

# A Mapping Strategy for Interacting with Latent Audio Synthesis Using Artistic Materials

SHUOYANG ZHENG, Centre for Digital Music, Queen Mary University of London, UK

ANNA XAMBÓ SEDÓ, Centre for Digital Music, Queen Mary University of London, UK

NICK BRYAN-KINNS, Creative Computing Institute, University of the Arts London, UK

This paper presents a mapping strategy for interacting with the latent spaces of generative AI models. Our approach involves using unsupervised feature learning to encode a human control space and mapping it to an audio synthesis model’s latent space. To demonstrate how this mapping strategy can turn high-dimensional sensor data into control mechanisms of a deep generative model, we present a proof-of-concept system that uses visual sketches to control an audio synthesis model. We draw on emerging discourses in XAIxArts to discuss how this approach can contribute to XAI in artistic and creative contexts, we also discuss its current limitations and propose future research directions.

## ACM Reference Format:

Shuoyang Zheng, Anna Xambó Sedó, and Nick Bryan-Kinns. 2024. A Mapping Strategy for Interacting with Latent Audio Synthesis Using Artistic Materials. In *Proceedings of Explainable AI for the Arts Workshop 2024 (XAIxArts 2024)*. ACM, New York, NY, USA, 4 pages.

## 1 INTRODUCTION

Research in the field of eXplainable AI for the Arts (XAIxArts) examines how AI models could be made more understandable for creative activities [1]. One approach to increasing the explainability of an AI model is to provide real-time interaction and feedback [2]. This approach allows the user to learn how the model works and find ways to interact with it by exploring how their actions affect the model’s output [11]. In practice, it involves mapping a human control space to the parameter space of the AI model and providing real-time audio feedback. For example, mapping a user’s spatial coordination in an immersive environment to the latent space of a deep audio synthesis model so that the user learns the effect of specific latent space axes by ‘walking through’ the environment [13].

The design of this human control space is an open question in the field of XAI [3], especially in a creative and artistic context. While typical approaches have explored parametric control mechanisms such as sliders and 2D pads [2, 15], it is important to consider other interaction mediums that may be more expressive in creative and artistic contexts [5]. These interaction mediums can be the artistic materials from the artists’ performance space [1], for example, the movements of dancers [11], sketches [9], or electromyography data of performers [14]. However, materials in an artistic performance space can be high-dimensional sensor data, which are difficult to map to the AI model and use as a medium for interaction [12]. To tackle this, we propose a mapping strategy that encodes a human performance space and maps it to the latent space of an AI model to help users interact with the model’s generation process. We use a sketch-to-sound controller we developed in a previous work [16] to demonstrate how this mapping strategy can turn high-dimensional sensor data into control mechanisms of an AI model.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

## 2 MAPPING STRATEGY

Materials in an artistic performance space can be high-dimensional sensor data such as sketches [9] or movements of dancers [11]. In order to map artistic materials to the latent space of an AI model, a two-step approach is needed: (i) feature extraction techniques are required to encode these data into representative features [12], (ii) these representative features need to be meaningfully mapped to the latent space of the AI model. Therefore, as shown in Figure 1, we propose using latent mapping [11] and Interactive Machine Learning (IML) [7] to achieve these two steps.




Fig. 1. Proposed mapping strategy

**Step 1** uses latent mapping [11], an unsupervised feature learning approach, to learn representative features from a large corpus of unlabelled data. A trained latent mapping model encodes high-dimensional inputs into lower-dimensional latent representations, holding representative features of the artist’s current movements, sketches, or other artistic practices, which can be mapped to sound controls. Commonly used unsupervised feature learning methods include Variational AutoEncoders [11] and unsupervised dimensionality reduction [12]. In practice, training a latent mapping model requires the artist to create a training dataset for their artistic materials. In other words, they need to record a full set of possible outputs that they might produce in the performance space [11]. For example, training a latent mapping model for a dancer may require recording body skeletal data from the dancer of all their possible dance movements, including intermediate transitions between two movements.

**Step 2** uses Interactive Machine Learning (IML) to connect the extracted features and the AI model’s parameters. IML allows users to construct mappings between two spaces using a few personalised training examples [7]. In this way, it creates a mapping between the human performance space and the AI model’s output, allowing artists to interact with the AI model using data from the performance space. In practice, the artist manually defines a few paired data points in the performance space and the AI model’s parameter space as training examples. Tools such as *Wekinator* [7] can be used to facilitate this process.

