 SOM grid, with color intensity indicating cluster density.
- **Draw FO:** Opens a graphical representation of the Factor Oracle's current state structure, showing forward transitions and suffix links among the states.

Two text windows provide real-time logging: one displays system events as they occur (onset detection, cluster assignments, oracle state changes), while the other records general messages and warnings. A digital timer tracks elapsed performance time in minutes and seconds.

Performance Controls: The bottom section contains two rows of MIDI-mappable controls. The first row consists of nine knobs controlling respectively:

- Levels for three fixed-frequency resonators;
- Level for the adaptive spectral centroid-based filter (based on the SC implementation [42] of the work by [44]);
- Three FO controls: generation interval, sequence length, and continuity parameter;
- Level for oracle-generated segments re-injected into guitar;
- Overall level for room speakers;

The second row features eight buttons. The first three control system state:

- **Start/Stop Listening:** Toggles the whole system listening capability, that is the onset detection, feature extraction, and SOM training;
- **Start/Stop Generation:** Activates/deactivates oracle sequence generation and playback;
- **Start/Stop FO Update:** Enables/disables automatic oracle update every N detected events.

The remaining buttons can be used to activate distinct **constraint modes**, each reconfiguring the three resonator synths with different parameter sets optimized for specific sonic behaviors. Such behaviors can be decided by the performers to accommodate specific customizations.

3.5 Interaction Design Philosophy

A critical design decision separates mouse-accessible controls from MIDI-only controls. Setup, initialization, and visualization buttons can be activated with the mouse for pre-performance configuration. In contrast, all knobs and performance buttons in the bottom rows are *exclusively* controllable via MIDI, preventing accidental parameter changes during performance when the performer's attention is focused on listening and playing rather than screen interaction.

This dual-access approach emerged from iterative development with H[t] Duo, balancing the need for quick setup against the risk of disrupting performance flow with unintended mouse clicks. The MIDI-only constraint for performance parameters also encourages tactile, muscle-memory-based interaction rather than visual attention to the screen.

4 Performance Practice and Evaluation

OTIAC was developed through an iterative, practice-led research approach in collaboration with the H[t] Duo[9] (Pierpaolo Dinapoli, classical guitar, and Matteo Tundo, electronics). Over a period of 9 months, the duo participated in multiple rehearsal sessions during which the system and interface underwent successive refinements based on their feedback and observed interactions.

After the third practice session, where the duo had gained substantial familiarity with the system, Author 1 conducted a one-hour semi-structured interview via video call. The interview took place a few hours after the duo's weekly rehearsal, ensuring the experience was still fresh while allowing time for initial

