xCoAx 2026
Torino, Italy
2026.xCoAx.org

Hostile Interferences Against Improvisational Predictability

Kevin Blackistone
kevin@blackistone.com
Independent Scholar

Within the world of the built environment, there is a term known as *hostile architecture*. This refers to designs purposed with preventing individuals from utilizing a space beyond desired intention, or in manners deemed unappealing. Often this manifests as forms of physical barriers or protrusions, hindering common uses such as rest or activity for all out of apprehensions of the perceived misbehaviors of few. A permutation of this adversarial approach is as well used as a means to train machine learning (ML) algorithms by providing a negative to prove the quality of a positive. Just as the communities around these hostile architectural designs find ways to move in and around them, so to do these algorithms find their way to a relative/desired truth. This research envisions an array of hostile/adversarial software architectures as a design foundation to shape performance and force or direct performers into new patterns by preventing, modifying or penalizing those behaviors common to them through the use of interactive machine learning (IML) strategies. By integrating this adversarial history with taxonomic analyses of musical performance, it envisions an array of potential hardware and software interventions to musical improvisation.

1 Background

"Hostile architecture" is an urban design strategy that uses elements of the built environment to purposefully guide behavior through restriction or interference. Typically, this is in regards to a perceived undesirable behavior or public nuisance. The interventions manifest in a number of formats that can include spikes, barriers and angled structures designed for discomfort. In its penalization of behavior it neither takes account of root causes, nor attempts any address thereof. As something which takes its forms within the realm of public (or public-use private) space, it tends to extend its effects beyond those targeted. This can be found in wildly uncomfortable park benches, an intentional lack of protective bus shelters, or the removal of seating from train stations – each exists to prevent the comfort of loiterers or sleep of the homeless, but no less prevents the civic enjoyment of those for whom they are ostensibly designed, while further interfering with elements of accessibility. Often, through their attempt to remove a given activity, they instead accentuate it through their visual presence, even in the absence of the given behavior.

1.1 Hostile Architecture as Artistic Approach

These interventions are at best a nuisance to those targeted, if not a direct frustration to all others. Yet, it is easy to see how these same at-

Keywords: Improvisation, Adversarial Design, Machine Learning, Interactive Training, Hostile Architecture.

DOI 10.34626/2026_xcoax_017

tempts to shape behavior frequently fail at their purpose through the inherent improvisations of the public in their interest in common use of public space. Many novel solutions can be found where-in these designs have been overcome. This includes bedding or planks over spikes (Sorrel 2015), artistic designs such as Sarah Ross's *Archisuits – fashion for hostile architecture* (Newman 2024), or the movement improvisations of *Fakir's Rest* by Stéphane Argillet and Gilles Paté (Argillet and Paté 2003).

This paper takes this human improvisational re-contextualization as a metaphor expanded to software architecture. There are already coined terms for contrarian design – "dark patterns" and "nag-ware" being two examples – to express software designed to interfere with comfortable and intended uses. These antagonistic designs form the basis from which to explore the use of software to interfere with musical intent, thus forcing improvisational explorations.

While performance often must adapt to any number of hindrances – from equipment failures to physical ailments, to unpredictable weather – this research proposes interference as intent. By presenting a barrier against the performer, it is required that their technique must change. By then incorporating modern computational approaches, the natural inclination to develop new habits can be discouraged by a progressive computational learning of new patterns in order to thus update encumbrances. While the incorporation of computationally controlled interference has already been explored in works such as *The Sabotaging Piano* (Dannemann and Bryan-Kinns 2023), they do not respond to the performer. They instead simply provide new hindrances to overcome. By incorporation of machine learning, one can produce a cyclical relationship between training, improvisation and interference.

1.2 Teaching the Student: Training the Machine

In isolated approaches to performance learning, we deal only with our cognitive layers and its various states of self-awareness. When we incorporate an instructor, things become more complex. What then, if the instructor is the instrument – and further, what then when the instrument is itself learning from the performer? As discussed in the next section, several attempts at producing instruments that aid as instructional tools have been designed, but these differ not only in the intent of aid rather than impediment, but in that they do not restructure themselves based on their use. Thus, there is no variation in their relationship to the hierarchies of learning any more than a traditional, non-responsive instrument would provide simply through its sound.

