TECHNO-UTOPIA: Music Emerging from Colliding Embedded AI Instruments with Radio Orchestras and their Archives

Robert Laidlow
robert.laidlow@kcl.ac.uk
King's College London
London, United Kingdom

Nicola Privato*
nprivato@hi.is
Intelligent Instruments Lab
University of Iceland
Reykjavik, Iceland

Victor Shepardson*
victorshepardson@hi.is
Intelligent Instruments Lab
University of Iceland
Reykjavik, Iceland

Figure 1: Image of Zubin Kanga performing on the Stacco with BBC Philharmonic. Photographer: Robin Clewley

Abstract

This practice-led paper explores the musical work *TECHNO-UTOPIA*, composed for orchestra and soloist performing on traditional acoustic instruments, electronic samplers and embedded AI instruments. *TECHNO-UTOPIA* was co-commissioned by two radio symphony orchestras and marked the first time either has worked with NIMEs.

We examine how these embedded AI instruments functioned as compositional tools that shaped structural, harmonic, and textural decisions across the entire work, extending beyond moment-to-moment sound generation. The paper details the technical development of several RAVE models trained on orchestral archive data and other methods of exploring archives compositionally.

The paper reports on the rehearsal process with professional orchestral musicians, examining how the soloist developed individual performance techniques distinct from the composer's approaches. We address practical challenges including notation systems for NIME performances, strategies for maintaining technological credibility in traditional orchestral settings, and methods for integrating archive-based instruments into live performance. This research contributes insights into ethical creative AI design using archival datasets obtained with express permission, and which pertain to the ensemble at hand and the relationship between NIMEs and established musical communities.

*Both authors contributed equally to this research.

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

Keywords

Embedded AI, Creative Process, Orchestra, Rehearsal, Neural audio synthesis, archives

1 Introduction

TECHNO-UTOPIA is a practice-led research project which explores the creative consequences of introducing new interfaces for musical expression (NIMEs) into music for symphony orchestras. The project, which was undertaken primarily from 2024-25, involved collaboration with two European radio symphony orchestras who acted as co-commissioners and performers: the BBC Philharmonic and the Rundfunk-Sinfonieorchester Berlin.

Its major creative output was a concerto of the same name, which lasts approximately 35 minutes in performance, for soloist Zubin Kanga - performing on piano, keyboards and embedded AI instruments - and orchestra. The NIMES used throughout the project allowed for real-time expressive control of sound archives through machine learning models and DSP algorithms. They are referred to as embedded AI in this paper due to how these algorithms were controlled through tactile, physical, interfaces and instruments, as opposed to operating a computer keyboard or other traditional computing interface directly.

Research, both recent and stretching back several decades, has focussed on a diverse range of intersections between embedded AI (or machine learning) and instrument design, including descriptions of new tactile instruments and interfaces [36], interviews with artists who have designed and worked with their own such interfaces [9] and pipelines for supporting the development of neural networks using specific hardware [27] among others.

This paper presents an exploration of the creative insights resulting from introducing embedded AI instruments into symphony orchestras, an account of the technical details of the NIMEs and embedded AI used, and a report on the rehearsal and performance process.

2 Background

There are numerous examples of valuable work undertaken between the NIME community and non-musicians [24] [18] [26] [14] and non-professional musicians [20], where the focus has been on original means for enabling wider participation in music-making. There has also been productive exploration of the creative benefits of using embedded AI and NIMEs with classical instrumentalists who are technology specialists [35] [22] and artists who are "actively engaged with precisely such questions [...] [who] bring with them years of professional artistic experience, strong opinions on the use of technology in performance, and commitments to public presentations" [11].

There has been comparatively less focus on the interaction between NIME research and the varied ecosystem of professional music performers who are not technology specialists. As more NIMEs begin to engage creatively with embedded AI models, this gap becomes increasingly important. Professional musicians are often wary of "AI" [12] and are generally not aware that AI and machine learning in music is a diverse field, encompassing both experimental creative work in addition to the now well-known commercially-driven prompt-based generative AI paradigm (see: [6] for discussion of such paradigms).

This is a problem for the wider credibility of our research field in today's cultural context and must be overcome if we intend to have a positive effect on that context. This credibility problem is more acute if we (that is to say, the community of researchers and artists interested in exploring human creativity through its intersection with technology) wish to continue to use the term "AI" despite its recent monopolisation by a narrow slice of Silicon Valley enterprises.

