SoundTraveller: Exploring Abstraction and Entanglement in Timbre Creation Interfaces for Synthesizers

Zefan Sramek
zefanS@iis-lab.org
IIS Lab, The University of Tokyo
Tokyo, Japan

Arissa J. Sato
arissa@iis-lab.org
IIS Lab, The University of Tokyo
Tokyo, Japan

Zhongyi Zhou
zhongyi@iis-lab.org
IIS Lab, The University of Tokyo
Tokyo, Japan

Simo Hosio
simo.hosio@oulu.fi
Center for Ubiquitous Computing,
University of Oulu
Oulu, Finland

Koji Yatani
koji@iis-lab.org
IIS Lab, The University of Tokyo
Tokyo, Japan

ABSTRACT

Timbre exploration and creation are key tasks in electronic music composition. Modern synthesizers can produce thousands of unique timbres, but this complexity hinders musicians' ability to explore these timbres effectively. We contribute SoundTraveller, an interactive timbre exploration system aimed at fostering electronic musicians' creative processes. SoundTraveller allows the user to explore the timbral space using two modes: evolutionary and morphing, with which they can generate hundreds of unique timbres without the need to edit individual parameters. Our user study confirmed that SoundTraveller supported participants' exploration, decreased their cognitive load, and increased their perceived creativity. Through analysis of our interview study, we contribute design considerations for timbre exploration systems, and for exploring large aesthetic parameter spaces more generally. Finally, we discuss how systems like SoundTraveller can fit into the existing workflows of electronic music composers, and how shared agency with creative technologies can impact the creative process.

CCS CONCEPTS

• Human-centered computing → Human computer interaction (HCI); **• Applied computing → Sound and music computing**.

KEYWORDS

Creativity support, synthesizers, sound authoring, timbre exploration, agency, entanglement

ACM Reference Format:
Zefan Sramek, Arissa J. Sato, Zhongyi Zhou, Simo Hosio, and Koji Yatani. 2023. SoundTraveller: Exploring Abstraction and Entanglement in Timbre Creation Interfaces for Synthesizers. In *Designing Interactive Systems Conference (DIS '23), July 10–14, 2023, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3563657.3596089

DIS '23, July 10–14, 2023, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-9893-0/23/07…$15.00
https://doi.org/10.1145/3563657.3596089

1 INTRODUCTION

Musicians composing and producing in a wide variety of contemporary genres regularly use synthesizers and electronic instruments in their musical practice. Synthesizers are instruments that produce sound by electronic means and provide users with the ability to control not only the volume and pitch of a note but also its *timbre*. Timbre is the property of a tone that distinguishes it from other tones played at the same volume and frequency, e.g., the difference between the sound of a guitar and a piano playing the same note. Unlike other instruments, synthesizers provide users with many parameters (e.g., oscillator wave shape, filter cutoff frequency) to manipulate their timbre. Consequently, this allows for a musician to control and explore nearly unlimited options when creating music.

The synthesizer's timbre creation interface supports systematic explorations by allowing users to manipulate individual *synthesis parameters*. Modern synthesizers typically combine physical knobs, digital menus, sliders, and many other types of input means for such explorations, offering great flexibility in timbre creation. However, timbre creation is a fundamentally technical task within the larger creative process, and musicians' ability to explore timbres effectively is hindered by both the extremely large size of the timbral space, and the complexity of understanding how synthesis parameters contribute to the production of a given timbre. This technical burden may hinder the composition process, and results in many musicians relying on preset timbres, which constrains their creativity.

Music creation and performance is a complex and abstract activity involving the interaction of human performers and music technology, making it a fruitful area of study for HCI [33]. Synthesizers, in particular, are of great interest because they differ from acoustic instruments in that their timbre control interface and sound generation methods are decoupled, and free from the physical constraints imposed by typical acoustic sound generation methods. As a result, the user's experience of timbre creation is mediated by an interface whose design can be independently explored. Furthermore, the operation of such timbre interfaces requires domain-specific technical knowledge, on top of creative or aesthetic judgements. This makes synthesizers and other forms of electronic music creation an intriguing research topic in HCI because they represent a space where interface designs can be freely explored and their impact on the creative process can be studied. Prior work has examined interface designs for electronic music synthesis or performance, but has not investigated how supporting timbre exploration can contribute to musicians' composition process, nor how such a system can integrate with musicians' current practices.

In this research, we set out to: 1) understand how electronic musicians use their current tools for timbre exploration and in the broader context of their music-making; 2) design an interactive timbre creation system that can be integrated into musicians' existing workflows and assist timbre exploration; 3) evaluate how the system supports timbre exploration and augments their already existing workflows. We first conducted a formative study with electronic musicians to understand their creative processes. Following the design considerations derived from this study, we conducted an iterative design process, which resulted in *SoundTraveller*. SoundTraveller allows the user to explore the timbral space using two modes: evolutionary and morphing, with which they can generate hundreds of unique timbres without the need to edit individual parameters. Finally, we conducted a user study where participants used SoundTraveller for timbre exploration during music composition, and we evaluated it from a variety of angles, including timbre creation support, cognitive load, and creativity support.

The results of our studies indicate that SoundTraveller significantly increased participants' perceived creativity, and significantly decreased their cognitive load. While for some participants, SoundTraveller was a tool to acquire new timbres, for others it acted as a source of inspiration by allowing them to experiment with timbres outside their habitual sound palettes. Furthermore, shared agency with SoundTraveller during the timbre creation process had unique impacts on our participants experiences, which can inform future work with creative technology systems. Combined, these results contribute to the literature on exploration support for large aesthetic parameter spaces and creativity in HCI.

The contributions of this paper are as follows:

- A set of design considerations for timbre exploration support systems, based on our formative study
- The design and implementation of a novel support system for timbre exploration: SoundTraveller
- An investigation of how the system supports musicians' timbre exploration and integrates into their existing workflows

2 RELATED WORK

2.1 Entanglement & Shared Agency in Creative Systems

All musical instruments can be thought of as technologies for augmenting the human body's ability to produce musical sound. As such, any type of musical instrument imposes certain requirements and limitations on users, which, according to theories such as Actor Network Theory, imbues them with agentic characteristics, despite being non-living objects [21, 45]. That being said, the type of agency we might ascribe to an acoustic musical instrument is quite different from that envisioned in human-computer interaction literature, where typically non-human agency would only be granted to technological systems, such as *"computers, robots, and virtual agents"* [49]. But regardless of how we view the agency of a musical instrument, we can understand it as mediating the relationship between the musician and the music or performance, or as Verbeek puts it, creating *"specific relations between its users and their world, resulting in specific experiences and practices"* [78].

Frauenberger's concept of *entanglement* [21] can be useful for considering the agential qualities of digital systems designed to aid in technical creative tasks. Using this concept, the creative interface, rather than being an object with which the user simply interacts, can be considered a non-human actor with which the user works to *"enact certain phenomena"* [21]. Whether or not they adopt this theoretical framework, many recent HCI papers on creativity support systems have directly or indirectly taken this approach to supporting creative work, for instance by using AI to generate new design examples [51], developing agentic systems that work on creative tasks collaboratively with users [46, 88], or providing tools that share some of the burden of design space exploration [39, 42, 81]. While some of these systems explicitly make use of virtual agents that collaborate with users, all shift the balance of agency between the user and the system. This has substantial implications for thinking more generally about creativity support systems in HCI. Resnick et al.'s 2005 paper [58] outlines twelve design principles for creativity support, which appear in previous work exploring creativity support tools in a broad range of fields, such as creative writing [13], architectural design [4], drawing [56, 70, 80], and computer graphics [30, 32, 67]. However, Resnick et al. do not address the possibility of shared agency, instead only discussing how the choice of black boxes represents a trade-off between ease of application and flexibility.

Although AI agents are one obvious direction for exploring shared agency, many technologies can be considered agentic in Frauenberger's sense. For future work on digital creativity systems, understanding not only how agentic systems can support creativity, but also how shared agency impacts the creative process and human actors' experiences will be essential, particularly given that such systems can impact users' thinking styles [22] and sense of artistic identity [72]. Indeed, Rodger et al. [59]'s re-evaluation of musical instruments in the HCI context asserts that, *"musicians are not users, but rather agents in musical ecologies."* As such, we must explore how musicians experience their current tools, and how new systems impact their workflows and creative processes, not only in terms of creativity, but also in terms of agency.

2.2 Parameter Space Exploration in the Creative Process

Creating and selecting timbres is a key aspect of many electronic musicians' creative processes. With traditional synthesizer interfaces, this is achieved either by manipulating the instrument's synthesis parameters, or by selecting *preset* timbres provided by the manufacturer. In this sense, timbre exploration can be regarded as a part of a class of creative problems that involve the exploration of aesthetic parameter spaces.