---

[8] https://github.com/musikinformatik/JSONlib

[9] https://www.pierpaolodinapoli.com/en/ht-duo-2/

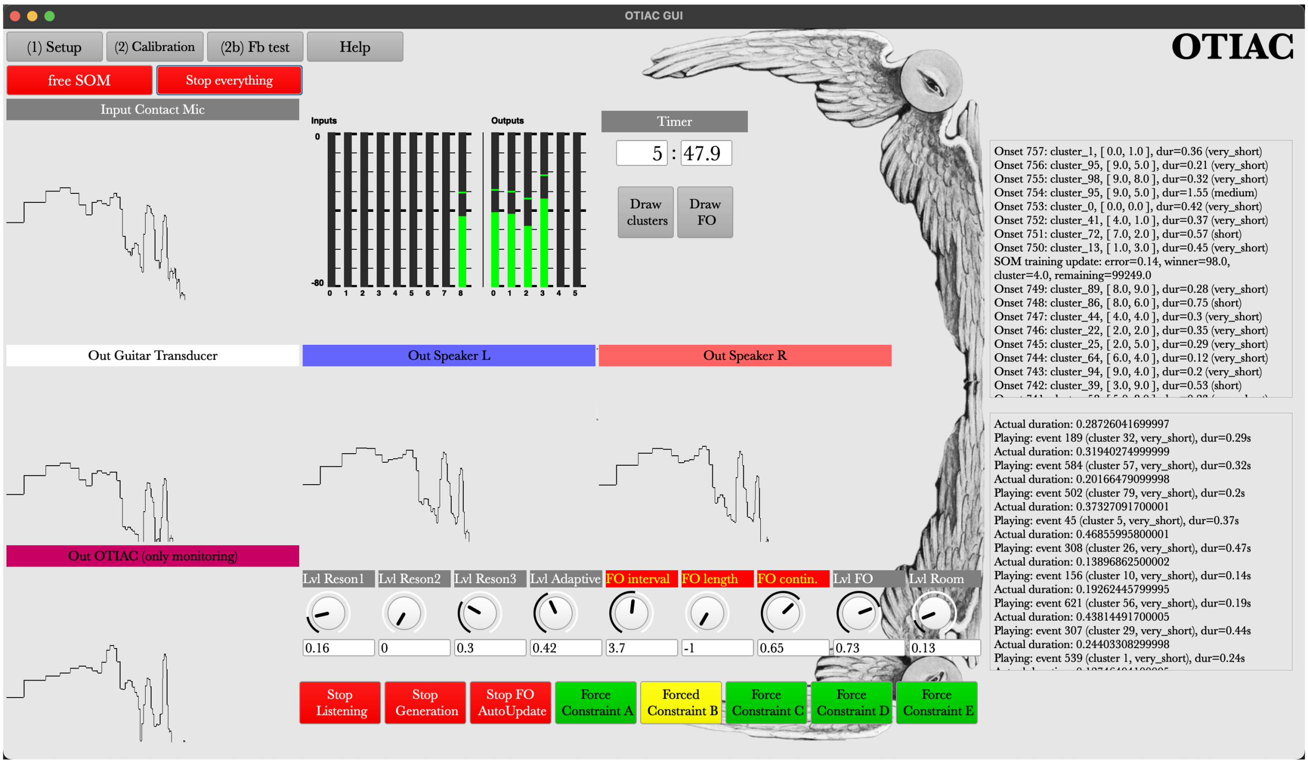

Timer: 5 : 47.9

Lvl Reson1: 0.16
Lvl Reson2: 0
Lvl Reson3: 0.3
Lvl Adaptive: 0.42
FO interval: 3.7
FO length: -1
FO contin.: 0.65
Lvl FO: 0.73
Lvl Room: 0.13

```
Onset 757: cluster_1, [ 0.0, 1.0 ], dur=0.36 (very_short)
Onset 756: cluster_95, [ 9.0, 5.0 ], dur=0.21 (very_short)
Onset 755: cluster_98, [ 9.0, 8.0 ], dur=0.32 (very_short)
Onset 754: cluster_95, [ 9.0, 5.0 ], dur=1.55 (medium)
Onset 753: cluster_0, [ 0.0, 0.0 ], dur=0.42 (very_short)
Onset 752: cluster_41, [ 4.0, 1.0 ], dur=0.37 (very_short)
Onset 751: cluster_72, [ 7.0, 2.0 ], dur=0.57 (short)
Onset 750: cluster_13, [ 1.0, 3.0 ], dur=0.45 (very_short)
SOM training update: error=0.14, winner=98.0,
cluster=4.0, remaining=99249.0
Onset 749: cluster_89, [ 8.0, 9.0 ], dur=0.28 (very_short)
Onset 748: cluster_86, [ 8.0, 6.0 ], dur=0.75 (short)
Onset 747: cluster_44, [ 4.0, 4.0 ], dur=0.3 (very_short)
Onset 746: cluster_22, [ 2.0, 2.0 ], dur=0.35 (very_short)
Onset 745: cluster_25, [ 2.0, 5.0 ], dur=0.29 (very_short)
Onset 744: cluster_64, [ 6.0, 4.0 ], dur=0.12 (very_short)
Onset 743: cluster_94, [ 9.0, 4.0 ], dur=0.2 (very_short)
Onset 742: cluster_39, [ 3.0, 9.0 ], dur=0.53 (short)
```

```
Actual duration: 0.28726041699997
Playing: event 189 (cluster 32, very_short), dur=0.29s
Actual duration: 0.31940274999999
Playing: event 584 (cluster 57, very_short), dur=0.32s
Actual duration: 0.20166479099998
Playing: event 502 (cluster 79, very_short), dur=0.2s
Actual duration: 0.37327091700001
Playing: event 45 (cluster 5, very_short), dur=0.37s
Actual duration: 0.46855995800001
Playing: event 308 (cluster 26, very_short), dur=0.47s
Actual duration: 0.13896862500002
Playing: event 156 (cluster 10, very_short), dur=0.14s
Actual duration: 0.19262445799995
Playing: event 621 (cluster 56, very_short), dur=0.19s
Actual duration: 0.43814491700005
Playing: event 307 (cluster 29, very_short), dur=0.44s
Actual duration: 0.24403308299998
Playing: event 539 (cluster 1, very_short), dur=0.24s
```

Figure 5: The main window of OTIAC when in action. The background image is a detail from the Hand of the Mysteries by John Augustus Knapp.

