Explainable AI in Music Performance: Case Studies from Live Coding and Sound Spatialisation

Jack Armitage
Intelligent Instruments Lab
Iceland University of the Arts
Reykjavík, Iceland
jack@lhi.is

Nicola Privato
Intelligent Instruments Lab
Iceland University of the Arts
Reykjavík, Iceland
nicola@lhi.is

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

Celeste Betancur Gutierrez
Center for Computer Research in Music and Acoustics
Stanford University
Stanford, CA, USA
celes@stanford.edu

Abstract

Explainable Artificial Intelligence (XAI) has emerged as a significant area of research, with diverse applications across various fields. In the realm of arts, the application and implications of XAI remain largely unexplored. This paper investigates how artist-researchers address and navigate explainability in their systems during creative AI/ML practices, focusing on music performance. We present two case studies: live coding of AI/ML models and sound spatialisation performance. In the first case, we explore the inherent explainability in live coding and how the integration of interactive and on-the-fly machine learning processes can enhance this explainability. In the second case, we investigate how sound spatialisation can serve as a powerful tool for understanding and navigating the latent dimensions of autoencoders. Our autoethnographic reflections reveal the complexities and nuances of applying XAI in the arts, and underscore the need for further research in this area. We conclude that the exploration of XAI in the arts, particularly in music performance, opens up new avenues for understanding and improving the interaction between artists and AI/ML systems. This research contributes to the broader discussion on the diverse applications of XAI, with the ultimate goal of extending the frontiers of applied XAI.

1 Explainable AI for the Arts (XAIxArts)

The relevance of Explainable Artificial Intelligence (XAI) to music is multifaceted and extends beyond the technical realm. Public interest and discourse around AI in music has been growing, with increasing curiosity about how these systems work and how they can be used creatively. From a pedagogical perspective, XAI can provide valuable insights into the underlying mechanisms of AI/ML-based music systems, thereby enhancing learning and fostering a deeper understanding of these tools. Furthermore, XAI can contribute to artistic immersion and intuition by making AI/ML systems more transparent and interpretable, allowing artists to better understand and control the creative process.

As authors of artistic tools, there is a desire to create systems that are self-explanatory and simple to use. This is not only for practical reasons but also for aesthetic and materiality considerations. The design of these tools should be intuitive and engaging, allowing artists to focus on the

XAI in Action: Past, Present, and Future Applications @ NeurIPS 2023.

creative process rather than the technicalities of the system. XAI can play a crucial role in achieving this, by making the workings of AI/ML systems more understandable and accessible.

In the following review, we will discuss key articles that have contributed to the exploration of XAI in music, highlighting the potential of XAI in enhancing the understanding and interpretation of AI/ML-based artistic practices.

Llano et al. 2020 and Bryan-Kinns, Banar, et al. 2021 were two early attempts to directly address XAI in the arts, specifically in music. In particular, Bryan-Kinns, Banar, et al. 2021 propose a method to make a latent variable model for music generation more explainable. They extend the MeasureVAE model, which generates measures of music, by using latent space regularisation to map specific dimensions of the latent space to meaningful musical attributes. They also provide a user interface feedback loop and a visualisation of the musical attributes in the latent space. This approach bridges the gap between the latent space and the generated musical outcomes, making the model and its outputs more explainable and "debuggable".

Inspired by the first article, Bryan-Kinns, Ford, et al. 2023 organized the first international workshop on XAI for the Arts at the ACM Creativity and Cognition Conference 2023 (ACM C&C 2023). The workshop aimed to bring together researchers from various fields to explore the role of XAI in the arts. The focus was on how complex AI models, such as deep learning techniques, can be made more understandable to people, particularly in creative domains. In a submission to this workshop, Privato and Armitage 2023 proposed a context-sensitive approach to XAI in music performance. The authors argue that there is no universal approach to explainability and that the context and audience are crucial in developing explainability requirements. They propose an *Explanatory Pragmatism* (EP) framework for XAI in music performance, which tailors explanations to specific audiences and refines them based on feedback. This approach offers a promising direction for enhancing the transparency and interpretability of AI systems in broad artistic applications.

Finally, Bryan-Kinns, Banar, et al. 2024 expanded on their previous work in a book chapter, reviewing one hundred AI and music papers to illustrate how AI models are being explained, or more often not explained. The chapter explores the potential of XAI for music generation and demonstrates how a latent space model for music generation can be made more explainable. The authors conclude with four key challenges for XAI for music and the arts: the nature of explanation, the effect of AI models, features, and training sets on explanation, user-centred design of XAI, and interaction design for explainability.