Beyond any creative or technical benefits, which are the focus of this paper, working with non-technology specialist musicians and embedded AI in this context will therefore raise awareness of the variety of technological research and how it relates to their practice. It may also positively affect the case for regulating "generative AI" through demonstrating that alternative methods to AI in music do exist, and do not require unethical data scraping practices – a direct rebuttal of common commercial arguments [1]. Cultural organisations are frequent participants in discussions around creative industries and technology and are trusted by their audiences; audiences who frequently share misgivings surrounding the term "AI" [2]. They make ideal partners for those in the community working towards the goal of creative, equitable and ethical AI.

Orchestras, which can be understood as a museum of musical technologies, have particular resonances with NIMEs. These orchestral technologies include both tactile interfaces (i.e. musical instruments) and proficiency with underlying technologies which influence, structure and capture musical thought, such as music notation (see: [25]). Previously, orchestras have readily engaged with new musical interfaces affording creative expression – such as the ondes martinot [1], synthesizers [32] and player pianos [23]. There are also instances of orchestras working with machine learning, including improvising machines [19] and neural synthesis models [16] – which has been called "performance-time" machine learning [8] – and symbolic-generative methods [10] – or "design-time". These interactions, however, have tended to be "composer-like" or "performer-like", in that they generate fixed audio or musical information, rather than "instrument-like" [15], designed specifically to allow a human performer to express themselves in real-time through the interface.

Radio orchestras are particularly well suited for the kind of project described in this paper. These orchestras were generally established internationally in the first half of the 20th century to respond to new technology [28] and, as will be discussed in Section 3, having built up an archive of radio sound recordings which can be integrated into embedded AI design. Finally, orchestras rehearse and perform together, aiming toward a clear and effective performance goal. There have been project write-ups and artistic/technical reflections on rehearsed performances and improvisations using NIMEs [9], but fewer when the performer(s) is not also the person(s) who originally designed the NIME.

3 Personal Artist Motivation

I[2] devised and composed *TECHNO-UTOPIA* as a response to a number of artistic, technical and ethical questions that arose from my previous creative practice with classical musicians and emerging technologies. I have composed for symphony orchestra several times previously, but this marked the first time I had integrated NIMEs into that practice.

[1] e.g., "it would be impossible to train today's leading AI models without using copyrighted materials" - OpenAI; Deposition to UK Parliament Select Committee (2024)

[2] This paper is presented by three authors. The first author is the composer of the work, while the other two authors collaborated to provide technical expertise, instrument fabrication and machine learning model training. "We" refers to work we did together. "I" refers to the first author and its use is intended to avoid the implication that the other two authors exactly share my aesthetic and conceptual opinions.

One question concerns has been previously termed "musical structuralism" [15], in which the dominance of machine learning in our daily lives erases the kinds of music-making that cannot be effectively captured in data [17]. I am interested in exploring the gulf between the act of music-making and the record of music-making – that is to say, the difference between music as process and music as product.

I also intended *TECHNO-UTOPIA* to join, and respond to, the lineage of orchestral works in which the orchestra interacts with a simulacrum of itself or its past performances. Works such as Steen-Andersen's *TRIO* [34] make effective use of sampling technology to do this, while works such as Berio's *sinfonia* [4] do so through structural and notational means.

4 *TECHNO-UTOPIA*'s Data and Models

4.1 Overview of NIMEs

TECHNO-UTOPIA uses three newly- or recently-developed instrumental interfaces. This paper describes each, with focus on the Stacco, an instrument designed for the intuitive control of neural audio synthesis latent parameters through the displacement of magnetic objects on a wooden board concealing four magnetometers [29]. Dramatically, I was interested in featuring a moment where the work's soloist could deconstruct orchestral sounds through physical movement. Having considered other options, including the MiMu gloves [13] and motion capture via video camera, I chose to work with the Stacco for this purpose. We, the three-person author team, spent a week at the Intelligent Instruments Lab at the University of Iceland in November 2024 developing the Stacco and the RAVE models it controls for this project. Our collaboration extended across 2025, up to and beyond the first performance of *TECHNO-UTOPIA* in July 2025.

4.2 Data Provenance

The project took as its sole source of audio data an archive of BBC Philharmonic Radio broadcasts featuring the BBC Philharmonic to which I was granted access by BBC Radio 3 for this project. This was important to the project for two reasons: first, all data was obtained with express permission from the rights-holder; second, the data was specific to the ensemble that would be performing the music and was therefore site-specific in a way that posed aesthetic questions. The full archive stretched back some thirty years and contained over 2600 hours of music. We used stereo bounces which were originally recorded at 48 kHz using microphone setup suitable for professional radio orchestra.