Exploring large aesthetic parameter spaces can be a challenging task in many domains [83]. Researchers in HCI and beyond have sought to harness machine learning and other technologies to support users in these processes. Outside of timbre exploration, researchers have explored support for font selection [11, 53], color palette selection [87], computer graphics [43], fashion [39, 76], photo filters [69, 73], game design [38], music composition [6], and audio effects [15]. These support tools present abstractions of the parameter space, for instance by organizing font libraries based on perceptual characteristics [11, 53] or learning the relationships between fashion descriptions and clothing items [76], or recommending options to users based on their initial inputs, such as photo filters for social media [73] or audio processing effects chains [15]. While such systems can liberate users from being overwhelmed by the entire parameter space, we argue that they do not necessarily support the full extent of users' creative potential because users are constrained by the systems' pre-defined solutions, instead of being encouraged to explore their own choices.

SoundTraveller takes an approach similar to that of Yuan et al.'s *InfoColorizer* [87], which supports the exploration and specification of color palettes for infographics. Rather than simply producing suggestions of entire color palettes, with *InfoColorizer* color palettes can be iteratively refined by adjusting individual 'vague' and 'specific' color preferences, and generating updated suggestions. This type of design supports one of Resnick et al. [58]'s principles: *"Low Threshold, High Ceiling, and Wide Walls"*, which means that the systems present a low barrier to entry for novice users, but allows advanced users to explore in more nuanced ways. We argue that this approach is critical to supporting users' creativity, particularly in tasks like aesthetic parameter space exploration. However, *InfoColorizer* uses machine learning trained on an existing library of infographics, and can therefore make appropriate recommendations, but may constrain users' explorations to regions of 'best practice,' rather than allow for open experimentation. In the domain of music, Scurto et al. [64] instead use collaborative AI to support users' exploration. With SoundTraveller, we demonstrate a different approach, where the user can freely locate themselves in a region of the parameter space using preset timbres, and thus explore according to their own personal aesthetic taste, creative needs, or experimental ideas.

For timbre in particular, perceptual timbre spaces constructed based on listening test results have been used for several synthesis applications [26, 40, 66, 77, 82], but timbral spaces have been designed in a variety different ways, such as parameter relationships [47], and audio latent spaces [19, 20]. The burden presented by large parameter spaces [83] means that unique timbre spaces can be particularly advantageous when they reduce the number of parameters presented to the user [14, 44, 62, 79]. However, many such systems require users to master technically complex interfaces with interaction paradigms that differ substantially from traditional synthesizer interfaces, hindering their practical application. SoundTraveller similarly makes use of a timbral space, but by using presets as a basis for exploration, presents an interface that is more familiar to musicians.

2.3 Sound Synthesis and Timbre Exploration Support

Synthesizers present users with numerous parameters that can be adjusted to vary the instrument's timbre. In a traditional synthesizer architecture, the mapping between the interface controls and the timbre parameters is one-to-one. However, as Hunt and Wanderley [37] note, this is rarely the case for traditional acoustic instruments, which typically have only a small number of inputs that map non-linearly to simultaneously influence the timbre. Furthermore, Seago et al. [65] argue that the technical language used by synthesizer interfaces does not map well onto subjective experiences of timbre, meaning that musicians must learn to create timbres using a non-intuitive process different from how they experience timbre with traditional instruments.

Researchers have explored alternative interfaces for sound synthesis and music composition by drawing connections between sound and other domains, such as physical shape [52, 75], body movements [5, 50, 89], touch gestures [61], and writing, drawing, or sketching [9, 25]. Schatter et al. [60] approach the problem by creating an abstract 3D visualization for timbre manipulation, while Jordà et al. [41] use the *ReacTable* to visualize the synthesizer's internal signals. Approaches like these can aid users re-conceptualize or visualize hidden aspects of the synthesis process.

That being said, as Palani et al. [55] found in their analysis of the adoption of creativity support tools, integration into a user's existing workflow ranks only second to functionality in determining their adoption of new tools. Studies with digital music composition systems, like those of Garcia et al. [23], have also highlighted the necessity of understanding musicians' diverse practices. While much previous work on electronic music aligns with the principles outlined by Resnick et al. [58], our work extends previous literature by situating our design in musicians' existing practices, and by contributing qualitative results that provide guidance for future investigations of timbre exploration and other facets of electronic music creation.

Another typical timbre creation approach is to reproduce the parameter set of an existing timbre automatically. This is challenging, particularly with complex synthesis engines like FM synthesis [12], and has led researchers to explore the use of genetic algorithms [34, 48, 74, 85] and deep neural networks [86]. Systems based on the WaveNet model [54] have also been developed for exploring timbre [27, 36]. Using this model, Google's NSynth [18] can also generate realistic instrument timbres and their novel combinations, but the model is too computationally expensive to run in real time, and the playable instrument is limited to pre-computed models [17], thus limiting its potential for free exploration.

While SoundTraveller also employs evolutionary concepts, it is designed to support musicians' own timbre explorations, instead of automatic timbre reproduction. Our work also contributes an interface design that can be integrated with musicians' existing workflows, which is critical for user adoption. This work thus uniquely expands the landscape of research on interactive timbre exploration and music creation in HCI.

3 FORMATIVE INTERVIEW STUDY

3.1 Procedure and Participants

Understanding the existing practices and workflows of potential users is key to developing effective creative tools. To situate our design process within the practices of electronic musicians, we conducted an interview study to understand musicians' experience using synthesizers during composition and performance, as well as their processes for selecting and creating timbres. The semi-structured interviews were conducted online, and lasted between 45 and 90 minutes. The pre-determined interview questions are listed in Appendix B. The interview protocol was approved by the Institutional Review Board of the authors' university.

Ten participants (P1–10) were recruited for the interview study through the authors' personal networks using a snowballing approach, and were from North America and Asia. All participants aside from P3 are part of a loose global scene of underground ambient and dance music. Participants were recruited on the basis of being active musicians who use synthesizers. Although their level of experience and training varied, all participants were self-trained with regard to electronic music and synthesizer programming. All participants except P3 had performed and released music semi-professionally. Participants were compensated approximately $18 USD. All participants are listed in Table 2 in Appendix A.

Each interview was audio recorded and transcribed. We then conducted a deductive thematic analysis [7]. The interviews were initially coded by the first author. The codes were then validated through discussions and negotiation with another author, informed by the purpose of revealing themes that translate directly into design considerations for timbre exploration. As we were using a deductive approach, only codes agreed upon by both authors to be relevant for establishing design considerations were included for analysis and report in this paper. Finally, two authors grouped the codes into themes, which are presented below.

3.2 Results

3.2.1 More sounds but not more work. It was essential to differentiate between the participants' desire to acquire new sounds for their music and their desire to undertake timbre programming. This was reflected in the responses of five participants (P3, P5, P6, P7, P9). P7 described how her process was not so much about understanding how synthesizers work, as it was about listening to sounds:

> *My knowledge of sound is not that advanced, so when I'm using a software instrument...I just listen to the presets and when I think one is nice, I'll try loading that sound or using that synth...* [P7]

3.2.2 Presets as a starting point for exploration. Another key insight was that, while many of the participants did not do substantial timbre creation from scratch, they did in their own ways undertake timbre exploration. Eight participants (P1, P2, P3, P4, P5, P6, P7, P9) mentioned using presets as a starting point for editing or exploration:

> *I usually start with a preset. So I try to go in and find a preset that I like and then I'll try to modify it in different ways.* [P1]

Five participants (P1, P2, P6, P7, P9) in particular talked about using presets as their main method of sound exploration. However, they mentioned that presets were often inadequate to fulfill their desires for unique and appropriate sounds. P1 discussed the Yamaha SY-35 synthesizer in this context, explaining how the lack of useful presets made him value sounds he managed to program himself:

> *The SY-35 is this weird, esoteric sort of strange thing that's out there that people don't know about. And it's hard to use it. And most of the presets [are] really [bad]; they're terrible presets... So when you can really produce a cool preset on an SY35, it's a really unique thing.* [P1]

3.2.3 External equipment support. While many complex synthesizers provide vast capabilities for sound design, eight participants (P1, P3, P5, P6, P7, P8, P9, P10) mentioned further customizing their sounds using external effects. In fact, for some participants, such as P6, this was their main method for creating their own unique sounds. This suggests that timbre exploration support systems should function effectively with musicians' own external effects processors or techniques to achieve integration with their existing workflows:

> *I get a rough... preset, I'll tweak it a little bit and then I just record that... And then 90 percent of the sound design is either with... three or four [pieces of software], usually some pedals, a sampler... So that's how I've been able to develop my own sounds...* [P6]

3.2.4 Fast and immediate user experience. Four participants (P1, P2, P5, P6) specifically mentioned speed or fluidity as a key factor of their experience with synthesizer programming. P2 expressed a desire to be able to work quickly:

> *Immediacy - that's a good way to put it... intuition, something that allows me to just foster the speed of creativity, because that's a really big thing for me - being able to go fast.* [P2]

Three participants (P2, P6, P8) also noted the desire for having some key parameters always readily available. These were often parameters that have immediate and clear effects on the sound.

