A Live-learning Punitive Interface for Improvisational Performance Dynamics

Kevin Blackistone
Martin Kaltenbrunner
kevin.blackistone@kunstuni-linz.at
Tangible Music Lab, University of Arts Linz
Linz, Austria

Abstract

There have been many permutations in methods to train musicians in the performance technique and the technical skills required for an instrument. Among these, many modern computational approaches have been developed which incorporate some level of feedback from the student. These often assess mastery based on a response, dynamic or learned, to the practiced result. This research presents an interface expanding these concepts, which guides the performer to novel improvisation through punitive techniques. As opposed to models using prior composition or predetermined reference points, it live-learns in situ to assess predictability without pretraining. It provides a framework and interface which seeks to be implementable for a variety instruments and styles analyses of sound and pose.

CCS Concepts

• **Human-centered computing → Interaction devices**; • **Applied computing → Performing arts**; • **Computing methodologies → Machine learning approaches**; • **Computer systems organization** → Real-time systems.

Keywords

Machine learning, improvisation, interference, penalty, performance

ACM Reference Format:
Kevin Blackistone and Martin Kaltenbrunner. 2026. A Live-learning Punitive Interface for Improvisational Performance Dynamics. In *Proceedings of International Conference on New Interfaces for Musical Expression (NIME '26).* ACM, New York, NY, USA, 6 pages.

1 Introduction

While punitive measures are neither useful, nor ethical in the education of children[10] they might still be have a place in playful contexts of music. The use of antagonism can be recontextualized as less the direct educational tool, and more an interference strategy to overcome. One can already look to works such as *The Sabotaging Piano*[7], in which the pitches of performed keys are shifted, for inspiration on how the implementation of adversity can be a tool for novel musical output. Rather than that work's randomized variation, one might consider an interface that makes performance difficult using predictive modeling. While not broadly applicable to composed works, this has potential to push *improvised* performance in unexpected directions. The live-learning antagonist approach here provided expands on ideas such as the prior piano's forced variation, by providing an performance-learning algorithm. It thus creates an interaction in which the performer must learn to compete against a rigged game, while the algorithm itself forever adjusts to any such learned adjustment. Instead of learning *what to do*, the musician must instead devise *how to adjust*. This interactive machine learning (IML) technique might then produce a mutual relationship of training, torture and inspired improvisation.

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

1.1 Prior Computer-Aided Learning and Instruction

Computation has regularly been used as a tool in music education and one can readily find examples of such NIMEs. PianoTouch [16] provides a haptic interface to guide the student's knowledge of which finger should be used. MoveMe [12] takes takes the approach of physically positioning the performer's hands, while Zhang, et al propose physical guidance of individual digits to respective keys[26]. Machine learning has quite expectably been incorporated in the role of tutor to both detect performance error and provide live feedback [4]. These approaches are thus-far based on preprogrammed compositions, and have no basis to intervene in any improvisational context. Neither do they incorporate responses learned from the performer to affect the instruction given.

1.2 Learning Through Vague Antagonism

These prior teaching strategies attempt a fairly direct feedback in order that the musician understand requirements to better perform. This research attempts instead to give an undefined punishment to the performer. It seeks only to inform the performer *that* their behavior is predictable, but not *how*; neither hinting to changes to make it less so. This can lead to a cyclical tension, as the performer attempts to please the algorithm, which is meanwhile learning the musician's new strategy, and all the while in an attempt to please an audience. One must as well include in this mix that the algorithm is fully capable of, and with some regularity, predicting correctly only by chance, rather than algorithmic robustness, and thus push to change, even when unwarranted. In an inversion to NIMEs designed to improvise with a computer analysis[6] or a co-creative AI [25], the musician here improvises against the NIME.

2 Development

This research began with an analysis of the various methods of performance and the physicalities for an array of instruments, then applied to potential algorithmic and hardware approaches. Many of those hardware approaches involved functional manipulation of specific instruments. with interventions including blocking keys, pitch shifting strings or muting drums. These presented significant difficulties for realization in that, not only is there a wide array of instruments from which to start, but the list of potential modifications available to many instruments can become quickly daunting. Instead, it was decided to produce an interface capable of learning directly from the performance and guide the performer via simple punitive measures should the result register as too predictable. This approach is thus simple to implement and less restricted in instrument selection.