4.3 RAVE Models

RAVE [5] is a variational autoencoder for real-time neural audio synthesis, trained to reconstruct a corpus of sounds via a multi-dimensional latent space that can be audibly explored and manipulated. It has been used widely in new instrument design, improvisation practices and music composition [36] [29] [7]. Subsets of the BBC Philharmonic archive were used to train a range of experimental RAVE models, of which five were considered for the final project:

(1) Full orchestra audio "Model 1" (training data approximately 18.5GB)
(2) Strings-only audio (training data <1GB)
(3) Brass-only audio (training data approximately <1GB)
(4) Woodwind-only audio (training data approximately <1GB)
(5) Percussion-only audio (training data approximately <1GB)

Models were trained using a causal RAVE v3 architecture at 48kHz on an Nvidia A4000 or A5000 GPU. ""Model 1"" was trained first, which took approximately two weeks, with the remaining models trained utilising transfer learning [30], which took approximately 3 days each.

Data for models 2-4 was compiled by selecting broadcasts from the archive only involving certain orchestral sections or by manually segmenting longer symphonic works. Data for model 5 was recorded for the project with the BBC Philharmonic's principal percussionist. Music composed by living composers was removed from all training data for models used in the final work, as were spoken segments by BBC radio presenters. These RAVE models were controlled through the Stacco instrument via Bela and Max 8. Models were navigated directly (magnet location mapped to RAVE latents) and indirectly through two methods: capturing and looping orchestral audio, encoding it into latent space, then perturbing the decoder further through magnetometer input; and sending eight waveforms to the eight latents before adjusting their (de)coherence, amplitude and frequency via magnetometer input.

Investigation of latent representations in "Model 1" through sustained author interaction revealed unexpected characteristics (Table 1). Unlike typical RAVE models where the first latent determines dynamic [29], ""Model 1"'s" first latent seemed to encode timbral "thickness" (number of "instruments" playing). Seven of eight latents mapped to timbral elements in some way (i.e., richness or instrument type) with most representing liminal interpolations between colours rather than specific instruments. Latent 7 was an exception, representing a specific continuum from string timbres (low values) to brass (high values). Dynamics were entangled across latents 2, 4, and 8, requiring simultaneous magnet manipulation for intuitive control and despite the dataset's enormous pitch range, pitch correlated strongly only with latent 6 and weakly with the overtones represented in latent 2.

4.4 Concatenative Synthesis

During the first movement, the soloist performed on a ROLI Seaboard 2 (a keyboard with MPE capabilities) using a preset synth voice. A Max patch mapped finger pressure to the balance between preset synth and live, granular, concatenative reconstruction using approximately 5 hours of the BBC Philharmonic archive. The concatenation was modelled using FluCoMa [21] machine learning tools before being re-implemented with a commercial plugin [3] to maximize stability and consistency between composer and performer machines.

4.5 Archive Dreamer

BBC Philharmonic archive data was also used as part of what was dubbed the "Archive Dreamer". The Archive Dreamer initially used FluCoMa machine learning tools to identify an input pitch in real-time before matching that pitch automatically with a randomly selected 10-second snippet of the archive which also featured the same pitch according to a prior analysis of the corpus. Input was provided by the soloist's piano, transforming the piano into a site of live archive exploration. Initial real-time lookup in FluCoMa was replaced with a sampler-based approach selecting randomly between pre-selected snippets due to latency constraints and the discovery that live lookup across a much larger archive didn't meaningfully increase performer expressivity.

Table 1: Author Exploration of Latent Representation

Latent	Representation
1	Timbral "thickness" of sound
2	Dynamic and harmonic richness (overtones)
3	Timbre (i.e., which "instrument" forms the central sound output)
4	Timbre and dynamic
5	Thickness and timbre
6	Pitch
7	Timbre – specifically strings and brass
8	Harmonic warmth, pitch and dynamic

4.6 Where Is The Interface?

None of these interfaces is strictly "new"; Stacco was presented at NIME Conference in 2024, the ROLI Seaboard 2 was released in 2022 and the piano is, of course, centuries old.

To an orchestral community, however, each of these interfaces (excepting the piano) and means of expressive control over sound were, generally, truly new. Furthermore, we suggest that it is the artistically fertile relationship between this community and the expressive and adaptive existing NIMEs that provides creative interest and novelty. This relationship is the new interface for musical expression that really matters in the project. The archive itself is also presented as a major site of interaction. Each NIME interacts with it, or parts of it, in one or more ways. Modern technology, however, is only one method of interfacing with an archive of sounds to create new music. The orchestra's museum of technologies can do the same, especially when working with the musicians who created the archive in the first place.

5 NIMEs in the *TECHNO-UTOPIA* Composition Process

Working closely with the NIMEs described above affected the entire compositional process, including parts of the composition that did not involve performing using a NIME at all. This section details demonstrative examples from each of *TECHNO-UTOPIA*'s three interconnected movements.

