


KEVIN BLACKISTONE



**HOSTILE ARCHITECTURES AS
ALGORITHMIC INTERVENTION IN
MUSICAL IMPROVISATION**

Universität für künstlerische und industrielle Gestaltung
Kunstuniversität Linz Institut für Medien
Postdigital Lutherie

Hostile Architectures as Algorithmic Intervention in Musical Improvisation

Kevin Blackistone

Masterarbeit zur Erlangung des akademischen Grades
Master of Arts

Betreut von:

Univ. Prof. Dr. Martin Kaltenbrunner

Univ. Prof. Dr. Volkmar Klien

Univ. Dr. Enrique Thomas

Datum der Approbation: 23.01.2025

Unterschrift der Betreuerin/des Betreuers:

Jan. 2025, Linz

Table of Contents

Synopsis

Background

- 4 *Hostile architecture*
- 8 *Teaching the student : training the machine*
- 14 *Cultural histories of acoustic & algorithmic antagonisms*

Concepts & Thought Experiments

- 20 *Taxonomies of performance*
- 22 *Performance learning strategies*
- 23 *Software architectures*
- 27 *Instrument-centric interventions*
- 31 *Sample instrument modifications*
- 33 *Performer-centric intervention*
- 38 *Generalized interference*

Project Development

- 42 *Background*
- 45 *Software design*
- 53 *Hardware design*
- 56 *Aesthetics*

Output

- 63 *Results*
- 69 *Further research*

References and credits

Synopsis

Using music & performance as metaphor to explore the intrinsic societal complications imposed through unfettered use and cross-pollinations of hostile architecture and algorithmic control systems.

Alternatively:

An exploration of algorithmic approaches to prevention of intentional habits, there-by forcing new personal approaches to musical performance.

Abstract:

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. This adversarial approach is as well used as a means to train machine learning 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 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.

Background

Conceptual histories of hostile architecture

Machine and human learning

Adversarial design

Performative antagonism

Hostile architecture

History

Hostile architecture is an urban design strategy that uses elements of the built environment to purposefully guide or restrict behavior.



Figure 1: An anti-bandito hump or "Pissotte" in Venice. (Photo by author)

It might be more honestly stated that its specific intent is to wholly prevent behavior as it could hardly be said the elements provide any guidance to an alternative. For this reason it has acquired the current view of such interventions as *hostile*. Typically, this is in regards to a perceived undesirable behavior- often [rather problematically stated] as preventing an unwanted public nuisance, such as skateboarding or homelessness. Moreover, this form of design seeks to penalize a behavior without

any understanding of root causes or any attempt at address to said root causes. As something which takes its forms within the realm of public (or, quite often, overlapping public-use private) space, it tends to extend its effects beyond those targeted by the design. 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 prevent the civic enjoyment of those for whom they are ostensibly designed. A further critique can be found in that, in their non-address of root causes, their use merely displaces the action from specific locations (public and private) and into alternative (public and private) locations,

although many novel solutions can be found where-in these designs have simply been overcome. Often they, through their attempt to hide-through-prevention a given activity, they accentuate it through their perpetual presence even in the absence of the given behavior.

Examples

1. Spikes – One of the most common examples used. The inclusion of an array of spikes or studs on rails, ledges and ground surfaces is employed to prevent common uses such as sitting, laying or sliding.
2. Bench dividers – The separation of the long space is intended to minimize use to exclusively the upright seated position
3. Angled seating – These designs operate to minimize seated comforts of long-term use. This has been extended by corporate offices to toilet designs such as the *Slanty* to minimize the break time of workers.(Pinkser, 2019)
4. Corner deflectors – In use at least since the 19th century in England, and further back in Italy – these corner nubs serve a primary purpose as preventative measure against public urination. In Italy they have as well been considered a design to minimize the ability of muggers to wait hidden in corners.




Figure 2: Ledge spikes. (Paydah)



Figure 3: Bench rails to prevent horizontal relaxation. (Pinki, H.)

Obstacles overcome

Such designs have not escaped scrutiny. Neither have they been wholly useful. Several groups have taken to guerrilla approaches to simply, illegally remove these where possible through sawing off rails or unfastening of spikes. Many of those for whom these interferences are designed have been quite capable of navigating means to enjoy utilization of the space, whether it be through something as simple as placing a mattress upon spikes to sleep. Meanwhile, skaters have found designs targeting them as much a challenge as an opportunity to inspire new styles, demonstrating that designs discouraging one undesired behavior may inspire another. These same skaters having taken readily to grinding on angled benches intended to prevent long-term seated comfort.

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 attempts to shape behavior frequently fail at their purpose through the inherent improvisations of the public in their interest in

common use of common space. Here, I use this very aspect of human jazz-recontextualization as a metaphor within the field of software architecture. There are already coined terms for this form of contrarian design – “dark patterns” and “nag-ware” being two examples¹ – among other such terms that exist to express software design that exists to prevent the user from being able to comfortably use the software in the manner intended.

Given these software designs – intentionally contrarian to user intent – it is perhaps interesting that the synonymous term of the *adversarial* is used in an array of methods used to train machine learning algorithms. As will be explained in greater detail, this incorporates both attempts to find faults in learning models as well as using error for improvement of model training.

Between these, this work considers utilization of the concepts of the architectural through software-based interventions as a means to inspire alternative performative practice. Artists, musical and otherwise, frequently and understandably find patterns they like. Even within improvisation this comes from a learned response of which elements, sound and techniques engage both artist and audience.

There are often hindrances within any artistic approach. They require frequent, improvised work-arounds. Performatively, these can include equipment failures, physical ailments, or the unpredictable weather of an outdoor function. The responses can range from mild reworking of the planned performance to full cancellation of the function. What this work proposes is that the work’s intent *is* the interference. By presenting a barrier against the performer, it is required that they change their technique. What’s more, by incorporating modern computational approaches, the natural habit of defining new habits can be discouraged by a constantly updated learning of the new variations and providing constantly updating

1 These are listed here merely as example concepts. Much can be written about them and their further relations to the idea of the physically hostile structure. I leave them here without further explanation as each, especially “dark patterns,” require a more robust background than should be included in this overview to adequately address. A perhaps too-brief explanation is that dark patterns obfuscate intent through design to drive intended behavior, where-as nag-ware interferes with interaction the interruption. The reader is encouraged to explore further.

encumbrances. These can result in a mutual relationship of training, torture and inspired improvisation.

Teaching the student : training the machine

In a basic approach to learning we have multiple layers of cognition, often categorized in a hierarchical fashion ranging from those with which we are fully aware of to levels for which we lack adequate knowledge or consciousness to express in any but the most abstract of terms. Through each of these (non-definitive) levels, one can attribute specific learning elements. At the most low-level we are not (to our general awareness) learning to make our hearts beat, or our lymphatic system to flow. As we proceed through the hierarchy we can become more aware of what we are doing – and thus more able to adapt those levels. In, "The Emotion Machine," Marvin Minsky writes in references a prior diagram[Figure 4]

...self-reflection has limits and risks. For any attempt to inspect your thoughts is likely to change what you're thinking about. It is hard enough to describe a thing that keeps changing its shape before your eyes-and surely it is harder yet to describe things that change when you think about them. So you're virtually certain to get confused when you think about what you are thinking now-which must be one of the things that make us so puzzled about what we call "consciousness," (Minsky, 2006. p.145)




Figure 4: Summary of the organization of human mind (Minsky, 2006. p.29)

These complications expand when we incorporate any instructor into the paradigm, as this complexity becomes intertwined. 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 paragraphs, several attempts at producing instruments that aid as instructional tools have been designed (and infinitely more software packages, not discussed), 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 not 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. The feedback cycle is simple – relying on either the performer’s own auditory assessment in rehearsal, or the assessment (direct or interpreted) of the audience. There is no response from the device as to what the performer should or should not do. In the inclusion of such a layer, there is an addition point of interpretation, which itself yields a further point of failure. With a four point interaction, there is no long a pure cycle as the exponent of any additional point(s) require(s) intersections or dimensions. [Figure 5] 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,” (Zicarelli, 1999, from Cascone, 2002) this tension might be used to produce a more intriguing performative structure through the cross-instruction between the performer and the performed.




Figure 5: **[A]** A simple feedback cycle as per performer, instrument, audience. **[B]** An additional node, such as an instructor, doubles the edge count. **[C]** Represented then in tetrahedral 3-space. Further influences operate as additional variables/dimensions, increasing edges exponentially.

Computer aided learning & instruction

Learning to play an instrument is intrinsically multimodal. It usually involves learning music notations via the visual system, memorizing the tones via the auditory system, and mastering the performance skills via the motor system. Though visual and auditory interfaces, such as sheet music and recordings, have long been used to assist music learning, haptic interfaces that are able to reproduce motion expressions have just been invented in the recent years. ... In general, haptic interfaces offer guidance via tactile or kinesthetic perception: tactile perception is vibrations or pressure conveyed through the skin, while kinesthetic perception is receptors in muscles and tendons that allow us to feel the pose of our body. (Zhang, 2019)

Several attempts have been made to using a combination of computational and interface-based approaches to improve a learner's skill at an instrument. These designs have included those that work only to measure the student's performative skill, such as on the piano (Huang, 2008), those providing guiding force to the user (Fujii, 2015. Zhang, 2019) and those which use a combinatory approach to detect performance error and use haptic and visual cues to inform, correct, or hint to prevent possible error. (Chin & Xia, 2022) These approaches are thus-far based on interpretations of a composition. They 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 – or even one which is pre-composed, but for which it has not in some way ingested the score – leaving out further considerations that these prior references as-yet lack considerations of complex instruments, let-alone multi-instrumentation. For such, a machine 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, as will be noted in the section on instrument-based interventions.

Machine learning

Machine learning is a branch of computational science that focuses on algorithms which produce a result that varies based on some form of learning from assigned inputs. These are best known, at present, for the creation of several variants of text-to-image generation and chat bots, but the field is far more varied and goes back at least as far as research in the 60's using punch-tape to attempt to analyze sonar signals, electrocardiograms, and speech pattern². While generally considered as a multi-dimensional neuronal modeling approach to prediction, it may also incorporate any number of learning, regressions and analysis models.




Figure 6: Representation of a simplified ANN network. (Glosser.ca) Such designs would include several more hidden layers, and likely a number more inputs and/or outputs.

A more recent and common approach is the artificial neural network (ANN). These attempt to (in various capacities) simulate the networks of neurons in an organic system.³ This primarily views those neurons in an interlinked nodal configuration.⁴ In this field several layers of inter-connected nodes take a number of inputs to predict a number of outputs.(Figure 6) The weighting of these nodes connections at each layer changes as the network is trained. As these layers and node counts increases it eventually becomes what is known as *deep learning*. One of the known difficulties of the machine learning field, especially as this depth increases, is that once these intermediary nodes are trained and weighted, it is nearly impossible to look inside to determine how the model

² Machine Learning, Wikipedia 2023

³ More modern approaches to simulate the organic brain and some of the efficiency benefits come in the form of spiking neural networks (SNNs) but they still follow the same general nodal structure.

