

SKETCHING MAGNETIC INTERACTIONS FOR NEURAL SYNTHESIS

Nicola Privato*

Intelligent Instruments Lab
Reykjavík, Iceland
nicola@lhi.is

Thor Magnusson

Intelligent Instruments Lab
Reykjavík, Iceland
thor.magnusson@lhi.is

Giacomo Lepri*

Intelligent Instruments Lab
Reykjavík, Iceland
giacomo@lhi.is

Einar Torfi Einarsson

Iceland University of the Arts
Reykjavík, Iceland
et@lhi.is

ABSTRACT

The application of neural synthesis for sound generation has grown significantly in recent years. Models such as RAVE offer real-time control by mapping sounds to and from numerical vectors in an abstract latent space whose features of entanglement, arbitrariness, and continuity pose novel challenges to musicians and composers in encoding and interpreting the inscription. In this paper, we introduce Stacco, a magnetic score system that addresses these challenges by functioning as a performative and compositional support for neural synthesis models through embodied sketching. We describe the system and present insights from a workshop aimed at exploring the compositional potential of this platform. We conclude by reflecting on how, during the workshop, Stacco’s playfulness and magnetic materiality translated into the participants’ scores and on the broader implications of embodied sketching in notating music for neural synthesis.

1. INTRODUCTION

Over the past few years there has been notable advancement in neural synthesis techniques for music generation, with models such as RAVE [1] offering real-time control and the direct navigation of their latent representations.

As neural audio synthesis for real-time applications is a very recent technology, little work has been done in developing specific performative and compositional strategies for such tools, whose inherent features of parameter entanglement, arbitrariness, and continuity pose new challenges for composers and performers working with digital audio.

Within this domain, and more generally in sonic interaction design, one of the risks is to focus on the abstract layers characterizing software models, where the development of the interface often prioritizes technical and sonic

*Nicola Privato and Giacomo Lepri contributed equally to the development of Stacco. Nicola Privato led the activities, collected and elaborated the data used in this paper.

factors, with little attention to the bodily experience of the performer [2].

Based on these premises, the contribution of this paper is twofold. On the one hand, we introduce Stacco, an *instrument-score* based on magnetic sensing that functions as an ideal performative and notational mediator for neural synthesis models, thus enabling us to reflect upon the unique compositional challenges that this novel synthesis technology brings forth. On the other hand, we address these challenges by reporting on the outcomes of a workshop with Stacco, where participants creatively notated musical scores on the body of the instrument as they experimented with a series of neural synthesis models.

In the next section, we discuss relevant literature and existing practices to situate our work, with a focus on the overlapping domains of notation and instrument design.

2. BACKGROUND

In the past few decades the field of *Digital Musical Instruments* (DMI) has witnessed a flourishing of innovative interdisciplinary methodologies, as musicians, designers and technologists have integrated diverse fields into their research practice. This concoction of viewpoints, ranging from media studies [3], social sciences [4] and political studies [5] represents a shift, and at times a dissolution, of the boundaries that traditionally separate musical categories and roles.

Gurevich describes a practice-based attitude to instrument design and performance which promotes an “ecological view of music-making” challenging the traditional model of music as communication, the rigid distinction between composer-performer-listener as well as the locus of agency across human and non-human factors [6]. Media theorist and sociologist Jonathan Sterne argues that “[i]n music technologies, media collapse into instruments – or, rather, the line between instruments and media grows fuzzy” [7]. Similarly, Battier and Schnell introduce the term *composed instrument* to underlay the fact that “computer systems used in musical performance carry as much the notion of an instrument as that of a score” [8].

Our work resonates with such attitudes and perspectives, and, more specifically, it builds on the notion of *instrument-score*, bringing into the debate new method-




Figure 1. Nicola Privato's Thales.

ological and conceptual tools. Of particular relevance is here the work of Tomás and Kaltenbrunner who, drawing on Alvin Lucier's interviews [9], introduce the concept of *inherent scores* to describe the progressive embedding of inscriptions within the instrument [10] itself. Building on this, Tomás develops his *tangible scores*: a particular kind of inherent score allowing sound production through the tactile interaction with drawings engraved on a surface [11].