These articles collectively highlight the potential of XAI in enhancing the understanding and interpretation of AI/ML-based artistic practices, particularly in music. However, they also underscore the need for further research and discourse in this area, particularly in terms of the nature of explanation, the role of context and audience, and the design of explainable interfaces.

2 Case Studies

The case study research methodology combines a first-person, practice-based (see Candy and Edmonds 2018 and Johnston 2016) approach with an autoethnographic journaling of the experience. The analysis of the second case study is carried out through the framework offered by Latour's Actor-Network Theory (ANT) (Latour 2005). The advantages provided by ANT in the analysis of complex sociotechnical systems such as the one described in this case study are manifold: since ANT views agents as temporary assemblages of other agents, complex networks can be simplified within the level of abstraction that better allows to respond to the research question. Furthermore, ANT considers agency as a feature of any active entity, being it a human being, an inanimate object or even a concept. It does so by shifting the focus to the semiotic domain, in which "agents are continually coming into being, fading away, moving around, changing places with one another" (Pickering 1996). Finally, ANT allows to understand a network by describing its connections and accounting for the ways in which actants negotiate their interactions through dynamic processes of translation. This principle, rather than referring to physical transformations of actants within the network, accounts for their semiotic reformulation within language, i.e. to their interpretation (Janicka 2023).

3 Case Study 1: XAI in Pandora's Mycophony

This case study focuses on artistic live coding (Blackwell et al. 2022), that is, programming onstage as performance with dedicated arts programming systems, to produce an improvised audiovisual work.

3.1 Context

Pandora's Mycophony was premiered at New Interfaces for Musical Expression (NIME) 2023 in Mexico City, Mexico (see program note in Figure 1, and concert photo in Figure 4), and is an audiovisual dreamscape that re-imagines Pandora's story (see concert program note in Figure 1, and Betancur and Armitage 2023). A second iteration was performed at Mengi in Reykjavik, Iceland in August 2023 (see concert photo in Figure 5). The performance features two performers who combine their own live coding systems, Tölvera and Pandora, respectively.

Pandora's Dream is a live coding playground that opens up a world of possibilities for performers. As detailed in Betancur Gutierrez 2023, Pandora's Dream empowers artists to harness the perspective of multiple tools, languages, and approaches to create audiovisual experiences in real time integrating AI/ML algorithms in live coding environments. Pandora's Dream is a C++ application that hosts interpreters for Lua, OpenGL, ChucK and more (see system diagram in Figure 3). These languages share global variables, so that the artist-coder can read and write from any language to any other. Finally, Pandora's Dream features a re-implementation of the Wekinator (Fiebrink, Trueman, and Cook 2011) and also ChucK AI (CHAI) (Mulshine et al. 2023), enabling the implementation of small AI models (Vigliensoni, Perry, and Fiebrink 2022) that can be trained using the on-the-fly paradigm (Wang, Cook, and Salazar 2015).

Tölvera, is a software library for artistic exploration of artificial life and self-organising systems, as described in Armitage and Magnusson 2023 (see example scenes in Figure 2). This system emphasises the unique challenges that emerge from engaging with musical scores and notations as real-time agents, promoting fluidity of form in notation and prompting performers to identify with, mirror, and attune to them. In this work, Tölvera was paired with an interactive machine learning library also inspired by the Wekinator (Fiebrink, Trueman, and Cook 2011). Our mapping system[1] breaks down into three modules: embedding, search, and interpolation. Given a dataset of input-output pairs, an embedding module maps inputs to feature vectors. A search module maps a query feature vector to set of relevant examples v from the dataset with relevance scores d. Finally, a smooth interpolation module combines the output examples according to their scores with linear weights:

$$w_i = \frac{1}{d_i} - \frac{d_i^2 - 3d_i d_k + 3d_k^2}{d_k^3}$$

where $d_i = ||v_i - q||_2 + \epsilon$ and $d_k = \max(d)$.

In Pandora's Mycophony, inputs are point-cloud data from Tölvera and outputs are uniformly distributed random vectors in $\mathbb{R}^n$. Inputs are embedded by random projection to $\mathbb{R}^m$ followed by sorting, such that the L2 distance between vectors implements the sliced optimal transport distance Bonneel et al. 2015 between point clouds. A nearest-neighbor search is used to find the closest features in the dataset, which are used to smoothly interpolate output vectors.