> *[My ideal] would be like the [Elektron] Digitone except that it would have more dedicated knobs for filters and stuff like [that] that are playable... I can menu dive within two button pushes or whatever to get to the filter... but I don't want to do that...* [P2]

In other words, even if a system provides new methods for sound design and exploration, musicians will still need easy access to some key parameters.

3.2.5 Tactile input. All participants made use of both hardware and software, but in particular, five participants (P1, P2, P4, P6, P8) noted that having a tactile interface was important to their experience. Some participants stressed that a tactile interface was more fun, exciting, or engaging. As P1 put it:

> *I think programming [software synthesizers] is... not as fun... There's something [about] the lack of tactile sensations [of] turning a knob or... moving your hands around on a table.* [P1]

3.2.6 Overwhelming parameter access. Seven participants (P1, P2, P3, P4, P6, P8, P9) relayed that too many parameters, even if they are easily accessed, can be overwhelming or frustrating. P6 conveyed that he did not like having many parameters easily accessible, because this would distract him from the process of sound creation:

> *I don't try to figure it out and I don't even want to figure it out... When I know what I'm actually doing with a piece of gear, I feel like I'm not actually listening as closely to the sounds... and then I go in a route that's really predictable for me.* [P6]

Overall, participants seemed to find large numbers of parameters frustrating or counter-productive. This suggests that new interface designs cannot work by simply giving musicians improved direct control over every parameter.

3.2.7 Meaningful limitations. Similarly, three participants (P1, P4, P8) expressed the importance of limitations for their creative process. Understanding that a level of constraint is important for the creative process is thus a key insight when designing new sound authoring tools.

> *I find that that helps me too, with limitations. I just have to set the preset that I've already created and then already know this is a cool sound. It kind of makes it easier for me because if I'm constantly tweaking endlessly I just can't decide on anything.* [P1]

3.3 Design Considerations

Based on the seven themes presented above, we derived the following nine design considerations. These design considerations guided the development of SoundTraveller and can inform the development of other creative interfaces, particularly those related to music or other exploratory processes. C2 and C3 are both derived from the theme *presets for exploration* and C5 and C6 are derived from *fast and immediate user experience.*

C1. Musicians' desire for new sounds should not be equated with a desire to do more timbre programming.
C2. Explorations can be initiated with presets rather than by tweaking individual parameters.
C3. Presets should be viewed as a starting point for timbre exploration, rather than as a solution for it.
C4. Timbre exploration designs should be conscious of and support the use of external effects processing.
C5. The exploration process should be fast, fluid, and immediate.
C6. A few key parameters or controls should always be readily available.
C7. A tactile, physical interface should be considered for the exploration process.
C8. Direct access to all parameters should not be equated with ease of exploration.
C9. The timbre exploration process should be guided by meaningful limitations.

4 INTERFACE DESIGN ITERATION

The design considerations in Section 3.3 can work as a validation tool for designers to evaluate whether a proposed prototype can potentially meet the requirements of electronic musicians. However, the guidelines are still implicit in that they do not translate directly into concrete interface designs. In practice, creating a design prototype for supporting timbre exploration is non-trivial. In order to develop our final interface implementation, we conducted an iterative design process, starting from low-fidelity sketches and holding design sessions with multiple participants.

4.1 Intuitions for the Initial Prototypes

Creating the first design concepts to initialize an iterative design process is never an easy task. Our initial concepts for SoundTraveller emerged from a reflection on the considerations developed in the formative study and the first author's expertise in the field of electronic music composition. Here we describe the intuitions that inspired these initial concepts.

Our first intuition came from prior research that has treated timbre as a vector, leading to the notion of timbral space as a way to conceptualize the relationships between timbres [26, 40, 66, 77, 82]. Considering timbres spatially allows us to perceive similar timbres as forming regions, and of presets as landmarks, or points of interest, in a multidimensional space. This thinking inspired our exploration ideas that would allow a user to situate themselves in the timbral space using presets (C3), and then explore a limited region of interest, which led to the development of our initial morphing and region exploration designs.

Our second intuition came from work on genetic algorithms, such as Horner et al. [34]'s, which suggests that the timbre parameter vector can be viewed as a gene, to which we can apply processes of mutation and genetic cross-over. Work on interactive evolution [16] and Scurto and Bevilacqua's work on human-AI collaboration [63] also points to the possibility of user-guided evolution of timbres, starting from known landmarks, or presets.

4.2 Iterative Design Process

Developing these ideas into a series of low fidelity user interface prototypes [68], we conducted an iterative design process consisting of three sessions to develop the features and interface of SoundTraveller. Four participants (P1, P2, P7, and P9) volunteered for this part of the study. The first two sessions were organized virtually with all participants, and the final session where we showcased an early hardware prototype took place in person with P7 and P9. Participants were compensated the equivalent of approximately $18 USD per session. All participants are listed in Table 2 in Appendix A.

Participants' questions, comments, and suggestions were recorded as notes by the interviewer, including drawings directly on the sketches to indicate modifications suggested by the participants. Although participants did not sketch new ideas themselves, they did verbally describe ideas for modifications or novel combinations of features. In all three sessions, participants were encouraged to use the presented prototypes as a basis for thinking about and suggesting changes or new features, and in all three sessions ideas emerged that were not originally part of the respective prototypes.

4.2.1 Design Session 1. We presented 17 sketches showcasing the results of our initial design process. These included early representations of our timbre exploration concepts and possible graphical representations of them, and different methods for interacting with the system, such as a joystick, knobs, and slider configurations. When developing these initial sketches it was important to consider the trade-off between the flexibility of novel interaction approaches and the familiarity of typical synthesizer

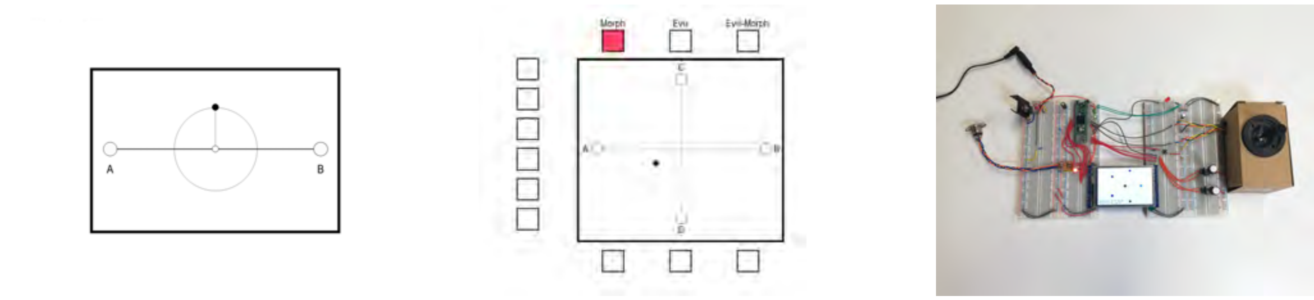

(a) Example Sketch from Design Session 1 **(b) Example Sketch from Design Session 2** **(c) Prototype from Design Session 3**

Figure 1: Examples of the materials used in the participatory design process.

interface elements, particularly given the conservative nature of existing synthesizer interface designs. Therefore, the sketches tend towards using interface elements that would be familiar to synthesizer users, but in novel or unconventional ways. The sketches include both variations on our timbre exploration concepts, and multiple variations of their associated control schemes.

Participants' suggestions during this session resulted in our choice of the morphing and evolutionary modes, as well as confirming support for including a physical interface. It also gave us insights into which types of graphical representations were most intuitive. P1 and P2 both described features for exploring limited sets of parameters, which led to the development of the parameter group controls, described in Section 5.3. P7 suggested that a joystick better supported exploration, but that knobs would be better for nuanced adjustments. Participants also preferred the clarity offered by a screen.

4.2.2 Design Session 2. We presented an animated visual prototype to test the clarity of our system's workflow. This confirmed that the participants found the core workflow intuitive. P9 argued that the interface should include some standard synthesizer controls, which led to the direct parameter controls described in Section 5.3. P9 also suggested that the user should be able to control the degree of randomization in the Evolutionary Mode, which led to the development of the *Mutation Rate* and *Mutation Range* controls described in Section 5.2.1. P2 and P7 emphasized the need for a quick and simple timbre saving feature, while P1 argued for including the undo feature.

4.2.3 Design Session 3. We presented an early hardware prototype, which allowed the participants to create timbres for the first time, and to test the physical experience. Letting participants hear the timbres they produced with the prototype confirmed that the algorithms could be useful for exploration and composition. Participants' suggestions from this session also led to refinements of the exploration algorithms and the final feature set described in Section 5, such as P7's suggestion to limit the range of certain parameters in the Evolutionary Mode, and both participants' assertion that five offspring provided diverse results without overwhelming the user.