TECHNO-UTOPIA's first movement (titled "In My Image") features the soloist making extensive use of live concatenative synthesis using the orchestral archive. This allowed the soloist to "emerge" subtly from the orchestra's sound and enabled compositionally "mirror" moments where the orchestra copied the soloist's concatenated material to create an uncanny echo. This NIME offered a genuinely new perspective on traditional orchestral textures in a (piano) concerto through its ability to modulate how far the soloist is sonically impersonating the ensemble.

The movement is structured as a closed circle: it ends with the orchestra performing the soloist's original music, which becomes the looped material that opens the second movement's Stacco cadenza. This looped material therefore functioned as both compositional endpoint and starting point. Since the looped audio is reconstructed in the "Model 1" latent space rather than preserved as original orchestra sound (sharing DNA with instruments such as the Living Looper [31], the swap between real orchestra and the Stacco is not perfect. There is an element of uncanniness that arises from the subtle differences between original and reconstruction that are heard side-by-side, even before the latent space is further perturbed by Stacco magnet manipulation. The orchestral material needed to be simple enough for these imperfections in reconstruction to be audible.

The movement's thematic structure involves five distinct and complex orchestral soundworlds each transforming into that final looping material through processes of timbral, melodic and harmonic distortion. The soundworlds themselves were composed after the looping material was decided, so were a response to it (a kind of negative – I imagined what five "opposites" of the looped material might be). The affordances of the Stacco and RAVE therefore directly influenced the structure, harmony and textural decision-making throughout the first movement, even when the Stacco was not being played.

TECHNO-UTOPIA's second movement is an extended cadenza for the Stacco instrument, lasting 5-8 minutes. Aesthetically, the intended effect was for the cadenza to begin as though the soloist had "stolen" the music of the orchestra before deconstructing that music into its constituent sonic parts.

The Stacco cadenza itself involved exploration of three RAVE models ("Model 1", "Percussion" and "Brass") and an FM synthesizer. MIDI faders were used to control which sound(s) were present in the mix. The cadenza was notated as a text score for the soloist, giving a framework for the cadenza but leaving details and timings only partially defined. This is partly because a cadenza is often semi-improvised in orchestral music, but also because traditional staff notation was unable to rise to the challenge of notating music for the Stacco. The instrument design deliberately lends itself toward experimentation and exploration, which in turn introduced those elements into the composition where they would not otherwise occur; exploration and improvisation normally happens in my orchestral work in the composition process, not during performance.

"Model 1" is extremely sensitive and, subjectively, a joy to perform with that seems to offer more each time one returns to it. I found while making the text score, I was interested in hearing the soloist luxuriating in sound and enjoying the liminality of sounds found in the latent space. This is the opposite of the hard-edged music of the first movement.

Because none of the latents were unimportant (Section 4.3) and none were directly linked to loudness, there was no "principal" magnet for "Model 1". This luxuriation could happen when focusing on any individual magnet.

Specific points in the latent space relating to specific sounds discovered during experimentation were notated on the instrument itself (Figure 2) so the soloist could easily navigate to them during performance. This form of notation is similar to how starting string players might mark particular notes in chalk on their fretboard or to how pianists might mark where to press lightly on a string to produce a specific harmonic when performing inside the piano.