In the performance of Pandora's Mycophony, one performer trains an interactive machine learning model that takes the 2D XY positions of every particle in a Tölvera scene and creates a random input-output mapping pair. The length of the output vector is decided by the other performer, who uses it to control musical parameters in their CHmUsiCK custom library. They also use ChucK AI (CHAI) to live code small data models that take the musical event stream from CHmUsiCK and retrain every bar. These models output a vector of arbitrary length, which controls the flocking and swarming high-level behaviour parameters of Tölvera. This process creates a feedback loop between the two performers' systems, resulting in a unique, dynamic performance (see Figure 5).

[1] `https://github.com/Intelligent-Instruments-Lab/iil-python-tools/tree/master/iml`

3.2 Reflections

Reflecting on the intersection of live coding and Explainable AI (XAI), we delve into the nuances of code and understanding. The act of writing code inherently seeks to make processes explainable, translating complex ideas into a language that machines can interpret. This translation process, in turn, often clarifies the concept for the coder themselves, as they must break down the idea into its most fundamental parts.

Live coding, as outlined in the TOPLAP Manifesto[2], is not just about creating a program, but also about explaining or narrating the process in real-time. The manifesto demands access to the performer's mind and the transparency of their screens, emphasising that programs are instruments that can change themselves and that the artificial language of code is the way to transcend the program. This transparency and visual representation of code and underlying algorithms make live coding a self-explanatory practice.

The original manifesto acknowledged that it is not necessary for a lay audience to understand the code to appreciate it, much like one does not need to know how to play guitar to appreciate a guitar performance. However, live coding practice still provides a level of explainability that is not typically present in other forms of performance involving computing technology. While the audience's understanding may not be as deep or detailed as the practitioner's, the live coding practice still provides a level of explainability that is not typically present in other forms of performance.

Knotts and Paz 2022 discusses the challenges and opportunities of combining machine learning with live coding. They examine how this combination affects core principles of live coding, such as transparency, liveness, and the visibility of algorithmic processes. The "small data mindset" advocated for by Vigliensoni, Perry, and Fiebrink 2022 aligns with their views that simplicity is a gateway to transparency for ML in live coding, along with techniques for visualisation that are often part of an XAI approach.

Despite that introducing ML into live coding does bring with it issues of explainability, our perspective is that making something "live codable" inherently makes it (more) explainable, compared to not. In reality, most ML practitioners and audiences do not experience it as a live coder does, they rather experience it in an "offline" capacity, with a slow feedback loop, lacking in the tactile material dialogue that live coding provides. Therefore, live coding has something unique to offer to the discussion around XAI. In live coding, explanations are embedded in code, but also in the live material that the code creates, which remains true in the case of machine learning. The process of live coding allows the practitioner to understand the inner workings of the code as they write it, and it also offers the audience a glimpse into this process. The unfolding experimental process is also written on the face of the performer as well as the screen, giving an emotional and psychological insight to audiences.

From this perspective, we propose that incorporating interactive and on-the-fly machine learning processes into live coding practice is a promising avenue for artistic investigation of explainability.

4 Case Study 2: XAI and Sound in Space

This case study focuses on sound spatialisation, and explores the possibility for this musical dimension to inform the performer's understanding and real-time navigation of autoencoders' latent dimensions.

4.1 Context

The research question arose from playing with an instrument we developed at Intelligent Instruments Lab[3](IIL), collaborating with the director Majella Clarke[4]: the sonic baton (see Figure 6). The sonic baton combines an accelerometer with a real-time audio variational autoencoder (RAVE). (Caillon and Esling 2021) RAVE is a fine-grained latent variable model, encoding a

[2] `https://toplap.org/wiki/ManifestoDraft`

[3] `https://iil.is/`

[4] `https://www.majella-clarke.com/`

stream of raw audio sampled at 48 kHz to a stream of latent vectors sampled at 23 Hz before decoding it back to audio. It combines a variational autoencoder with an adversarial reconstruction term, so that the coarse qualities of the sound tend to be represented in the latent space while realistic details are sampled by the decoder. Depending on the dataset and hyperparameters, the dimensionality of a RAVE latent space is usually between eight and thirty-two dimensions.

We trained three RAVE models and built three batons. We then performed with the batons at Limbó Finissage[5], inviting the audience to play with them after the exhibition. During the activity, we observed how people would point the batons towards different objects in the space, to access similar sound features through gestures anchored to the three-dimensional space. This led us to question whether the performer's projection of the model's internal mechanics could be further enhanced by using a spatialised system.