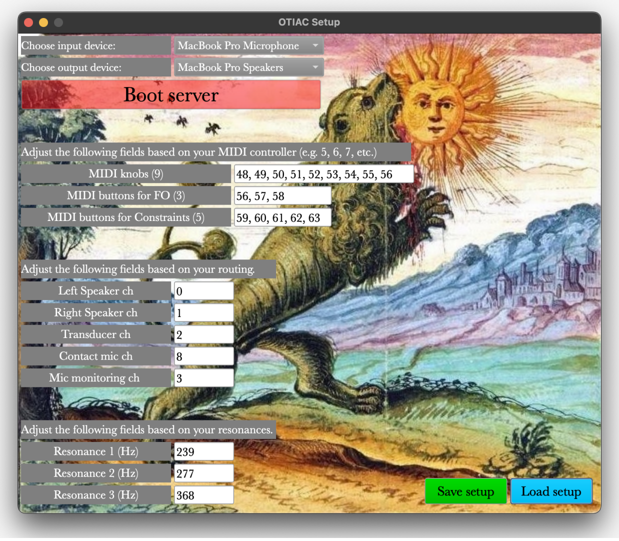

Choose input device: MacBook Pro Microphone
Choose output device: MacBook Pro Speakers

Adjust the following fields based on your MIDI controller (e.g. 5, 6, 7, etc.)
MIDI knobs (9): 48, 49, 50, 51, 52, 53, 54, 55, 56
MIDI buttons for FO (3): 56, 57, 58
MIDI buttons for Constraints (5): 59, 60, 61, 62, 63

Adjust the following fields based on your routing
Left Speaker ch: 0
Right Speaker ch: 1
Transducer ch: 2
Contact mic ch: 8
Mic monitoring ch: 3

Adjust the following fields based on your resonances.
Resonance 1 (Hz): 239
Resonance 2 (Hz): 277
Resonance 3 (Hz): 368

Figure 6: The setup window. The background image is the "Leo Viridis, the Green Lion Devouring the Sun", an illustration from the Alchemical and Rosicrucian Compendium.

reflection. The conversation was conducted in Italian (the participants' native language), recorded, transcribed, and analyzed for recurring themes and critical insights. The guitarist is labeled as P1 and the electroacoustic performer as P2.

The interview explored the following themes: the distribution of agency among the three performing agents (guitarist, electronics performer, autonomous system), strategies for controlling and responding to system behavior, the cognitive demands of the interface, temporal and formal considerations for performance, and the performers' subjective experience of collaboration with the musical agent.

4.1 Distributed Agency

One notable observation concerned a marked difference in how the two human performers experienced agency and awareness during performance. P2, operating the electronics and GUI, reported feeling "in full control" since he could activate or deactivate system components, adjust parameters, and shape the overall form of the piece. In contrast, P1 described moments of being *"suonato"* (played by the system), uncertain whether the sounds he was hearing originated from his own actions, P2's parameter changes, or the oracle's autonomous generation.

This asymmetry appeared to stem from access to visual information. As P1 articulated: *"P2 speaks like this because he has the screen in front of him. I don't have it, so my perception is completely different because I'm unaware of the system."* Without visual feedback about system state, P1 relied primarily on auditory cues and physical sensations (vibrations in the guitar body) to distinguish between different sound sources—a challenging task given the timbral similarity of feedback-based material regardless of origin.

The asymmetry became particularly pronounced when the oracle activated and P2 began making rapid parameter adjustments. P1 reported difficulty determining which feedback he was actively controlling versus which segments were being generated and re-injected by the oracle. This lack of clarity was described as frustrating when seeking more conscious control over the musical discourse, though P1 acknowledged it might resolve through extended practice and the development of *"diachronic mastery"* [35], that is the dynamic expertise enacted in improvisation within electroacoustic ecologies.

When asked whether they experienced the piece as a trio or quartet, responses diverged in ways that further illuminated this asymmetry. P2 positioned himself as a conductor coordinating three elements (*"Io mi sento di comandare!", "I feel I am leading here!"*): P1's guitar, the feedback system, and the oracle. P1, however, experienced a more complex relationship with the musical agent. While acknowledging P2's structural control, P1 conceptualized the performance as involving four distinct agents: *"I would say [we are] even a quartet... because in reality the feedback does things... it's true that I act on the feedback but when you open everything, the feedback also moves a bit as it wants. There's you [P2], there's the oracle, there's me, and the feedback".*

4.2 Emergent Control Hierarchies

Through their exploration of the system, P1 articulated a three-tiered hierarchy of controllability that emerged from practice:

1. **Fixed resonances:** Nearly completely controllable once systematically mapped. Pressure points on the top plate using the fingers or an ebow,[10] pressure points, and string tensions produce predictable results (harmonics jumps, microglissandi, sustained tones). With dedicated study, these could function as reliable compositional material.
2. **Adaptive spectral centroid:** A middle ground which moves in response to input, requiring the performer to track and adapt to its behavior, but remaining broadly predictable in character.
3. **SOM and FO generation:** The most complex and least predictable layer. While high-level behavior follows expected patterns (generating sequences based on learned material), the specific timing, cluster selection, and resulting timbres introduce genuine unpredictability that resists simple mastery.