The proposed solution analyzed both audio output and performer movement for prediction followed by punitive electroshock in the case of performance-prediction matching. Heavier processing is handled by laptop processor, while the motion and punishment interface manifests as a discrete wearable using wireless data transmission. This allows low-latency of the training/prediction models with freedom of movement for the performer and required circuit isolation discussed in section 4. The software selected to handle predictions was MaxMSP [5], due to the availability of a variety of packages capable of handling the analysis and prediction algorithms. Within these, the primary driver for analysis was the Fluid Corpus Manipulation(FluCoMa) toolkit [9], which features an array of nodes capable of handling the spectral analysis, onset detection, dimensional reduction and clustering addressed below. Markov and regression models were also used for training and predictions [3, 24]. A simplified diagram of the software's prediction strategies is shown in figure 1.

3 Software Design

3.1 Audio Analysis and Prediction

3.1.1 Rhythm. This section predicts the likelihood of something that might count as a new event. This is accomplished using FluCoMa's *fluid.onsetslice* with the *energy* metric as detector. Rather than amplitude thresholding, it responds to spectral variation and phase deviation to determine relative entropy [14]. An auto-threshold function is included to minimize the need for calibration to on-site volume levels or energy variation between musical styles. The time differences between each new beat are collated into a sequence used to train a Markov model (5th order) to determine the likely onset timing of the next event. To compensate for the possibility of very slow/fast tempos, the sequence treats the time logarithmically rather than linearly. A prediction is deemed accurate if three sequential predictions are within a ten percent time difference to the result.

3.1.2 Pitch. Upon detection of onsets via *Modified Kullback-Leibler* metric[1], this chain makes a pitch detection using FluCoMa's *fluid.pitch* with the *YinFFT3* algorithm[2], with the MIDI note integer output simplified to pitch class. The pitch algorithm, while not designed for polyphonic interpretation, was still reasonably reliable in prediction, even with complex sounds as demonstrated via predictive result analysis (see Methods). Each new note is added to a sequential queue of pitches. The determined sequence then trains three Markov models (1st, 4th and 8th order) alongside one regression model based on the latest 32 pitch sequence. With each new pitch, predictions are made by each model and a consensus from their collected results determines if a prediction matches that which was performed via 3:4 predicted results matching current pitch.

3.1.3 Novelty. The prior techniques are poorly suited for ambient or slow-variation music, and would require long time-frames to train models. In order to compensate, a general novelty analysis is included, determining whether there exists a general change in sonic character within a variable time-frame. It uses a sliding window approach of recent input that determines not specific onsets, but longer period variation using FluCoMa's *novelty* resource with the *spectrum* algorithm option - combining spectral domain analysis and self-similarity matrices [11]. Intended to allow for longer period slicing of audio into similar regions, it is here used to determine sufficient general variation. The resulting value is compared against a self-calibrating threshold which resets a slow timer if exceeded. Should the time run out, it indicates insufficient novelty and punishment is triggered.

3.2 Pose Analysis and Prediction

3.2.1 Planning. In addition to sound analysis, the interface attempts to surmise variation in performer motion. There exist several techniques to obtain positional information. Video-based pose estimation [15], while minimally cumbersome for performance was ruled out as models were untrained for the possible visual interferences of a variety of musical instruments. While worn gyroscopic sensors were a potential solution, electromyography (EMG) inference was instead selected, for it's physiological relation to the punitive electroshock, and it's potential to, with few signals, infer poses (potentially including digital motion) in a low profile. Such interfaces have been incorporated into performance such as in *Xth Sense* [8], although usually as signal for direct audio interpretation. Prior literature has demonstrated the possibility of EMG signal in determining such gestures [22] and have been effective in defining up to seventeen pose classifications using only two or three signals [17, 19]. Pre-trained EMG placement techniques such as SparseEMG [18] could not be implemented as relevant poses change between different instruments and this system uses no pre-training, allowing gestural coupling to a performed instrument in situ via frequency analysis, dimensional reduction and clustering; followed by a training and prediction approach similar to that used for pitch prediction. Experimental result demonstrated sufficient correlation of prediction to suggest the suitability of this approach.