In the interactions generally built between a performer and instrument there is a direct, if non-linear, relationship from practice to result and a simple feedback cycle – first in private rehearsal via on one's own self-assessment, next in instructed practice and finally in performance via the assessment of the audience. There is no response from the device as to what the performer should or should not do. The addition of any such layer yields another node of interaction. Within this four-point interaction, there is no longer a pure cycle; as each additional point expands the exponent and affiliated intersections or dimensions (Fig. 1). This specific extra node exerts itself as a negative influencer, however. Instead of reinforcement towards success, it actively fights towards failure. Perhaps if, as David Zicarelli writes, “failure tends to be far more interesting than success” (Cascone 2017), this tension might be used to produce a more intriguing performative structure through the cross-instruction between the performer and the performed.

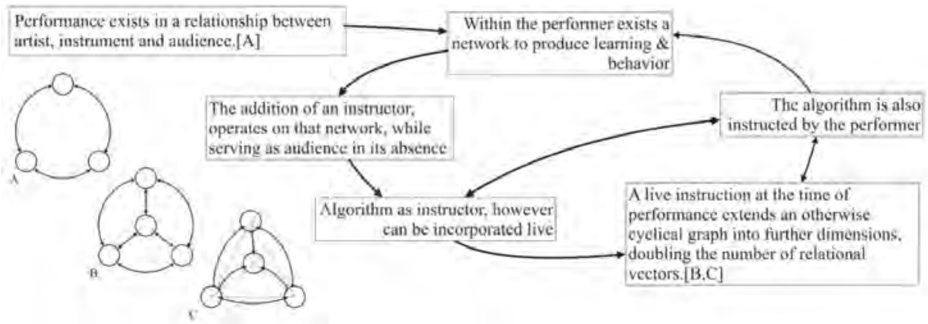

Fig. 1. Networks of performance and learning interaction when instruction is included concurrent to performance.

2 Computer Aided Learning and Instruction

Computational and interface-based approaches have been repeatedly implemented to improve a learner’s skill at an instrument such as those that work measuring only the student’s performative skill (Huang et al. 2008), those providing guiding force to the user (Fujii et al. 2015; Zhang et al. 2019) and those which use a combined approach to detect performance error while providing haptic or visual cues to inform on possible error (Chin and Xia 2022). These approaches are thus-far based on interpretations of a composition and can only detect how correct a performer is in the domains of timing and note accuracy. As such, they are incapable of learning to interpret an improvised performance. These further lack considerations of complex instruments, let-alone multi-instrumentations of collective performance. For such, an algorithmic learning approach is required. These prior studies do inform of two useful components to this research. The first is means to measure and interpret performed notes. The second is inspiration on potential mechanisms to guide the performers movement. (On this second point, the devices thus-far produced to aid learning might even be better suited to this antagonistic approach given their frequent physical design interferences.)

2.1 Machines Learning Music

Much of the uses to date within the realm of sound and music IML has sought to replicated that for which it has been used within the visual and text fields. This can easily be seen in the close analogs between text-to-image generation and equivalent acoustic mimicries of music by genre and text description. As well, the large language model approaches that have been used to produce blocks of writing can be alternatively tweaked, trained and redeployed to produce generative musical scores.

Sound / song recognition algorithms (such as rights management, song identifiers and word recognition) equivocate themselves with with visual identification algorithms both in style and often in use-cases. Sound style transfers such as Qosmo's Neutone (Fyfe and Mitcheltree 2022) take an audio input and translate it to sounds of trained instruments through Rave (Caillon and Esling 2021) in a manner not unlike image style transfers. The intent of these works is rather not to use prior invention to assist in the creation of like content, but to take the prior efforts and prevent their continuation within live performance.

2.2 Antagonism as Play

While machine learning is regularly promoted as a positive technology; with claims to erase biases and prevent human error (and as many critiques as to how it exacerbates both) this research aims for the reverse, by exploring it as direct aggression against the performer, allowing a novel means to explore improvisation in a manner resembling competitive sport. Performance may have its built-in feedback mechanism in the form of the audience, but the nature of courtesy prevents its common use in the negative. The strategic yawn or booing by the audience is rarely witnessed to provide such – more often it is only the visual disinterest or exit of the crowd. The IML in this case provides a stern judge that can give a direct response to the lack of creative variation in a given performance. All this again belies the fact that it is the musician who is training the machine at all times on what and how to judge. It is in this (im)balance of teaching:training that a tension can occur. The performer wants to play what is unexpected by the machine, as the machine learns how the performer tries to avoid prediction. Adding to these, is that predictions will yield false positives (and/or correct predictions, but with incorrect reasoning), while the musician, unable to determine an accurate or inaccurate match, might attempt to further shift their performance style. This offers the difficult possibility that both performer and algorithm may, over time, neither converge nor diverge, but produce merely an incoherent dialogue of data.