This schema suggests a gradient of agency from human to machine, and importantly, demonstrates how performers conceptualize the system's different components through the lens of control rather than purely technical function. The hierarchy also suggests different performance strategies: the resonances can be "played" with intention, the adaptive filter requires responsive listening and adjustment, while the oracle invites a more improvisational stance of reacting to and incorporating its contributions.

P2's interaction with the interface paralleled this hierarchy. He reported adjusting resonance levels frequently and responsively, while FO's parameters (generation interval, sequence length, continuity) were set less often, typically established at the beginning of a section and left to run. The constraint buttons occupied an intermediate position, being used strategically to shift the sonic palette or create formal delineations, but not requiring constant adjustment.

4.3 Performance strategies, temporality and formal concerns

P1 engaged fully in the practice-led research methodology. He dedicated regular solo practice sessions to methodically exploring the feedback system's behavior, mapping which ebow positions trigger which resonances, where feedback can be interrupted or sustained, how string tension and damping affect timbral evolution. This approach, which he explicitly connected to Melbye's concept of diachronic mastery, treats the augmented guitar as an instrument requiring dedicated technical study rather than purely intuitive exploration.

Specific discoveries made include:

- Ebow positions that function as "switches," immediately killing or triggering feedback;
- Locations on the guitar body (particularly near the bridge and on the sides) that produce less stable but more complex feedback behaviors;
- Techniques for inducing harmonic jumps (4th, 5th harmonics) through controlled string damping;
- Methods for creating microglissandi between pitches by subtle tension adjustments.

On his side, P2 praised the possibility of having the possibility of activating **constraint modes** in the GUI, allowing the construction of more structured musical form. The five following modes that the duo used in their rehearsals were partially suggested by Author 1 and partially developed together:

- **Constraint A:** Basic resonant feedback at user-defined frequencies with minimal processing;
- **Constraint B:** Frequency-shifted feedback (2 Hz upward shift across all resonators);
- **Constraint C:** Granular synthesis mode with bandpass filtering and impulse train excitation;
- **Constraint D:** Complex mixed mode combining granulation, frequency shifting, and phase manipulation;
- **Constraint E:** Spectral displacement via frequency multiplication (1.2× multiplier) to force new harmonic regions;

Both performers felt that OTIAC would benefit from extended durations (20+ minutes) to allow thorough exploration of its sonic possibilities. The limited gestural palette, currently focused primarily on ebow techniques and resonance activation, risks redundancy in shorter performances. As P1 noted: *"If we limit ourselves to this type of exploration we're doing now, clearly the instrument sounds in a fairly similar way [for] ten minutes...".* Therefore P1 proposed expanding the guitar's vocabulary of techniques currently employed in OTIAC to sustain interest across varied durations, suggesting to adopt *"instrumental actions that interact with the instrument in a more, excuse the term, traditional way, so you can increase the duration aspects without any problems."*

A tension emerged between two approaches to formal structure: P2 leaned toward establishing a loose *"canovaccio"* (roadmap) outlining sections or constraint progressions, while P1 preferred to first achieve fuller instrumental mastery before committing to predetermined forms. This difference may reflect their different relationships to the system: P2, with fuller awareness and control, can more easily conceive of planned trajectories, while P1, still mapping the instrument's possibilities, prefers to preserve improvisational freedom.

4.4 Interface Affordances

Both performers expressed satisfaction with the current interface iteration, particularly compared to earlier versions. The ability to load complete configurations from JSON files was perceived as reducing setup friction, making the system feel more practical for concert use. P2 appreciated that the cognitive load remained manageable: the nine knobs and eight buttons provided sufficient

---

[10] Interestingly, the use of a switched-off ebow as a prop to exert pressure on the guitar's top plate nodes has proven to be very effective in controlling the feedback.

expressivity without overwhelming his attention, allowing him to focus primarily on listening rather than interface manipulation.

However, P2's satisfaction with the interface contrasted with P1's expressed need for greater awareness. He suggested that a secondary display showing spectral information, fundamental pitch tracking, and oracle state might help him distinguish sound sources and make more informed musical decisions. During practice, he had begun using a tuner app on his phone to track pitch and spectral content in real-time, which is an improvised solution to the awareness gap. The question of whether to integrate such feedback into the system remains open: it could enhance conscious control but might also shift attention from embodied listening to visual monitoring.

5 Discussion

The development and practice-led study of OTIAC points to several tensions in designing co-improvisation systems where human performers and computational agents negotiate shared musical agency. The following discussion offers exploratory findings, intended to generate hypotheses for future systematic study rather than definitive conclusions.