We further explored this question during a residency at Tangible Music Lab[6] [7], in Linz, working on the Dodekaotto system[8], a 20+1 speakers dodecahedron with the diameter of ca 2 Meters with the performer's head at the centre, using ICST Ambisonics algorithm (Arteaga 2023) for spatialisation (see Figure 7). For this purpose, we developed and performed with a novel interface capable of extending the control of the model to a higher dimensionality than the batons: the magnetic discs (see Figure 8). (Privato, Magnusson, and Einarsson 2023) The discs are hand-held by the performer, and combine the navigation of three absolute spatial dimensions with the three-dimensional position relative to the field of magnets embedded in the instruments themselves or within the performative space. Through this, it is possible to achieve a high degree of control over multiple latent dimensions at once, with each disc controlling all eight latent dimensions of one RAVE model. (Privato, Magnusson, and Einarsson 202e)

4.2 Actants

For the purpose of this analysis we define a series of actants interacting within the network:

- **Performer**: The musician controlling the instrument in its dual role of performer and researcher, placed at the center of the sound space.
- **Sound Space**: This actant combines the physical space, the Dodekaotto's structure, and the representational space offered by the system through sound spatialisation. From the perspective of the performer the physical space and the sound space coincide, as any sound displaced in a given area of the dodecahedron is to be perceived as coming from a particular three-dimensional coordinate.
- **Composed Instrument**: We use the notion of Composed Instruments (Schnell and Battier 2002) to describe the actant resulting from the entanglement between the model and the interface it is navigated with, mediated by an arbitrary mapping. This reduction seems in contradiction with the initial research question, aimed at exploring the interpretability of the model's latent spaces. Nonetheless, it is not possible for the performer to directly interact with the latent spaces per se. Notably, once the model is embedded into an interface affording a defined set of performative gestures, some areas within the latent representation become more accessible than others.
- **Research Aim**: Finally, because it is known *a priori*, the research question itself affects the network by informing the practice of the performer and should be considered as an actant. The research aim frames the research question within the aesthetic boundaries of the performance, as there would be little artistic use in an interpretable yet non-musical configuration.

4.3 Translations and Reflections

In ANT, translations account for the ways in which the different actants negotiate and adapt to each other's agency. A first process of translation may be detected in the research aim itself: in the attempt of exploring the interpretability of the model in action an interface needs to be provided, and with that a particular mapping. Within this process, the affordances of the instrument as well as those of the performer's physicality redefine the navigation of the latent

[5] https://www.nylo.is/en-us/events/17002

[6] https://tamlab.kunstuni-linz.at/

[7] https://youtu.be/y09Gg5LoUkQ

[8] Structural design by Martin Kaltenbrunner and Ben Wesch for the OTTOsonic project -open platform for immersive sound-. OTTO Kulturgenossenschaft Ottensheim and the Tangible Music Lab, University of the Arts Linz, Austria.

dimensions by facilitating or restricting the access to particular representational areas. As a consequence, the initial research aim of understanding the model's representation per se translates in this network into the understanding of how this particular AI-based composed instrument is experienced through sound in space.

In the process of testing and redefining the mappings, the material agency of the interface became apparent. We realised that the different shape of the discs compared to that of the baton was having a critical impact on the playing. As ANT suggest, agents act ambiguously within the material and the semiotic domain: indeed, the baton's shape was outward-projected, with the representation emanating from the discs to a particular point in the sound space as the trajectory of an arrow. On the contrary, the circular shape of the discs was inward-projected, with the latent dimensions and the spatial ones meeting halfway, thus collapsing their respective agentialities in the performer's hands.

A third process of translation involves the practice of the performer within the sound space. During the week of practice we realised that the way we learnt to use the instrument and to find sounds within the Dodekaotto system is sensibly different from how we understand it in a stereophonic environment. The spatial projection of the sound allows for a higher degree of sensitivity and control, and the positioning of the sound in the space becomes both and aesthetic and artistic expedient, and a cue for finding a particular timbre trajectory. Through sustained practice, RAVE, the discs and the Dodekaotto became a single agent or large-scale musical instrument, interacting with the performer in a dynamic negotiation between the inner workings of the discs and the spatial projection of the sound. At the same time, such negotiations between the components have found a synthesis that could be hardly described through language, but could be instinctively and progressively understood and navigated by the performer through sustained practice.

To conclude, through a practice-based method we experienced how sound spatialisation may facilitate the understanding of an auto encoder's latent dimensions as they're experienced through the materiality of a particular interface. We also realised that the high dimensionality of the model as well as the complexity emerging from its negotiation with the network's agents can be intuitively understood through practice.