⁴ There are several attempts to further this metaphoric approach through simulations of the actual electrochemical spiking patterns produced by the neurons, rather than purely the volume of signal per neuron, but such research and references are beyond the scope of this thesis.

makes its decisions.⁵ Even in comparatively simple attempts to explore the layers of the AI blackbox, it can be incredibly difficult to understand what one is looking at. In *What Neural Networks See* (Kogan, G. 2017) one is provided software to do just this on a convolution neural network (CNN) to assess the nodal layer determining a simple image classification system. As can be seen, while some identifiable patterns might emerge, the relevance of those is difficult if not impossible to determine let alone there cumulative meaning.(Figure 7) As such, any biases introduced into the system, mathematical or human, becomes baked into the algorithm with little capacity to audit besides via analysis of inputs and presumed outputs.



Figure 7: Sample feature detection maps. (Kogan, G.)

Adversarial design

One of the methods that has been built on top of the ANN training model is the capacity to not just train the model on known values, but to *adversarially* train it against its results in order for it to prove that it can perform correctly – refining the model’s fit if it can not. This technique best gained recognition in the field of generative adversarial networks (GANs) (Goodfellow, 2014) used to train text-to-image models such as OpenAI’s Dall-e.⁶ But, beyond its use as a means to train the model, it also defines a research field of attacks

5 There are several organizations working in the realm of AI auditing and transparency, but most large existing models were not built with these needs in mind.

6 These model have largely been supplanted by diffusion models in the case of image generation and large language models in the case of chatbots, but these algorithms are again out of scope of this research as they have little bearing on the theory or design.

on machine learning models to demonstrate failures and weaknesses. These have most popularly been used against image recognition systems through works such as Adam Harvey's *CV Dazzle* (Harvey, 2010) which used hair and make-up styles to prevent facial recognition or stickers that while easily discernible to a human, cause an image recognition algorithm to hallucinate a false physical identity (Wei, 2022). Such attacks are constantly being updated and deployed by researchers as those producing the algorithms retrain to defeat them.

In the following experiments, this will serve not only as part of the training functionality for the machines, but as well an extension of the hostile architecture metaphor – as they work adversarially to the human creative intelligence and move to prevent the actions and shape the reactions intended by the performer.




Figure 8: Saliency map for CV Dazzle Look 6. Model: Jason. Saliency evaluation from New York Times Op-Art photoshoot. (Image: © Adam Harvey)

Machines learning music

Much of the uses to date within the realm of sound and music machine learning 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 mimics 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 alternative tweaked, trained and re-deployed to produce generative musical scores. Sound / song recognition algorithms (such as rights management, song identifiers, shot spotter, translation, and word recognition for wire-tapping) equivocate themselves with with visual identification algorithms both in style and often in use-cases. Sound style transfers such as Qosmo's *Neutone* (Shibata, Et Al. 2022) audio plugin take an audio input and attempt to translate it to the sound of specifically trained instruments through *Rave* (Callion, 2021) in a manner not unlike image style transfers.

The intent of these works is rather not use prior invention to assist in the creation of like content, but to take the prior efforts and prevent their continuation within live performance.

Cultural histories of acoustic & algorithmic antagonisms

Within the realms of sound, there is already a long history of its use as a disruptive tool. We can assume the shofar's ability to crumble the walls Jericho were likely not a literal tale, but its clear from the story's existence that sound has a long held place as a force weapon in the cultural zeitgeist. Several modern and real examples can be found in this direction. As a torture device it has been used in locals ranging from Chile under Pinochet to Auschwitz-Birkenau – from Afghanistan's Bangram to Guantanamo bay.⁷ These include intentional uses of music such as heavy metal and

7 Avramova, Nina. The dark side of music: Using sound in torture. Published by CNN Health, February 20, 2019. Retrieved from <https://edition.cnn.com/2019/02/08/health/music-in-torture-intl/index.html> October 2024

against Manuel Noriega in 1989⁸ and *Baby Shark* as prisoner torture in the United States more recently⁹. Offensive or directed sound, rather than music, can be found in the use of the LRAD sound canon, deployed against protestors in Ferguson, MI (USA).¹⁰ Perhaps the most direct intersection of sound with hostile architecture is in the *Mosquito*,¹¹ a device designed to play a high-pitched tones most adults can no longer hear to deter teenagers from using public space.

Meanwhile, machine learning has as already presented itself in many formats as an antagonistic force. Beyond the prior-mentioned algorithmic designs built into the dark patterns of app usage, one can look to law enforcement's willful ignorance of years of science fiction and legal precedents and speculations on the concept of innocent until proven guilty to extend into the realms of pre-crime, exceeding even the errors already presented by algorithmic sentencing mentioned later. This interpretation of *justice* has been already deployed in military use with devastating consequences. In his essay, *The Future of Death: Algorithmic Design, Predictive Analysis, and Drone Warfare*, Anthony Downy discusses just such concerns. Having the opportunity to see him speak at the *How an Image Matters* panel of Transmediale 2023 (Downey, 2023) he goes in depth discussing how in Kabul in 2021 a US aid worker, Zemari Ahmadi, was algorithmically targeted based on movement pattern profiling and a drone strike authorized without human intervention or review. This pattern-based interference underscores the extremes of attempts to interfere based on intended movement. These systems show little sign of being recalled to better analyze their efficacy or externalities – as can be seen by the present-day

-
- 8 Music torture: How heavy metal broke Manuel Noriega. Published by BBC News, May 30 2017. Retrieved from <https://www.bbc.com/news/world-latin-america-40090809> June 2024
 - 9 Minnyvon Burke, Oklahoma inmate who sued over alleged 'Baby Shark' torture tactic is found dead in his cell. Published by NBC News, Sept 13, 2022. Retrieved from <https://www.nbcnews.com/news/us-news/oklahoma-inmate-sued-alleged-baby-shark-torture-tactic-found-dead-cell-rcna47486> May 2024
 - 10 Lily Hay Neuman. This Is the Sound Cannon Used Against Protesters in Ferguson. Published in Slate Aug 14, 2014. Retrieved from <https://slate.com/technology/2014/08/lrad-long-range-acoustic-device-sound-cannons-were-used-for-crowd-control-in-ferguson-missouri-protests.html> Aug 2024
 - 11 Gary Crystel, Is the Mosquito Alarm an Infringement on Human Rights?. Published online in Civil Rights Movement, Feb 28 2024. Retrieved from <https://www.civilrightsmovement.co.uk/mosquito-alarm-infringement-human-rights.html> Aug 2024

use and expansions of systems such as “Gospel.” (Brumfield, 2023) Such might also be interpreted through the use of ML sentencing software used by courts to determine likelihood of recidivism (Nikolskaia, 2020) but in this case, on similarities between the defendant’s life circumstances and those of an accumulation of strangers (Ho, 2023).

Antagonism as play

While the above uses of ML present several detriments, they are sold as positive technologies. They are claimed to erase biases and prevent human error. This research aims for the reverse. I sell it here as a direct aggression against the performer, but in doing so, hopefully provide a platform that in fact allows a novel means to explore improvisational performance in a manner closer to competitive sport. It has already been demonstrated that pain may exist as an element of play – most clearly in *Painstation*¹² merging the classic computer game, *Pong*, with electroshock punishment, but that lacks the performative level of play I address here. Returning to the acoustic warfare of this section’s open, each example given is the use of sound as antagonist, but not as any direct interference with a musical performer. This is perhaps because the act of performance has a built-in feedback mechanism in the form of the audience. Due to the nature of courtesy, however, this feedback is rarely direct when in the negative. The performer must assess engagement, but the use of a strategic yawn or boo-ing¹³ by the audience is rarely (outside of Hollywood fantasy) witnessed to provide the negative response – more often it is but the visual disinterest or exit of the crowd. The ML in this case provides a stern judge, that while not assessing the performance for audience interest, can at least give a direct (if occasional – or perhaps even frequent) 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. In this (im)balance of teaching:training that a tension can occur. The

12 Volker Morawe and Tilman Reiff, 2001. Between its similarities to the function of a device in the James Bond movie, *Never Say Never Again* (1983) and the legal concerns of Sony, it is now known as “*The artwork formerly known as PainStation*”

13 I did in fact witness boo-ing of *Le Compte de Hoffman* at Salzburger Schauspiele, so in some formats (re-interpreted classical production, ie) this is more regular than my understanding.

performer wants to play what is unexpected by the machine, as the machine learns how the performer tries to avoid prediction. Adding to these 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. They may, over time, neither converge nor diverge, but produce merely an incoherent dialogue of data – as the audience watches on eating popcorn.

Concepts & Thought Experiments

Performative taxonomies

Algorithmic strategies

Combined methodologies

Experimental designs

Taxonomies of performance

These classifications exist to provide a basic framework for exploration and should only be regarded as highly simplified ones at that. They are not intended to be exhaustive. The first explores the motions from the perspective of the performer.

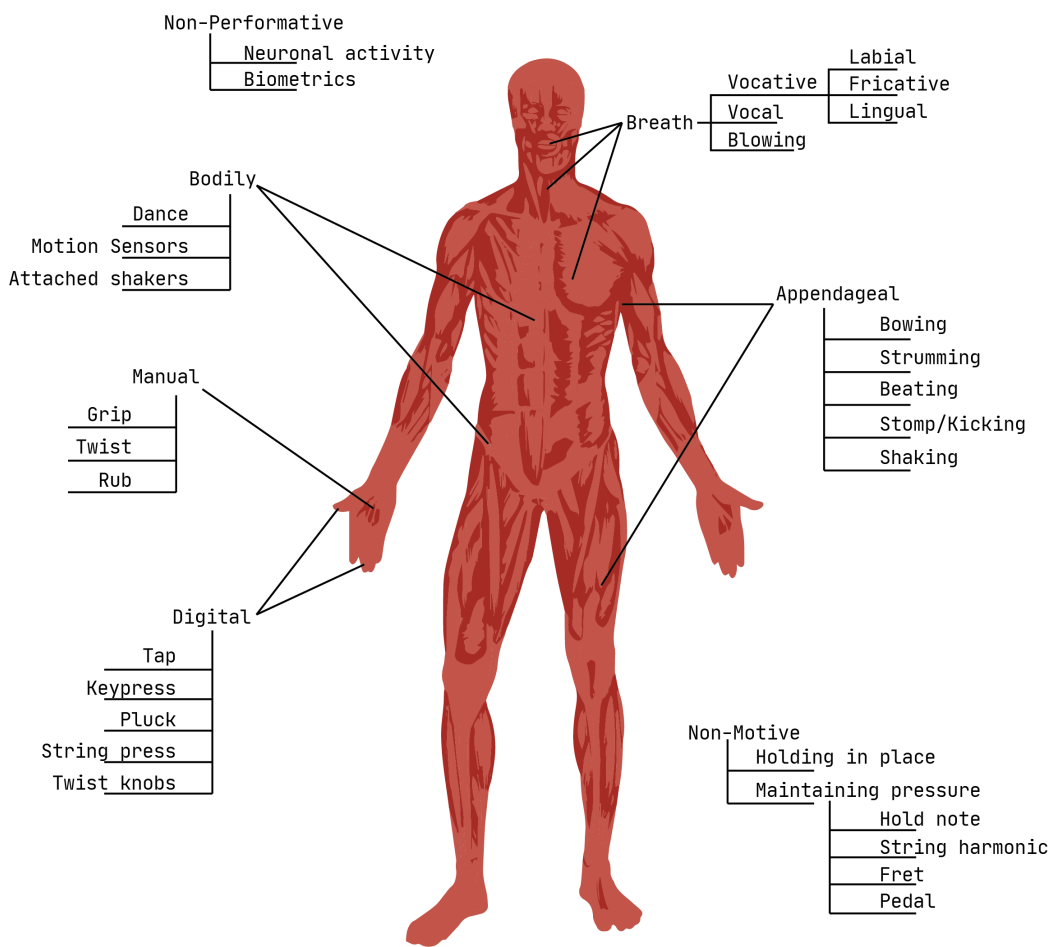


Figure 9: Division of performance by anthro-mechanical action

The second is instrument-centric, addressing the form of actions performed upon them.¹⁴