4.3 Finalizing the Design

Our participatory design process produced rich results from which to finalize the design of SoundTraveller. However, since it is impractical to ask participants for inspiration about every detail of the design, our final system necessarily contains design aspects that were informed primarily by the first author's over ten years of experience as an electronic musician and synthesizer user. These were carefully chosen to be familiar to musicians with experience using traditional synthesizer interfaces, to support integration with their existing workflows. For instance, the vertical and horizontal layout in the morphing mode was chosen to mimic a familiar Cartesian grid, and the layout of the knobs for the direct parameter controls was chosen to align with the standard synthesizer signal flow. These controls were included above the screen in order to place emphasis on the novel exploration modes offered by SoundTraveller. The exploration controls were laid out to both support two-handed and single-handed operation while simultaneously playing the synthesizer's keyboard.

5 SOUNDTRAVELLER: A NOVEL TIMBRE EXPLORATION SUPPORT SYSTEM

Through our iterative design process, we developed SoundTraveller, a timbre exploration support system in the form of a physical controller that connects with a synthesizer via MIDI[1]. SoundTraveller augments the traditional programming interface by giving the user new tools for timbre exploration that do not rely on individual parameter editing, but instead inspire experimentation and intuitive manipulation, based on musicians' existing sound authoring workflows (C1).

5.1 Hardware and Software Implementation

SoundTraveller was implemented as a standalone hardware device using the Teensy 4.1 microcontroller, chosen for its support of the Arduino environment, relatively high clock speed of 600MHz, and large amount on I/O in a small package. The screen is a 3.5" color TFT LCD display module from Adafruit Industries, chosen for it's relatively large size, color graphics support, and ease of implementation. The joystick is a standard 2-axis resistive type (Figure 2, C). All knobs use endless rotary encoders that

[1]Musical Instrument Digital Interface: https://midi.org/

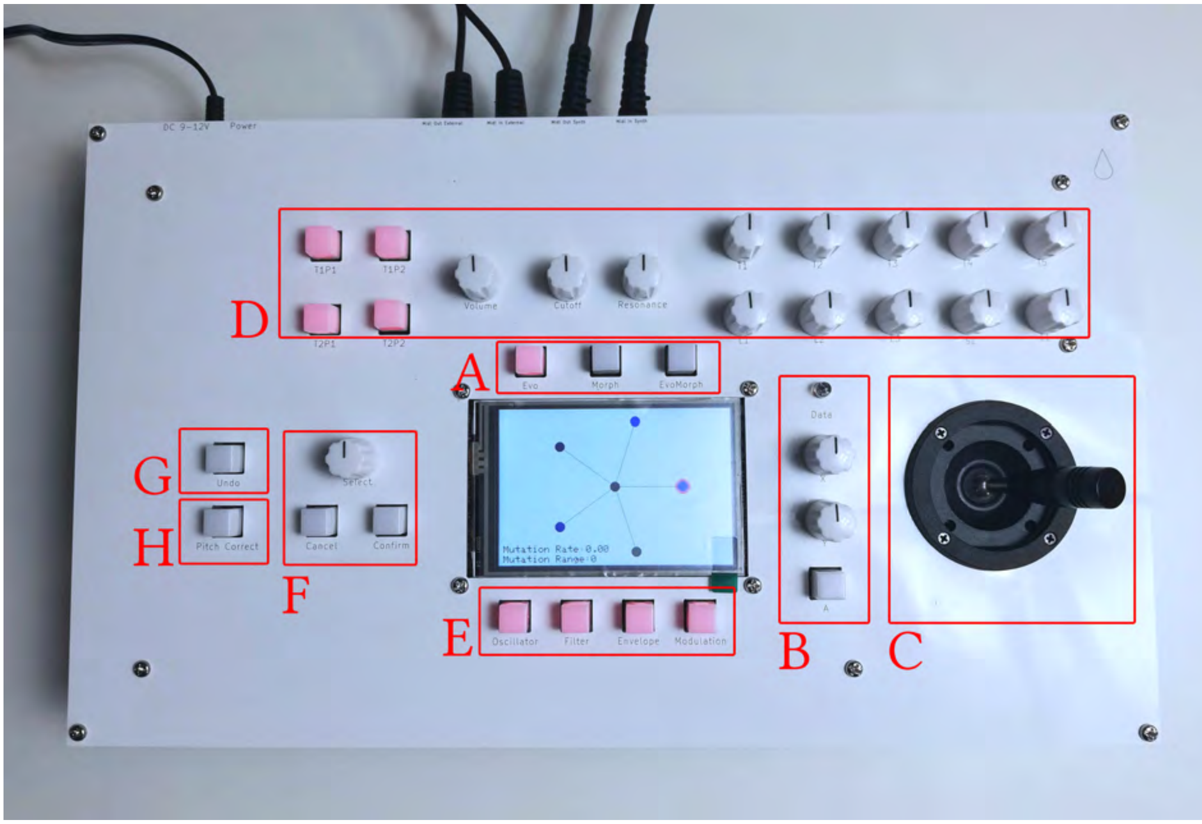

Figure 2: The SoundTraveller interface resultant from our iterative design process, in the evolutionary mode. A: Exploration mode selection buttons. B: Data transfer LED, X Y and A controls used primarily in the evolutionary mode. C: The joystick. D: Direct parameter controls. E: Parameter Group selection buttons. F: Menu controls. G: Undo button. H: Pitch Correct button.

allow parameter adjustment from their current values without any sudden transitions. SoundTraveller features four 5-Pin DIN MIDI connectors to allow for pass-through of MIDI messages to and from external equipment. The entire prototype is implemented on a single printed circuit board and is powered by a standard 9V DC power adaptor.

SoundTraveller is a timbre exploration interface and not a synthesizer in and of itself. As such, we required a synthesizer for it to control. We chose the Roland *D-50 Linear Synthesizer* because it remains a popular instrument among today's electronic musicians despite the original hardware not being manufactured anymore. More importantly, Roland has recently released both software and hardware recreations of the D-50, which ensures the open nature and replicability of research conducted with it. Although the D-50 was first released in 1987, its synthesis engine with over 300 parameters [1] makes it one of the most complex and flexible widely-used synthesizers produced to date. This vast timbral space of the D-50 makes it an excellent candidate with which to evaluate SoundTraveller. SoundTraveller's connection with a standard synthesizer also means that it can easily be integrated into an existing signal chain that may include external audio effects units or other devices (C4).

The firmware for SoundTraveller was developed in the Arduino environment and handles all operation of the system. Communication with the D-50 is achieved using MIDI System Exclusive (Sysex) messages according to the D-50 Midi Implementation [1]. However, MIDI's data transfer rate is limited to 31.25 kbits/s [2], and updating the timbre of the D-50 requires sending approximately 350 Sysex messages[2], each of which consists of at least 11 bytes. This creates an unavoidable bottleneck, meaning that updating the timbre takes approximately 1 second, during which the synthesizer cannot produce sound. This limitation could potentially be overcome using alternative communication protocols, such as OSC [84] or MIDI 2.0 [3], although these still have yet to be widely adopted. Alternatively, a system such as SoundTraveller could be built directly into a synthesizer, eliminating the reliance on an external communication protocol.

5.2 Exploration Modes

The final design of SoundTraveller features two exploration modes, evolutionary and morphing, both of which allow the user to start from one or more preset timbres and explore the neighboring

[2]This is equivalent to one message per parameter, and is a limitation of the D-50's MIDI implementation [1].