5 Discussion

Through our case studies, we have explored the application of Explainable AI (XAI) in music performance, specifically in the contexts of live coding and sound spatialisation. Our findings suggest that XAI can significantly enhance the understanding and interpretation of AI/ML-based artistic practices, providing artists with a deeper insight into the workings of their tools and systems.

In the first case study, we found that live coding inherently promotes explainability, as the process of writing code in real-time provides a transparent view into the underlying algorithms and processes. The integration of interactive and on-the-fly machine learning processes into live coding practice further enhances this explainability, offering both the artist and the audience a deeper understanding of the AI/ML systems in use.

In the second case study, we discovered that sound spatialisation can serve as a powerful tool for understanding and navigating the latent dimensions of autoencoders. Through sustained practice with a novel interface and a spatialisation system, we were able to intuitively understand and navigate the high-dimensional latent space of the autoencoder, demonstrating the potential of XAI in this context.

However, our research also highlighted several challenges associated with the application of XAI in the arts. The nature of explanation, the role of context and audience, and the design of explainable interfaces all emerged as key areas for further investigation. Additionally, the complexity of AI/ML systems and the high dimensionality of their latent spaces present significant challenges for explainability.

The case studies we presented propose two different approaches to XAI in music, each one tailored to a particular performative context and aesthetic need. This testifies the importance of a pragmatic approach to XAI, one that defies any essentialist view of explanations in favour of a diffractive (Barad 2014) and nuanced analysis of a particular sociotechnical assemblage. In future research, we envision integrating live coding and sound spatialization into a unified performance and practice-based investigation, and to explore through ANT the novel associations, translations, and emergent properties of the network in comparison to those delineated in this paper.

6 Conclusion

In conclusion, our research suggests that XAI holds great potential for enhancing the understanding and interpretation of AI/ML-based artistic practices. However, further research is needed to fully realise this potential and to address the challenges associated with explainability in these systems. We believe that the exploration of XAI in the arts, particularly in music performance, opens up new avenues for understanding and improving the interaction between artists and AI/ML systems. This research contributes to the broader discussion on the diverse applications of XAI, with the ultimate goal of extending the frontiers of applied XAI.

Acknowledgments and Disclosure of Funding

Celeste Betancur and Jack Armitage carried out Case Study 1, Nicola Privato carried out Case Study 2, and Victor Shepardson contributed software design for the interactive machine learning package. All authors made an equal contribution to the writing and editing, and the author names are not ordered in terms of contributions.

Jack Armitage, Nicola Privato and Victor Shepardson are supported by the Intelligent Instruments project (INTENT), which is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

Celeste Betancur is supported by The Hume Graduate Fellowship Fund in the Arts.

7 Supplementary Material

Pandora hears her own dreams, they talk to her in mysterious voices, unknown languages.

You find yourself standing alone, in the middle of her darkness.

You don't know how you got there.

Are you one of Pandora's dreams?

Talk to her, maybe she will answer you.

In this audiovisual dreamscape lies a re-imagining of Pandora's story, where the contents of her jar are bioluminescent swarming spores that seek to fill the world with hope instead of evil, and life instead of death. The spores want to get out, their evolutionary powers are hidden, and the whole universe is waiting to be explored. Meanwhile, Pandora is dreaming, condemned to keep the box closed. Life waits to be released.

Figure 1: Program notes for Pandora's Mycophony.

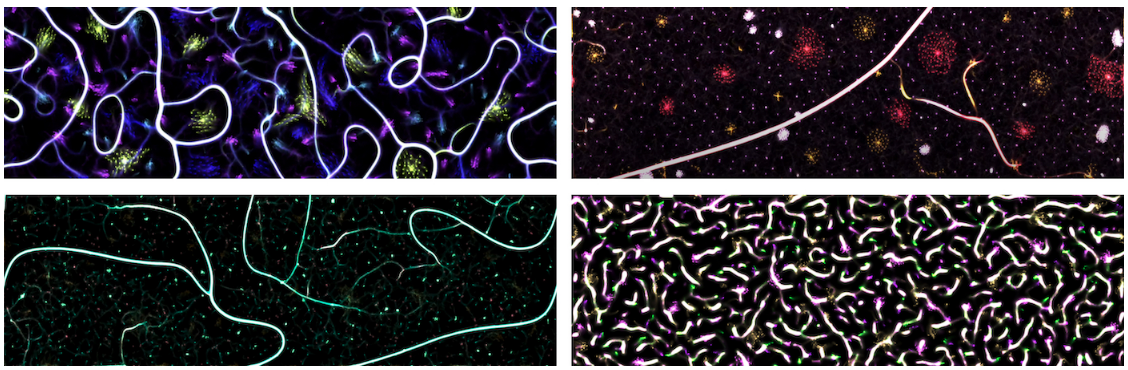

Figure 2: Case Study 1: Tölvera example scenes.