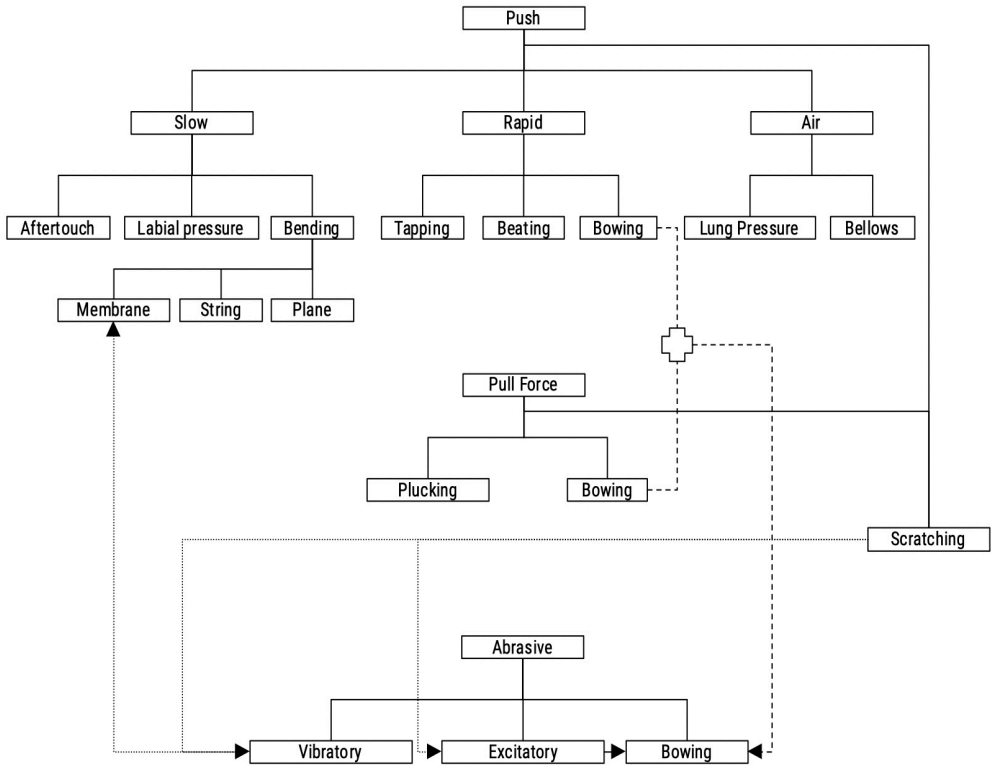


Figure 10: Divisions of performance by applied force characteristic

I will not further describe the decisions in producing these taxonomies. Their primary function was that of an early analysis for potential design creation, rather than their own work. One can find techniques in the following sections that easily might be referenced back, if one is inclined to return to this section.

14 The basic separations made here are primarily going to, by some necessity, be defined by the format of instrument and will thus somewhat follow a typical instrument breakdown. Sub-branches are based on the physical motion rather than the produced sound.

Performance learning strategies

(all may be combined)

Live response

The primary presumed means of interaction with the learning model would be one where-in, with some pre-training, the model continues to train during the performance, such that there is an interplay between the performer and the algorithm, each 'trying' to thwart the other. The algorithm, of course, 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 means of performance which do not trigger the antagonistic response. As such a performance, however, would continue to train the model, it becomes harder for the performer to improvise something that is neither pure randomness, nor too predictable. This general concept of the live response is what the final project and many of the thought experiments are based on. The following are just as suitable and adaptable, but in this instance were not pursued as further avenues of exploration.

Per-performance evolution

In this alternative approach, the model could be trained only on prior performance, and thus there might be the opportunity for the artist to learn what it predicts through 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 someone lesser dynamics of interplay that the live learning would possess, this would require playing out over a much longer time frame and could actually become quite interesting as a series if done well.

In composition

The prior strategies extend to another intriguing concept that could function with either, although likely be better suited for the per-performance model (in either case out of the scope of this present research). This is the capacity for one to attempt to compose a score ahead of time which presumes the algorithms predictions. Of course, such a score, if perfectly composed in this regard, would be more like beating a video game than providing an interesting performance, because if perfect, the algorithm would never respond, at which point one just witnesses a normal performance. It could also lend itself to cheating if viewed this way as one could, for example, compose once in a minimalist capacity, and then go for a complex jazz arrangement the next (and nothing besides honesty would prevent the same improvisationally) Depending on the means of interference, however, the more imperfectly planned the score, the more interesting it could become, because if the musician cannot improvise, then all errors are pre-baked and all predictive responses could well overtake the intended actions.

Software architectures

The following are software only approaches to the learning strategies. Being software, they are thus the simplest to implement, but the forthcoming hardware suggestions will be built upon these core concepts. These are expanded in the discreet and variable interference paradigms below, but are given initially as these software considerations inspired those paradigms as well as the other sections below.

Performance blocking

(expected notes can not be performed)

As has already been lightly tried, it is simple enough to train a system to learn a sequence of notes and mute those which are predicted as next in the sequence. This might be expanded, with some complication, to consider polyphonic input. Many techniques for predictive scoring already handle polyphony in a manner not too

dissimilar from single note input, with chords just counting as a sequence with minuscule durations. This produces quite reasonable results.

Expansions on this theme include the use of weighting (explored frequently in further concepts). Two implementation strategies come to mind. The first is to use the weighting as a randomizer, so the prevented input is not solely the one deemed most likely, but allows variation that might be expected in performance. A cut-off threshold for top-weighted can be included to avoid penalizing the performer unnecessarily based on random flukes. For example, of twelve notes, none below the three most likely would be muted, with each of those three having a mute probability based on its predicted chance. As such, the most likely note is most likely muted but the second or third may be as well. This also allows for more resiliency should two notes be equally determined to follow. The 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. This latter approach seems more optimal, or at least interesting, for the analog possibilities inherent in a physical realization, rather than software, however.

Performance altering

(expected notes are changed)

In the preceding format, the software is relying solely on prevention (or potentially at least, diminishing sound). Alternatively, it might, upon performance of a given note, swap that note to the one least expected. Or one of those least expected based on a probability function of the lowest weighted. This could lead to interesting live recomposition (and the possible confusion of the performer).

Time domain

Neither of the above so-far consider the time domain, but purely pitch and specifically pitch as note in a sequence. The time between each note is not yet in play. An alternative algorithm might instead use the above techniques of silencing or altering based on

a time expectation rather than sequential selection. A more advanced approach still could incorporate both time and sequential expectation into the equation.

Beyond the algorithm's intent, it should be considered that the longer time domain features of song structure may require longer training periods. This may be partially mitigated by the sufficient pre-training. It might be further aided through an expectation that each section maintains a structure that can be recognized, and that many performances keep to a key, and so a range of notes can be generally deemed unlikely (although this latter is less consistent in experimental contexts).

The other area in which the time domain becomes the primary concern is in the tracking of rhythmic performances, although something that incorporates the sequential approach might still be relevant to the series of drums in a kit.

Live-generative [anti]scoring

Another approach still, which I do not intend to explore, but will discuss for completeness is one that is only hostile to the performers capabilities to perform what they know. It is in a situation where the performance is analyzed as the musician plays, but rather than a real-time approach, it could play with the idea of pre-composed sheet music. As each sheet's performance nears completion, the next page is generated for the musician to attempt to play. While this may seem antagonistic to the musician only in-as-much as they are playing something unfamiliar, and potentially anti-familiar, this would only hold true for the first page of generation. Each following page would actually provide the counter-structure to that produced prior plus any account for errors of the performer. In such a situation, the algorithm becomes primarily self-hostile, and in fact may end up writing pages that are only against itself, as it would compose against the performer's errors, which may in fact make an easier score for the performer. The more perfect the performer, the more self-effacing the algorithm.

In the following section I will demonstrate and categorize the application and expansion of these concepts to physical devices as well as provide some design concepts that, while software-based, cannot exist without the hardware component.

Instrument-centric interventions


Expansion into adversarial instrument design with sketches.

Discreet interferences

Blocking operations per-event.


i. Prevention (timing)

Simple blocks or mutes prevent the possibility of moving any performative element of the instrument. This would be most applicable to rhythmic considerations, when taken in isolation. Anything pressed or hit could in some manner be blocked. In the case of the drum or drums, other methods may be better suited such as muting. Prevention in this case would likely use an interference method as below.



ii. Prevention of multiple (melody, sequence)

The first logical step building upon the prior interference is to grow this to an array of blocked elements. Keys, for example, could be incorporated with the above performance blocking algorithm to prevent the playing of any note deemed sufficiently likely. Beyond piano-style keys, this could also be applied to other instruments such as a



woodwind's pads. As discussed in the background research section, there have already been developed strategies to suggest movement to aid in performance learning [Figure 11]. Such a device could easily be modified at the algorithmic level to instead suggest against or prevent a motion.





Figure 11: Illustration of clutched mechanism to shift finger position on holed instrument. (Zhang, 2019)

iii. Muting (timing, melody, sequence)

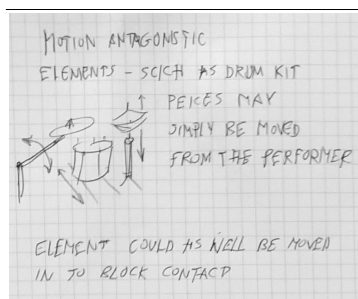
Some instruments do not lend themselves to a means of fully stopping the motion. Stringed instrument and drums would require clunky devices for such. Mutes offer the compromise of

hindrance over prevention. In these cases the system applies pressure through a padded device minimizing 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



iv. General interference

In this scenario, the instrument itself does not change per se. Instead the instrument is either providing interference or is interfered with. Such formats might include a drum kit that moves the kit pieces in, out, up, down or not at all 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 *distractive response* section below in discussion of performer-based, rather than instrument-based, interferences.