2.1 Magnetic Scores

Building upon Tomás' contribution, we describe Stacco, the musical interface introduced in this paper, as a *magnetic score* [12]: a particular type of inherent score featuring embedded magnets and a visual interpretation of their magnetic fields. To perform with a magnetic score means to interact with its magnetic fields, usually by means of other magnetic elements, as the variation in the magnetic field is detected and converted into sound.

One notable example of magnetic scores is Thales [13] (Fig. 1), which combines an engraved board embedding a series of magnets with two controllers allowing the tangible navigation of the magnetic fields. As the musician moves the controllers on top of a score with embedded magnets, the magnets in the controllers and those in the score repel or attract each other, suggesting various performative gestures.

Another relevant artwork is the Chowndolo [14] (Fig. 2), a magnetic pendulum whose trajectories are altered by magnets placed underneath the device. The chaotic oscillations of the pendulum are then translated into sound. The magnets below the pendulum can be arranged to compose new shapes: different configurations will produce unstable sonic variations, articulating music that evolves based on the pendulum's dance.

The instrument-scores mentioned above have a distinct embodied character, as they encode the inscription within the materiality of the interface, with the performer interpreting it through touch and gestures. This materiality is apparent in the score's resistance and action in response to the performer's gestures, empowering creativity and boosting the communication between the performer and the audience [15]. Our research aims to leverage such material



Figure 2. Giacomo Lepri's Chowndolo.

features, exploring embodied ways to embed articulated compositional ideas, and how this might help composing and performing with neural synthesis.

2.2 Embodied Sonic Sketching

The investigation of embodied musical interactions encompasses different fields, including digital instrument design, sound design, dance, and theatre performance. Drawing on *soma design* [16], Avila et al. describe a set of workshops for the design of digital musical interfaces emerging from bodily explorations of guitar performance [2]. Delle Monache and Rocchesso thoroughly investigate embodied approaches to sound generation and manipulation and propose a set of interactive tools to sketch sounds using the voice [17, 18]. Based on long-lasting collaborations with dancers and musicians, Camurri et al. develop computational tools and frameworks to extract and interpret meaningful gestural information to be exploited while designing interactive experiences [19, 20].

Particularly relevant for our methodology is Kristina Andersen's *Magic Machines* workshop [21]. Drawing on diverse performance and theatre practices [22], Andersen proposes a playful design fiction approach to technology ideation, which involves the creation of mock-up objects that work *as if by magic*. Andersen exploits the notion of the "magical unknown" to free a participant's imagination and generate manifestations of original technologies and interactions. Crucial to the Magic Machine workshop is the notion of "thinking with the hands" [23], where the physical making and manipulation of artefacts allows participants to sketch ideas through the spontaneous and intuitive assemblage of materials. Lepri and McPherson have also exploited Andersen's Magic Machines workshop to facilitate the emergence of subjective aesthetic values and priorities of artists engaged in the design of novel musical interactions [4].

The activity presented in this paper may be viewed as a variation of such work, as participants were invited to augment Stacco with mundane objects, by sketching physical scores to be placed on the instrument itself and guiding the performance with an interface controlling a neural synthesis engine.

2.3 Neural Synthesis

The recent introduction of deep learning methods has brought forth exciting technologies for the generation of raw audio [24, 25]. Whereas the slow responsiveness of neural audio synthesis models has initially limited their use in performative scenarios, the introduction of RAVE [1], a Real-time Audio Variational Autoencoder (VAE) that performs fast and high-quality audio synthesis, has drastically facilitated the use of neural synthesis in interactive contexts [13].

RAVE is a fine-grained latent variable model. It encodes a stream of raw audio to a stream of latent vectors, before decoding it back to audio. It combines a variational autoencoder with an adversarial reconstruction term, where realistic details are sampled by the decoder and the coarse qualities of the sound are represented in the *latent space*: a multi-dimensional, compressed representation, in between the model’s encoding and decoding functions.

RAVE may be used for style transfer, by forwarding raw audio to its encoding function or explored by navigating its latent representation through control signals. It is also possible to combine these two approaches, for instance by forwarding audio through the encoding phase to one or two latent dimensions whilst driving the others with dedicated signals, or by mixing audio and control signals on individual latent dimensions. In the study described in this paper, we apply the former method, mapping each of Stacco’s data points to a single latent dimension, and bypassing the encoding function.

3. STACCO

In this section, we provide a high-level analysis of neural audio synthesis’ affordances from the perspective of the composer, motivating the reasons that led us to develop our framework for embodied sketching. We continue by describing Stacco’s hardware and software, and the neural synthesis models we used in our study.