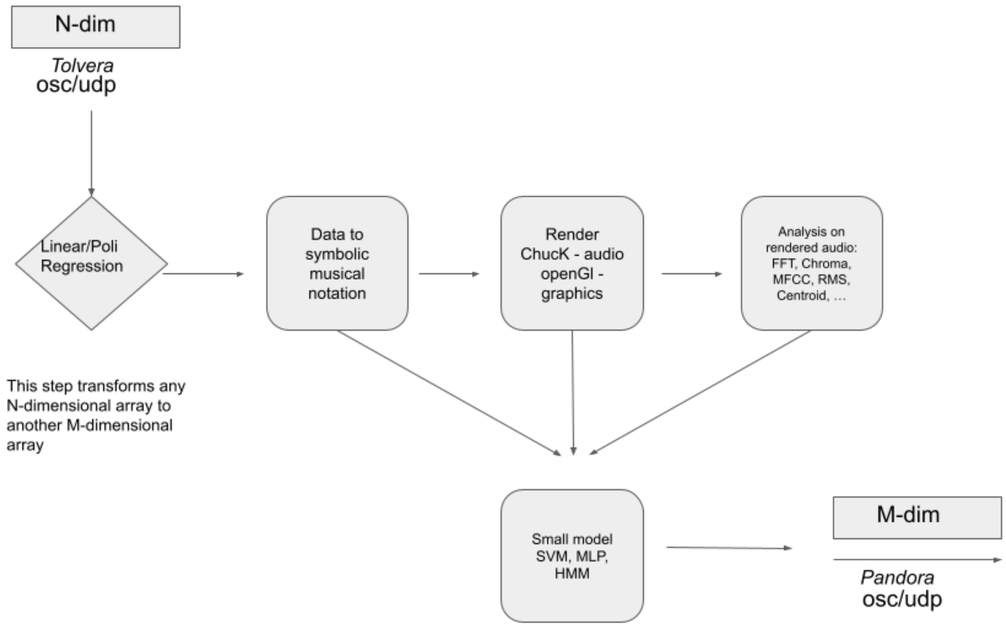

Figure 3: Case Study 1: Pandora's Dream implementation diagram.

Figure 4: Case Study 1: Premiere of Pandora's Mycophony at NIME 2023.

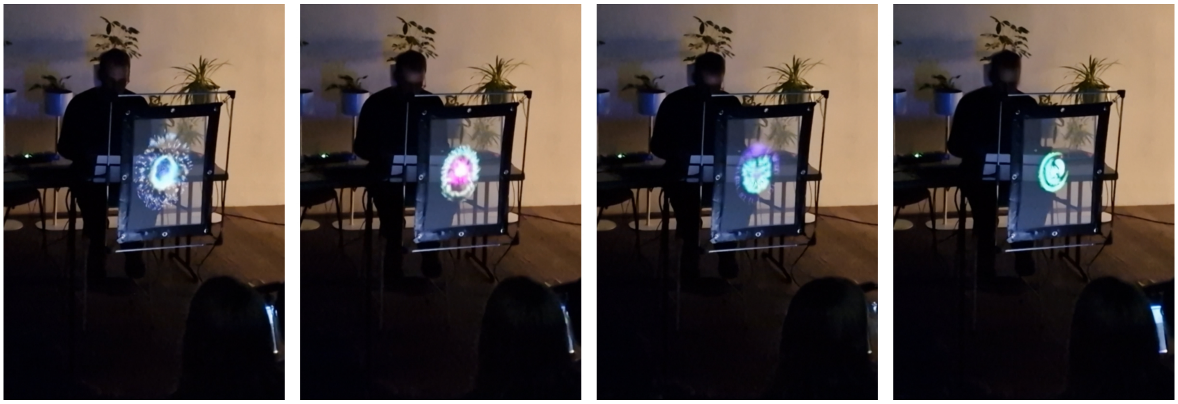

Figure 5: Case Study 1: Second iteration of Pandora's Mycophony at Mengi, Reykjavik, Iceland, using a custom hologauze display.