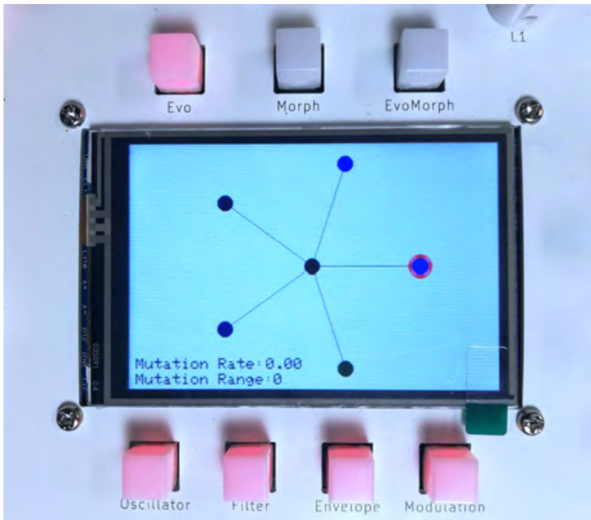

(a) Evolutionary Mode Detail

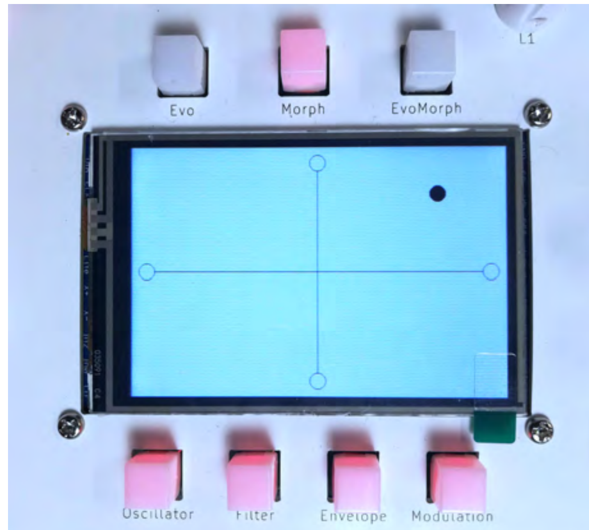

(b) Morphing Mode Detail

Figure 3: Details of the exploration mode graphics. (a) The center black circle represents the parent timbre, with the five offspring timbres generated by the system spaced around its perimeter. Each circle represents a different timbre, which the user can select and audition using the joystick. The values of the mutation rate and mutation range are displayed in the bottom left corner. (b) The two axes indicate the two-dimensional representation of the timbre space defined by the four presets chosen by the user. The circles at the ends of the axes each represent one timbre slot, and the influence of the timbre loaded there on the resulting timbre increases as the user moves closer to it. The black circle visible in the first quadrant is the cursor representing the user's current location and can be freely moved using the joystick. The resolution of the morphing is 100 x 100.

timbral space without having to manipulate individual synthesis parameters (C2, C3).

5.2.1 Evolutionary Mode. The *evolutionary mode* (Figure 3a) allows the user to select a single preset timbre as a starting point and, by iteratively generating offspring using different levels of randomization, traverse the timbral space without the need to edit individual synthesis parameters. By adjusting the *Mutation Rate* and *Mutation Range*, and choosing which offspring to use as the new parent, the user can guide the evolution in a direction of interest. Low *Mutation Rate* and *Mutation Range* values will produce timbres resembling subtle variations on the parent timbre, while higher values will produce increasingly unpredictable results, up to complete randomization (C9). By generating five offspring timbres in a single generation, the morphing mode allows the user to quickly experiment with timbral variations (C5). Pseudo code describing the full algorithm is available in Appendix C.1.

5.2.2 Morphing Mode. The *morphing mode* (Figure 3b) works by interpolating between four preset timbres chosen by the user. This allows the user to locate themselves in the timbral space, using the presets as landmarks. As the user approaches a preset timbre with the cursor, its influence on the resulting timbre increases. The morphing resolution is 100 x 100 points, resulting in 10,000 individual parameter combinations. In practice, however, the morph space is generally composed of several regions of similar timbres. Using the joystick (C7), the user can quickly discover these regions (C5) and then home in more specifically by making small adjustments. Pseudo code describing the full algorithm is available in Appendix C.2. Please refer to our accompanying video for demonstrations of the two modes.

5.2.3 Evolutionary-Morphing Mode. This was an imagined mode that combined features of the evolutionary and morphing modes, allowing users to morph between timbres generated using our genetic algorithm. Due to time constraints and the overlapping nature of these features, it was not implemented for this research.

5.3 Auxiliary Controls

The auxiliary controls give more advanced users deeper control of the timbre creation process (C6, C8). They include: direct parameter controls to change the parameters of the amplifier volume, filter cutoff frequency, resonance, and volume envelope (Figure 2, D); parameter group controls to include or exclude certain parameters from morphing and evolution processes (Figure 2, E); save onto and load from and an SD card (menu system accessed via Figure 2, F); undo (Figure 2, G); and the pitch correction button to align all the oscillators to a standard tuning and set pitch-related parameters to their defaults (Figure 2, H).

6 USER STUDY

We conducted a user study with SoundTraveller to understand its effect on timbre exploration. We had two primary objectives: 1) examine how SoundTraveller can reduce musicians' workload and support their creativity in explorations of new timbres through a comparative study; and 2) investigate the user experience of

SoundTraveller in realistic music composition tasks through a qualitative study. The following protocol was approved by the Institutional Review Board of the authors' university.

6.1 Participants

Ten participants (P7, P9, P11–18) were recruited through the authors' network using a snowballing approach. P7 and P9 continued from the interview study. P11–18 had not participated in the project until this point. All participants were part of a loose global scene of underground ambient and dance music. Although their level of experience and training varied, all participants were self-trained with regard to electronic music and synthesizer programming. All participants have performed and released music semi-professionally.

No participants had prior experience with the Roland D-50, other than P7 and P9 who had used it briefly during Design Session 3 (Section 4). Because the user study was conducted fully in-person, local participants were recruited to replace those who originally participated online, or were otherwise unavailable. As the participants were non-native English speakers, we conducted our interviews in their local language. Participants were compensated the equivalent of approximately $43 USD. All participants are listed in Table 2 in Appendix A.

6.2 Tasks

Previous research into support systems for music composition used structured observation of short, well-defined tasks to collect data from participants under consistent conditions [24, 35]. We followed this methodology and developed our own task design to simulate more realistic sound composition scenarios. In the first task, participants created timbres using SoundTraveller alone. However, artists rarely compose music with one instrument only. To ensure the ecological validity of our subsequent music composition tasks, the participants were encouraged to use SoundTraveller along with their own typical instruments they use to compose music, which they brought to the user study. To this end, participants in the two composition tasks explored SoundTraveller with their own instruments, such as laptop computer and Logic Pro digital audio work station, Eurorack format modular synthesizer, or Teenage Engineering OP-Z.

6.2.1 Timbre Creation Task. This task was intended to evaluate how and how much SoundTraveller aids in the creation of novel timbres. As such, it was designed to encourage the participants to orient themselves toward the goal of creating complete and usable timbres. The participants were instructed to create as many unique timbres as possible for use in their own music within 20 minutes. This was repeated first using the D-50 synthesizer with SoundTraveller, and then using the D-50 only. This fixed condition order ensured that participants were operating SoundTraveller without direct experience of the D-50's synthesis architecture. We chose to limit the comparison to the D-50 because it is one of the most complex synthesizers that is widely used by electronic musicians, unlike alternative interfaces that have only been explored as part of research projects. With both interfaces, participants were initially provided with four preset timbres[3], taken from the D-50 Creative Book [8]. The participants are also asked to rate the starting timbres for reference.

6.2.2 Focused Composition Task. Since musicians rarely create timbres in isolation, and instead integrate sound design into their larger music-making process, we wanted to evaluate the use of SoundTraveller in this context. The *focused composition task* was intended to evaluate how SoundTraveller fits into our participants' larger music-making workflows, as well as allow us to observe how this influences the way participants use it. Participants were instructed to create a short composition using their own equipment and timbres created with SoundTraveller, within 20 minutes. Participants were provided with the same four preset timbres to start with, to ensure they began by producing new timbres.

6.2.3 Exploratory Composition Task. While the *focused composition task's* time pressure meant we were able to observe a very directed timbre creation and composition process, this is not a realistic condition for music composition. The *exploratory composition task* gave the participants more leeway to explore and experiment during the compositional process, which reflects more closely how they would normally use SoundTraveller in their music-making. The participants were instructed to create a short composition using their own equipment and timbres created with SoundTraveller, within 40 minutes. During this task, the participants had access to a variety of presets created in advance by the first author from which to explore and create new timbres.

6.3 Procedure

Before starting the study, participants were offered explanations and time to become familiarized with the use of SoundTraveller and the test synthesizer. To standardize this process, we created two videos explaining the features and use of SoundTraveller and the D-50 synthesizer, which were provided to participants in advance. The study itself was broken into two sessions with a 10-minute break in-between: the first session for timbre creation tasks and the second session for music composition tasks.

After each task, we administered the questionnaires described in Section 6.4. After the completion of all tasks, we conducted a semi-structured interview to obtain qualitative data on participants' experiences with SoundTraveller. We video- and audio-recorded each session and task. We also took detailed notes on the participant's actions and progress. The entire study lasted approximately five hours. The complexity of navigating the D-50's menu system is such that a novice user, even after understanding the available synthesis parameters, may be unable to make meaningful adjustments to the instrument's timbre due to the difficulty of locating the relevant parameters within the menus. To mitigate this issue, the experimenter provided assistance to the participants in order to help them locate any parameters they wished to modify at any time during the experiment. For instance, if the participant wanted to modify a specific parameter, but could not locate it, the experimenter helped them navigate to the appropriate menu. Participants typically required such assistance several times during

[3]The four timbres (patches) were *Bohemian*, *Syn-Strings Hi 1*, *Chorus*, and *Poly-Synth*.

the study, either because they wanted to access a parameter but could not locate it, or because they had forgotten the location of a parameter they had previously adjusted.

6.4 Evaluation Metrics

6.4.1 NASA Task Load Index. The NASA Task Load Index (TLX), is a workload assessment tool first developed by NASA's Human Performance Research Group [28, 29]. The NASA TLX measures workload using six factors, for each of which a sub-score is calculated based on subject's responses to pair-wise comparisons. These sub-scores are weighted according to the subjects ranking of importance of each factor in the task, and summed to calculate a total score out of 100. We included this to measure participants' cognitive workload during each task.