3.1 Composing for Neural Synthesis

Neural Synthesis is a new, exciting synthesis method whose distinctive features open new compositional approaches and expressive possibilities. In composing with RAVE, the artist curates datasets of raw sounds. During the training phase, the model learns to reconstruct the datasets and distributes its most meaningful sound features in the latent space.

In our artistic explorations of different models’ latent space, we encountered three distinctive qualities that influence the interaction with the model and that, ultimately, redefine the compositional strategies:

- **Entanglement:** Latent dimensions are deeply intertwined; they appear as a complex and interdependent structure of relations, where any change in the state of one latent variable affects the behaviour of all the others. Consequentially, it is not possible to manipulate one single parameter, such as amplitude or frequency, without affecting the others.

- **Continuity:** Latent spaces represent continuous dimensions that remain active at all times and cannot be selectively deactivated. Consequently, systems employing RAVE are typically designed to offer continuous values rather than booleans.
- **Arbitrariness:** The distribution of the sound features within the latent spaces is autonomously performed by the model during the training phase, it differs with every dataset and initialisation seed, and is empirically explored *a posteriori* by interacting with the trained model. This implies that it is not possible to abstract gestural constants affecting the sound in similar ways with two different models.

Because of the aforementioned arbitrariness, composers need to spend time exploring every new model, looking for patterns, configurations, and gestures that lead to satisfactory acoustic results. This process is necessarily mediated by a physical interface, that influences the control of the system as well as the understanding of the algorithm [26].

Once the model has been thoroughly explored, and meaningful gestures and interesting areas within the latent space have been detected, the composer faces the challenge of defining an appropriate notational strategy. This is of course dependent on the overall aesthetic aim of the composer. For the case described in this paper, i.e. the navigation of RAVE through control signals, it is nevertheless possible to outline two general considerations.

First, because the latent distribution of a model is learned through a gestural exploration on a given interface, an immediate and intuitive approach to notation is the one that allows the reconstruction of those gestures on that interface; second, because sound features such as frequency or amplitude cannot be disentangled from the latent representation, it is not possible to notate them individually. This further binds the representation to the synthetic potential of the gesture.

As a consequence of the notation’s gestural reliance, the act of composing becomes intrinsically intertwined with the interface. By building upon the capability of magnetic scores to incorporate the notation into the instrument, our instrument-score allows artists to compose by sketching gestures on interchangeable musical scores placed on the interface itself. In addition, Stacco features a one-to-one, unlabeled and continuous mapping to latent dimensions, and the interaction of the magnets, their agency and resistance makes the playing engaging and enjoyable.

3.2 Hardware

Stacco (Fig. 3) is a novel type of magnetic score embedding magnetic attractors and sensors underneath an engraved surface, and whose design features are aimed at facilitating the interaction and composition with RAVE whilst providing a rich and playful musical experience.

Stacco consists of four miniaturised magnetic discs [12] combining a magnetometer with a permanent magnet each. Each sensor performs two-dimensional readings of nearby magnetic fields whilst actively attracting and repelling nearby magnets and ferromagnetic objects. The sensors



Figure 3. Stacco.



Figure 4. Workshop Material.

are displaced in four symmetrical points underneath an oval wooden board and connected to a Bela [27] for embedded synthesis. The 32 x 217 cm engraved board features a raised edge and is enclosed via a living hinge structure.

The performer interacts with Stacco by throwing and displacing on the board a series of magnetic spheres of variable dimensions. The four magnets under the board, each coupled to a magnetometer, actively engage the performer in a playful dance of agencies, enacting the inscription through the interactions of their magnetic fields with the spheres controlled by the performer.

The engraved board reveals the position of the four sensors and the magnetic field through a series of circular patterns. An additional smaller circle is engraved in the centre and may host magnets with opposed polarity for more complex interactions. A larger circle encloses all four attractors within the larger oval.

Stacco's design allows embedding sketches as tailored oval sheets with two or three-dimensional inscriptions. Since the spheres interact with the instrument through their magnetic field and do not need to touch the board, the presence of the oval sheets does not affect the instrument's playability.

3.3 Software

Stacco uses Bela [27] to forward the sensors' data to the laptop via OSC, for a total of eight data points, two axes per sensor.