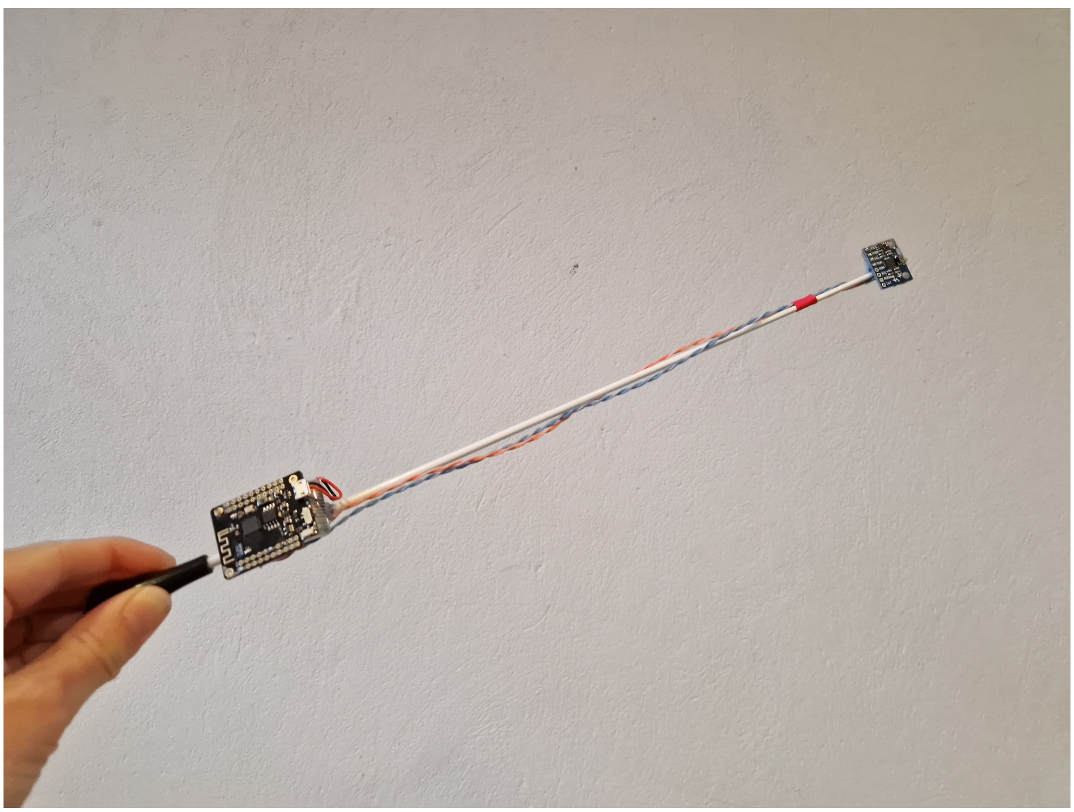

Figure 6: Case Study 2: Sonic Baton.

Figure 7: Case Study 2: Dodekaotto Spatialisation System.

Figure 8: Case Study 2: Magnetic Discs.

References

Armitage, Jack and Thor Magnusson (2023). "Agential Scores: Exploring Emergent, Self-Organising and Entangled Music Notation". In: *Proceedings of the 8th International Conference on Technologies for Music Notation and Representation*. Northeastern University, Boston, Massachusetts, USA.

Arteaga, Daniel (May 2023). "Introduction to Ambisonics". In: DOI: 10.5281/zenodo.7963105.

Barad, Karen (July 2014). "Diffracting Diffraction: Cutting Together-Apart". In: *Parallax* 20. DOI: 10.1080/13534645.2014.927623.

Betancur, Celeste and Jack Armitage (2023). "Pandora's Mycophony". In: *Proc. New Interfaces for Musical Expression*. Mexico City, Mexico.

Betancur Gutierrez, Celeste (Aug. 2023). "AI in Live Coding Environments: Pandora's Dream". In: *AI Music Creativity 2023*. University of Sussex, Brighton, UK: Zenodo. DOI: 10.5281/zenodo.8329742. (Visited on 10/02/2023).

Blackwell, Alan F. et al. (2022). *Live Coding: A User's Manual*. MIT Press. (Visited on 11/27/2022).

Bonneel, Nicolas et al. (Jan. 2015). "Sliced and Radon Wasserstein Barycenters of Measures". en. In: *Journal of Mathematical Imaging and Vision* 51.1, pp. 22–45. ISSN: 1573-7683. DOI: 10.1007/s10851-014-0506-3. URL: https://doi.org/10.1007/s10851-014-0506-3 (visited on 10/02/2023).