Variable interference paradigms

Operations with a non-binary inference or result variability.

i. Variation to tension (ie. string)

Rather than preventing or minimizing a note, string tension adjustments could create notes intentionally against the performers wishes. 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, if the affected element (string, i.e.) is played, the note will not be that which is intended. The degree to which such a strategy could shift the pitch would depend greatly on the instrument string tension and would be technically quite challenging.

ii. Cross-talk to alternative note

Most easily implemented in MIDI-style performance and comes from the prior software-based discussion – this approach could quite easily extrapolate a predicted likelihood for all notes. Upon the playing of a note, the one with a lesser likelihood would instead be played. Variations to how this is determined including cutoff threshold to play the performed note at and which alternative notes to play could lead to wide avenues of exploration. As with many of the purely algorithmic techniques, the need to predict and compare before sounding would introduce a small latency in performance, but with sufficiently simple predictive models the discrepancy is likely to be minimal. With little doubt, it would likely be less a concern for the performer than the actual change of note, especially to more practiced musicians. 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 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.

iii. Variable resistance

As opposed to purely blocking/muting a key, one might instead have some form of resistance applied. In this way the performer has feedback as to what the algorithm believes will be played and can still, through a certain degree of force, play any desired element. A less analog version could provide merely a minor, tactile resistance as indicator, so the performer has the feedback with less interference of pushing through variable resistance forces.

iv. Stochastic probability variation

This technique is not based on a specific note, or variable probability directly. The inferred note here is instead based on a random chance equation using the performed note as input. In a similar vein, the timing might be offset using a randomizer based on probability of the note at that time – the less likely the performer is to play at a time, the more on-time that performed note is. This could be determined by timing of any event, the timing of the specific note, 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. Unfortunately this would not allow performance of a note faster than performed. As a potential remedy, notes deemed unlikely by pitch or timing could be added without the performer intent.

Sample instrument modifications

Based on the above, a few simple instrument designs are given which might better present some of the ideas from before. Some of these would be relatively easy to implement, others less so.

Drum




Figure 12

Perhaps the simplest of interventions, A controlled arm(A) mutes the drum head when a hit is expected. As per the prior section, the degree of pressure can be used vary the head response. Due to the required placement away from the likely striking position, there would be expected variation to resonances and limited ability to fully mute the head.

Cello

In this representation, performance of strings(A) may be prevented by the intervention of mutes before the bridge(B). Tonal modification and interference, meanwhile, could be adjusted through controlled gearing of secondary tuning pegs(C) – while leaving the primary tuning intact.




Figure 13

Keys

This example suggests modification of a fairly standard electronic keyboard. As is standard, keys(A) send the note when depressed sufficiently to hit the receptors(B). However, a variably resistive pressure is instituted through a pneumatic device(C) which allows inflation response into chambers (i.e. 1,2,3) beneath the keys. In the illustration, the chamber(1) is fully inflated, indicating the algorithm

has determined it to be a likely next note, while the partial inflation of chamber(2) indicates some possibility and deflated(3) is unexpected.¹⁵ The greater the pressure in each chamber, the more effort required by the performer to play that given note, rewarding an unexpected progression.




Figure 14

15 The pressures here are illustrative. In a more realistic representation all chambers would have similar levels of inflation and the variation to internal pressure would not be so apparent.

Performer-centric intervention

Rather than affecting the instrument, these restrict or alter performer motion directly. Using the prior taxonomies, one can break down a few simple categories of which physical actions and gestures can be limited and what effect those might have.

Body physicalities to interpret

i. Large motions (waist, spine)

These include an array of full-body kinetics most frequently employed in the control of an instrument in whole. A swaying of the hip or full-body lean are most often dramatic carriers on stage. Less often do these actions produce direct sound, however in the case of some larger instruments (movement along a sufficient range on piano), or those with some heavier strength requirement (ie washtub bass) these larger dynamics are a required element to produce their tone or character.

ii. Medium movement (shoulders, neck, elbows, knees)

For many percussive instruments, this is the primary driver. Similarly the percussive or long motion elements of other instrument, such as the bowing of a violin, or strumming of a guitar. This category frequently carries a pitch character – i.e. through ascending/descending the length of a piano's keys or guitar's frets. More directly, a trombone's slide action directly translates arm to pitch. Naturally, the limbs can similarly add flare and style to a performance.

iii. Small movement (digital, labial, lingual)

These cover an array of musicalities. Certainly, digital motion is most commonly associated with the playing of a range of keyed instruments, from woodwinds, to brass, to keyboards and includes the pitch control of strings. Many simple, though often quieter, percussives may be produces as well.

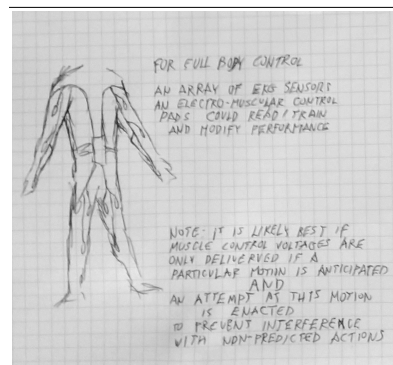
In the case of oral movements, the versatility of the mouth through voice performance is quite broad, and must be sectioned off as out-of-scope for this discussion as I'm focussing on instrumentation. The labial and lingual controls however provide many of the rhythmic and tonal qualities of breath-based performance instruments.

iv. Breath

I've separated breath as it can fall into a range from quite large, to imperceptible. Even when not directly producing an element of sound as it would in a brass or woodwind instrument, it is nonetheless a feature in the more obvious acts above. It is integral to all styles of performance and has heavy influences on tempo and musician temperament, even if little acknowledged in many cases.

Means to register

The array of techniques available to ingest various human actions is complex and broad. 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 wearables (*falling often into a category that might be defined as bulky*). Additional options could be incorporated on a per instrument basis with a variety of 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.




Additionally, while likely less sensitive to definitive motions, and more error-prone than many options, especially motion sensors, electro-muscular sensors can provide a somewhat detailed look at motion with relatively slim interface to the performer. Extrapolating motion from this type of component

become much more complicated on the algorithmic side and becomes a trade-off on the number of sensors feasible for any generalizable performance. Fortunately, these have been demonstrated capable of discernment of a range of motions from a fairly limited sensor count with reasonable accuracy. (Jiralerspong, 2017)

Pose estimation provide a more modern approach through analysis of video by trained machine learning to infer the locations of joints. This frees the musician from any physically worn devices. However, these algorithms are only recently becoming fast enough to collect realtime information from this, and even still with some delay. These models are also trained for specifics – ie one model responds the joints of the larger body, while another is required for (reliable) finger motion. They are also likely to falter when faced with the capacity of most instruments to get in the way. In the future, individually trained models could be produced that acknowledge the instrument (or possibly even track it as addition data points), but at present these complications prevent its viability for my uses.

Means to control



The greatest difficulty, when considering an adversarial component as a function affecting performer, rather than instrument, is that most design approaches require some elements of bulk. Any such machinery that is attached to the performer may well look interesting, but fail when one of the goals is generalizability to any

instrument. A number of bio-mechanical interfaces exist, and many have even been used in artistic performance artists such as Stellarc,¹⁶ among others. Pullies, motors, pneumatics and the like can all be rigged to control the potential human movement. The

¹⁶ Stellarc's *Parasite, Event for Invaded and Involuntary Body* (1997) utilized, among several other components, an array of muscle stimulators to control his movement involuntarily. These were however controlled by an online public rather than algorithm.




Figure 15: Early sketches of breath-limiting devices

only ideas for such interferences I could devise having *not* seen were either tensioning belts or some form of piston device that in both cases serves to function as an absurdist breath control (Figure 15). These would, of course, suffer no less as a hindrance to anything approaching a natural performance. (even if successfully amusing)

An alternative strategy would be to provide resistance through electromuscular stimulation. A properly supplied and controlled voltage, sent through a given muscle, results in contraction. An assortment of these has been already shown to performatively cause involuntary motions. [Figure 16] Here, it would be not driven by rhythm or intended facial control, but produce contractions deemed by the algorithm to be counter to the performer's intent forcing anti-desired motions (with potentially unsafe results)



Figure 16: Daito Manabe's "Face visualizer, instrument and copy" 2008 used an array of electrodes to stimulate facial muscles in sync with electronic beats. (Schwartzman, 2011)

Any of these options, mechanical or electromuscular, still have limitation in the amount of hardware one can attach to the body, especially at the level of small digits, and still have a reasonably capable performer. As discussed later, due to the interlinked nature of musculature, there are means to infer the motion of digital muscles indirectly.




Figure 17: Combination device sensing muscular movement and controlling finger motion via wire tension

Generalized interference

So far, these examples have been based on quite specific physicalities, especially in regards to those factors that are intended to limit or control the abilities of the performer. Here are presented more generalizable approaches that can be employed to the/any performance without specific reference to a particular discreet aspect of action.

Adverse response

This technique strips the response down to its most basic foundation of function. The system learns from the performer, makes a prediction, and punishes the performer should the prediction be correct. The deeper considerations here are more in the array of inputs which might be used as penalty. A number of methods of negative feedback can be used in this situation. The most gentle of these being that of visual indicator such as a light's activation or color shift. More belligerently, auditory indicators can be used – with the added benefit that these: a) interfere with the performed sound, b) can punish the audience. More assertive still are the traditional techniques of penal retribution – the physical. Any number of mechanical devices could be employed in this regard, such as a motorized stick or whip. Pokers on solenoids or other such mechanism could be other possible methods (although perhaps not so sharp as to risk injury [ie Figure 18]). The problems with such (beyond danger) is that their bulk can, once again, limit the performer motion to a great degree. A further option, then, is electroshock, as such devices are readily and easily attached discretely (although, no less requiring some level of safety precaution).

To expand this option, the level of feedback can be controlled based on a number of controls, but primarily the confidence level of the prediction. It might be further expanded by comparing multiple most-likely events, and sending a response at a level deemed appropriate for the confidence, so multiple results may simultaneously be deemed *wrong* or *predicted*, but those for which it has lower confidence would result in a lesser sanction.



Figure 18. Image from scene, "The Torturer's Apprentice," a farcical opera in The Baron Munchausen (Gilliam, 1989) in which the performer tortures humans with pikes to each then scream a different note. Here, I suggest a similar technique but to interfere with the performers themselves.

Distractive response

Taking a side-angled view of the above punishing response, there are some interesting, and performatively amusing, alternatives that develop. These fall under the category of responses to distract/confuse the performer, thus to cause irregular performance. These can introduce a certain sense of play back into the, perhaps ever-more-sterile approaches above. Instead of some form of painful whipping, a tickling feather in the ear might be one example. Taking the concept of the auditory indicator; a buzz or screeching sound may work well as a direct punishment, but one could as easily trigger vocal samples. A performer constantly being yelled at by the computer for their laziness and predictability could be quite the performance indeed! To take it a step further, and incorporate one more level of performance, why not involve volunteers who, given cues by the algorithm, visit any number of vocal or physical interventions upon the performer if they are incapable of sufficient novelty. As with most elements discussed here, a combinatory approach is certainly a possibility.