3 Taxonomies of Performance

In considering how to interfere with musical performance it was necessary to produce simplified taxonomic breakdowns of the means to produce sound. These classifications provided the basic framework for exploration and are not intended to be exhaustive. Their production was approached from two angles. The first (Fig. 2) considers types of motions typical to the human player. This framework was developed to understand both the means to register and the means to manipulate the performance from the standpoint of the performer's body.

The second (Fig. 3) is instrument-centric, addressing the form of actions performed upon them. The basic separations made here are primarily going to be defined by the format of instrument and will thus

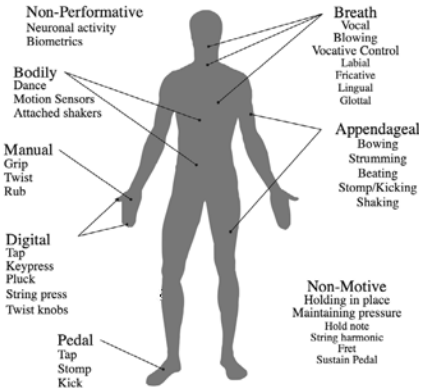

Fig. 2. Performance possibilities by anthro-mechanical action.

somewhat follow a typical instrument breakdown. This charting gives a simple overview of the network interactions as well as the kinds of universalize-able approaches might be used in control or registration of an instrument itself.

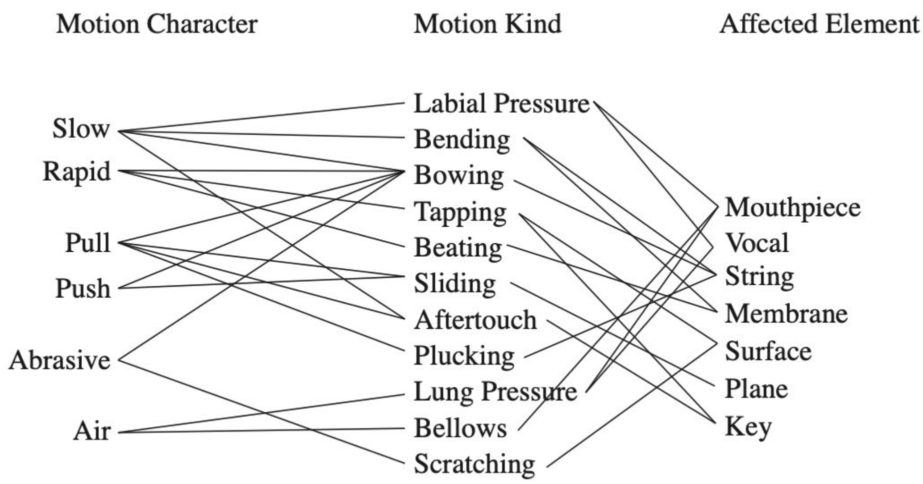

Fig. 3. Performance motion characteristics with type of motion and affected instrument elements

4 Performance Learning Strategies

The prior taxonomic investigations set the groundwork for what styles of algorithmic learning and intervention can take place. The primary, presumed means of interaction with the learning model would be one wherein the model continuously trains and responds through a performance, such that there is an interplay between the performer and the algorithm, each 'trying' to thwart the other. The algorithm, in that it is directly intended to prevent what it has learned as the performer's habits, and the performer, in an attempt to find styles which do not trigger the antagonistic response. As such, as a performance continues to train the model, it becomes harder for the performer to improvise something that is neither pure randomness, nor too predictable.

In an alternative approach, the model could be trained only on prior performance. This is the more common method for machine learning, but does not permit the interplay here desired, however there might be the opportunity for the artist to learn what it predicts over the duration of performance. In this way, the performer may start incapable of performing regularly but be able to play in a new fashion semi-fluidly by the end. To keep a versatile approach, each performance could be learned in succession, such that each sequential performance is iterative of the prior(s). Beyond the somewhat lesser dynamics of interplay that the live learning would possess, the dynamic here could be played out over a much longer time frame through a performance series. Additionally, a pretrained model is capable of being retrained in performance with sufficient planning and thus allow prior performance to affect the current while still adapting to present variation.