A notable observation concerns the emergence of asymmetric awareness between performer: P1 and P2 experienced agency differently during performance. P2, operating the electronics interface, reported feeling in full control, while P1 described moments in which he was unable to distinguish whether sounds originated from his actions, P2's parameter changes, or the oracle's generation. This asymmetry is not merely technical but epistemological: the performers do not share the equivalent perceptual access. P2 possesses both sonic and visual knowledge of system state, while P1 relies almost exclusively on auditory and haptic cues. Such uncertainty around source attribution in electro-acoustic performance is not isolated to OTIAC. Borgo notes that performers in electro-acoustic settings can find it genuinely difficult to distinguish live sounds from processed or replayed material [10], and more broadly that in group improvisation each performer may have a substantially different interpretation of what is occurring, *i.e.* intersubjectivity is intrinsic to collective performance [9]. This ambiguity may prove aesthetically productive, encouraging deeper listening, but can also frustrate when seeking decisive control.

The performers' disagreement about whether OTIAC constitutes a trio or quartet reflects genuine ontological ambiguity. P2's trio conception (guitarist, electronics performer, oracle) treats feedback as controllable material. P1's quartet formulation, on the other hand, grants feedback autonomy. Drawing on Barad's notion of intra-action, Tahiroğlu [49] argues that in human-AI musical contexts agency is not locatable in any single entity but emerges through mutually constitutive relationships between performers, instruments, and technological systems. From this perspective, both the trio and quartet framings may be insufficient: the disagreement would therefore point at the distributed, intra-active nature of the system, where boundaries between agents are not fixed but continuously renegotiated in performance. This might also resonate with Di Scipio's notion of ecosystemic agency [16]: agency is a property of the whole ecosystemic network in operation. Notably, P1 discovered that feedback could be interrupted at specific guitar's body locations, suggesting hidden control affordances within apparently autonomous processes.

Interestingly, P1 articulated a three-tiered hierarchy of controllability: fixed resonances (nearly fully controllable once mapped), adaptive spectral centroid filter (requiring continuous tracking), and oracle/SOM generation (most unpredictable). This phenomenological taxonomy reflects how the system *feels to play* for the musicians rather than its technical organization, *i.e. what each component does technically*: the performers experience the system through the lens of control and predictability rather than functional modules.

OTIAC operates without pre-existing corpora, learning exclusively from current performance. This proves particularly apt for feedback-based music, where each context (tuning, acoustics, transducer placement) produces substantially different sonic possibilities (hence pre-trained corpora from previous performances would likely prove ineffective). However, the corpus-free approach creates a significant aesthetic limitation: only feedback material becomes part of the learned vocabulary. Both P1 and P2 recognized this constraint, noting that the gestural palette risks redundancy over extended durations. P1 proposed expanding the instrumental techniques: by incorporating traditional guitar gestures alongside feedback, the oracle could learn from a richer sonic vocabulary, potentially enabling more varied generative output.

6 Conclusion and Future Work

In this paper we presented OTIAC, a SuperCollider-based musical agent for co-improvising in feedback-based guitar performances. The system integrates Self-Organizing Maps with Factor Oracle automaton to enable real-time learning and generation from live performance without relying on pre-existing corpora. Our research contributes to the NIME community in three ways: it provides a novel open-source Factor Oracle implementation for SuperCollider, enabling tight integration with real-time audio processing; it demonstrates corpus-free learning that constructs all knowledge from the current performance session, proving particularly suited for feedback-based music; and it documents the development of a feedback-augmented guitar prototype alongside empirical insights into the performative and collaborative dimensions of human-AI musical interaction.

While OTIAC has been specifically tailored for the feedback-augmented classical guitar setup described in this paper, its architecture is deliberately modular: the audio feature extraction pipeline, SOM clustering, and Factor Oracle are loosely coupled components, each configurable independently. Adapting the system to other feedback-based or acoustic instruments would primarily require adjusting the resonance frequencies to match the target instrument's modal characteristics and tuning the segmentation and feature extraction parameters to its timbral profile.

The practice-led research with the H[t] Duo suggested critical design considerations for co-improvisation systems. Awareness asymmetry between performers is not merely technical but epistemological: interface choices may create fundamentally different experiences of agency, suggesting these should be deliberately considered in the design process. In this context, predictability appeared to exist on a spectrum rather than as binary opposition, and successful systems may benefit from layering components with different degrees of controllability, allowing performers to navigate between stable and emergent behaviors according to musical needs. Embodied practice appeared central even in this highly mediated system, particularly with feedback-based instruments where performers can develop diachronic mastery through systematic exploration and mapping of physical-acoustic behaviors. Future research could further explore these preliminary observations through extended practice-led studies involving multiple performer pairs as well as different feedback instrument configurations, allowing for richer comparative insights into how awareness asymmetry and distributed agency manifest across different human-AI co-improvisation contexts.