We trained three RAVE models based on the following datasets, mapping each latent dimension to one of the eight data points forwarded by Stacco's magnetometers.

- Choir: model trained by the artist Jonathan Reus. The model is RAVE v1, with 16 latent dimensions mapped two-to-one with each of Stacco's sensor readings.
- Organ: recordings of open-source organ music. Small amounts of voice and other instruments are included and vinyl record noises are prominent. The model is RAVE v1 modified by Victor Shepardson, with 16 latent dimensions mapped two-to-one with each of Stacco's sensor readings.

- Magnets: one-hour recording of magnets of different dimensions interacting with each other or scratching wooden and metallic surfaces. The model is RAVE v1, 48Khz, 8 latent dimensions linearly mapped one-to-one with each of Stacco's sensor readings.

In addition to this, to compare neural synthesis with traditional synthesis methods we developed a stereo FM synth with eight control parameters, individually mapped to each of Stacco's sensor readings.

4. WORKSHOP

We explored the composed nature of Stacco in a workshop in Łódź (Poland) at Act In Out, a joint initiative by the Art Factory in Łódź, Slaturhusid Art Center (Iceland), Carte Blanche Dance Theatre and Visjoner Theatre (Norway). In this workshop, titled *Composing Magnetic Interactions*, we used Stacco as a support to the design of a series of musical scores, aiming to understand the compositional practices that may develop around neural synthesis and whether the concept of instrument-score might favour an embodied understanding of the model.

4.1 Description

The workshop involved five participants: three trained musicians (P1 to P3), one graphic designer (P4), and one visual artist and art curator (P5) (Fig. 4).

Before the event, we cut a series of cardboard ovals with a stylized version of Stacco's engraving printed on one of the sides. Thanks to the instrument's raised edge, the ovals fit the top of the board working as removable blank canvases for the participants' compositional ideas. We provided each of the participants with one oval, and the whole group with the following materials: paper straws and cups, magnets of different sizes and colours, pencils, pens, markers, clay, rubber bands, rope, scissors, chalks, wooden sticks, tape, wall gum.

In choosing the materials we avoided tool-kits, electrical components, sensors and software units as well as materials with distinct acoustic properties (e.g. boxes and wood), in order not to limit people's imagination and inventiveness. Following Andersen's approach, we privi-




Figure 5. P1's Score.

leged mundane materials and everyday objects to free participants and facilitate the emergence of subjective aesthetic views. The exploration of these relatable and familiar objects in novel creative technology-based projects empowers and asserts confidence in new and non-specialist audiences [28].

The workshop lasted three hours. During the first hour, we briefly demoed Stacco and invited each participant to try it out using the available set of magnetic spheres. The participants were free to play with and mix all the models plus the FM synthesizer. Each of them experimented with the instrument for about ten minutes, with P2 and P4 playing more than one session.

After the break, we invited the participants to “build” their scores. We avoided using the term “compose” for its strong semantic connotation to established traditions. This choice is in line with Andersen’s methodology, which avoids overly loaded terminologies to prevent participants from limiting themselves to pre-existing technological – in our case techno-musical – assumptions.

The time for the task was 50 minutes. This constraint is based on the guidance provided by Lepri and McPherson which found that a relatively fast-paced exercise helps participants sidestepping insecurity and avoid overthinking. During the remaining hour, the participants presented their works and performed by placing the ovals on top of Stacco. After the workshop, we submitted to the participants an optional online survey. The survey was completed by P1, P2 and P4.

The next paragraphs describe the strategies devised by the participants for notating their ideas and elucidate the rationales they put forward in support of their approaches.

4.2 Scores

P1 designed a one-shot composition by distributing small coloured magnets on the cardboard (Fig. 5).

The magnets were carefully distanced to not collapse into each other through their reciprocal attraction. Some of the magnets were placed on top of wooden sticks, others on top of magnetic stripes. As soon as the oval was placed on top of Stacco, its four attractors broke the balance and caused the system to collapse.



Figure 6. P2's Score.

The performance consisted of this initial chaotic event and an improvisation starting from the newly found balance, aiming to elicit new unpredictable interactions.

P1 chose not to title the score. The choice of sounds, FM synthesis and magnets model, was very consequential as the composition combined the materiality of magnets with the non-linear behaviour that is typical of FM.