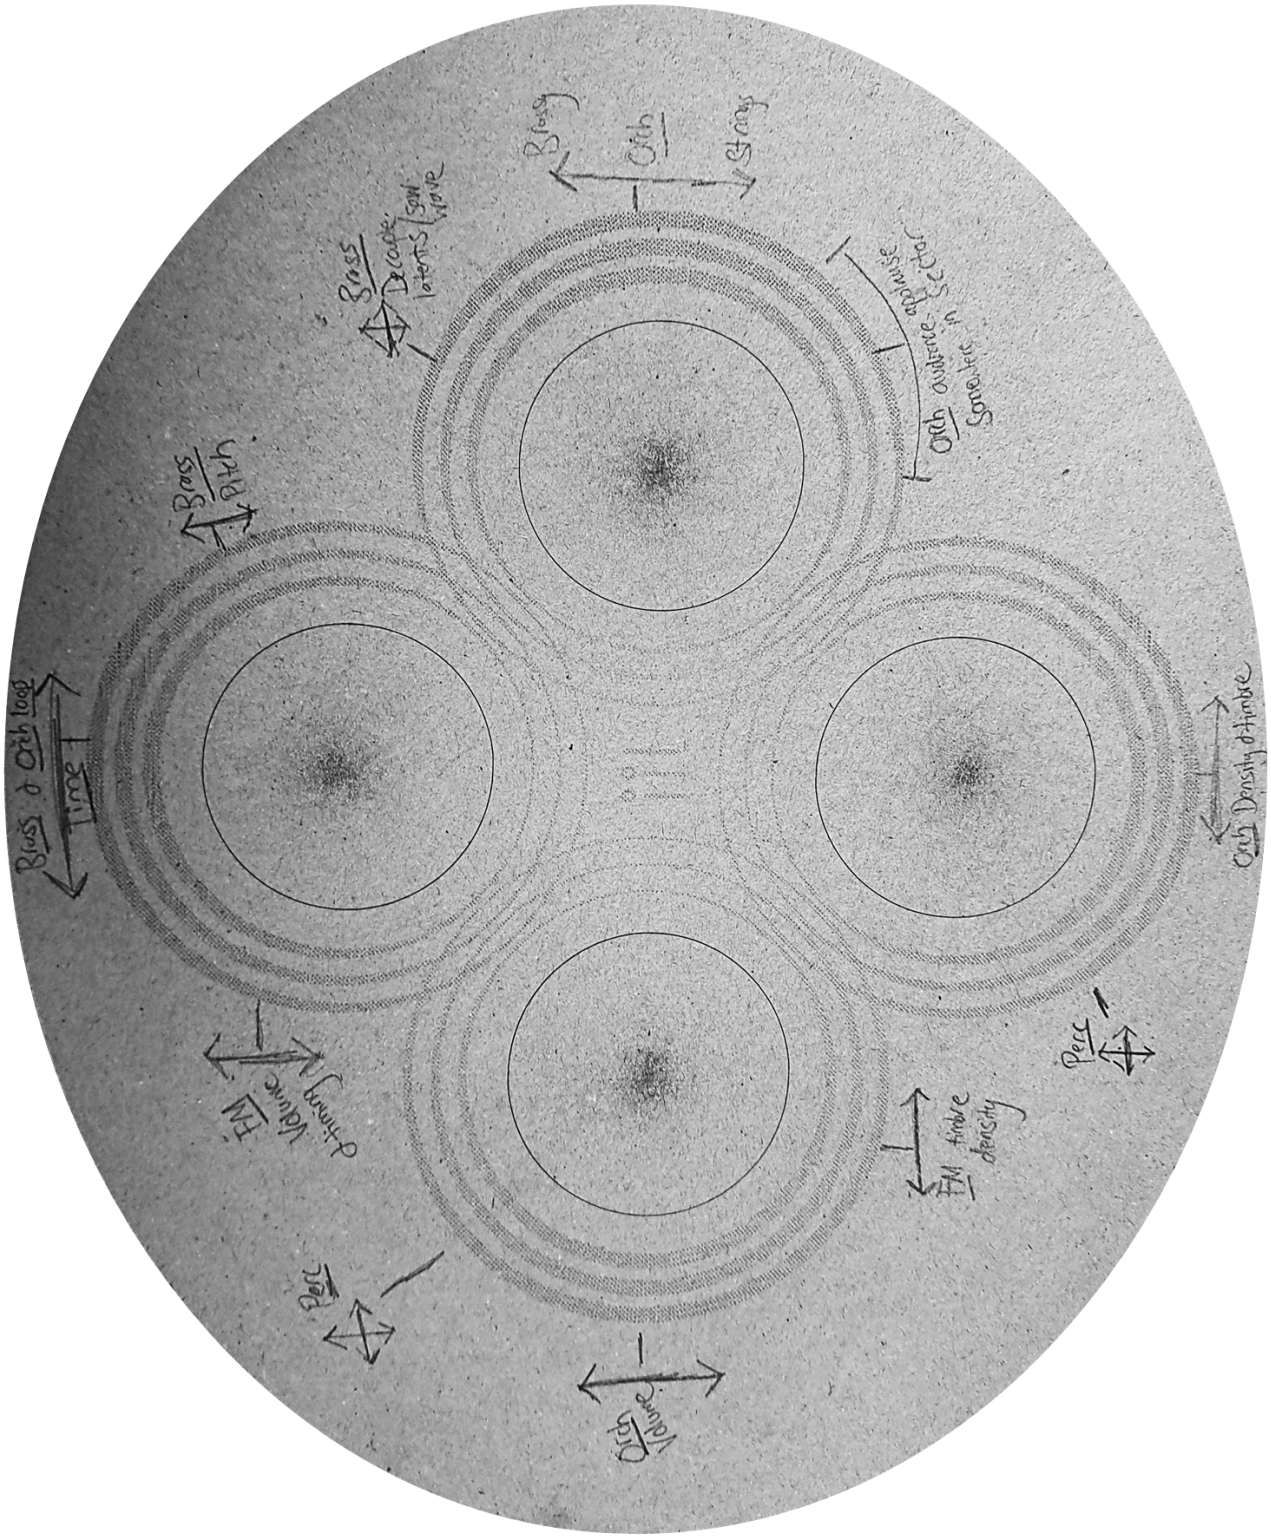

Figure 2: Scan of the surface of the Stacco

I was inspired by the way that "Model 1" interpolated seamlessly between different areas of the archive's latent space. I therefore composed the third movement (titled "Everything I Have Ever Heard Is Just Data") to translate this effect to the orchestra itself. The movement stitches together around 100 pieces of music from both the orchestra's archive and my own memories of music-making. Sometimes these references were notated (ie, I compositionally reconstructed moments from the archive) and sometimes they were literal archive samples. In this movement, music for the Stacco is notated with slashes (like jazz improvisation sections) with written instructions above the staff (Supplementary audio file 3). The end of the third movement features Archive Dreamer triggering archive fragments.

Interfacing with the archive remained forefront of the compositional process, even if its methods of doing so are very different to either previous movement.

To me, this is the advantage of colliding a NIME with an acoustic ensemble – it provides an opportunity to explore different ways of imagining what the act of making music is, what it means to make a record of music, and how an ensemble and technology might interact with the past. There was no distinction between the NIMEs themselves and the aesthetic questions they raised during the composition process. Their creative consequences were not limited to the moment-to-moment sounds they could make but rather were widespread across many approaches to interaction.

6 Notes on the Rehearsal, Performance and Recording Process

Ahead of its first performance and recording[3] for BBC Radio 3 by the BBC Philharmonic, there were eight hours of rehearsal scheduled with the orchestra and an additional tech setup day. New works of this length in the UK would generally be scheduled up to four hours. The additional time was granted by orchestral management on account of the work's technical complexities.

Before the first rehearsal, I gave a five-minute introduction to the work and its themes. I explained that I was interested in designing creative AI systems that were both interesting for artists and ethical and contrasted these to the commercially available generative AI models which had been trained without permission on copyrighted data (including many recordings of these same musicians). I demonstrated the NIMEs and explained how their orchestral recordings had been used to create these instruments. Finally, I discussed the aesthetic goals of the work.