The system's current focus on feedback-based material, while aesthetically coherent, was deemed restrictive over extended durations, suggesting future development. The system architecture already supports gestural expansion since any sound captured by the contact microphone enters the learning pipeline. However, performance practice has thus far privileged feedback exploration over incorporating traditional guitar techniques. The Duo's premiere of OTIAC in a public concert given at CIRMMT in Montréal (February 2026),[11] while the performance was well received, confirmed that the current focus on feedback-based material benefits from extended durations; longer future performances offer opportunities to develop more varied approaches integrating diverse instrumental techniques.

Future technical development could explore Reinforcement Learning (RL) approaches to enable the system to not only learn sequential patterns but also optimize its generative strategies based on performer responses and musical context. An RL framework could enable the oracle to learn higher-level performance strategies: while SOM and FO handle microform (event clustering and sequential patterns), RL would act as a decision-maker governing macroform. Practically, an RL agent could learn to adjust generation parameters (sequence length, continuity, timing) or modulate the oracle's navigation strategy through its state space based on explicit feedback from P1. For instance, the guitarist could use a foot pedal during performance to signal approval or disapproval of the system's current behavior, providing real-time rewards or penalties. By continuously updating its policy based on this feedback, the agent could adapt its behavior over time to better suit the performer's preferences and needs, developing sensitivity to successful versus unsuccessful interventions and gradually shaping its generative behavior toward patterns that meaningfully engage with P1's actions.

The complete system implementation and documentation are available as open-source resources, inviting other researchers and performers to adapt OTIAC for their own feedback-augmented instruments and co-improvisation contexts.

7 Ethical Standards

This work is made with the support of the French National Research Agency in the framework of the project MICCDroP (ANR-24-CE38-5860). The code developed for the project is released as open-source software, guaranteeing reproducibility. No human participants have been involved in the research besides the H[t] Duo. No conflicts of interest are reported by the authors.

References

[1] Cyril Allauzen, Maxime Crochemore, and Mathieu Raffinot. 1999. Factor oracle: A new structure for pattern matching. In *Proc. of the 26th Seminar on Current Trends in Theory and Practice of Informatics (SOFSEM'99).* 295–310.

[2] Jaime Arias, Myriam Desainte-Catherine, and Shlomo Dubnov. 2016. Automatic construction of interactive machine improvisation scenarios from audio recordings. In *Proc. of the 4th International Workshop on Musical Metacreation.*

[3] Gérard Assayag, Georges Bloch, Marc Chemillier, Arshia Cont, and Shlomo Dubnov. 2006. Omax brothers: a dynamic yopology of agents for improvization learning. In *Proc. of the 1st ACM Workshop on Audio and Music Computing Multimedia.* 125–132.

[4] Gérard Assayag, Laurent Bonnasse-Gahot, and Joakim Borg. 2022. Cocreative Interaction: Somax2 and the REACH Project. *Computer Music Journal* 46, 4 (2022), 7–25.

[5] Gérard Assayag and Shlomo Dubnov. 2004. Using Factor Oracles for Machine Improvisation. *Soft Computing* 8, 9 (2004), 604–610. https://doi.org/10.1007/s00500-004-0385-4

[6] Christian Blandhoel. 2021. Stadsbiblioteket 2021. https://myfirstrecordlabel.bandcamp.com/album/st-ybiblioteket-2021-pdf-book-catalog

[7] Benjamin D. R. Bogart and Philippe Pasquier. 2013. Context Machines: A Series of Situated and Self-Organizing Artworks. *Leonardo* 46, 2 (2013), 114–122. https://doi.org/10.1162/LEON_a_00525

[8] Joakim Borg, Gérard Assayag, and Mikhail Malt. 2022. Somax2 a Reactive Multi-Agent Environment for Co-Improvisation. In *Proc. of the Sound & Music Computing Conference.*

[9] David Borgo. 2006. Sync or swarm: Musical improvisation and the complex dynamics of group creativity. In *Algebra, Meaning, and Computation: Essays Dedicated to Joseph A. Goguen on the Occasion of His 65th Birthday.* Springer, 1–24.

[10] David Borgo. 2016. The Ghost in the Music, or the Perspective of an Improvising Ant. *The Oxford handbook of critical improvisation studies* 1 (2016), 91–111.

[11] Filippo Carnovalini and Antonio Rodà. 2020. Computational creativity and music generation systems: An introduction to the state of the art. *Frontiers in Artificial Intelligence* 3 (2020), 14.

[12] Nicolas Collins. 2021. Improvising with architecture: Pea soup and related work with audio feedback. *Resonance: The Journal of Sound and Culture* 2, 2 (2021), 168–181.