P2 used three-dimensional objects such as rope and straws to define separate areas of the oval through meticulous chromatic patterning (Fig. 6). The score was titled *Techno and Textiles*. The rope and the straws were chosen because of their alternating coloured and white diagonal lines. To reinforce the chromatic effect, P2 cut the straws into two-dimensional rectangles, obtaining an orange-and-white pattern, and juxtaposed multiple rope windings to create a chromatic continuity between the white and blue diagonals. In line with P1, P2 used FM and magnets’ sound.

As P2 explained, the title of the score is a reflection on the relationship between techno music and textiles, two themes that are rooted and in constant dialogue in Łódź’s culture, as the venues hosting techno events are for the most part repurposed textile industries.

P3 designed instead a two-dimensional graphic score using coloured markers (Fig. 7), a map sketching trajectories for four spherical magnets around Stacco’s four attractors. The patterns were very diverse and altogether gave the impression of electronic schematics. The choice of the sounds fell on the choir model. As opposed to the previous participants, P3 described the FM synth sounds as “too aggressive.”

P4 and P5, the ones in the group with no prior musical training, designed two three-dimensional scores (Fig. 8, 9). As P4 disclosed in the following survey, Stacco reminded him of Italian architecture, of Carlo Scarpa in particular and Nono’s work “A Carlo Scarpa, architetto, ai suoi infiniti possibili” (1984). Indeed, P4’s score was similar to the architectural model of a city’s square, with a main structural unit represented by a cup with a stick and a rubber band holding a magnetic pendulum. On the sides, wooden sticks functioned as rails on top of the attractors for two spherical magnets. The selected sounds were choir and FM, and the performance consisted of suc-



Figure 7. P3's Score.



Figure 8. P4's Score.

cessive reconfigurations of the architecture, moving the spheres along the paths of the wooden rails placed on top of the attractors. Finally, P5 developed a narrative around a magnetic persona, with a larger magnetic sphere as the body and head, and arms and legs made of smaller magnets. P5 built a series of structures with straws, clay and rubber bands around three of the attractors, and titled the score *A Day in the Park*. By wandering through the score, climbing and sliding on top of straws and sticks of different heights, the character modulated Stacco's magnetic fields. As P5 noted, the work was "more as a sculpture than a practical thing." In P5's open-world score, all interactions, and in particular those elicited by gravity, are redefined as magnetism. Coherently, the sole model chosen for the performance was the magnets one.

4.3 Performances

Once the scores were completed, the participants presented them and played for the group. In this section, we will observe how the scores redefined the gestural and sonic interactions with Stacco by comparing the initial open explorations with the final performance, and including insights from the conversation that followed. This section includes observations from P1 to P4, P5 chose to explore the score alone and during the group's break. For this reason, we chose to not record this performance.



Figure 9. P5's Score.

In his exploration of Stacco, P1 started with groups of two to three spheres on each attractor, using the choir model and focusing primarily on the interactions between the spheres rather than on their correlation with the sound. After a few minutes, P1 reduced the number of spheres to two, placed each one on a different attractor and began to systematically explore the sound by lingering on the gestures that had a bigger effect. Towards the end, P1 was holding the larger spheres a few centimetres above the board and rapidly switching from one attractor to the other. As P1 noted in the survey, performing was "very interesting but sometimes random." At the same time, the interface was "gesturally sensitive." During the final discussion, P1 also noted how playing Stacco is "like fighting with the instruments." This agency, he continued, is something that every instrument has, but with Stacco "you feel the physicality of the instrument, which in electronic music you often lose." Coherently with these insights, P1 used randomness as the defining feature of the piece, letting Stacco's magnetic fields autonomously reconfigure the score.

P1's performance with Stacco was a purposeful exploration of the instrument's agency. The entanglement of the latent dimensions in RAVE was mirrored through the unpredictable entanglements of the magnets, with the score conceptually and physically mediating the dialogue between the performer and the interface. The movements during the performance were not aimed at controlling specific sound features, since, as noted during the final discussion, P1 intuitively realised their entanglement. The performative gestures were aimed instead at breaking the system's homeostasis and letting the magnetic interactions affect the sounds.