6.4.2 Creativity Support Index. The Creativity Support Index (CSI) is a rating system developed by Cherry and Latulipe [10] based on the NASA TLX. Operationally, the CSI is the same as the TLX, except that the factors are about creativity support instead of workload. Remy et al. [57] suggest that the CSI may be a good contender for a standardized measurement of creativity support. We included this to measure how well SoundTraveller supported participants' creative processes throughout all tasks.

6.4.3 Implication Assessment Ratings. While the TLX and CSI are widely applied standardized measures, in order to assess more directly how our prototype adheres to the considerations we identified in the formative study, we developed nine corresponding questions using five-point Likert scale agreement statements (Strongly Disagree (-2) to Strongly Agree (2)), producing a total value between -18 and 18. The full set of questions can be found in Appendix D.

6.5 Semi-Structured Interviews

The first author conducted semi-structured interviews with each participant to gather qualitative data on the performance of SoundTraveller. Appendix E lists the pre-determined questions. Participants were also given the opportunity to freely express any thoughts and impressions. We made it clear that participants should express both positive and negative opinions, and asked explicitly about difficulties they experienced with SoundTraveller.

We performed a thematic analysis on the results. The first author coded the interviews first using a deductive approach guided by our design considerations in Section 3.3, followed by an inductive approach to categorize results that fell outside this scope. Two authors then derived 12 themes with 209 quotations. The themes were then reviewed with the other authors and combined or revised into the eight presented below. For example, the theme "changes to the musical process" contained quotations that overlapped with the other themes, so the relevant quotations were distributed among them. Likewise, "challenges and frustrations" and "future improvements" were combined into "areas for improvement."

7 RESULTS

7.1 Quantitative Results

For the timbre creation task, we sought to evaluate how well SoundTraveller supports timbre creation and how this compares with the original programming interface of the D-50 synthesizer. Table 1 shows the means and standard deviations of the evaluation metrics for both interfaces. We conducted a one-way repeated-measures ANOVA on the three metrics against the two interface conditions. We found significant differences for the NASA-TLX score ($F(1,9)$=20.12, p<.01, generalized η^2=.48), CSI score ($F(1,9)$=64.60, p<.001, generalized η^2=.82), and the implication assessment score ($F(1,9)$=146.18, p<.001, generalized η^2=.88).

Table 1 also includes the means and standard deviations of these metrics for the composition tasks. While our major objective of including composition tasks was to provide participants with more realistic experiences and thus obtain more ecologically-valid qualitative results, these quantitative results also suggest that SoundTraveller was positively received in these tasks.

7.2 Qualitative User Experience Results

7.2.1 Perceived Value of SoundTraveller. All participants had a positive experience with SoundTraveller, but they identified different value it provided them. P14, who often takes substantial time creating timbres for her music, felt that SoundTraveller would *"surely be useful for making a variety of sounds,"* and that its value lay in *"the fact that you can enjoy using it without having to overthink."* P18 explained that SoundTraveller is *"more intuitive than the equipment I normally use"* and felt she was *"playing around, seeing how it would turn out if I put this and that together."* P12 commented that SoundTraveller acted as assistant during the exploration process:

> *It feels like the device is thinking together with me, and I can [explore] in a way that's really human-like and fun. On the other hand, with the D-50, I have to think of everything and get there by myself. So it's an isolating process, and I think that could cause frustration. With [SoundTraveller], it feels like there's another person with me, so I can trust in it.* [P12]

Participants used SoundTraveller a variety of ways, particularly during the composition tasks. P9 and P17 primarily used the evolutionary mode, while P7 and P11 focused on the morphing mode. However, SoundTraveller supports iterating between both modes, and six participants (P12, P13, P14, P15, P16, P18) took advantage of this workflow. P15 used the evolutionary mode to sketch out some timbres before homing in with the morphing mode:

> *I think of course you can use the evolutionary mode as-is, but because it's a bit difficult to make sounds you like in one shot, I'd collect several sounds I liked and move them over, so that I could morph between several sounds I'd made using the evolutionary mode.* [P15]

7.2.2 Encouraging Exploration and Timbre Creation. Four participants (P7, P11, P12, P17) explained that they rarely create their own timbres. For these participants, SoundTraveller offered new possibilities for sound design despite their lack of experience or understanding of synthesis parameters. P17 stated that she valued SoundTraveller because it allowed her *"to come up with far more timbres than the number originally in the synth."* P7 emphasized that SoundTraveller allowed her to explore new timbres more intuitively than with presets:

Table 1: The average scores and standard deviations of NASA-TLX, CSI, and the implication assessment scores in the timbre creation and music composition tasks.

Task and Interface	NASA-TLX Score	CSI Score	Implication Score
Timbre Creation with SoundTraveller	46.0 (21.6)	77.7 (13.0)	8.7 (2.6)
Timbre Creation with D-50 Only	82.0 (17.8)	28.5 (18.9)	-8.7 (4.1)
Focused Composition with SoundTraveller	49.0 (15.8)	76.2 (21.3)	8.2 (4.8)
Exploratory Composition with SoundTraveller	40.0 (25.0)	80.2 (17.7)	8.3 (4.9)

> *It's more intuitive than going and looking at the [names of presets] on the computer, so it was really easy to search [for new timbres]. When it comes to choosing timbres, I feel like I can find them much faster than normal...* [P7]

7.2.3 Timbre Exploration as Inspiration. For six participants (P9, P13, P14, P15, P16, P18), the value of SoundTraveller lay not so much in the ease of accessing new timbres, but in how it allowed them to explore timbres beyond their habitual sound palettes. P14 emphasized the fun of making *"sounds that I wouldn't think of myself"*, and P18 similarly thought *"it was great to go places I hadn't imagined."* For P9, this ability to explore serendipity was really important to his music-making process:

> *When I want to use some somewhat coincidental elements that I wouldn't think to tweak that way myself, the evolutionary mode is really interesting; I like adding those coincidental elements when I'm making music, and without them I tend to lose interest, so the evolutionary mode was really just right in that way.* [P9]

As P16 commented, these types of coincidences can also be inspiring:

> *I thought it was great how timbres I wouldn't think of myself just pop out of [SoundTraveller], so from there, I already got inspiration, and maybe that brought out some new riffs I wouldn't have thought of.* [P16]

7.2.4 Quick Exploration. In our preliminary study, we identified that the speed of exploration would be an important factor. Five participants (P7, P12, P16, P17, P18) specifically expressed that SoundTraveller was a fast way to create sounds; P12: *"In any case it was fast;"* P17: *"The way the evolutionary mode produces five sounds feels quick; you can search for sounds quickly."* P18 compared the sound creation process with that using her analogue synthesizers, suggesting that SoundTraveller *"was super fast. It was more direct than with my analogue synthesizers..."*

However, as we noted in Section 5.1, there is an unavoidable delay when new timbre data is sent to the synthesizer. P9, P11 and P15 commented on this; as P15 said, *"I think it would be better if the loading was a bit faster; I can't wait..."*

7.2.5 Value of the Physical Interface. SoundTraveller was specifically developed as a hardware device, following the results of our interview study that suggested many participants appreciated the tangible, hands-on aspects of hardware. During the user study, six participants (P7, P9, P11, P12, P15, P18) referenced the value of having a physical interface. As P11 put it, *"it has a lot of meaning because it's visual."* P12 expressed how having all the controls laid out in front of him made SoundTraveller fast and easy to use:

> *...if you were doing this just on an iPhone app, for instance, you'd have to flip through multiple pages, and I don't think it would be as fast as [SoundTraveller]... It's really best not to have the work of flipping through [multiple pages].* [P12]

7.2.6 Integrating with Existing Workflows. Based on our interface design considerations, we were concerned with how SoundTraveller would integrate with users' existing music-making workflows. All participants brought some of their own equipment to the user study and created short compositions during the composition tasks, demonstrating SoundTraveller's ability to integrate effectively. Four participants (P12, P15, P16, P17) also commented on this explicitly. Particularly for participants who had workflows that started with preset timbres, SoundTraveller fit well by affording them a means to generate a variety of new timbres to process. P12 described how SoundTraveller worked in this way:

> *I pretty much always use presets to make music, and haven't really made sounds using actual synthesizers before, so the way [SoundTraveller] turns [the timbres] into presets for you really suited me I think.* [P12]

7.2.7 Level of Control. With SoundTraveller, we aimed to provide users with a flexible level of control without overwhelming them with too many parameters or options. Seven of the participants (P7, P9, P12, P13, P14, P16, P17) expressed that SoundTraveller provided them with sufficient control without being overwhelming. P9 felt that the level of control was *"just right"* and P13 thought *"the range of control was more than enough."* However, not all participants made use of all the controls provided. P7 reflected that she *"didn't really get"* the direct parameter controls, though *"if [I] could concentrate this much on composing, I don't feel like there are too many."*

While they similarly did not use all the controls provided, P12, P14, and P17 also felt that the interface provided more than ample control. P12 discussed in more detail the value of the direct parameters, suggesting that they played an important role in homing in on his desired sounds:

> *In the end, the sounds created by the randomization do need a bit of flavoring, so... I think it was good to have the envelope functions. If you didn't have that then... you [would] need to search for different sounds, and I think that would be quite a detour... I think they are probably really necessary knobs.* [P12]

However, six participants (P9, P11, P14, P15, P16, P18) also discussed what they were unable to do with the current prototype, such as removing the reverb effect, or changing other specific parameters.