[13] Arshia Cont, Shlomo Dubnov, and Gérard Assayag. 2007. Anticipatory model of musical style imitation using collaborative and competitive reinforcement learning. In *Proc. of the Workshop on Anticipatory Behavior in Adaptive Learning Systems.* 285–306.

[14] Tom Davis and Laura Reid. 2020. Taking back control: Taming the feral cello. In *Proc. of the International Conference on New Interfaces for Musical Expression.* 416–421.

[15] Agostino Di Scipio. 2003. 'Sound is the interface': From interactive to ecosystemic signal processing. *Organised Sound* 8, 3 (2003), 269–277.

[16] Agostino Di Scipio. 2022. A Relational Ontology of Feedback. *ECHO: a journal of music, thought and technology* 3 (2022).

[17] Shlomo Dubnov, Gérard Assayag, and Arshia Cont. 2007. Audio oracle: A new algorithm for fast learning of audio structures. In *Proc. of the International Computer Music Conference.*

[18] Ken Déguernel, Emmanuel Vincent, and Gérard Assayag. 2018. Probabilistic Factor Oracles for Multidimensional Machine Improvisation. *Computer Music Journal* 42, 2 (2018), 52–66. https://doi.org/10.1162/comj_a_00460

[19] Ludvig Elblaus and Gerhard Eckel. 2025. utrumque: A Careful Approach to Shared Acoustics and Site-Specific Feedback Music. *ECHO: a journal of music, thought and technology* 7 (2025).

[20] Alice Eldridge and Chris Kiefer. 2017. The self-resonating feedback cello: interfacing gestural and generative processes in improvised performance. In *Proc. of the International Conference on New Interfaces for Musical Expression.*

[21] Alice Eldridge, Chris Kiefer, Dan Overholt, and Halldor Ulfarsson. 2021. Self-resonating Vibrotactile Feedback Instruments ||: Making, Playing, Conceptualising :||. *Feedback Musicianship Network* (2021).

[22] Rebecca A. Fiebrink. 2011. *Real-time human interaction with supervised learning algorithms for music composition and performance.* Ph. D. Dissertation. Princeton University.

[23] Neville H Fletcher and Thomas D Rossing. 2012. *The physics of musical instruments.* Springer Science & Business Media.

[24] Jonathan Foote. 2000. Automatic audio segmentation using a measure of audio novelty. In *Proc. of the IEEE International Conference on Multimedia and Expo.*, Vol. 1. 452–455. https://doi.org/10.1109/ICME.2000.869637

[25] Liane Gabora. 2002. Cognitive mechanisms underlying the creative process. In *Proc. of the 4th Conference on Creativity & Cognition.* ACM Press, 126–133. https://doi.org/10.1145/581710.581730

[26] Jiffer Harriman. 2015. Feedback Lap Steel : Exploring Tactile Transducers as String Actuators. In *Proc. of the International Conference on New Interfaces for Musical Expression.*

[27] Dorien Herremans, Ching-Hua Chuan, and Elaine Chew. 2017. A functional taxonomy of music generation systems. *ACM Computing Surveys (CSUR)* 50, 5 (2017).

[28] Matthias Jung. 2023. Intelligent Music Performance Systems: Towards a Design Framework. *Studia Musicologica Norvegica* 49, 1 (2023), 28–44.

[29] Teuvo Kohonen. 2012. *Self-organizing maps.* Springer Science & Business Media.

[30] Otso Lähdeoja. 2016. Active acoustic instruments for electronic chamber music. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* 132–136.

[31] Thanos Polymeneas Liontiris. 2018. Low Frequency Feedback Drones: A non-invasive augmentation of the double bass. In *Proc. of the International Conference on New Interfaces for Musical Expression.* 340–341.

[32] Otso Lähdeoja. 2008. An Approach to Instrument Augmentation: the Electric Guitar. In *Proc. of the International Conference on New Interfaces for Musical*

---

[11] https://www.cirmmt.org/en/events/live-cirmmt/13-feb-2026

Expression. 53–56.

[33] Thor Magnusson, Chris Kiefer, and Halldor Ulfarsson. 2022. Reflexions upon Feedback. In *Proc. of the International Conference on New Interfaces for Musical Expression.* https://doi.org/10.21428/92fbeb44.aa7de712

[34] João M. Martins and Eduardo R. Miranda. 2006. A Connectionist Architecture for the Evolution of Rhythms. In *Applications of Evolutionary Computing*, Franz Rothlauf, Jürgen Branke, Stefano Cagnoni, Ernesto Costa, Carlos Cotta, Rolf Drechsler, Evelyne Lutton, Penousal Machado, Jason H. Moore, Juan Romero, George D. Smith, Giovanni Squillero, and Hideyuki Takagi (Eds.). Vol. 3907. Springer Berlin Heidelberg, 696–706. https://doi.org/10.1007/11732242_66