P2's test with Stacco began where P1's exploration had finished: by moving magnets from one attractor to the other and looking for emerging patterns. Once P2 found a drastic change in amplitude, he would stop and start moving one magnet around the board, looking for a gesture that could control the overall amplitude. He found a few, and one in particular producing silence, thus realising that at least one of the data points was affecting, among other parameters, the overall amplitude. P2 tried out Stacco a second time, experimenting with the cardboard before composing his piece. Through this intermediate phase he discovered a new gesture: by turning as a knob a spherical

magnet placed at the centre of one attractor and pushing it firmly on the cardboard, he could fine-control the sounds better than by moving the magnets on the board. He incorporated this approach into the following performance, by pushing spheres from the centre of one attractor and then rapidly moving them to another one.

The chromatic patterning in P2's score and the title of the performance found a correspondence in the rhythmic approach to the navigation of the latent space, with the magnets model combined with the FM synthesizer to generate noisy textures.

P2 purposefully built the score around the magnetic attractors, emphasizing through colours and textures the areas of finer control. As opposed to P1, P2 used Stacco's magnetic fields as a guide in exploring the instrument-score, and represented the close relationship of textiles and techno music in the area by rhythmically displacing a single sphere from orange to blue areas and vice-versa.

P3 used instead the organ model for his performance with Stacco. He placed one sphere on each attractor and moved them around with both hands, at first slightly, then more decisively, in search of areas on the board where the magnets' movement had the most effect on the sound. The score was then sketched around these areas, with the spheres rolling back and forth from the centre of the attractors, first one by one, then together. The closure of the performance was offered by Stacco itself: P3 was aiming to place the fourth sphere on the last free circle, but the magnet got attracted by a sphere in the nearby one and slipped from P3's hand. That particular configuration in the latent space, resulting from a sudden interaction with the magnetic inscriptions, produced complete silence and an unexpected, enjoyable finale to which the whole group reacted with a laugh.

If P3's approach to the writing was rigorous and methodical, and certainly more akin than the others to the traditional idea of a graphic score than in all other scores in the session, the performance was very playful and engaging: as P3 noted in the survey, playing with Stacco related to an "attitude to experiments" that "goes back to childhood."

P4 chose to experiment with the FM synthesizer and with the choir model. While playing, P4 compared the experience of playing with the FM patch to that of a no-input mixer, emphasizing the high responsiveness to small gestures. This is a distinctive feature of FM synthesis compared to RAVE: if both algorithms are similar in the high degree of entanglement between the parameters, RAVE seems to distribute them evenly within the latent representation. The interactions between the parameters of the FM synthesizer are instead less linear, with small transitions causing marked changes. In addition to this, the higher latency of RAVE contributes to reducing the overall responsiveness to subtle and rapid gestures.

P4 explored Stacco with two magnets, moving them rapidly above and between the attractors. The gestures afforded by the score were instead very different, with spheres moving on predetermined paths. The playful character of this score emerged clearly in the way P4 manipulated the main structure and the embedded pendulum to

interact with the spheres, much like a kid having fun with construction toys.

5. DISCUSSION

In this section, we describe and reflect upon themes that emerged from the workshop and the following survey, focusing in particular on how the notion of the instrument-score intersects to neural audio synthesis through embodied sketching.

5.1 Playfulness and Agency

During the demo, the participants were evidently surprised by the magnets spinning around the attractors, bumping into each other and combining into complex structures. After this initial surprise, during the practice sessions and the discussion the participants emphasised Stacco's agency as its most distinctive feature, with P1 even noticing that it felt like "fighting with the instrument," and the group laughing in witnessing Stacco's magnets calling P3's performance to an unexpected finale.

Notably, as the participants practised and acquired confidence with the system, the agency and unpredictability of the magnetic interactions, from a factor of surprise, became a compositional element, with the participants' scores leveraging it as a generator of randomness (P1), as an architectural component (P4), or as the main thread in constituting a narrative (P5).

The magnetic agency of the instrument-score, that P1's metaphor so effectively depicts, pairs with that of the algorithm in autonomously and independently re-modelling the dataset. To provide an effective way of accounting for, negotiating with, and incorporating these two features into the notation, we relied on embodied sketching.

5.2 Embodied Sketching

The approach of *Embodied Sketching*, that is, the practice of embedding the score onto the instrument itself has proven effective for different reasons.