5 Software Architectures

This section suggests code-based approaches to the learning strategies, as they are both simpler to implement than the later hardware approaches while also forming their foundation. It must be considered that the need to predict, compare (and in some cases train) before sounding or responding introduces latency. One must compromise algorithmic robustness against predictive delay, with sufficiently simple predictive models minimizing discrepancy.

5.1 Note Alterations

The principally considered means to intervene is via note sequence determination. The most basic of intervention with which is the prevention of an intended note by learning from the performed sequence on a MIDI scale to then disable note-on of the next predicted note or pitch class (with potential considerations for polyphonic input).

Expansions on this theme includes the use of weighting (expanded in section 6). In applying weighting as a randomizer, a prevented input may allow variation beyond that deemed *most* likely. A cut-off threshold for top-weighted predictions can be included to avoid excessive fluke penalization. For example, of seven base pitches, none below 10% probability would mute, with those above having adjusted mute probabilities (Fig. 4). As such, the most likely note is most likely muted, but others may be instead or as well. This allows for resiliency should two notes be equally determined to follow. A second implementation uses the weighting as a means of volume control such that, while the top option may be muted, the next most likely might be merely quieted as a function of its likelihood.

Fig. 4. Predicted most likely note probabilities (left) with cutoff for top results above 10% (dashed line) adjusted to next note mute probability (right).

note	**c**	**d**	**f**	**g**	**a**	**b**	**e**
predicted chance	63%	19%	11%	3%	2%	1%	1%

note	**c**	**d**	**f**
mute chance	68%	20%	12%

Rather than blocking/muting, upon performance of a given note, the algorithm could instead swap that note to the another note less or more expected. This could again use an inversion of the prediction weight. The result would then lead to live recomposition (and possible confusion of the performer) not unlike the Sabotaging Piano, but with constant variation to the reconfiguration of the notes.

5.2 Time Domain

Expanding from pitch, the rhythmic component can be used alone or in concert with concepts such as pitch. The clearest area in which the time domain becomes of primary concern is in the tracking of rhythmic performances, although something that incorporates the event/note sequence might still be relevant to the series of drums in a kit rather or tempo of melodic sequence. An algorithm here might use above techniques of silencing based on a time expectation rather than sequential selection. A more advanced approaches would incorporate time and sequential expectation.

6 Instrument-Centric Interventions

This section enumerates and categorizes the application and expansion of the algorithmic concepts to physical devices and provides some designs that, while software-based, cannot exist without the hardware component.

6.1 Discreet Interferences

Simple blocking or mutes provide the simplest option for any performative element of the instrument and a direct analogue to the software-based systems first discussed. Anything pressed or hit could in some manner be blocked. One can expand this to an array of blocked elements. Piano keys can easily incorporate mechanisms to prevent the playing of any note deemed sufficiently likely, while other instruments such as a woodwind's pads can be controlled with slightly greater complexity. There have already been developed strategies to suggest movement to aid in performance learning (Zhang et al. 2019). Such devices could easily be modified at the algorithmic level to instead suggest against, or prevent, these motions. Some instruments, such as strings or drums do not lend themselves to a means of fully stopping the motion, however applied pressure through a padded device may instead minimize vibratory potential on the relevant component. In many cases this can be as effective as prevention, although in others it merely diminishes the sound while changing its timbral characteristics.

It should be noted that instrument itself does not need to change, so much as its performance could be simply interfered with. Such formats might include a drum kit that moves the kit pieces in, out, up or down as determined by performance. Alternatively, elements might be designed to get in the way or annoy generally the performer. These dovetail in many regards with the later distractive responses (section 8).

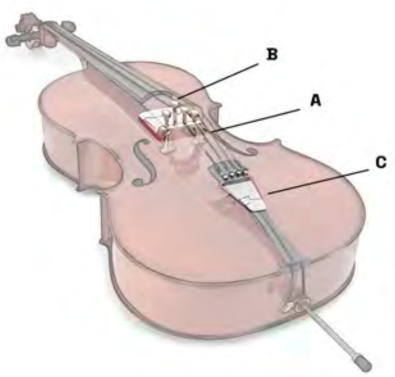

Fig. 5. Performance of strings (A) may be prevented by the intervention of mutes before the bridge (B). Tonal modification and interference could be adjusted through controlled gearing of secondary tuning pegs (C) – leaving the primary tuning intact.