[35] Adam Pultz Melbye. 2021. Resistance, Mastery, Agency: Improvising with the feedback-actuated augmented bass. *Organised Sound* 26, 1 (2021), 19–30. https://doi.org/10.1017/S1355771821000029

[36] Eduardo A. L. Meneses, Sérgio Freire, and Marcelo M. Wanderley. 2018. GuitarAMI and GuiaRT: Two independent yet complementary Augmented Nylon Guitar projects. In *Proc. of the International Conference on New Interfaces for Musical Expression.*

[37] Jérôme Nika and Marc Chemillier. 2012. Improtek: Integrating harmonic controls into improvisation in the filiation of OMax. In *Proc. of the International Computer Music Conference.* 180–187.

[38] Jérôme Nika, Marc Chemillier, and Gérard Assayag. 2017. Improtek: introducing scenarios into human-computer music improvisation. *Computers in Entertainment* 14, 2 (2017).

[39] Jérôme Nika, Ken Déguernel, Axel Chemla–Romeu-Santos, Emmanuel Vincent, and Gérard Assayag. 2017. DYCI2 agents: merging the "free", "reactive", and "scenario-based" music generation paradigms. In *Proc. of the International Computer Music Conference.* https://hal.science/hal-01583089

[40] Jérôme Nika, Diemo Schwarz, and Augustin Muller. 2025. Crafting Musical Agents: Interactive Generation as a Catalyst for Artistic Formalization. In *Proc. of the 6th Conference on AI Music Creativity.*

[41] Daniel Overholt. 2011. The overtone fiddle: An actuated acoustic instrument. In *Proc. of the International Conference on New Interfaces for Musical Expression.* University of Oslo, 4–7.

[42] Claudio Panariello and Chiara Percivati. 2023. "WYPYM": A Study for Feedback-Augmented Bass Clarinet. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* 174–179.

[43] Dario Sanfilippo. 2018. Time-variant infrastructures and dynamical adaptivity for higher degrees of complexity in autonomous music feedback systems: the Order from noise (2017) project. *Musica/Tecnologia* 11, 12 (2018), 121–31.

[44] Dario Sanfilippo. 2021. Time-domain adaptive algorithms for low-and high-level audio information processing. *Computer Music Journal* 45, 1 (2021), 24–38.

[45] Dario Sanfilippo. 2023. The Aesthetics of Musical Complex Systems. *Organised Sound* 28, 3 (2023), 381–391.

[46] Dario Sanfilippo and Andrea Valle. 2013. Feedback systems: An analytical framework. *Computer Music Journal* 37, 2 (2013), 12–27.

[47] Diemo Schwarz. 2007. Corpus-based concatenative synthesis. *IEEE signal processing magazine* 24, 2 (2007), 92–104.

[48] Benjamin D. Smith and W. Scott Deal. 2014. ML.* Machine learning library as a musical partner in the computer-acoustic composition flight. In *Proc. of the Joint Conference ICMC12-SMC14*, Vol. 2014.

[49] Koray Tahiroğlu. 2024. Musical intra-actions with digital musical instruments. *Journal of New Music Research* 53, 1-2 (2024), 126–138.

[50] Kıvanç Tatar and Philippe Pasquier. 2017. MASOM: A musical agent architecture based on self-organizing maps, affective computing, and variable Markov models. In *Proc. of the International Workshop on Musical Metacreation.* 5281.

[51] Kıvanç Tatar and Philippe Pasquier. 2019. Musical agents: A typology and state of the art towards musical metacreation. *Journal of New Music Research* 48, 1 (2019), 56–105.

[52] Notto J. W. Thelle and Philippe Pasquier. 2021. Spire muse: A virtual musical partner for creative brainstorming. In *Proc. of the International Conference on New Interfaces for Musical Expression.*

[53] Pierre A. Tremblay, Owen Green, Gerard Roma, James Bradbury, Ted Moore, Jacob Hart, and Alex Harker. 2022. *The Fluid Corpus Manipulation Toolbox.* https://doi.org/10.5281/zenodo.6834643

[54] Halldór Úlfarsson. 2018. The halldorophone: The ongoing innovation of a cello-like drone instrument. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* 269–274.

[55] Cathy Van Eck. 2017. *Between air and electricity: microphones and loudspeakers as musical instruments.* Bloomsbury Academic.

[56] Federico Visi. 2024. The Sophtar: a networkable feedback string instrument with embedded machine learning. In *Proc. of the International Conference on New Interfaces for Musical Expression.* 142–148.

[57] Cheng-i Wang and Shlomo Dubnov. 2014. Guided music synthesis with variable markov oracle. In *Proc. of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment.* 55–62.