- It favoured the transposition of the instrument's peculiarities into the notation. As participants practised, gaining a better understanding of the system, Stacco's features translated into the scores' semantics and subsequently, in P1, P2 and P5 in particular, into the choice of the sounds.
- It offered to the participants an open creative environment, in which they felt free to compose two-dimensional (P3) or three-dimensional scores (P4, P5), to adhere to traditional (P3) or novel compositional methodologies, often strongly related to their personal histories and backgrounds (P2, P4, P5) or with their personal views on composition (P1, P4).
- It allowed them to overlap latent exploration and gestural notation. This is of particular value in composing with neural synthesis, since, as we discussed in 3.1, this novel technology involves an embodied process of a posteriori exploration and a strong connection between gesture and notation.

Finally, during the discussion another theme has emerged: how different is composing and performing with neural synthesis than with a chaotic synthesizer based on a traditional synthesis technique? Even though this deserves further investigation, the participants' feedback and their consistent use of the FM patch provide some initial indications. Indeed, most of the scores combined the two methods, taking advantage of FM's higher sensitivity to small and rapid changes as well as of neural synthesis' potential for a smoother, organic sonic exploration. In future research, we aim to explore whether these two synthesis methods may suggest different notational approaches, and in what ways these might diverge.

6. CONCLUSIONS

Stacco is a new interface for a new synthesis method, a novel instrument-score that functions as a platform for composing with neural synthesis through embodied sketching, which we defined as the practice of embedding a notational layer on the instrument's body – i.e. removable oval cardboard sheets placed on top of Stacco's magnetic board.

As our workshop demonstrates, embodied sketching addresses some of the compositional and performative challenges posed by neural synthesis. Indeed, the tangible and playful interactions that Stacco entails allow to bridge gesture and notation and prioritise a holistic and embodied exploration of the model.

By "thinking with their hands" [23], rather than focusing on rationalising the model's intricate functioning, our participants intuitively developed original and highly subjective musical sketches. This was particularly clear, for instance, in P1's approach to the composition, leveraging Stacco's properties to elicit chaotic magnetic and sonic interactions, but also in P3's more traditional approach, searching for and drawing sonically satisfying trajectories around Stacco's attractors.

Despite the short duration of our workshop, the engagement with the instrument led to the creation of a variety of compelling sketches. In future workshops, we intend to investigate whether a more consistent and prolonged practice with our system might lead to a deeper embodied understanding of the model and to the creation of compelling musical works. To this end, we plan to expand Stacco's embodied musical sketches using transparent materials instead of cardboard sheets, so to stack multiple notational layers and create interchangeable scenes and modular compositions. Through this approach, and by experimenting with other inherent scores, we hope to contribute to the understanding of the creative practices and the novel semantics that the introduction of deep learning techniques into the musical and creative domains is bringing forth.

Acknowledgments

This research 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).

7. REFERENCES

- [1] A. Caillon and P. Esling, "Rave: A variational autoencoder for fast and high-quality neural audio synthesis," *arXiv preprint arXiv:2111.05011*, 2021.
- [2] J. M. Avila, V. Tsaknaki, P. Karpashevich, C. Windlin, N. Valenti, K. Höök, A. McPherson, and S. Benford, "Soma design for nime," in *Proceedings of the 2020 International Conference on New Interfaces for Musical Expression (NIME'20)*, 2020.
- [3] D. Holzer, H. Frisk, and A. Holzapfel, "Sounds of futures passed: Media archaeology and design fiction as a NIME methodologies," in *Proceedings of the Conference on New Interfaces for Musical Expression*, 2021, forthcoming.
- [4] G. Lepri and A. McPherson, "Making up instruments: Design fiction for value discovery in communities of musical practice," in *Proceedings of the Designing Interactive System Conference*, 2019, pp. 113–126.
- [5] F. Morreale, S. M. A. Bin, A. McPherson, P. Stapleton, and M. Wanderley, "A NIME of the times: Developing an outward-looking political agenda for this community," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, 2020, pp. 160–165. [Online]. Available: https://www.nime.org/proceedings/2020/nime2020_paper31.pdf
- [6] M. Gurevich, "Discovering instruments in scores: A repertoire-driven approach to designing new interfaces for musical expression," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, 2017, pp. 163–168. [Online]. Available: http://www.nime.org/proceedings/2017/nime2017_paper0031.pdf
- [7] J. Sterne, "Spectral objects: On the fetish character of music technologies," *Steintrager and Chow*, vol. 2019, pp. 94–109, 2019.
- [8] M. Battier and N. Schnell, "Introducing composed instruments, technical and musicological implications," 2002.
- [9] A. Lucier, "Origins of a form: Acoustical exploration, science and incessancy," *Leonardo Music Journal*, vol. 8, pp. 5–11, 1998. [Online]. Available: <http://www.jstor.org/stable/1513391>
- [10] E. Tomás and M. Kaltenbrunner, "Tangible scores: Shaping the inherent instrument score," in *Proceedings of the International Conference on New Interfaces for Musical Expression*. London, United Kingdom: Goldsmiths, University of London, Jun. 2014, pp. 609–614. [Online]. Available: http://www.nime.org/proceedings/2014/nime2014_352.pdf
- [11] E. Tomás, *The interface-score: Electronic Musical Interface Design as Embodiment of Performance and Composition*. Linz, Kunstuniversität, 2018.