Bryan-Kinns, Nick, Berker Banar, et al. (2021). "Exploring XAI for the Arts: Explaining Latent Space in Generative Music". In: *1st Workshop on eXplainable AI Approaches for Debugging and Diagnosis (XAI4Debugging@NeurIPS2021)*. Online.

– (2024). "Explainable AI and Music". In: *TBC (Forthcoming)*. Online: Taylor & Francis.

Bryan-Kinns, Nick, Corey Ford, et al. (2023). *Explainable AI for the Arts - XAIxArts*. https://xaixarts.github.io/. (Visited on 10/02/2023).

Caillon, Antoine and Philippe Esling (2021). "RAVE: A variational autoencoder for fast and high-quality neural audio synthesis". In: *CoRR* abs/2111.05011. arXiv: 2111.05011. URL: https://arxiv.org/abs/2111.05011.

Candy, Linda and Ernest Edmonds (Feb. 2018). "Practice-Based Research in the Creative Arts: Foundations and Futures from the Front Line". In: *Leonardo* 51.1, pp. 63–69. ISSN: 0024-094X, 1530-9282. DOI: 10.1162/LEON_a_01471. (Visited on 04/01/2021).

Fiebrink, Rebecca, Dan Trueman, and Perry R Cook (2011). "The Wekinator: Software for using machine learning to build real-time interactive systems". In.

Janicka, Iwona (2023). "Processes of Translation: Bruno Latour's Heterodox Semiotics." In: *Textual Practice 37, no.6: 846-66*.

Johnston, Andrew (2016). "Opportunities for Practice-Based Research in Musical Instrument Design". In: *Leonardo* 49.1, pp. 82–83. ISSN: 1530-9282. (Visited on 03/03/2016).

Knotts, Shelly and Iván Paz (Jan. 2022). "Live Coding and Machine Learning Is Dangerous: Show Us Your Algorithms". In: *Proceedings of the International Conference on Live Coding*. Valdivia, Chile: Zenodo. DOI: 10.5281/zenodo.5888356. (Visited on 03/03/2022).

Latour, Bruno (2005). *Reassembling the Social: An Introduction to the Actor-Network Theory*. Oxford, England and New York, NY, USA: Oxford University Press.

Llano, Maria Teresa et al. (2020). "Explainable Computational Creativity". In: *International Conference on Computational Creativity*. arXiv. DOI: 10.48550/arXiv.2205.05682. arXiv: 2205.05682 [cs]. (Visited on 11/21/2023).

Mulshine, Michael R et al. (2023). "Webchuck: Computer music programming on the web". In: *New Interfaces for Musical Expression*.

Pickering, Andrew (Nov. 1996). "The Mangle of Practice: Time, Agency, and Science". In: *Bibliovault OAI Repository, the University of Chicago Press* 38. DOI: 10.2307/3106908.

Privato, Nicola and Jack Armitage (Sept. 2023). "A Context-Sensitive Approach to XAI in Music Performance". In: *The 1st International Workshop on Explainable AI for the Arts (XAIxArts), ACM Creativity and Cognition (C&C) 2023*. New York, NY, USA: arXiv. DOI: 10.48550/arXiv.2309.04491. arXiv: 2309.04491 [cs]. (Visited on 10/02/2023).

Privato, Nicola, Thor Magnusson, and Einar Torfi Einarsson (2023). "The Magnetic Score: Somatosensory Inscriptions and Relational Design in The Instrument-Score". In: *Proceedings of the International Conference on Technologies for Music Notation and Representation – TENOR'2023*. Ed. by Anthony Paul De Ritis et al. Boston, Massachusetts, USA: Northeastern University, pp. 36–44. ISBN: 978-0-6481592-7-8.

Privato, Nicola, Thor Magnusson, and Einar Torfi Einarsson (202e). “Magnetic Interactions as a Somatosensory Interface”. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*. Universidad Autonoma Metropolitana, Mexico. DOI: inpublication. URL: https://www.nime.org/archives/.

Schnell, Norbert and Marc Battier (May 2002). “Introducing Composed Instruments, Technical and Musicological Implications”. In.

Vigliensoni, Gabriel, Phoenix Perry, and Rebecca Fiebrink (2022). “A Small-Data Mindset for Generative AI Creative Work”. In: *Generative AI and HCI - CHI 2022 Workshop*. New York, New York, USA: ACM.

Wang, Ge, Perry R Cook, and Spencer Salazar (2015). “Chuck: A strongly timed computer music language”. In: vol. 39. 4. MIT Press One Rogers St., Cambridge, MA 02142-1209, USA journals-info@ mit. edu, pp. 10–29.