## 3 DEMONSTRATION

To demonstrate how this mapping strategy works in practice, we build on our previous work on a sketch-to-sound controller that encodes visual sketches into sound controls. In this work, we use it as the latent mapping model to control a latent audio synthesis model. Technical details for the sketch-to-sound controller can be found in [16]. We use RAVE [4], a variational autoencoder that synthesises audio from a latent space in real-time. We use a RAVE model pre-trained on organ music recordings [8] that runs in a Max4Live<sup>1</sup> device. The pre-trained model has 16 latent space dimensions, and the sketch-to-sound model encodes sketches into a 32-dimensional latent space. Therefore, the IML model takes a 32-dimensional input and maps it to a 16-dimensional output. The communications between the sketch-to-sound model, the IML model, and the RAVE model are achieved through open sound control (OSC)<sup>2</sup>.

<sup>1</sup><https://maxforlive.com/library/>

<sup>2</sup><https://ccrma.stanford.edu/groups/osc/index.html>




Fig. 2. Two screenshots of user interfaces showing the sketch and the audio synthesis model’s latent spaces. A full demonstration video can be viewed at <https://tinyurl.com/xaixarts24-mapping>.

Figure 2 shows two screenshots of the user interface. In each screenshot, the left panel allows the user to interact with the RAVE audio synthesis model by sketching on the canvas, while the right panel allows users to train the IML model. When recording training examples for the IML model, users can start by either manually adjusting the latent parameters or randomising the RAVE model’s latent parameters using the “randomise” button, then clicking the “record” button and drawing sketches that depict the current sound produced by the RAVE model. After training the IML model, users can click the “run” button and interact with the RAVE model by drawing sketches.

#### 4 LIMITATIONS AND FUTURE WORK

This section draws on emerging discourses in XAIxArts [3] to discuss the proposed mapping strategy in terms of the temporal and cross-modal aspects. It also discusses the limitations and contributes with future research directions.

First, **the temporal aspect**. Real-time interaction and feedback in an XAI model can help artists discover and reuse unexpected phenomena for their artistic expression [1]. For example, in our demonstration video, movements such as quick transitions from short lines to full circles trigger interesting sounds. Users can reuse these sounds by memorising the sketch trajectories that trigger them, but is there a way to explain these trajectories to the users in the first place? Typical XAI approaches focus on explaining a static state of the model parameters [6], but explaining the trajectories of sketches focuses on explaining the transition, or the interpolation from one state to another, which involves taking account of the temporal dimension. This aspect of explanation contributes to the embodied knowledge about the AI model because it tells artists how to act with a system, which is an important component in artistic and creative contexts [10].

Second, **the cross-modal aspect**. Investigating interactions that involve cross-modalities, such as shape-sound associations, is useful to expand the explanation mediums in creative and artistic contexts [5]. In latent audio models, for example, future research can explore how human-interpretable features in non-audio domains can be applied to audio latent space. Typical XAI approaches aim to establish meaningful features in the audio domain, and a number of these features already rely on cross-modal metaphors to express the sonic characteristics, such as “brightness”, “noisiness”, and so on [9]. Our mapping strategy can be used to study other cross-modal mappings that might be useful in XAI. For example, future work on our sketch-to-sound system can explore programming and animating sketches using interpretable features such as rotation, zoom, or other affine transformations, and then use our mapping strategy to connect these features with the audio latent space to discover meaningful shape-sound associations that can be used to explain the audio synthesis model.

The motivation for the mapping strategy presented in this paper was to encourage explorations on explanation mediums that are salient in the creative and artistic contexts [5]. However, we only use sketch-to-sound to demonstrate

how the mapping strategy works. It would be necessary to explore other forms of creative practices and go beyond sketch and sound to generalise the discussion into wider creative and artistic contexts.

## 5 CONCLUSION

Real-time interaction and feedback increase the explainability of an AI model by supporting users in exploring how their actions affect the audio generation. Following this direction, this paper proposed a mapping strategy for interacting with latent audio synthesis models using data from an artistic performance space. We have described how this mapping strategy works in practice and demonstrated it with a system that controls an AI audio synthesis model using visual sketches. We discussed that future work using this mapping strategy should take account of the temporal and cross-modal XAI aspects, and it is important to experiment with how this mapping strategy can be used in other forms of creative practices in future work.

## ACKNOWLEDGMENTS

Shuoyang Zheng is a research student at the UKRI Centre for Doctoral Training in Artificial Intelligence and Music, supported by UK Research and Innovation [grant number EP/S022694/1].