6.2 Variable Interference Paradigms

Analogue Responses

There is an array of potential means to manipulate in less binary formats than the above. Rather than preventing or minimizing a note, string tension adjustments could create notes intentionally against the performers wishes (Fig. 5). Such a device would not so easily be able to perform a note based on its lack of expectedness, but could instead interfere with the performers melodic sense as an affected played element will not be that which is intended. The degree to which such a strategy could shift the pitch would depend greatly on the instrument and would be technically quite challenging.

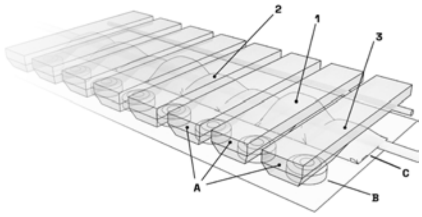

Fig. 6. Variable resistance blocking. Keys (A) predicted with higher probability require greater pressure to play depending on inflation (C) of numbered pockets to reach contact triggers (B). Here pocket 1 is most inflated, indicating greatest likelihood of performance and thus greatest difficulty in playing.

Simpler than pitch adjustment can be volume control. By expanding the use of muting, a variable pressure might be applied either to the sounding element or the performed key (Fig. 6) providing a tangible feedback as to what the algorithm believes will be played. Through additional effort, the performer may still force a desired note. Tactile feedback, such as vibration, is another option, permitting the performer awareness to change without pushing through variable resistance forces.

Stochastic Responses

Extending prior software-based discussion (section 5.1) – via the predicted likelihood for all notes and upon the playing of a note, the one with an inverse likelihood would instead be played. In theory this could be implemented in hardware as well, but only for very few instruments. The hilarity implicit in the complexity of a piano set to regularly shift hammers to alternative strings is exceeded only by the comedic simplicity of a trombone slider constantly trying to move to its least likely position.

The inferred notes may as well incorporate random chance with the trained weighting with timing offsets using a randomizer based on probability of the note at that time (ie. the less likely the performer is to play at a time, the more on-time that performed note is). This could be determined by event timing, specific note timings, or the timing in relation to a melodic/harmonic sequence. When considering the performance feedback elements of this approach, the existence of a low-chance note being adjusted, when the performer would not expect it, could lead the performer to infer a level of predictive dimensionality to the algorithm based on surprise – especially if the offsets are allowed to be especially large.

7 Performer-Centric Interventions

7.1 Body Physicalities to Interpret

Rather than affecting the instrument, one might restrict or alter performer motion directly as extrapolated from Fig. 2. A range of motion scales might be affected from the large-body movements of the waist/spine (commonly carrying dramatic action rather than sound) to limbic to digital. Movement of the limbs are often a primary driver for percussive instruments, as well as the rhythmic or long motion elements of other instrument (i.e. violin bowing, guitar strumming). Limb-movement frequently carries a pitch character (i.e. ascending/descending piano keys or guitar's frets) as well. Thus, manipulation of the limbs can create a wide range of affects, dependent on the instrument in use. Smaller motions (including digital, labial and lingual) cover an array of musicalities. Digital motion is commonly associated with the playing of a range of keyed instruments and includes the pitch control of strings.

Many simple, though often quieter, percussive motions are thus produced as well. Interference with the movements of the fingers is the most direct option for keyed instruments, although not without technical difficulty.

The versatility of oral control is quite broad. Labial and lingual controls provide many of the rhythmic and tonal qualities of breath-based performance instruments. Breath, even when not directly producing an element of sound, as in brass or woodwinds, is integral to all styles of performance, having heavy influences on tempo and musician temperament. It's potential for manipulation is general limited to how much it might be restricted.

7.2 Means to Register

The array of techniques available to ingest various human actions is complex. There is not yet any full-body action determination which is not in some manner unreliable or a burdensome hindrance to a live show. Some basic inputs include the simplistic if unreliable, such as stretch sensors, or a wide array of existing movement sensors, gloves, and other (often bulky) wearables. Per-instrument option may be incorporated through touch or pressure at the point of contact and, although such designs might be the most relevant to the sound, they ignore the range of motions by other elements of the body bringing that motion to bear. Electro-muscular sensors can extrapolate performed motion in a fairly slim footprint (Fig. 7) with greater algorithmic complexity. Such have demonstrated reasonable discernment of a range of motions with limited sensor count but require pre-training to learn specific motion. (Jiralerspong et al. 2017).