The rehearsals themselves focussed purely on musical questions, led by the conductor Jack Sheen (Figure 3). Despite our relatively long rehearsal period, there remained a lot to get through and there was no time at all to waste in serious or lengthy technological failures. From my past experience, such failures would result not only in wasted time but also a loss of credibility in the work from the orchestral musicians (similar to if there were mistakes in the parts, or information missing from the conductor's score). Working with an orchestra is like a spell; magical, but once the enchantment begins to unravel for any reason, it can vanish into thin air. It was crucial that this credibility was maintained, since the project was just as much about introducing orchestral musicians to the world of embedded AI as it was about studying the creative effects of NIMEs. Avoiding live technological failures was one reason that I chose to recreate the Archive Dreamer rather than scan the archive live (Section 4.5).

[3]https://on.soundcloud.com/7dBc94yMOi0ewbrC7O

Figure 3: Jack Sheen conducting the BBC Philharmonic. Photographer: Robin Clewley

The orchestral musicians were very receptive towards the NIMEs and our approach, and particularly towards the Stacco which many of them wanted to learn more about, try out for themselves and make videos to send to their colleagues. I speculate they were particularly interested in the Stacco because it looks like a new instrument, whereas the other NIMEs only sound new. It was repeatedly commented that the instrument reminded players of the theremin. The BBC Philharmonic had in recent years performed Kalevi Aho's Eight Seasons (Concerto for Theremin and Chamber Orchestra) at the BBC Proms. For the Stacco to be immediately understood as an analogue to a historical "NIME" was heartening. Some of them, forgetting the instrument's name, took to calling it the "magic ball instrument".

Rehearsing with the soloist, pianist Zubin Kanga, was an illuminative process. Zubin had no hand in designing any of the NIMEs but has previously performed with an array of advanced technologies [33]. An obvious, but important, point was that his methods of interacting with the NIMEs were different to mine. The way that he chose to glide between original pitch and concatenated archive synthesis on the ROLI Seaboard 2 (through finger pressure), the length of time he left between triggering Archive Dreamer samples and, above all, the way that he interfaced with the Stacco in the cadenza and third movement were highly individualistic. All of these are rooted in the act of listening; listening to the NIME and the orchestra and responding.

During the Stacco cadenza, I found that Zubin's style was more exploratory and contemplative than mine. He would concentrate on a single magnet for an extended period, moving manipulating it with tiny motions and letting the result sit. Supplementary audio files 1 and 2 demonstrate the cadenza performed by the first author (demo) and by the work's soloist (first performance) respectively. This supports the existing evidence that the Stacco is a robust musical instrument that can support diverse playing techniques (see: [29] Section 6).