3.2.2 Implementation. The EMG signals are sent as open sound control (OSC) packets via WiFi to the computer. Beyond the native filtering built into the AD8232s, they are additionally filtered in software to remove the signals below 20Hz known to be a principle source of EMG amplitude noise [23]. The signals are then analyzed using seven band mel-frequency cepstrum coefficients (MFCC) with the first coefficient removed to exclude amplitude interference from the spectral variation. Detection of motion through onset features (using *spectral flux* metric) sample current MFCC values after a brief ($\tilde{2}0$ ms) signal clarity delay. Those values are then reduced to two dimensions using principle component analysis (PCA). K-means clustering is applied with a set cluster number of 10. Once sufficient points (n=100) have been collected to maintain stability in the results, the cluster of each new point is added to a sequence which is sent to three Markov models (1st, 4th and 8th order). The predictions are deemed to match here if greater than three accurate predictions are made in the sum of the three models over two sequential predictions. The clustering caps the number of data points to 750 after which any new data points overwrite the oldest to prevent observed processing lag of larger data sets after lengthy performance, while maintaining a balance between older and more recently observed data for prediction.

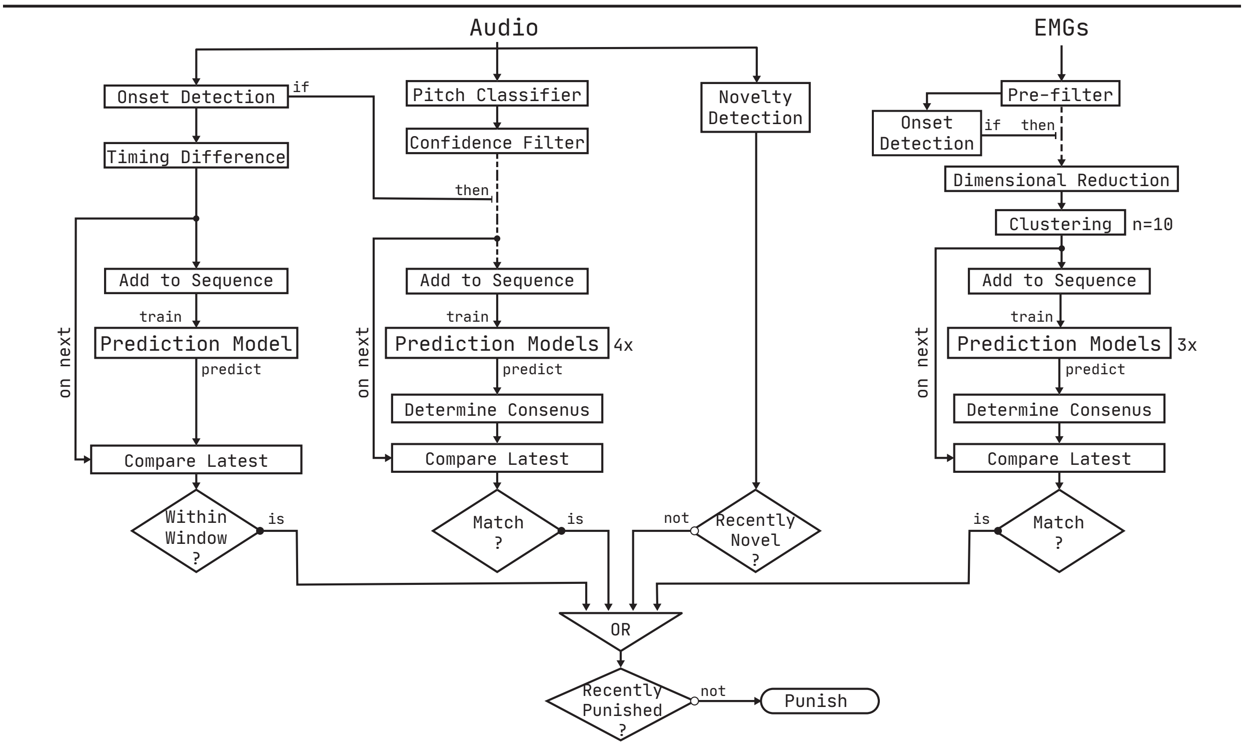

Figure 1: Simplified flowchart overview of the four predictive algorithms.