Project Development

General design decisions

Software design

Physical interface

Background

In realizing what might be a conceivable project design using the prior concepts, a few specific choices had to be made based on 1) feasibility (within budget & timeline) and 2) creative performativity.¹⁷ To the first end, a combinatory audio and body-signal-based interpretation approach was decided for generalizability to any musical performance, relative simplicity of components, and permitting a failsafe (audio only) variation should greater complexity designs not be realized. The intervention strategy then involving electromuscular stimulation (EMS) to cause intentional (counter performance) muscular motion, with the likely fallback being the use of the EMS as punitive shock in the case of performance-prediction matching. This combination then satisfies the second question of performative interest in that:

- It can function with multiple styles of musical instrumentation.
- There is a clear visual component in the form of connected electrodes
- There are likely visual and auditory component which become part of the performance as result from the performers response to the EMS
- Additional visual information may be added for the audience and/or performer to see how close to registering as predicted the performance is and to stimulate audience inquiry

17 An early experiment in this realm existed in the form of a simple pitch classifier/predictor/mute. Each time a key was struck, the note was simplified to a single octave twelve note scale and using a neural network, retraining on each note, the next note would be muted if the prediction was played. While seemingly functional, it did not hold up under testing, certain errors like the same note or two always being predicted would occur. On the other side, playing only one note repeated often would not even be caught. Without sufficient additional programming knowledge or explanation of the base algorithms/libraries I was unable to debug and this experiment was abandoned. Due to the very limited result of the experiment, I felt it best relegated to a footnote rather than a significant exploration on the path to final production.

Design Order

1. Develop auditory analysis structure, test of various forms of music and sound
2. Source EMG components, test functionality
3. Produce intercommunication between EMG and software layer
4. Develop and test strategies for analyzing and classifying EMG signals
5. Build EMS and test circuits and interface
6. Build pretrained models for audio and EMG pose predictions
7. Test models with visual indicator (no EMS)
8. Calibrate for reasonably expected performance
9. Fine-tune and stylize wearable components.
10. Test with real performance and calibrate further
11. Perform

Software design

The following diagrams and explanations will show the general algorithmic plan that was designed through this process. On the software side, a combination of MaxMSP, FluCoMa toolkit, SP-Tools and simpler Markov and regression models were used for the analysis, training and predictions.

Largely, these are divided into two main sections:

1. Audio analysis determining tonal, rhythmic, and general variation in sound characteristics
2. Biosensorial (electromyographic) analysis, determining physical movement

These two are then combined to create an overall determination of prediction vs detection. It should be noted that as this text and the software are being developed concurrently and the software regularly updated, that some elements of the result may vary from what is presented here, but only in fairly minimal considerations.

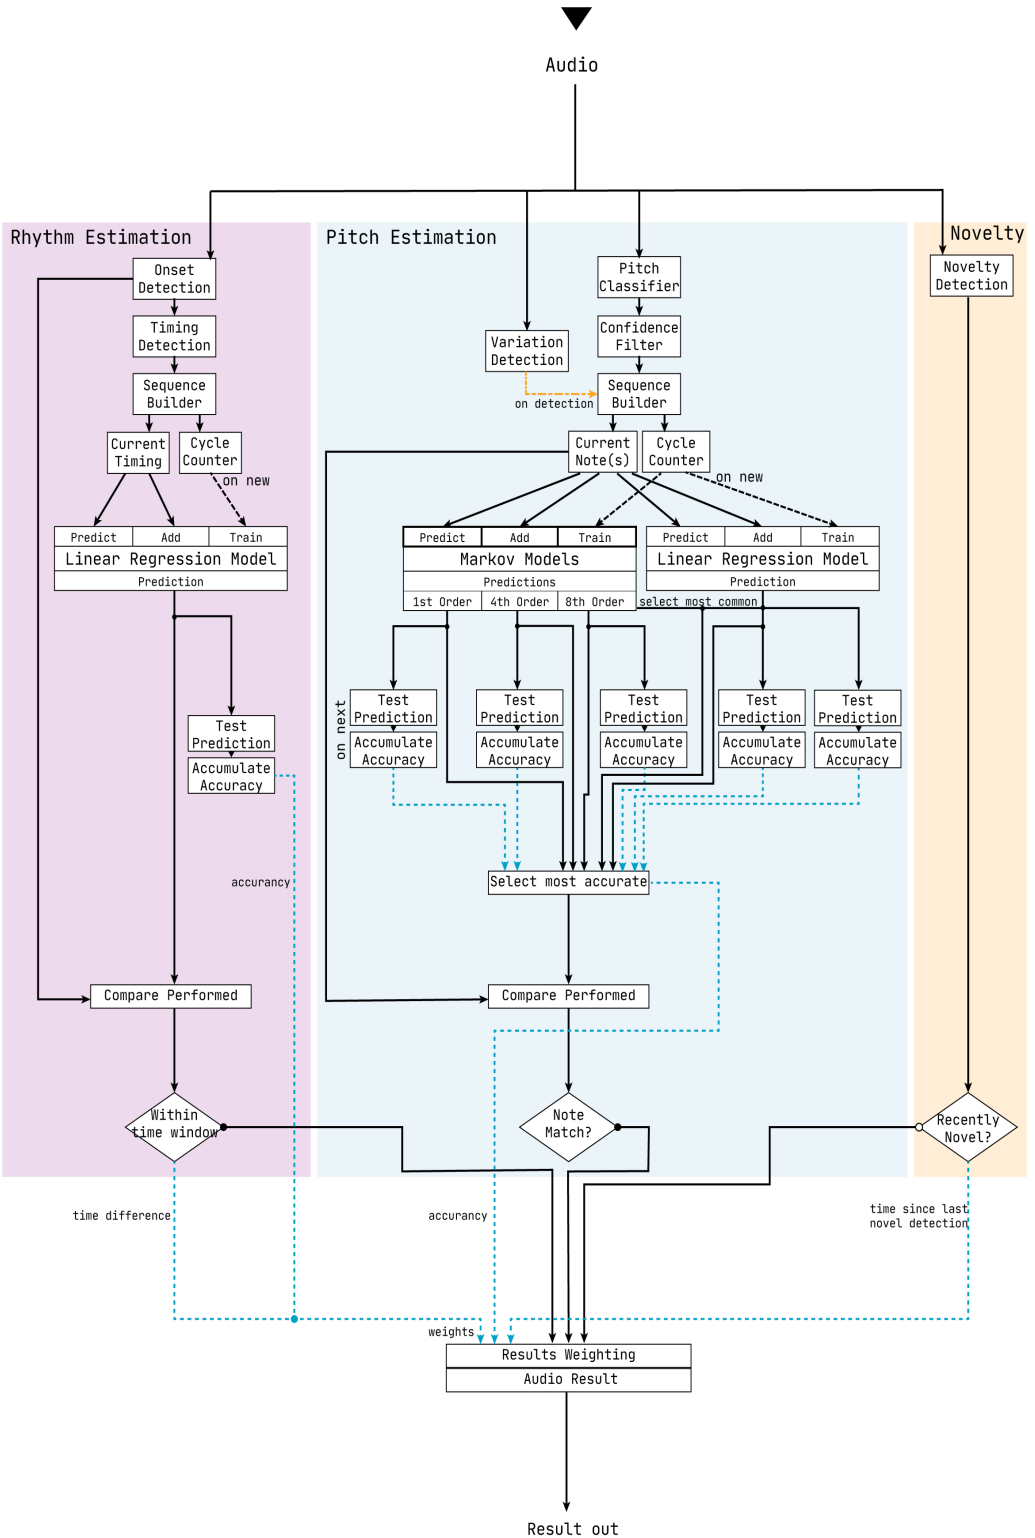


Figure 19: Data path of audio learning software

Audio analysis

The first thing one might note when analyzing this section is there are multiple paths of analysis and comparison. While there are several factors at play, and some of this was by original intent, the sub-paths within pitch-based system were the emergent result of development testing as will be discussed below.

Rhythm

This section is tasked with predicting the likelihood of something that might count as a new hit (note, beat, etc). This is accomplished using FluCoMa's "onsetfeature" algorithm. Rather than a pure amplitude thresholding, it can respond as well to components such as spectral analysis and phase deviation.¹⁸ The complex work being thus handled, the time between onsets is then built into a sequence of time differences. These are fed into a simple (regularly retrained) regression model which then determines the likelihood of the next onset from the prior sequence. If the time between performed onset and predicted time to next onset is within a threshold then this determined to be a match. Weighting, based on how close to the exact prediction, may be then applied.

Pitch

While the design for this section started out as direct as the above, the testing procedure resulted in much greater complexity. This chain starts with a pitch detection using the YinFFT¹⁹ algorithm with additional input from the onset estimation to determine a new value. The YinFFT algorithm is unfortunately not well suited for polyphonic interpretation, yet was still quite reliable in prediction even with complex polyphonic sounds. Alternative methods exist such as Mel-Frequency Cepstrum Coefficients (MFCC) to determine spectral character, followed by dimension reduction, but beyond the added complexity layers, this also is not as well suited to tone based performance. Alternatively one could use MFCC on a per-note, per-instrument basis to train for polyphony, but beyond

18 The algorithms used can be explored in Hainsworth, Stephen and Macleod, Malcolm. 2003. Onset Detection in Musical Audio Signals. *Proceedings of the 2003 International Computer Music Conference, ICMC 2003*, Singapore, September 29 - October 4, 2003. Michigan Publishing 2003

19 https://essentia.upf.edu/reference/streaming_PitchYinFFT.html

the extensive labor of such an approach, it would not be suited for any instruments untrained. Again, in testing, the monophonic pitch estimator was sufficient in most cases.

The frequency was then simplified into a pitch class (0-11 corresponding to c-b(+1 octave)) and the sequence fed into predictive models similar to the above rhythmic estimation. This is where this section became more complicated. In the interest of determining what the optimal settings were, I built out multiple models to play against each other. In this testing, I found that different musical styles resulted in different methods producing improved results. As such, all tested formats were kept and two extra layers added. First, it was noted that quite often, multiple model resulted in the same prediction. As such, I included an extra prediction based on the most common prediction. Additionally each prediction (including the *most common* prediction) was compared with the performed result. The percentage of matching values is then used to determine both an overall (all-time) accuracy and, using a ring buffer, an accuracy based on the most recent values. These accuracy values can then be used to select which value should be used for upcoming determinations.

A complete merger and weighting of all predictions may have been used, but beyond the unnecessary complexities of such a system, it also gives two results which are undesirable. In one, the values are all so low weighted that the performance is never determined to match. In the other, with so many potential predictions, it gives an unfair chance the algorithm finds one that matches by chance and punishes the performer excessively.

In tests, the performance ranges were typically 1.5-2x better than chance and sometimes as high as 4-6x. Occasionally and interestingly, the individual results came in as low as one half as accurate as predicted by chance suggesting the predictions had a specific offset bias. This is quite an interesting result, but did not occur with sufficient frequency to determine what might cause the deflation. The 1.5-2x range seems a relatively comfortable area to fall in as 2x implies an accuracy of one in six notes, and excessive accuracy risks over-punishment of the performer.