- [12] N. Privato, T. Magnusson, and E. T. Einarsson, “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*, A. P. D. Ritis, V. Zappi, J. V. Buskirk, and J. Mallia, Eds. Boston, Massachusetts, USA: Northeastern University, 2023, pp. 36–44.
- [13] —, “Magnetic interactions as a somatosensory interface,” pp. 387–393, May 2023. [Online]. Available: http://nime.org/proceedings/2023/nime2023_54.pdf
- [14] N. Merendino, L. G. Lepri, A. Roda, and R. Masu, “Redesigning the chowndolo: a reflection-on-action analysis to identify sustainable strategies for nimes design.” 2023.
- [15] T. Cochrane, “7475On the resistance of the instrument,” in *The Emotional Power of Music: Multidisciplinary perspectives on musical arousal, expression, and social control*. Oxford University Press, 07 2013. [Online]. Available: <https://doi.org/10.1093/acprof:oso/9780199654888.003.0006>
- [16] K. Hook, *Designing with the body: Somaesthetic interaction design*. MIT Press, 2018.
- [17] D. Rocchesso, D. A. Mauro, and S. D. Monache, “Mimic: The microphone as a pencil,” ser. TEI ’16. Association for Computing Machinery, 2016, pp. 357–364. [Online]. Available: <https://doi.org/10.1145/2839462.2839467>
- [18] S. Delle Monache, D. Rocchesso, F. Bevilacqua, G. Lemaitre, S. Baldan, and A. Cera, “Embodied sound design,” *International Journal of Human-Computer Studies*, vol. 118, pp. 47–59, 2018.
- [19] A. Camurri, G. Volpe, S. Piana, M. Mancini, R. Niewiadomski, N. Ferrari, and C. Canepa, “The dancer in the eye: towards a multi-layered computational framework of qualities in movement,” in *Proceedings of the 3rd International Symposium on Movement and Computing*, 2016, pp. 1–7.
- [20] A. Camurri, S. Hashimoto, M. Ricchetti, A. Ricci, K. Suzuki, R. Trocca, and G. Volpe, “Eyesweb: Toward gesture and affect recognition in interactive dance and music systems,” *Computer Music Journal*, vol. 24, no. 1, pp. 57–69, 2000.
- [21] K. Andersen and R. Wakkary, “The magic machine workshops: Making personal design knowledge,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–13.
- [22] A. Boal, *Games for Actors and Non Actors*, tr. A. Jackson. London, Routledge, 1992.
- [23] K. Andersen, “Making magic machines,” in *10th European Academy of Design Conference*, 2013.
- [24] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *CoRR*, vol. abs/1609.03499, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03499>
- [25] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. Courville, and Y. Bengio, “SAMPLRN: An unconditional end-to-end neural audio generation model,” 2017.
- [26] C. Dobrian and D. Koppelman, “The ‘e’ in nime: Musical expression with new computer interfaces.” 01 2006, pp. 277–282.
- [27] G. Moro, S. Bin, R. Jack, C. Heinrichs, and A. McPherson, “Making high-performance embedded instruments with bela and pure data,” 06 2016.
- [28] G. Lepri, J. Bowers, S. Topley, P. Stapleton, P. Bennett, K. Andersen, and A. McPherson, “The 10,000 instruments workshop - (im)practical research for critical speculation,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, 2022. [Online]. Available: <https://doi.org/10.21428%2F92fbeb44.9e7c9ba3>