3.3 Prediction and Assessment

The models made predictions from the above four analysis with reasonable quality. Unsurprisingly, across musical genres the lowest accuracy tended to be in experimental noise and highest in clear, repetitive music. Statistical comparisons were performed with results averaging 1.5-2.5x accuracy over chance. For pitch class and clustering predictors, this compared resulting match percentages to those predicted by chance (see methods). While there was no data analysis to verify that resulting clusters corresponded to poses, that the predictions performed well suggests this to be the case. Assessing the accuracy of the beat prediction was more difficult as the time between any two onsets was not discrete and could have infinite range. A better implementation might take into account a beat estimation and meter as discussed in future planning.

The prior mentioned consensus models and requirements of consecutive results minimize the likelihood of false positives. If a majority of the models predicted correctly, this was considered a match event and an electric shock was triggered to punish the performer. As these sequences would sometimes be predicted in series, a lock-out timer was added for performance adjustments before any next prediction could trigger punishment (usually 10 seconds).

4 Hardware Design

The worn hardware consisted primarily of an ESP-32 to wirelessly send EMG signals and receive shock triggers activating an off-the-shelf transcutaneous electrical nerve stimulation (TENS) device. These were packaged in a worn vest with elastic armbands to hold the electrodes (figure 4).

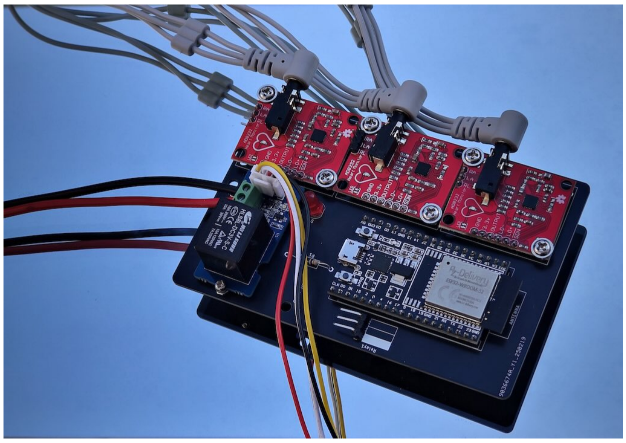

Figure 2: The sensor and control board.

The EMG-based pose detection evaluates only muscle movements of one arm with electrode placements at positions of mid-arm, elbow, and wrist as suggested by the prior referenced literature in section 3.2.1. This single-arm approach was primarily to keep the electro-shock components safely on the other arm, away from the highly sensitive EMG pre-amps. For this application, two (later three) EMG project boards were used (AD8232 module) featuring signal processing, ground reference and filtering to the typical range of EMG signal (5-450Hz) [20]. The resulting data is then ingested by a battery-powered ESP32 and transmitted as OSC. Since little processing power was needed, the ESP32 was selected for it's low cost and wireless transmission,

Figure 3: [Above] The vest interface. Interior breast pockets left and right hold pre-amp/transceiver module and the TENS device. Electrodes for TENS (left) and EMG (right) are held by elastic straps. [Below] Video still of the vest in use. (from video: https://www.dorftv.at/video/45716).

permitting freedom of movement and preventing interference from connection to other equipment.

The TENS device operates in a always-active state and is controlled by in-line relay located between itself and electrodes via a break in the transmission cable. Thus, the ESP32 could effectively control release of voltage without concern for cross-talk to the low power EMG amplification boards. The TENS electrodes were kept to a single arm for the reasons above as well as preventing dangerous current-flow crossing through the heart. For aesthetic reasons and to minimize waste, it was decided to use hardware-based electrodes rather than medical pads for this design. This provides a replicable system for future development and permits easy on-site repair if needed. A simple system of elastic straps featuring electrodes made of commonly available nut/bolt/washer combinations was used with crimp leads soldered to the preamp input cables. These electrodes' signals are often unstable when dry, but use of electro-conductive gel resulted in consistent and reliable performance, lasting for over forty minutes in practice.

5 Results

This work has now been presented in several performances, included its premiere as part of *Tangible Music Club*, STWST, Linz and the opening night of the *International Computer Music Conference* (ICMC, 2025), Boston. These performances involved electronic instrumentation, primarily in the form of a sequencing drum machine with additional instruments or noise generators changing per-performance. The drum machine was prepared with six to ten channels, each with an array of sequenced loops, or sound triggers. In additional to triggering, stopping, and layering these loops, variation was achieved through manipulation of tempo and filters. As such, many of the musical elements followed certain patterns more-so than what might be expected of single-instrument improvisation. The additional instrumentation would thus aid when primary manipulations felt too simplistic. Rehearsals were implemented without the use of the interface in order to avoid excessive awareness of algorithmic response or punishment acclimation.