During the third movement, there is a passage where the soloist performs on the Stacco simultaneously with a sample of birdsong and orchestral material quoting archive material. Zubin understood his role here to be in dialogue with the birdsong, which had not been my original intent while composing. Nonetheless, he was able to achieve his vision, finding sounds which imitated birdsong, within the constraints of the selected model. To me this demonstrates the instrument's diverse utility and can be heard in Supplementary audio file 3 (from 20s).

7 Evaluation and Further Questions

This project demonstrates how embedded AI systems can function as compositional tools that extend beyond moment-to-moment sound generation to shape structural, harmonic, and textural decisions across an entire work. The successful integration of these technologies into traditional orchestral rehearsal and performance contexts, evidenced by musicians' enthusiastic reception and the soloist's distinct performance approaches, suggests that the "newness" of such interfaces lies as much in their capacity to inspire creatively versatile relationships between technology and established musical communities as in their technical novelty.

A straightforward area for further work is in integrating NIMEs more fully into the fabric of a classical ensemble. In *TECHNO-UTOPIA*, the orchestral musicians performed alongside NIMEs but not with them. It was originally suggested that we might try the latter but later agreed with the orchestral management to restrain the scope of new technology to a soloist who had plenty of time in advance to work on the project. This is a difficult problem to overcome. Learning how to perform on a NIME requires time and practise. It is extremely rare to have time with orchestral musicians before the first rehearsal, by which point the music needs to be complete. Issues of notation would need to be solved to better balance between what an orchestral player expects to see in front of them with the inherent characteristics of the NIME itself. The problem of notation is one that is more broadly applicable to the NIME field because it is inherently entangled with the question of how these interfaces might be adopted by people who did not design them.

Further directions for future work include a wider investigation of the insights professional musicians can bring to understanding latent space dimensions in audio models, expanding an instrument like the Archive Dreamer beyond simply pitch-class matching, and further documenting how performer-specific techniques emerge through sustained engagement with these instruments in ensemble settings. We would like to see NIMEs interacting with many more types of non-technical communities, as we believe this can only be positive for their impact on the wider musical ecosystem and further strengthening their design methodologies and, ultimately, will result in exciting, relevant and unique artistic work.

8 Ethical Standards

This work was done following the university's policies and did not require formal ethical approval. It did not involve personal data collection. Orchestra members consented to be photographed ahead of the rehearsal period. All neural synthesis models were trained with data obtained with permission from the rights holder. Musicians were remunerated according to their contracts with the relevant radio orchestra(s).

Acknowledgments

This project would not have been possible without funding from Jesus College, University of Oxford and in-kind funding and support from BBC Radio 3 and the Rundfunk Sinfonieorchester Berlin. The Intelligent Instruments Lab is supported by the European Research Council (ERC) as part of the Intelligent Instruments project (INTENT), under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848) This research was also supported by an NVIDIA hardware grant. I would like to thank every musician who took part in this project.