Pose estimation provides a more modern approach through analysis of video by trained algorithms to infer the locations of joints, freeing the musician from physical devices. These algorithms are only recently becoming fast enough to collect near real-time information and their models are typically trained for specifics scales (i.e. limbs vs fingers). They are also likely to falter when faced with the capacity of most instruments to get in the way visually. In the future, individually trained models could be produced that acknowledge or track the instrument.

Fig. 7. A proposed performer-based combination device sensing muscular movement through electrode-based EMG and controlling finger motion via wire tension

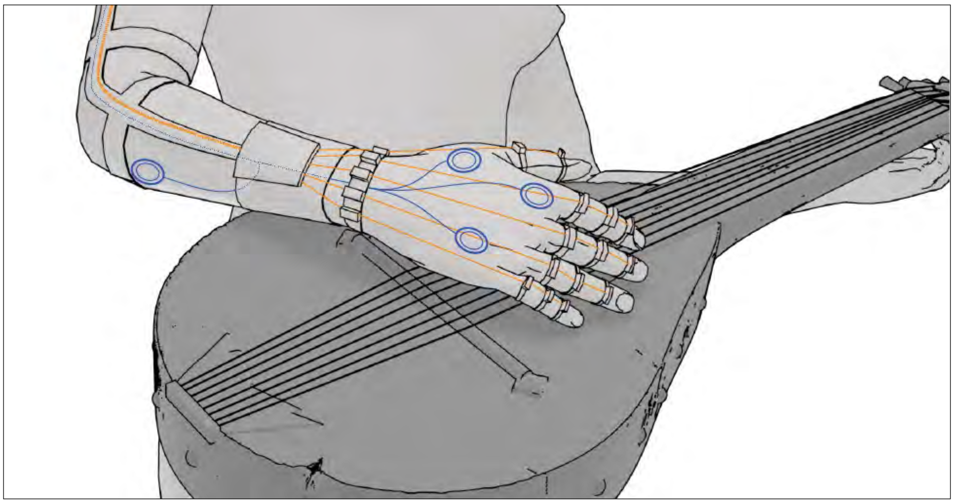

7.3 Means to Control

The greatest difficulty in considering components affecting performer, rather than instrument, is most require some elements of bulk. A number of bio-mechanical interfaces exist, having been used in artistic performance such as those of Stellarc (Garoian and Gaudelius 2001). Pullies, motors, pneumatics and the like can all be rigged to control the potential human movement. Machinery attached to the performer may look interesting but fail at generalizable performance across any array of potential instrument.

One strategy involves resistance through electro-muscular stimulation. A properly controlled voltage, sent through a given muscle, results in contraction. An assortment of these have been use performatively for involuntary motion and expression (Elsenaar 2009). This strategy would then produce contractions deemed by algorithm as counter to performer intent (with potentially unsafe results). All options, mechanical or electro-muscular, have limitations in the amount of hardware one can attach to the body, especially at the scale of small digits, and still have a reasonably capable or unencumbered performer.

8 Generalized Interferences

Prior examples were based on specific physical motions or instruments. One of the longest implemented forms interference in pursuit of behavior modification is the use of a punitive response. This may be doled upon determination that the performer behaves as presumed by the system. The deeper considerations in this frame are more in the array of outputs which might be used to inflict penalty (and the ethical considerations of each), rather than which performance inputs might be interfered from. The most gentle of these being that of visual indicator such as a light's activation or screen-based text informing of what elements are predictable. More belligerently, auditory indicators can be used – with the added benefits that these interfere with the performed sound and can affect the audience. More assertive are penal techniques through electronic or mechanical devices such as whips, pokes or electroshock. Such devices can be controlled with variability based on a number of factors, primarily prediction confidence.

There exist performatively amusing alternatives that include the distraction or confusion of the performer, expanding the sense of play. Instead of punishment, a tickling feather might be one example. Taking the concept of the auditory indicator; a buzz or screech may work well as punishment, but vocal samples yelling at the performer for laziness and predictability could induce any number of anxieties. To expand further one can engage volunteers who, given cues by the algorithm, might impose vocal or physical interventions upon the musician.