7.2.8 Areas for Improvement. While the participants all expressed that they enjoyed using SoundTraveller, there were some frustrations and challenges during the user study. As we discussed above, P9, P11 and P15 were frustrated by the data transmission delay. Due to the unpredictable nature of both algorithms and the numerous parameter combinations that can result in little or no sound being produced by the D-50, sometimes the morphing and evolutionary modes do not produce usable timbres. P15 commented on the difficulty this caused at times:

> *I thought it might be best if the scope of the randomization of the volume and attack was limited a little.* [P15]

P15 noted that the randomization allows you to just keep making new timbres over and over, but P12 suggested that this can also be frustrating:

> *You can't get to a specific sound with absolute certainty, so there's nothing to do except leave it up to [SoundTraveller] until by chance you come across a sound you were imagining. You have to repeat it over and over until you get [a timbre] where you go 'Ah, like this!'.* [P12]

But, P12 followed up: *"If you use it while thinking of it that way, maybe it's not a problem."*

Three participants (P13, P14, P15) also noted the fact that the direct parameter controls, being implemented with endless encoders, do not provide visual feedback regarding the current parameters' values. P14 explained how this forced the user to make *"rather intuitive adjustments"* and P13 stated that this led to some confusion because *"you can't see visually what state it's in."*

8 DISCUSSION

8.1 Aesthetic Choices in Creative Processes

SoundTraveller supports a variety of use strategies (Section 7.2.1) and allows participants to explore different options quickly, starting from preset timbres (Section 7.2.4). This approach is in line with our participants' existing workflows, and successfully encouraged them to explore and create new timbres during the user study's tasks.

Creative and technical systems with high dimensional parameter spaces share a number of key characteristics, including that within an overwhelmingly large design space, *"usually only a small subset of all possible settings is of interest and has practical value"* [83]. All domains of creativity require *"the identification of appropriate building blocks"* [31], and, as a result, working with presets is a typical approach for design explorations, particularly for novice users [53, 83, 87]. However, the difficulty of exploring effectively, both with and without presets, can effectively limit users to *"a few familiar, but poor, choices"* [53].

A key difficulty in parameter space exploration then, is how to go about finding the subsets of settings, or *"building blocks,"* that are of interest to the user. With SoundTraveller we started from the insight that presets, while not necessarily representative of users' end goals, provide starting points or reliable landmarks in the parameter space and are a familiar way of working for many musicians. SoundTraveller's exploration modes are then about being able to meaningfully constrain the parameter space being explored. While it was acceptable, and even desirable, for there to be a level of unpredictability in the system's output, by using known preset timbres to constrain the exploration, participants were able to create new timbres in *"much faster"* (P7) and *"more direct"* (P18) ways. Our results suggest that SoundTraveller's exploration modes and overall paradigm helped participants overcome the difficulties of large parameter space exploration for timbre creation.

SoundTraveller is designed to support users' creativity through exploration; however, our participants stated that they also benefited from SoundTraveller's *"suggestions"* (P16) and the way that they could explore various potential timbres *"without having to over-think"* (P14). We also found that SoundTraveller placed a significantly lower workload burden than a traditional timbre creation interface on participants, suggesting that it alleviated some of the difficulty associated with exploring a large parameter space. The design of SoundTraveller can thus complement existing aesthetic parameter space explorations by providing quick access to design variations. This suggests that both our design considerations in Section 3.3 and our user study results in Section 7.2 are generalizable beyond electronic music and that SoundTraveller's design could be used as a basis for further exploration in other creative domains.

8.2 Timbre Exploration: Creativity Through Sound Design

The results of our interview study suggest that traditional synthesizer interfaces do not adequately support timbre creation even though many electronic musicians want access to new and unique timbres. Research in HCI for electronic music has typically focused on musical performance and composition, rather than sound design or timbre exploration. Our interview results demonstrate that timbre is also a key element of electronic musicians' creative process and this motivated us to design SoundTraveller. A major challenge in timbre exploration and creation is not only the manipulation of potentially hundreds of synthesis parameters but also that *"the core language of [said parameters] has no well-defined mapping onto the task languages of subjective timbre categories used by musicians"* [65]. This, as we have seen in our interviews, results in many musicians relying on preset timbres, despite expressing dissatisfaction with their quality or quantity.

Our results positively support that giving musicians access to new timbre exploration tools encourages them to develop their own unique timbres, leading to better support for their creativity. Furthermore, participants described SoundTraveller as offering them suggestions for timbres (P16) or even acting as an assistant (P12), supporting the notion suggested by Scurto et al. [63, 64] that systems for computer music can best support users' creativity by acting as collaborators. We also observed that the unpredictability of SoundTraveller's exploration modes acted as a source of inspiration for some participants, and encouraged them to use timbres from beyond their typical sound palettes. These results suggest that interfaces like SoundTraveller can simplify timbre explorations as well as contribute to users' creative processes in a variety of ways.

We therefore argue that future synthesizer interfaces should integrate human-computer collaboration to support multiple approaches to this aspect of musical creativity. SoundTraveller demonstrates that the abstraction of parameters within musicians' existing practice of using presets is a promising direction for supporting creativity in future synthesizer interfaces.

8.3 Shared Agency in Creativity Support

Supporting creativity has been named one of HCI's grand challenges [71], and prior work has explored its key principles [31, 58], placing substantial emphasis on users' ability to explore their design space. However, notions of shared agency do not play a role in these theories of creativity. As we discussed Section 2.1, recent work on creativity support has explicitly or implicitly approached the problem from this agentic perspective, but largely without discussing its implications for creativity or interactive systems in HCI in general.

Considering synthesizers and electronic instruments from the perspective of shared agency is particularly interesting when we consider the timbre creation interface. While traditional acoustic instruments typically only support indirect and non-linear mappings between inputs and their resulting timbre [37], a synthesizer's timbre creation interface gives users controls for manipulating the instrument's timbre that are independent from its performance interface. In this sense, the timbre interface forms an added layer of mediation between the musician and the instrument, which can help or hinder their ability to complete sound design tasks, an essential part of the composition process.

Traditional synthesizer timbre interfaces give users access to all synthesis parameters, which can be manipulated to change the timbre of the instrument. Although the interface imposes some restrictions, from a design perspective, agency for manipulating timbre is vested entirely in the user. However, because the synthesis engine and interface are decoupled, novel timbre creation interfaces can shift the balance of agency between the user and the system, which can aid users in their creative tasks, but also poses questions for the role of agency in the creative process.

Recent work on creativity support using AI agents suggests that explicitly creating intelligent agentic systems is a powerful way to aid users achieve their goals in technical creative tasks. For example, Scurto et al.'s work on timbre creation with AI collaborative agents [63, 64] suggests that shared agency can be most meaningful when applied to idea generation or creative exploration tasks. SoundTraveller is different in that it does not use any sort of artificial intelligence, nor was it explicitly designed to act as an agent. However, SoundTraveller's exploration modes take over the task of manipulating synthesis parameters and introduce a level of unpredictability into the exploration process, which led our participants to experience SoundTraveller as though it were *"another person"* (P12) that gave them *"various suggestions"* (P16).

This sense of SoundTraveller's agency seemed to come from the fact that the user's inputs to the system influence, but do not completely determine the resulting timbres. This suggests that the agentic nature of SoundTraveller emerged from this balance of control between the system and the user. Offering the user too much control would have resulted in SoundTraveller appearing as simply an alternative control surface, while too little would make it an automatic timbre generation system. Likewise, too little unpredictability would rest all agency in the user, while too much would make the system practically unusable. This balance seems particularly pertinent in technical creative systems, because when it is off, users will either feel frustrated because they cannot access the control they want, or overwhelmed by too much control, both of which distract them from their creative goals. Furthermore, SoundTraveller's different modes and granularities of control provided participants with the ability to find an appropriate *personal* balance of agency between themselves and the system, where some participants relied exclusively on the exploration modes, while others used the direct parameter controls and parameter groups to hone their timbres more directly. This led to not only easier timbre exploration, but also to inspiration or unique ideas for composition.