Novelty

This layer operates quite differently from the others. It is more of a sliding window approach of recent input that determines not specific onsets, but variation in the character of the input overall. This is accomplished with the novelty resource from FluCoMa which uses a combination of spectral domain analysis and self-similarity matrices.²⁰ The result is meant to allow for longer period slicing of audio into similar regions. Here it is instead operating as a determinant of variability overall. Based on the input, there is a general weight of current novelty based on frequency of novel feature discovery. This final layer is useful for, among other things, types of performance the other two layer might not be able to predict. In fact, this is largely because this section is not what one would typically define as predictive per se. It operates instead on variations over time. The clearest examples of where this might best be able to detect what the others would miss being those styles within ambient or drone genres as these would have very few onsets to predict rhythms and quite slow tonal progression to predict next tones. While these could be compensated for by dialing in the thresholds for variation detection, this method allows them to operate within reasonable parameters.

Combination

These three layers are merged with their weighting to produce a global comparison of predictive accuracy. Each layer can, above a certain threshold, trigger independently, while also at lower thresholds trigger if the predictive accuracy of others suggests a performance quality that was generally predicted on all counts.

Thresholding

Each of these has a built in self-adjusting threshold level in an attempt to have an out-of-the-box system for any performance. In reality some adjustments will usually need to be made based on a given style and also the musician's personal thresholds for pain level and regularity.

20 https://ccrma.stanford.edu/workshops/mir2009/references/Foote_00.pdf

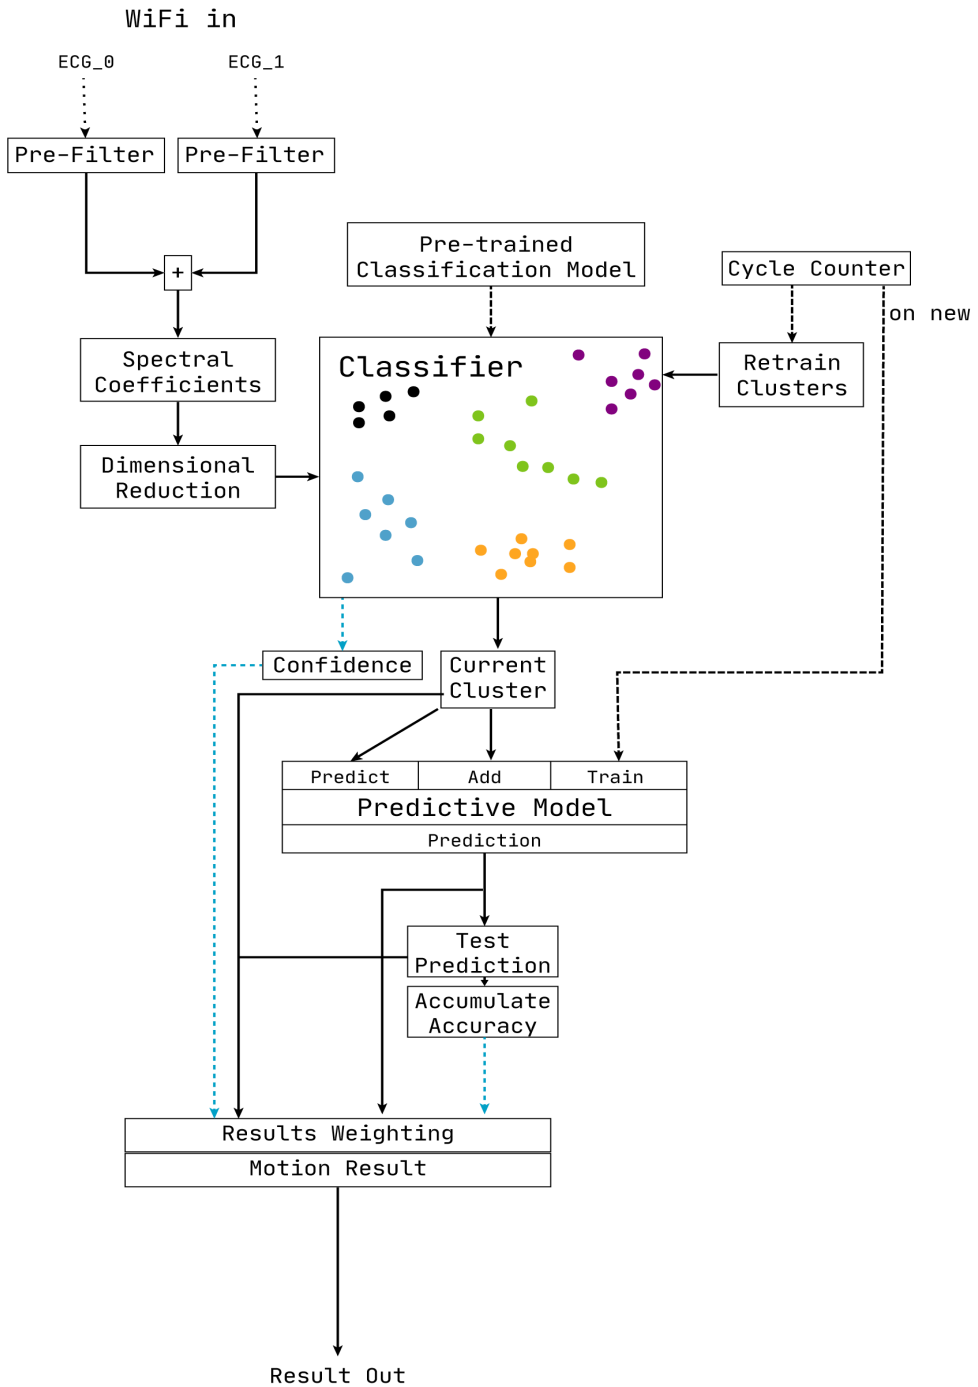


Figure 20: Data path for electromuscular software learning

Electromyography analysis

Electromyography (EMG), or the measurement of the electrical signals produced through muscle movement, can be used to determine an array of physical motions using deep learning networks as has been explored in the literature.²¹ Although many of these setups are more complex than I could produce in this work and more detailed in their process, there exists systems that have demonstrated reasonable ($n=17$) pose classifications using fewer signals ($2x^{22}$, $3x^{23}$). This is easily sufficient for my task. Additionally, while prior research needed to train for specific poses, either for interpolation for prosthetics or as a means of gestural control, my project requires only that I have classifications of whichever gestures are common to a given performance. It need not determine what those postures entail as specific gestures, only the commonalities in signals between them.



Figure 21: Testing of EEG pad placements

This provides a freedom that a pre-trained model is not only unrequired to begin a performance, it may in fact be a hindrance, as training from one motion set to another, especially on different performers, is unlikely to hold sufficient and regular clustering. The periodic learning can and does reliably build the classification clusters during live performance.

The system for prediction is algorithmically more complex than the audio predecessors, but as much of this is handled by the Flu-CoMa toolkit it is not so complicated on the production side. The signals from the two EMG signals uses variation detection to determine changes in posture. The signals are then analyzed to a multidimensional field using the pre-described MFCC method. the

21 Pizzolato S, Et Al. 2017

22 Jiralerspong, Trongmun, Et Al. 2017.

23 Li, Q. and Li, B. (2013)

full array of coefficients from both signals are then reduced to two dimensional clusters which are given a per-cluster value assignment. Clustering takes place using principle component (PC) analysis. More modern perceptron based reductions were too unstable with each new data point having a chance to completely reorder the two dimensional plot resulting in a complete rearrangement of clusters. Each determined change is also fed back into the algorithm for periodic re-analysis and rebuilding of the clusters, so the stability of PC is significant.

As these retraining steps are more intensive than retraining of the simple linear models, a balance must be struck between coefficient counts, trigger regularity and retraining frequency. For this reason, the number of coefficients is kept relatively low. Given the minimal frequency range (around 5-450hz) compared to a typical audio signal, this reduction is not of dramatic concern and higher coefficient quantities tend to have significantly diminished returns. An additional benefit of the MFCC method is that by stripping the first coefficient, we remove base amplitude, which is generally not associated with gestural response, but with signal error.

After the clustering stage, a much simpler regression model is again employed to learn the sequence of cluster labels, predict the following cluster and compare to the performed result. As the analysis is relatively fast, the latency from the prediction is primarily determined by the sample time required to collect sufficient frequency data. Onset for prediction is only used to produce the training, and prediction is based on the change from one cluster to another.

In the case that this project is further expanded, these signal could be used to send the prediction's reverse signal to trigger muscle motion, but in its current iteration, the result is merely a weighted predictive match. This is once again combined with the tonal matching and the overall confidence used to determine whether to send a punitive signal.

Hardware design

The physical production, while seemingly comparatively simple, had several issues of both hardware and personal safety that had to be considered. At its base, the system uses two EMG sensors and one EMS stimulation device controlled by the software side. For the microcontroller, an ESP32 was selected for its WiFi capability, thus isolating it from the computer in case of circuit issues arising from the high voltage stimulation device. The air-gap further prevents possible crosstalk with the EMGs as their circuits rely on detection of low-level signals and voltage interference over USB was regularly visible, especial in the event the computer relied on 50Hz AC supply. The EMG boards are relatively inexpensive AD8232 boards. These are effectively op-amps with rectification, smoothing and filtration to typical EMG signal range (5-450Hz). EMS was generated by an off-the-shelf transcutaneous electrical nerve stimulation (TENS) intended for pain/pleasure recreational application. Its only modification was to the output leads cable to be able to control its output via ESP32-activated relay. Unfortunately, for this revision this also meant there could not be confidence-based electrostimulation. Replacing the manual level control with a digital potentiometer is a possible future consideration, although given the other limitations of the current device, it is likely preferable to first look at new options for the EMS generator.




Figure 22: [left] First functioning assembly of components, [right] Analysis of electrostimulation signal




Figure 23: Signal diagram of physical components

Electrode Placement

Considerations in the realm of postural/gestural recognition required considerations beyond placement at the muscular ends, especially considering the limited lead count I would be able to measure. As a starting point, the suggestions in prior literature (esp. Jiralerspong, 2017) provided sufficient. Given that my algorithm is non-gesturally specific, I was able experiment with the placements and settled more on what gave the most reliable signals. This was important in the exploration of producing a wearable device. An array of attachment strategies exist, and while the most common and reliable is the one-time sticky pad, the wearable is superior in many respects. These include not only aesthetics, but the repeat re-usability which is especially important if used in rehearsal as repeat application of disposable pads is time consuming, mildly frustrating and wasteful. Prior exploration in this vein exist (incl. H. Wu, 2021). Drawing from these, simple metal pads seemed the most accessible technique. There is signal loss from this form of contact, but given the non-specificity of classification as discussed above and certain mitigation techniques, such as application of electroconductive gel, these losses seemed minimally problematic. This bore out in testing, even without the use of the gel (although the gel did minimize spiking). The greatest issue being the need to maintain contact. With good design practice, even this was not fully necessary for satisfactory results, as the detection and clustering of lost contact tended to have equivalencies with specific forms of motion, but again, full contact and gel providing quite improved response.