## REFERENCES

- [1] Nick Bryan-Kinns. 2024. Reflections on Explainable AI for the Arts (XAIxArts). *Interactions* 31, 1 (Jan. 2024), 43–47. <https://doi.org/10.1145/3636457>
- [2] Nick Bryan-Kinns, Berker Banar, Corey Ford, Courtney N Reed, Yixiao Zhang, Simon Colton, and Jack Armitage. 2021. Exploring XAI for the Arts: Explaining latent space in generative music. In *1st Workshop on eXplainable AI approaches for debugging and diagnosis*.
- [3] Nick Bryan-Kinns, Corey Ford, Alan Chamberlain, Steven David Benford, Helen Kennedy, Zijin Li, Wu Qiong, Gus G. Xia, and Jeba Rezwana. 2023. Explainable AI for the Arts: XAIxArts. In *Creativity and Cognition*. ACM, Virtual Event USA, 1–7. <https://doi.org/10.1145/3591196.3593517>
- [4] Antoine Caillon and Philippe Esling. 2021. RAVE: A variational autoencoder for fast and high-quality neural audio synthesis. (2021). <https://arxiv.org/abs/2111.05011> arXiv: 2111.05011.
- [5] Michael Clemens. 2023. Explaining the arts: Toward a framework for matching creative tasks with appropriate explanation mediums. In *Proceedings of The first international workshop on eXplainable AI for the Arts (XAIxArts)*.
- [6] Shipi Dhanorkar, Christine T. Wolf, Kun Qian, Anbang Xu, Lucian Popa, and Yunyao Li. 2021. Who needs to know what, when?: Broadening the Explainable AI (XAI) Design Space by Looking at Explanations Across the AI Lifecycle. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference (DIS '21)*. Association for Computing Machinery, New York, NY, USA, 1591–1602. <https://doi.org/10.1145/3461778.3462131>
- [7] Rebecca Fiebrink, Dan Trueman, and Perry R. Cook. 2009. A meta-instrument for interactive, on-the-fly machine learning. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Zenodo, Pittsburgh, PA, United States, 280–285. <https://zenodo.org/record/1177513>
- [8] Intelligent Instruments Lab. 2023. rave-models (Revision ad15daf). <https://doi.org/10.57967/hf/1235>
- [9] Sebastian Lobbers and George Fazekas. 2023. SketchSynth: A browser-based sketching interface for sound control. *Proceedings of the International Conference on New Interfaces for Musical Expression (May 2023)*, 637–641. [http://nime.org/proceedings/2023/nime2023\\_95.pdf](http://nime.org/proceedings/2023/nime2023_95.pdf)
- [10] Tim Murray-Browne and Panagiotis Tigas. 2021. Emergent interfaces: Vague, complex, bespoke and embodied interaction between humans and computers. *Applied Sciences* 11, 18 (2021). <https://doi.org/10.3390/app11188531>
- [11] Tim Murray-Browne and Panagiotis Tigas. 2021. Latent mappings: Generating open-ended expressive mappings using variational autoencoders. In *International Conference on New Interfaces for Musical Expression*. <https://doi.org/10.21428/92fbeb44.9d4bcd4b>
- [12] Gerard Roma, Owen Green, and Pierre Alexandre Tremblay. 2019. Adaptive mapping of sound collections for data-driven musical interfaces. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Marcelo Queiroz and Anna Xambó Sedó (Eds.). UFRGS, Porto Alegre, Brazil, 313–318. <https://doi.org/10.5281/zenodo.3672976> ISSN: 2220-4806.
- [13] Hugo Scurto and Ludmila Postel. 2023. Soundwalking deep latent spaces. In *Proceedings of the 23rd International Conference on New Interfaces for Musical Expression (NIME'23)*. Mexico, Mexico. <https://hal.science/hal-04108997>
- [14] Hugo Silva, Michael Zbyszynski, Ataru Tanaka, and Federico Visi. 2020. Interactive machine learning: Strategies for live performance using electromyography. (2020). <https://research.gold.ac.uk/id/eprint/28215/>
- [15] Gabriel Vigiensoni and Rebecca Fiebrink. 2023. Steering latent audio models through interactive machine learning. In *In Proceedings of the 14th International Conference on Computational Creativity*. Ontario, Canada. [https://computationalcreativity.net/iccc23/papers/ICCC-2023\\_paper\\_100.pdf](https://computationalcreativity.net/iccc23/papers/ICCC-2023_paper_100.pdf)
- [16] Shuoyang Zheng, Bleiz M. Del Sette, Charalampos Saitis, Anna Xambó, and Nick Bryan-Kinns. Accepted. Building sketch-to-sound mapping with unsupervised feature extraction and interactive machine learning. *Proceedings of the International Conference on New Interfaces for Musical Expression (Accepted)*, Utrecht, Netherlands.