In addition to personal observations, audience feedback was collected in the form of intentional, albeit informal conversations with those present at performances This included both unguided general inquiry about audience impressions as well as a few guided questions - most commonly along the lines of, "How clear were the moments of punishment?", "Were improvisational variations apparent?" and "If apparent, were variations interesting?".

- Personal observations
  - In practice, the interface could be worn with minimal physical interference to the performance, with the connecting wires as only mildly invasive. For some instruments, these wires might be prove more cumbersome.
  - Even at the lowest setting, the TENS device's shocks were readily apparent.
  - The shocks felt to result in significant jerking motions.
  - Shocks had a performance affect frequently resulting in a moment of psychological and physical recovery to perform following actions.
  - The above performance lag diminished over the performance even with intentional increase in voltage over time.
  - Early and end-performance attempts to create change resulted from shocks, rather than attempting to prevent the shocks. Mid-performance tended towards attempted evasion.
  - The punishments resulted in an impetus to manipulate these patterns to a much greater extent than typical.
  - It has been very difficult to implement any rehearsed plan for progression or evasion once in performance with the interface.
  - The added pressure to please the algorithm resulted in a lack of awareness in regards to audience response.
  - After performances there were sensations of phantom shocks for the next hour.
  - Performance of this acoustically complex music made it difficult to assess quality of prediction or avoid punishment with clarity, however it was clear that there were false positives in prediction, including a punishment received after a musical equipment failure and corresponding reset.
- Audience responses
  - The jerking motions I felt did not readily translate visually. Co-triggered strobes and/or glitched visuals cues aided, but were not always clear. This has been largely mitigated through brief explanation before performance.

- A general response to the performance was that it was fairly chaotic and thus difficult to surmise actual performance improvisations, however some moments of larger change were clear.
- Audiences largely did not care about the quality of prediction, but acknowledged the likely difficulty in avoiding predictability in a quasi-structured and beat-driven improvisation style.
- Most of those spoken to considered the performance to be musically interesting, beyond the elements of the interface.
- In general, audience members responded very positive afterwards from a musical/sonic perspective.
- The idea of live performance punishment was universally amusing amongst all respondents.

In general, it can be summarized that, from a performative side, the interface was not cumbersome physically but did affect the performance in several ways beyond the intent. This included focus on rapid variation as strategy for prevention as opposed to hoped strategic evasive learning. As above, the focus on shaping performance to algorithmic response unfortunately resulted in a lack of focus towards the live audience. While the responses of the audience were nonetheless positive, this is something to consider. Somewhat surprisingly, both of the last points on strategy and attentiveness to the audience have not changed in spite of repeated performances and awareness of these issues. Although difficult to state definitively, it seems there has been a greater result in personal performative variation to those occurring *without the interface*, but *since* its repeated use - the personal results being broader consideration to alternative variation even in non-improvised performance. So while this work is intended to be rather antithetical to a teaching interface, it yet seemed to train my performative direction. This may suggest that performative adjustments after use may be more interesting than performative result, and may connect with the practice-led methods of Green [13].

From the audience perspective, these concerns were of limited relevance. Their responses can can be summarized to suggest that this interface, no matter the quality of prediction, proved entertaining and garnered interest in further research in the topic, with several instrumentalists having expressed interest in attempting performance thus encumbered.

6 Conclusion

This research has been undertaken to consider unique explorations in interactive machine learning, improvisation and performative autonomy in the face of the technosphere's growing influence in both learning and control. It is hoped that this investigation into the use of IML 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 algorithmic control structures in regards to machine learning and its influences on our lives and behaviors.

One of the interesting components of this project is its cyclic feedback nature which, when taken to its logical conclusion, suggests that nothing designed here could ever be allowed to be fully accurate. It pits the human mental prediction of musical progression as discussed by Pearce [21] against an algorithmic one. The closer any predictive model gets to perfection, the less performance is possible, as a perfect model would not only predict the performed, but infer the next attempts to circumvent the prediction. Should a model work this well, it collapses, as the performer can play nothing (or is under constant pain in the current production). Fortunately, no such predictive algorithm is so perfect and part of calibration is tuning exactly what threshold requires a performer's efforts of avoidance, without prevention of execution. In this way it echos many of the predictive algorithms of punitive systems, in that the behavioral manipulation occurs as much, if not more-so in the threat of the punishment, than the accuracy of the prediction.