Aesthetics

Physical presentation

The first, and perhaps easier of the two presentation elements to discuss is the worn, physical elements. The design is primarily built around function. Wiring and device containers are designed to be a point of visible interest, naturally, but also to be relatively inconspicuous and out of the way for the needs of the performance.



Figure 24: Sleeve test

On the EMG side, the first instinct was the elastic compression sleeves present a simple and re-useable method for wearing the electrodes. By using snap fasteners, I could both secure the wire contact and use the fastener metal as skin contact. (Figure 24) Additionally, by drilling through the buttons, I could provide a means to apply electroconductive gel through the sleeve. Unfortunately, the amount of compression provided by these sleeves was insufficient to maintain good contact at the wrists. This design was

discarded to try the updated version using a series of adjustable elastic straps using a bolt/washer/spike combination. (Error: Reference source not found) These demonstrated their ability to maintain contact quite well, were visible more interesting, generally more secure and provided easier access for gel application.

For the TENS/electroshock side I simply produced two wide-band elastic straps with velcro fasteners. The TENS device was then connected to two—each of 25mm spikes with interior washers for improved contact. (Figure 25) Gel could be easily applied before to increase the level of punishment.




Figure 25: Spiked arm-straps

Tying the two together is a high-visibility vest. This was selected in part due to the aesthetics implications of danger, as well as general style choices. Interior pockets left and right contain the controller circuits and the TENS device, allowing some visual representation of the operation through their LEDs and quick access in case of emergency need to shut down.(Figure 26)



Figure 26: Vest fit test. [right arm(L)]: EMG sensor straps. [left arm(R)]: electroshock sleeves. [vest center-L]: microcontrollers. [vest center-R]: TENS device

Supplementary visual devices

There is a difficulty in presentation when using something such as electroshocks. Depending on the level to which the voltage / pain level is set, there may not be much to engage with from the audience perspective as to realizing when the performer has been punished – without which they are left to wonder when/if the musician is adjusting strategies. To remedy this a few plans were considered and any or all may be implemented for a given performance.

The simplest version is to use some form of lighting to give an indicator that a shock is taking place – using the common language of a red light for example, or momentary high-frequency strobe. This can be expanded depending on the algorithmic approach to give multiple light cues to evidence which element(s) of predictability has been violated (rhythmic, melodic, etc) or also the current probability risk for anything that has a temporally growing risk threshold. The latter could again use common lighting language such as green→yellow→red progressions common to street lights.

The second technique in development is providing a data feed through projection or other large, visible display. Having already implemented visual testing of the algorithms for training, I was able to make an arrangement that shows some of the behind-the-scenes data that is entering and leaving the system. The look is constantly being refined to make the aesthetic more interesting, but one element I consistently, intentionally avoid in this is labeling or any exposition to give clear meaning to the data. The reasons are two-fold. First, if the audience is spending their time reading and comparing the data to the performance, they are less focused on that performance, and the visual is meant as an accentuating element of interest, rather than focus. The other is that it creates a kind of game for those who are intent on focusing to try to look at the data displayed and attempt to figure out what each of the datagraphs and numbers mean.

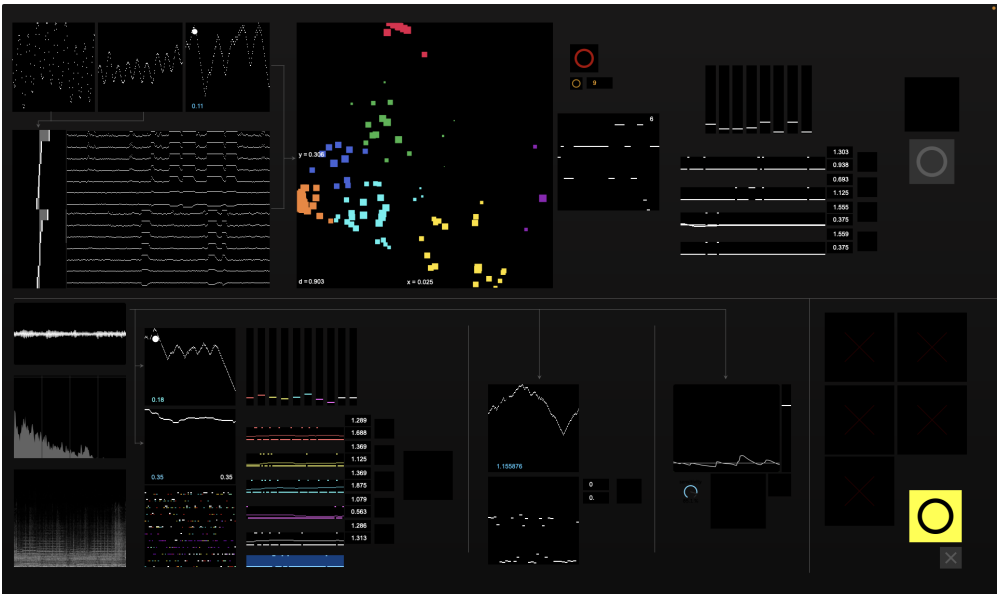


Figure 27: Visual data presentation. [Top, left to right] The progression of EMG sensor input to analysis to prediction. [Bottom] The same for audio but in three sections for pitch prediction, beat prediction and novelty. Additionally a final section (bottom right) includes indicators for trigger punishment and emergency re-set controls.



Output

Performance planning

Performance results

Discussions and further research



Photos: Bogi Nagy @bophotoadventures

Results

First Performance

***Sonic Saturdays - Lecture Performance. September 8, 2023.
Anton Bruckner University, Linz***

In order to assess this project in its goal as performative interference, an initial short demonstration was scheduled for the *Sonic Saturdays* series of the *Ars Electronica Festival 2024* to take place at Anton Bruckner University. The originally intended structure includes four parts.

1. A short performance by myself without explanation to the audience.
2. A brief lecture explaining what was happening in the performance.
3. An additional performance by a volunteer.
4. Audience feedback.

The intention of this structure is to allow an audience to view a performance without understanding what was going on (although with some visual feedback, to give some clue of what was occurring). Giving the explanation after then allows the second performance to occur with an understanding where they can better judge (and completes the audience ↔ performer ↔ computer feedback cycle). In the feedback I could then document the reaction of the audience both when unaware and when fully informed

Due to a variety of factors, the above structure was not fully adhered to, neither did it function as intended. While in the above I planned the introduction to focus on theoretical underpinnings, allowing the audience space to attempt analysis and determine what was happening with the system. Due to time constraints, I realized that:

1. I would not have time for proper rehearsal with my intended instrument and
2. This would lead to even less clarity than intended.

As such, my introduction shifted, somewhat late, into a merger of the original with a too-condensed version of the technical aspect. Between this overly-simplified explanation, both of the technology and background, with the poorly rehearsed performance, the results were not ideal. It can be stated that, from conversation with attendees, the concept in general was quite interesting. Yet, many of the comments were inquiries into those elements I did a poor job explaining. There was also some intrigue in the idea – specifically in the context of the electronic music device with which I performed, as many styles of electronic music default to quite repetitive structures. Audience members liked the attempt at forcing a breakage from this tradition.

Regarding the setup itself, as few difficulties became apparent. First was that I did not have the time to actually correctly wire the electro-shock device, so as a stand-in the punishment became annoying buzzing sounds. This had two advantages: it incorporated the audience in the punishment, and different sounds were used that could indicate to me which of the types of predictability I was violating (motor, pitch, rhythm, etc). The downside is that the element of risk would have made the performance better in general. The second issue, which is solved in later iterations, is that, having run out of EMG pads for the sensors, I used the electrostimulator pads. These, it turned out, had a different connection gauge than the EMG pads. This is why you can see my arm wrapped in black tape [Figure 28] rather than the electro-pads seen in testing [Error: Reference source not found]. As might be noted, the intent to have an audience volunteer performer did not happen. The reality of the

structure would have made it difficult, but beyond this without the electro-shock functioning at that time, the drama of requesting a volunteer would have been quite diminished. Neither would it have been likely a random volunteer would immediately have been able to produce an engaging performance with only the Moog Mother-32. In any case, the realization that the EMG and EMS pads were not connected through the same standard prevented that it could have happened as no volunteer could have been wired.



Figure 28. The tape-as-fastener approach is here seen during the lecture section from *Sonic Saturdays*

Lessons

While EMG sensors were selected for pose and gesture estimation based on existing literature, this was also with an intention to be able to counter some of those movements through the electro-stimulation. As this became too complex to produce, and punitive electro-shock was used, an array of other sensor might have provided a much simpler and potentially more consistent means to analyze and predict motion, especially given the long history of such wearables. That said, given that using EMG for gestural determination is still a field in active development by many large research players (including those with infinite budgets such as

Meta), it was certainly more interesting to explore the challenges inherent in their use as signal. This project seems unique in de novo, rather than pretrained gestural determination and there could likely be further research in that path, as this technique requires no foreknowledge of predetermined discreet motions.

Regarding the EMGs boards, the AD8232s used are cheap, readily available and reasonably documented. On the other hand, they lack protective circuitry and I found myself burning through the first two with simple wiring error (or in the second case, the interference of another in the lab). While not dramatic, these setbacks each meant both a stoppage of the ability to work while awaiting the shipment of new boards. The electronics-based delays also appeared in my realization that none of my MOSFETs or relays were reactive to the voltage levels of the ESP32s pinouts. (ie. the reason I made the change to sound penalty over shock during the performance). Moreover, while the direct line is between EMG and muscle contraction is no longer there, there use nevertheless works in this regard conceptually as, to the audience, there is still a biomuscular signal and a biomuscular response. In future development of this project, with the possibility to use higher quality voltage stimulators, it can also be of benefit that I have already done this side of the work and further expansion to greater electrode counts would require little new developmental overhead.

The thresholding at time of performance was set manually and auto-thresholding functionality has since been added, but the experience demonstrated that it is likely that some manual level of threshold testing/setting will always be required. These primarily affect the interpretation of the audio, as the EMG signals are reasonably consistent, especially once sent through the MFCC analysis.

Design updates

The following have been addressed in the hardware/software design sections, but to cover with acknowledgement of production timeline, they are briefly restated here. On the software side, a few elements were added. Among these were:

- Auto-threshold functions to limit the amount of time in setup would be required to fine tune note, beat and motion detection.
- A time-out function to prevent multiple shock triggers without time to adjust performance
- An array of visual improvements to the data display
- Various tweaks and calibration updates

By the prior performance none of the wearable components had yet been completed (or in some cases started). This test-execution at *Sonic Saturdays* aided in their creation by giving a preview of some of the issues that might be faced, such as cables and skin contact maintenance.

Preparations for second performance

Tangible Music Club. Upcoming January 22, 2024. STWST, Linz

With the production of all layers of the device successfully completed, it was able to be demonstrated in a very preliminary context at the *Tangible Music Lab*. (Figure 29) As the equipment and performance were quickly assembled without rehearsal, the performance itself lacked in any particular interest musically, but through it the hardware/software combination showed itself to be able to function within performative context.

In the intervening time between this writing, the project will continue to be updated for improved aesthetic and function. Ideally the software will be updated to control lighting as discussed in the “Supplementary visual devices” section prior. In the meantime the musical elements are being assembled and rehearsed to create an improvisational electronic set that can play off of the design predictors in a manner engaging the audience.