9 Example Result

The *Algorithmic Antagonist* is here presented in demonstration of one direction for generalized performance intervention. This attempted to learn from the performer via both audio and bio-signal analyses to make predictions on musical direction and participant motion. Should the predictions prove sufficiently accurate, impetus to change tactics was given via punitive electro-shock. The physical interface was fairly simple, consisting primarily of a wireless device to send electromuscular signals and to receive shock triggers. Triggers activated an off-the-shelf transcutaneous electrical nerve stimulation (TENS) fetish shock device. These were packaged in a worn vest with elastic armbands to hold the electrodes (Fig. 7). The analysis software consisted of MaxMSP (Puckette and Zicarelli 1991) with the Fluid Corpus Manipulation (FluCoMa) toolkit (Tremblay et al. 2021), Markov models via ML.* (Smith 2019) and regression models from ml.lib (Bullock n.d.) for the analysis, training and predictions. As this paper is focused on its theoretical framework and background, the specifics of the algorithms will not be described in-depth, however full details of its specific and technical development are to be available (Blackistone 2026).

9.2 In Operation

This realization has been presented several times including events with the Tangible Music Club (Linz), and the International Computer Music Conference (Boston) in a performance of pre-programmed electronically sequenced motifs manipulated by the performer, during which the following were observed. This musical style was not ideal for the interface as its complexity makes determination of accuracy difficult. While the punishments resulted in greater impetus to manipulate these patterns to a much more-so than typical, it was clear that there were false positives in prediction, such as a moment when the audio equipment failed registering as predicted.

Regarding responses from the audience, that performances were somewhat sequenced and beat driven made them more interesting from the audience perspective, due to their understanding that due to their inherent baked-in patterns, such improvisation would be likely more difficult to vary sufficiently against predictability. While the shock response felt to result in large, apparent motion, it was regularly stated that it was not visually significant from the audience perspective. This was compensated for via audio and visual signifiers such as tones, strobing lights and glitched visuals (iterated between performances), which were generally, understood as signifiers of punishment.

The device was successful in causing minimal physical interference from the worn components, however the performance of synthesizers is not as engaged as might be found in performance of instruments such as guitar or saxophone. In general, the result further successful in that audiences were very interested in the background theory and the operational techniques of the interface. From the perspective of producing an

Fig. 8. The vest and interface before premiere performance at STWST, Linz. (video available at: https://www.dorftv.at/video/45716) Further performances include The Red Room, Baltimore (video available at: https://www.youtube.com/watch?v=NTmER1qCT5E)

enjoyable performance, while seemingly successful, it should be considered that the audience at the selected experimental venues in which it was performed might more appreciate such masochism as manifested by machine intelligences than a *regular* audience.

10 Summary

This research presents a conceptual thread that began with considerations of the affects of hostile architecture on behavior to then make a game of forced improvisations in a musical, performative context. Over the course of its development, original connections to interactions with and improvisations around hostile architecture became less direct than at the project's inception. They are maintained here to better tell the story. That metaphor's extension to adversarial software design, provides a lens through which later threads were considered. These tie into what Selena Savić and Gordan Savičić suggest as, "The Algorithm of Unpleasantness," and indeed the presented result connects directly with their discussion of "Repellant systems" (Savić and Savičić 2014). As the ever-growing surveillance states and data accumulations are turned into tools of applied control – corporate, civil and military – we find ourselves forever in a tug of war over whether and/or how to change our behaviors to adapt to, push against or circumvent these new systems. Should such algorithmic self-antagonistic interfaces then become part of the self-educational toolkit for considered responsiveness to such systems? This does not go far enough to answer any such fundamental questions. Perhaps against a gamified system, a punishing play at performance and the affiliated conversations can be one small piece to draw attention to the problems and help identify novel means to address them.

10.1 Expansions

This research need not limit itself to the domain of music. Interfaces can be envisioned to interfere with any number of art forms – from painting, to sculpture, to, perhaps most amusingly, dance. Further, such interferences need not be applied only in performative or artistic contexts, while punishment as habit-breaking tool has a long-standing (if oft ineffective) tradition, it is little used outside of known-detrimental habits. There may be a place for learning applications to get in the way of any number of daily processes, from walking cadence to scribing journal publications.

Acknowledgments. This work is both condensation and expansion of the theoretical elements of thesis research produced at the University of Arts and Design Linz (Blackistone 2025). The author would like to thank Univ. Prof. Dr. Martin Kaltenbrunner, and Univ. Dr. Enrique Tomás for their knowledge and assistance in the realization of this research.

Ethical Standards. The used administration of electroshock via the TENS device for the production of this work was primarily upon only the author. Limited exposure to the device was granted colleagues only upon their request and only with full disclosure of any known health concerns.