Figure 4: Performance at Red Room, Baltimore. Data visualization behind performance gives the audience a peak at predictor function, while feed disturbance signals electroshock.

6.1 Further Development

This interface is under constant refinement. The techniques used for EMG clustering and prediction are ever-shifting to provide improved response. In consideration of the ethical standards, the performances thus far been limited to myself, and using the collection electronic instruments and sequencers familiar to me, rather than any singular instrument (such as piano or guitar) for which the interface was ostensibly designed. As such, more experimentation is required with instrumentalists to determine its performance on any other styles and individual instruments.

Among recent algorithmic improvements, addition of peak frequencies to MFCC data for dimensional reduction seems to help with more natural cluster formation, while clustering to three dimensions is being explored to improve definition. Beat prediction is currently the least stable (compensated for by requiring sequential prediction stability) of the predictors and may be significantly improved with the inclusion of beat estimation but, as it behaves sufficiently well in its current state, this has not yet been implemented.

While focused on singular performance, many of these technique could be expanded to multi-performer interaction. This research could lend itself, in the vein of motion prediction, to other movement-based performative art forms such as dance as well.

7 Ethical Standards

The used administration of electroshock vie the TENS device for the production of this work was primarily upon only the author. Limited exposure to the device was granted colleagues only upon their request and with informed consent. Any attempt to replicate this research should take into consideration all safety precautions when dealing with the application of high voltage to the skin.

Acknowledgments

Thanks are given to the Tangible Music Lab of the University of Arts and Design Linz for their aid and insights into the thesis work through which this project developed.

References

[1] Juan Pablo Bello, Laurent Daudet, Samer Abdallah, Chris Duxbury, Mike Davies, and Mark B Sandler. 2005. A tutorial on onset detection in music signals. *IEEE Transactions on speech and audio processing* 13, 5 (2005), 1035–1047.

[2] Paul M Brossier. 2006. *Automatic annotation of musical audio for interactive applications*. Ph. D. Dissertation. Queen Mary University of London.

[3] Jamie Bullock and Ali Momeni. [n. d.]. *ml.lib*. https://github.com/irllabs/ml-lib

[4] Daniel Chin and Gus Xia. 2022. A Computer-aided Multimodal Music Learning System with Curriculum: A Pilot Study. In *NIME 2022*. PubPub.

[5] Cycling'74. [n. d.]. *MaxMSP*. https://cycling74.com/products/max

[6] Palle Dahlstedt. 2015. Mapping strategies and sound engine design for an augmented hybrid piano. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 271–276.

[7] Teodoro Dannemann, Nick Bryan-Kinns, et al. 2023. The Sabotaging Piano: key-to-pitch remapping as a source of new techniques in piano improvisation. (2023).

[8] Marco Donnarumma. 2011. XTH SENSE: a study of muscle sounds for an experimental paradigm of musical performance.. In *ICMC*.

[9] FluCoMa Project, Department of Music, University of Huddersfield. [n. d.]. *FluCoMa Toolkit*. https://www.flucoma.org/

[10] Sarah A Font and Jamie Cage. 2018. Dimensions of physical punishment and their associations with children's cognitive performance and school adjustment. *Child abuse & neglect* 75 (2018), 29–40.

[11] Jonathan Foote. 2000. Automatic audio segmentation using a measure of audio novelty. In *ieee international conference on multimedia and expo. icme2000. proceedings. (cat. no. 00th8532)*, Vol. 1. IEEE, 452–455.

[12] Katsuya Fujii, Sophia S Russo, Pattie Maes, and Jun Rekimoto. 2015. MoveMe: 3D haptic support for a musical instrument. In *Proceedings of the 12th international conference on advances in computer entertainment technology*. 1–8.

[13] Owen Green. 2014. Nime, musicality and practice-led methods. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 1–6.

[14] Stephen Hainsworth, Malcolm D Macleod, et al. 2003. Onset detection in musical audio signals. In *ICMC*.