Figure 29: Brief demonstrative performance at the Tangible Music Lab, Linz

Further research

This project has been an enjoyable exploration of what, I believe, is a unique take on some old concepts regarding machine learning, improvisation and performative autonomy in the face of the technosphere's growing influence. Many of the algorithmic techniques discussed here are not new. Neither are many of the ideas of performative measurement, musical analysis, or punishment based interaction. It is hoped, however, that this combination – used against the performer – can provide not only a novel framework for exploring originality and creativity in musical performance, but can also work to draw attention to the metaphorical connections of hostile architectures, both those literally within the built environment, and figuratively in regards to presumptions both glorious and terrifying in the realm of machine learning and its influences on our lives and behaviors.

Algorithmically, this might have used more modern ML techniques, such as perceptrons (as was done in the first experiment with pitch class prediction). But these are not particularly necessary in this capacity and, moreover, the CPU overhead of newer techniques certainly does not equal the returns. One of the interesting components of this project is its now oft-mentioned cyclic feedback nature, which when taken to its logical conclusion mean that nothing designed here could ever be allowed to be fully accurate – as if that were possible. The closer any predictive model gets to perfection in this cycle, the less *performance* is possible. Any perfect model would not only predict the performed, but infer the performer's likely next attempts to circumvent the prediction. Should a model work this well, it collapses, the performer can play nothing (or is under constant pain in the current production). As with Dr. Doolittle's pushmi pullyu (Loffing, 1920), it would cease to be a performative cycle and would instead just be two heads pulling from center and going nowhere – the anti-Ouroboros (or uber-Ouroboros, perhaps?). Fortunately, no such predictive algorithm is so perfect and, as mentioned, part of the algorithmic calibration is trying to tune exactly what threshold of predictive accuracy must

require a performer's efforts of avoidance, without prevention of execution.



Figure 30: A pushmi pullyu as represented in the 1967 adaptation of *Doctor Doolittle* (Fleischer, 1967)

Perhaps hostile architecture is not the right metaphor here, so much as evolutionary response. Although in both it could be said that there is no activity on the current generation, whether architectural or biological, but in the following generation's response to the response we see a similar cycle to that intended with this performance. Nature's self selection for those things which survive²⁴, whether it be antibiotic resistant bacteria or those plants defined as weeds against which our agriculture struggles. A recent article addressing the later in the MIT Technology Review has the title, "The weeds are winning." (Main, 2024) This positions our technological systems as that which the weeds overcome. If we turn this on its head and view the prior-mentioned control systems, both architectural and algorithmic as instead the technological weeds, and produce systems exploring how to improvise against them, perhaps we can manage to keep our own weeds from winning. Or perhaps this can just be a musician's exercise in improvisation through frustration and punishment built on an inversion of the common intent of the learning algorithm.

24 There are multiple forms of genetic selection beyond the natural selection implied here, especially in more complex species. For the sake of metaphor, these will be here-ignored.

Afterword:

Expansions to other movement-based performance artforms

While this research is focused on actions as relate to musical practice, it lends itself as well other performative artforms, most specifically dance, but also in some cases of experimental theater. The taxonomies above could easily be expanded to include style of personal motion. The means of interference, especially those which interface with the body directly could have quite interesting use-cases – performers that must stay upright as the devices decide typical upright is not allow for example. The breath-work integral to many movement-based practices as well being as much a part of the performance tempo as on the musical stage.

References and credits

References

- Brumfield, Geoff. "Israel is using an AI system to find targets in Gaza. Experts say it's just the start", Dec. 2023. Morning Edition, National Public Radio, <https://www.npr.org/transcripts/1218643254>
- Cascone, Kim. "The Aesthetics of Failure: 'Post-Digital' Tendencies in Contemporary Computer Music." *Computer Music Journal* 24, no. 4 (2000): 12–18. <http://www.jstor.org/stable/3681551>.
- Caillon A., Esling, P. (2021) "RAVE: A variational autoencoder for fast and high-quality neural audio synthesis". arXiv 2021. <https://doi.org/10.48550/arXiv.2111.05011>
- Chin, D., & Xia, G. (2022). A Computer-aided Multimodal Music Learning System with Curriculum: A Pilot Study. *NIME 2022*. <https://doi.org/10.21428/92fbeb44.c6910363>
- Downey, A. 2023. Algorithmic Anxieties and the Future of Death. from How an Image Matters with Alaa Mansour, Anthony Downey, Lesia Kulchynska moderated by Zoé Samudzi. *Transmediale*. Feb 2023. Video: https://www.youtube.com/watch?v=e_yh6cHakOc Accessed Mar 2024
- Fleischer, Richard. *Doctor Doolittle*. APJAC Productions, 1967. 152 Minutes.
- Fujii, K., Russo, S. S., Maes, P., & Rekimoto, J. (2015, November). MoveMe: 3D haptic support for a musical instrument. In *Proceedings of the 12th International Conference on Advances in Computer Entertainment Technology* (pp. 1-8).
- Gilliam, Terry. *The Adventures of Baron Munchausen*. Allied Filmmakers, 1989. 126 Minutes.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Hainsworth, S. and Macleod, M.. 2003. Onset Detection in Musical Audio Signals. *Proceedings of the 2003 International Computer Music Conference, ICMC 2003*, Singapore, September 29 - October 4, 2003. Michigan Publishing 2003

- Harvey, A. CV Dazzle. Accessed 14 Nov. 2023.
<https://adam.harvey.studio/cvdazzle/>
- Ho, Yi-Jen (Ian) and Jabr, Wael and Zhang, Yifan, Ai Enforcement: Examining the Impact of Ai on Judicial Fairness and Public Safety (August 6, 2023).
<http://dx.doi.org/10.2139/ssrn.4533047>
- Huang, K., Do, E. Y.-L., & Starner, T. (2008). PianoTouch: A wearable haptic piano instruction system for passive learning of piano skills. *2008 12th IEEE International Symposium on Wearable Computers*, 41–44
- 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, no. 11: 1163. <https://doi.org/10.3390/app7111163>
- Kogan, G. What Neural Nets See. 2017. Retrieved July 2024 from
<https://experiments.withgoogle.com/what-neural-nets-see>.
- Li, Q. and Li, B. (2013) Online Finger Gesture Recognition Using Surface Electromyography Signals. *Journal of Signal and Information Processing*, 4, 101-105. doi: 10.4236/jsip.2013.42013.
- Loffing, Hugh. *The Story of Doctor Doolittle*. US: The Stokes Company, 1920
- Main, D. "The weeds are winning." 2024. MIT Technology Review. Retrieved Nov 9, 2024, from
<https://www.technologyreview.com/2024/10/10/1105034/weeds-climate-change-genetic-engineering-superweeds-food/?src=longreads>
- Minsky, Marvin, "The Emotion Machine". USA: Simon & Schuster, 2007
- Nikolskaia, K. and Naumov, V., "Artificial Intelligence in Law," *2020 International Multi-Conference on Industrial Engineering and Modern Technologies (FarEastCon)*, Vladivostok, Russia, 2020, pp. 1-4, doi: 10.1109/FarEastCon50210.2020.9271095.
- Pinsker, Joe. "Slanted Toilets Are the Logical Extreme of Hyperproductivity" 2019. The Atlantic. Retrieved Dec. 2024 from
<https://www.theatlantic.com/health/archive/2019/12/slanted-toilet-standardtoilet-productivity/603898/>
- Pizzolato, S, Tagliapietra, L, Cognolato, M, Reggiani, M, Müller H, et al. (2017) Comparison of six electromyography acquisition setups on hand movement classification tasks. *PLOS ONE* 12(10): e0186132.
<https://doi.org/10.1371/journal.pone.0186132>

Shibata, A., Tokui, N., & et al. (2022). Neutone. Accessed 14 Nov. 2023.
<https://www.qosmo.jp/en/r-and-d/neutone>

Schwartzman, Madeline. "See Yourself Sensing: Redefining Human Perception."
UK: Black Dog Publishing Limited. 2011. ISBN 978-1-907317-29-3

Wei, X., Guo, Y., & Yu, J. (2022). Adversarial sticker: A stealthy attack method in the physical world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3), 2711-2725.

H. Wu, G. Yang, K. Zhu, S. Liu, W. Guo, Z. Jiang, Z. Li, Materials, Devices, and Systems of On-Skin Electrodes for Electrophysiological Monitoring and Human–Machine Interfaces. *Adv. Sci.* 2021, 8, 2001938.
<https://doi.org/10.1002/advs.202001938>.

Wikipedia contributors. (2023, November 14). Machine learning. In *Wikipedia, The Free Encyclopedia*. Retrieved 07:36, November 14, 2023, from
https://en.wikipedia.org/w/index.php?title=Machine_learning

Zhang, Y., Li, Y., Chin, D., & Xia, G. (2019). Adaptive Multimodal Music Learning via Interactive Haptic Instrument. In M. Queiroz & A. X. Sedó (Eds.), *Proceedings of the International Conference on New Interfaces for Musical Expression* (pp. 140– 145). UFRGS. <https://doi.org/10.5281/zenodo.3672900>.

Figure Credits:

All figures by author except:

Figure 2: Paydah. Own work, CC BY-SA 4.0,

<https://commons.wikimedia.org/w/index.php?curid=104103324>. Accessed Aug 2024

Figure 3: Pinki, Herzi. Own work, CC BY-SA 4.0,

<https://commons.wikimedia.org/w/index.php?curid=56625576>. Accessed Aug 2024

Figure 6: Glosser.ca. Own work, Derivative of File:Artificial neural network.svg, CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=24913461>. Accessed Aug 2024

Figure 7: Kogan, Gene, 2017. From: <https://experiments.withgoogle.com/what-neural-nets-see>. Accessed July 2024

Figure 8: Harvey, Adam. 2014. *Saliency evaluation from New York Times Op-Art photoshoot*. From: <https://adam.harvey.studio/cvdazzle/>. Accessed: April 2024

Figure 30: Fleischer, 1967

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. THIS ADVERSARIAL APPROACH IS AS WELL USED AS A MEANS TO TRAIN MACHINE LEARNING 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 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.

ABOUT THE AUTHOR

Kevin Blackistone (US/AT) is an transdisciplinary media artist and researcher using immersive, tangible, participatory and performative elements as tools for exploratory engagements. His work investigates the networks of cross-interactions between our human organism(s), its habitats & inhabitants, and their technological interrelations from cultural, medical and ecological perspectives. His research background includes a BA in Intermedia and Digital Arts (US), post-bac research with the Laboratory of Neurogenetics (US), an MA in Interface Cultures (AT), as well as recent participation in the Institute for Digital Sciences Austria x Ars Electronica Founding Lab. He has shown, performed and exhibited works at festivals and venues including Ars Electronica Festival (AT), Siggraph Asia (JP & AU), Zeiss-Großplanetarium (DE), Miraikan (JP) and City Digital Skin Arts (CN/SG/IT/DE) and, with this pioneering text, is completing his Postdigital Lutherie MA.