References

[1] Thomas Adès. 2016. The Exterminating Angel. Musical Work.
[2] Alessandro Ansani, Friederike Koehler, Lisa Giombini, Matias Hämäläinen, Chen Meng, Marco Marini, and Suvi Saarikallio. 2025. AI Performer Bias: Listeners Like Music Less When They Think it was Performed by an AI. *Empirical Studies of the Arts* 43 (7 2025), 1137–1161. Issue 2. https://doi.org/10.1177/02762374241308807
[3] DataMind Audio. 2025. *Concatenator.* Retrieved February 12th, 2026 from https://datamindaudio.ai/concatenator-v1/
[4] Luciano Berio. 1968. Sinfonia. Musical Work.
[5] Antoine Caillon and Philippe Esling. 2021. RAVE: A variational autoencoder for fast and high-quality neural audio synthesis. (11 2021).
[6] Luca Casini, Laura Cros Vila, David Dalmazzo, Anna-Kaisa Kaila, and Bob L. T. Sturm. 2025. Data-Driven Analysis of Text-Conditioned AI-Generated Music: A Case Study with Suno and Udio. (9 2025). http://arxiv.org/abs/2509.11824
[7] Ricardo Climent. 2024. Sonic Tales of a Windsurfer (2024)-17: 45| 4-channel Fixed media. Musical Work.
[8] Rebecca Fiebrink and Baptiste Caramiaux. 2018. *The Machine Learning Algorithm as Creative Musical Tool.* Oxford University Press, 181–208. https://doi.org/10.1093/OXFORDHB/9780190226992.013.23
[9] Rebecca Fiebrink and Laetitia Sonami. 2020. Reflections on Eight Years of Instrument Creation with Machine Learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[10] Mark Gotham, T U Dortmund, Kunpeng Song, Nicolai Böhlefeld, and Ahmed Elgammal. 2022. Beethoven X: Es könnte sein! (It could be!). In *Conference on AI Music Creativity (AIMC 2022).*
[11] Ian Hattwick and Marcelo M Wanderley. 2017. Design of Hardware Systems for Professional Artistic Applications. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[12] Jonathan Herington, Raffaella Borasi, Benjamin J. Guerrero, David E. Miller, Blaire Koerner, Yu Jung Han, Zenon Borys, and Rachel Roberts. 2025. Musicians' ethical concerns about AI: an interview study. *AI and Society* (2025). https://doi.org/10.1007/s00146-025-02601-6
[13] Zubin Kanga. 2023. The Cyborg Hand: Gesture, Technology, Disability and Interdisciplinarity in Whatever Weighs You Down. *Contemporary Music Review* 42, 3 (2023), 319–338.
[14] André Knörig, Boris Müller, and Reto Wettach. 2007. Articulated Paint: Musical Expression for Non-Musicians. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[15] Robert Laidlow. 2022. *Artificial Intelligence Within the Creative Process of Contemporary Classical Music.* Ph. D. Dissertation. Manchester Metropolitan University. https://e-space.mmu.ac.uk/631170/
[16] Robert Laidlow. 2022. Silicon. Musical Work.
[17] Robert Laidlow. 2024. *Classical Music Futures: Practices of Innovation.* Open Book Publishers, Chapter Artificial Intelligence and the Symphony Orchestra.
[18] Jeppe Veirum Larsen and Hendrik Knoche. 2017. Hear You Later Alligator: How delayed auditory feedback affects non-musically trained people's strumming. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[19] George E. Lewis. 2024. The Reincarnation of Blind Tom (Double Concerto for Human Soloist, AI-Pianist and Orchestra). Musical Work.
[20] Anders Lind. 2018. ANIMATED NOTATION IN MULTIPLE PARTS FOR CROWD OF NON-PROFESSIONAL PERFORMERS. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[21] Fluid Corpus Manipulation. 2026. *Fluid Corpus Manipulation.* Retrieved February 12th, 2026 from https://www.flucoma.org/
[22] John A Mcdowell and Dermot J Furlong. 2018. Haptic-Listening and the Classical Guitar. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[23] Ensemble Modern. 2011. Nancarrow: As fast as possible. CD.
[24] Lucía Montesinos, Halfdan Hauch Jensen, and Anders Sundnes Løvlie. 2025. Enabling Embodied Music-Making for Non-Musicians. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[25] Walter J. Ong. 2011. *The Linguistics of Literacy.* John Benjamins Publishing Company, Chapter Writing is a technology that restructures thought.
[26] Dale E Parson. 2009. Chess-based Composition and Improvisation for Non-musicians. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*
[27] Teresa Pelinski, Rodrigo Diaz, Adán L Benito Temprano, and Andrew Mcpherson. 2023. Pipeline for recording datasets and running neural networks on the Bela embedded hardware platform. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* https://onnxruntime.ai/
[28] Joseph E. Potts. 1955. European Radio Orchestras: Western Germany. *The Musical Times* 96 (9 1955), 473. Issue 1351. https://doi.org/10.2307/938221
[29] Nicola Privato, Victor Shepardson, Giacomo Lepri, and Thor Magnusson. 2024. Stacco: Exploring the Embodied Perception of Latent Representations in Neural Synthesis. In *Proceedings of the International Conference on New Interfaces for Musical Expression.*

[30] Victor Shepardson. 2025. *Intelligent Instruments Lab RAVE Fork*. Retrieved February 12th, 2026 from https://github.com/victor-shepardson/RAVE
[31] Victor Shepardson and Thor Magnusson. 2023. The Living Looper: Rethinking the Musical Loop as a Machine Action-Perception Loop. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.
[32] Matthew Shlomowitz. 2019. Glücklich, Glücklich, Freude, Freude. Musical Work.
[33] Cyborg Soloists. 2026. *Cyborg Soloists*. Retrieved February 12th, 2026 from https://www.cyborgsoloists.com/
[34] Simon Steen-Andersen. 2019. TRIO. Musical Work.
[35] Halla Steinunn Stefánsdóttir and Thor Magnusson. 2025. Of altered instrumental relations: a practice-led inquiry into agency through musical performance with neural audio synthesis and violin. *Frontiers in Computer Science* 7 (2025). https://doi.org/10.3389/fcomp.2025.1578595
[36] Federico Visi. 2024. The Sophtar: a networkable feedback string instrument with embedded machine learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression*.