Online Resources. The files used in the ongoing productions from this research can be found in the github repository at: https://github.com/blackistone/ImprovisationalAntagonist

References

Argillet, Stéphane, and Gilles Paté.
2003. "The Fakir's Rest." YouTube video. Accessed May 16, 2026. https://www.youtube.com/watch?v=Qv3M7FxJqtM

Blackistone, Kevin.
2025. "Hostile Architectures as Algorithmic Intervention in Musical Improvisation." Universität für künstlerisch und industrielle Gestaltung Linz.

Blackistone, Kevin.
2026. "A Live-Learning Punitive Interface for Performance Dynamics." In *New Instruments for Musical Expression (NIME) 2026*. In production.

Bullock, Jamie.
n.d. "ml.lib." GitHub repository. https://github.com/irllabs/ml-lib.

Caillon, Antoine, and Philippe Esling.
2021. "RAVE: A Variational Autoencoder for Fast and High-Quality Neural Audio Synthesis." arXiv preprint arXiv:2111.05011.

Cascone, Kim.
2017. "The Aesthetics of Failure: 'Post-Digital' Tendencies in Contemporary Computer Music." In *Electronica, Dance and Club Music*, 97–103. Routledge.

Chin, Daniel, and Gus Xia.
2022. "A computer-aided multimodal music learning system with curriculum: A pilot study." In *Proceedings of the International Conference on New Interfaces for Musical Expression*. https://doi.org/10.21428%2F92fbeb44.c6910363

Dannemann, Teodoro, and Nick Bryan-Kinns.
2023. "The Sabotaging Piano: key-to-pitch remapping as a source of new techniques in piano improvisation." In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 574–578. http://nime.org/proceedings/2023/nime2023_82.pdf

Elsenaar, Arthur.
2009. "Electric Eigen-Portraits." *Leonardo* 42 (4): 366–67.

Fujii, Katsuya, Sophia S. Russo, Pattie Maes, and Jun Rekimoto.
2015. "MoveMe: 3D Haptic Support for a Musical Instrument." In *Proceedings of the 12th International Conference on Advances in Computer Entertainment Technology*, 1–8.

Fyfe, Andrew, and Christopher Mitcheltree.
2022. "Neutone—Real-Time AI Audio Plugin for DAWs." https://qmro.qmul.ac.uk/xmlui/handle/123456789/88475

Garoian, Charles R., and Yvonne M. Gaudelius.
2001. "Cyborg Pedagogy: Performing Resistance in the Digital Age." *Studies in Art Education* 42 (4): 333–47.

Huang, Kevin, Ellen Yi-Luen Do, and Thad Starner.
2008. "PianoTouch: A Wearable Haptic Piano Instruction System for Passive Learning of Piano Skills." In *2008 12th IEEE International Symposium on Wearable Computers*, 41–44. IEEE.

Jiralerspong, Trongmun, Emi Nakanishi, Chao Liu, and Jun Ishikawa.
2017. “Experimental Study of Real-Time Classification of 17 Voluntary Movements for Multi-Degree Myoelectric Prosthetic Hand.” *Applied Sciences* 7 (11): 1163.

Li, Quang, and Bo Li.
2013. “Online Finger Gesture Recognition Using Surface Electromyography Signals.” *Journal of Signal and Information Processing* 4 (2): 101–5.

Newman, Nick.
2024. *Protest Architecture: Structures of Civil Resistance*. RIBA Publishing.

Puckette, Miller, and David Zicarelli.
1991. *MAX Development Package*. Opcode Systems, Inc.

Savić, Selena, and Gordan Savičić.
2014. “Unpleasant Design: Designing Out Unwanted Behaviour.” In *Proceedings of the 5th STS Italia Conference: A Matter of Design*. *Making Society through Science and Technology*, 975–88.

Smith, Benjamin Day.
2019. “ml.*.” https://www.benjamindaysmith.com/ml-machine-learning-toolkit-in-max/.

Sorrel, Charlie.
2015. “This Guerrilla Mattress Turns Evil Antihomeless Spikes into a Force for Good.” Fast Company, July 2015. Accessed March 1, 2025.

Tremblay, Pierre Alexandre, Gerard Roma, and Owen Green.
2021. “Enabling Programmatic Data Mining as Musicking: The Fluid Corpus Manipulation Toolkit.” *Computer Music Journal* 45 (2): 9–23.

Zhang, Yian, Yinmiao Li, Daniel Chin, and Gus Xia.
2019. “Adaptive Multimodal Music Learning via Interactive-Haptic Instrument.” arXiv preprint arXiv:1906.01197.