[15] Robert M Haralick, Hyonam Joo, Chung-Nan Lee, Xinhua Zhuang, Vinay G Vaidya, and Man Bae Kim. 1989. Pose estimation from corresponding point data. *IEEE Transactions on Systems, Man, and Cybernetics* 19, 6 (1989), 1426–1446.

[16] Kevin Huang, Ellen Yi-Luen Do, and Thad Starner. 2008. PianoTouch: A wearable haptic piano instruction system for passive learning of piano skills. In *2008 12th IEEE international symposium on wearable computers*. IEEE, 41–44.

[17] Trongmun Jiralerspong, Emi Nakanishi, Chao Liu, and Jun Ishikawa. 2017. Experimental study of real-time classification of 17 voluntary movements for multi-degree myoelectric prosthetic hand. *Applied sciences* 7, 11 (2017), 1163.

[18] Anand Kumar, Antony Albert Raj Irudayaraj, Ishita Chandra, Adwait Sharma, and Aditya Shekhar Nittala. 2025. SparseEMG: Computational Design of Sparse EMG Layouts for Sensing Gestures. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–20.

[19] Qiang Li and Bo Li. 2013. Online Finger Gesture Recognition Using Surface Electromyography Signals. *Journal of Signal and Information Processing* 4, 2 (2013), 101–105.

[20] Roberto Merletti and PJJEK Di Torino. 1999. Standards for reporting EMG data. *J Electromyogr Kinesiol* 9, 1 (1999), 3–4.

[21] Marcus T. Pearce and Geraint A. Wiggins. 2012. Auditory Expectation: The Information Dynamics of Music Perception and Cognition. *Topics in Cognitive Science* 4, 4 (2012), 625–652. https://doi.org/10.1111/j.1756-8765.2012.01214.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-8765.2012.01214.x

[22] Stefano Pizzolato, Luca Tagliapietra, Matteo Cognolato, Monica Reggiani, Henning Müller, and Manfredo Atzori. 2017. Comparison of six electromyography acquisition setups on hand movement classification tasks. *PloS one* 12, 10 (2017), e0186132.

[23] Mamun Bin Ibne Reaz, M Sazzad Hussain, and Faisal Mohd-Yasin. 2006. Techniques of EMG signal analysis: detection, processing, classification and applications. *Biological procedures online* 8 (2006), 11–35.

[24] Benjamin D. Smith. [n. d.]. *ml.**. https://www.benjamindaysmith.com/ml-machine-learning-toolkit-in-max/

[25] Notto JW Thelle, Rebecca Fiebrink, et al. 2022. How do Musicians Experience Jamming with a Co-Creative "AI"? (2022).

[26] Yian Zhang, Yinmiao Li, Daniel Chin, and Gus Xia. 2019. Adaptive multimodal music learning via interactive-haptic instrument. *arXiv preprint arXiv:1906.01197* (2019).

A Research Methods

A.1 Algorithm Selection and Testing

Onset threshold algorithms used for determination of new beats, notes and poses, while originally attempted through research of algorithms [1], the most effective option was better found via trial and error. Similarly, the delay between onset and determination of new stable pose state was based on observed maximum durations of unstable low-frequency or clipped data. PCA dimensional reduction was selected due to computational simplicity and stability of resulting plots compared to other options. K-means clustering with a set cluster number was selected as well for it's consistency in resulting clusters.

Validation of predictive quality was achieved via comparison of resultant true vs false positives against expectation in random chance. In the case of pitch class determination this is expected to be 1 in 12 (8.3%) and for the clustered bio-signals 1 in 10 (10%) success rates for random chance. The algorithms both separately and as a whole tended to approximately 1.5 to 2.5 times this expectation from chance. While this still implies very regular false positives, the consensus approach from before reasonably diminishes their impact. Exceeding these accuracy values proved difficult in production with higher accuracy results often resulting from anomalous data, such as when sequential points fell excessively or redundantly into specific pitches or clusters.

The cluster count (n=10) was determined based on experimental result, with the count needing to be high enough to minimize regular false positives, but low enough to differentiation pose. Clusters could roughly be assumed to represent a certain pose of the arm and hand, but are difficult to prove, as there is no data to correspond specific poses to specific clusters. That the predicted results were consistently above chance thresholds suggests cluster-pose correlation.

B Online Resources

Files used in the production of this work are available online at github.com/blackistone/ImprovisationalAntagonist.