This suggests that systems that abstract from technical parameters and therefore have some level of unpredictability have deeper impacts than those envisioned by Resnick et al. [58] in their discussion of black boxes. Balancing the level of control and unpredictability, and therefore fostering a sense of shared agency between the user and the system can be an effective way to support creativity, particularly in exploration tasks. But this also suggests a need for more exploration around the boundaries of this phenomenon. If a system like SoundTraveller can be felt by users to be agentic, then what are the necessary and sufficient conditions for this effect? And given that the agentic nature of SoundTraveller had explicit impacts on our participants' creative experiences, how can such agentic systems be harnessed to better support particular creative outcomes? Future work with creativity support should explicitly consider the entangled nature of users and systems, and specifically how this can impact users' creative workflows and sense of artistic authorship.

8.4 Limitations

There were several limitations in the design and evaluation of SoundTraveller. The limited number of participants and the highly idiosyncratic nature of their music-creation processes may constrain the generalizability of our findings. Our study participants had different levels of experience in electronic music creation (see Appendix A) and composed in various styles, but future work should study a broader set of musicians to further validate how common the music creation processes observed in this work are. Although we argue, in Section 5.1, that the D-50 is a good candidate synthesizer for testing SoundTraveller, it is only one example among many different synthesis methods and control schemes. As such, the design of the D-50 likely influenced our participants' experiences, and thus exploring the use of SoundTraveller with other synthesizers would broaden our results. SoundTraveller is in some ways tailored to the architecture of the D-50, and how best to adapt it to other synthesizers or creative domains remains an open question. The controlled nature of our experimental design allowed us to systematically uncover the benefits of SoundTraveller. However, the effects of SoundTraveller during longer-term use are still unknown. Future work should investigate how the long-term use of systems like SoundTraveller impacts musicians' creative processes, experience of agency, and sense of artistic authorship.

9 CONCLUSION

Timbre exploration and creation are key elements of the creative process when composing electronic music. Through our interview study and iterative design process, we developed SoundTraveller, a novel interactive timbre exploration prototype. Our evaluation confirmed that SoundTraveller supported participants' exploration, decreased their cognitive work load, and increased their perceived creativity. By abstracting away from technical parameters and providing flexible timbre explorations, SoundTraveller offers users new and inspiring ways of exploring timbre, supporting their creativity in electronic music composition. But our results also suggest that the entangled nature of SoundTraveller shifted the balance of agency between the user and system, pointing towards the need for more research into the nature of shared agency in creativity support systems. We encourage future work to explore how the insights gained from SoundTraveller can transform future synthesizer interfaces, but, more generally, to investigate explicitly how shared agency in creative systems can be harnessed to aid users in their creative processes across domains.

ACKNOWLEDGMENTS

We thank Carla F. Griggio, Kengo Shibata, Shoko Sano, Keitaro Shimizu, Ginshi Shimojima, Takuma Masuda, and the other members of IIS Lab for their support during this work. We also thank Naotake Masuda for his comments and insights. Finally, we thank all our participants for their time and willingness to share their experiences. This research was supported by the Japanese Government MEXT Scholarship.

REFERENCES

[1] 1987. *D-50 MIDI Implementation.* Technical Report. Roland Corporation. https://archive.org/details/synthmanual-roland-d-50-midi-implementation. Accessed: 2022-02-18.

[2] 2014. *MIDI 1.0 Electrical Specification Update.* Technical Report CA-033. The MIDI Association.

[3] 2020. *MIDI 2.0 Specification Overview.* Technical Report M2-100-U. The MIDI Association.

[4] Dzmitry Aliakseyeu, Jean-Bernard Martens, and Matthias Rauterberg. 2006. A computer support tool for the early stages of architectural design. *Interacting with Computers* 18, 4 (2006), 528–555.

[5] Daniel Bernhardt and Peter Robinson. 2008. Interactive Control of Music Using Emotional Body Expressions. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (Florence, Italy) *(CHI EA '08)*. ACM, New York, NY, USA, 3117–3122. https://doi.org/10.1145/1358628.1358817

[6] Louis Bigo, Jérémie Garcia, Antoine Spicher, and Wendy E Mackay. 2012. Papertonnetz: music composition with interactive paper. In *Sound and Music Computing*.

[7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3 (01 2006), 77–101. https://doi.org/10.1191/1478088706qp063oa

[8] Jeff Burger. 1988. *Roland D-50 Creative Book.* Roland Corporation.

[9] Michael Bylstra and Haruhiro Katayose. 2005. Painting as an Interface for Timbre Design. In *Entertainment Computing - ICEC 2005*, Fumio Kishino, Yoshifumi Kitamura, Hirokazu Kato, and Noriko Nagata (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 303–314.

[10] Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools through the Creativity Support Index. *ACM Trans. Comput.-Hum. Interact.* 21, 4, Article 21 (June 2014), 25 pages. https://doi.org/10.1145/2617588

[11] Saemi Choi, Kiyoharu Aizawa, and Nicu Sebe. 2018. Fontmatcher: font image paring for harmonious digital graphic design. In *23rd International Conference on Intelligent User Interfaces*. 37–41.

[12] John M Chowning. 1973. The synthesis of complex audio spectra by means of frequency modulation. *Journal of the audio engineering society* 21, 7 (1973), 526–534.

[13] Elizabeth Clark, Anne Spencer Ross, Chenhao Tan, Yangfeng Ji, and Noah A Smith. 2018. Creative writing with a machine in the loop: Case studies on slogans and stories. In *23rd International Conference on Intelligent User Interfaces*. 329–340.

[14] Giovanni Costantini, Massimiliano Todisco, and Massimo Carota. 2007. A neural network based interface to real time control musical synthesis processes. In *Proceedings of the 11th WSEAS International Conference on CIRCUITS, Agios Nikolaos, Crete Island, Greece.* Citeseer, 41–45.

[15] Paulo Mateus Moura da Silva, César Lincoln Cavalcante Mattos, and Amauri Holanda de Souza Júnior. 2019. Audio plugin recommendation systems for music production. In *2019 8th Brazilian Conference on Intelligent Systems (BRACIS)*. IEEE, 854–859.

[16] Palle Dahlstedt. 2001. Creating and exploring huge parameter spaces: Interactive evolution as a tool for sound generation. In *ICMC*.

[17] Jesse Engel. 2017. Making a Neural Synthesizer Instrument. https://magenta.tensorflow.org/nsynth-instrument. Accessed: 2022-12-05.

[18] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan. 2017. Neural audio synthesis of musical notes with wavenet autoencoders. In *International Conference on Machine Learning*. PMLR, 1068–1077.

[19] Philippe Esling, Naotake Masuda, Adrien Bardet, Romeo Despres, and Axel Chemla-Romeu-Santos. 2019. Flow synthesizer: Universal audio synthesizer control with normalizing flows. *Applied Sciences* 10, 1 (2019), 302.

[20] Philippe Esling, Naotake Masuda, and Axel Chemla-Romeu-Santos. 2021. FlowSynth: simplifying complex audio generation through explorable latent spaces with normalizing flows. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 5273–5275.

[21] Christopher Frauenberger. 2019. Entanglement HCI the next wave? *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 1 (2019), 1–27.

[22] Jonas Frich, Midas Nouwens, Kim Halskov, and Peter Dalsgaard. 2021. How digital tools impact convergent and divergent thinking in design ideation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–11.

[23] Jérémie Garcia, Thibaut Carpentier, and Jean Bresson. 2017. Interactive-compositional authoring of sound spatialization. *Journal of New Music Research* 46, 1 (2017), 74–86.

[24] Jérémie Garcia, Theophanis Tsandilas, Carlos Agon, and Wendy E. Mackay. 2014. Structured Observation with Polyphony: A Multifaceted Tool for Studying Music Composition. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) *(DIS '14)*. Association for Computing Machinery, New York, NY, USA, 199–208. https://doi.org/10.1145/2598510.2598512

[25] Nan-Wei Gong, Mathew Joel Laibowitz, and Joseph A Paradiso. 2009. Musicgrip: A writing instrument for music control. New Interfaces for Musical Expression.

[26] John Michael Grey. 1975. *An Exploration of Musical Timbre Using Computer-Based Techniques for Analysis, Synthesis and Perceptual Scaling.* Stanford University.

[27] Lamtharn Hantrakul and Li-Chia Yang. 2018. Neural Wavetable: a playable wavetable synthesizer using neural networks. *arXiv preprint arXiv:1811.05550* (2018).

[28] Sandra G Hart. 1986. NASA task load index (TLX). (1986).

[29] Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (2006), 904–908. https://doi.org/10.1177/154193120605000909 arXiv:https://doi.org/10.1177/154193120605000909

[30] Erin J Hastings, Ratan K Guha, and Kenneth O Stanley. 2008. Interactive evolution of particle systems for computer graphics and animation. *IEEE Transactions on Evolutionary Computation* 13, 2 (2008), 418–432.

[31] Thomas T Hewett. 2005. Informing the design of computer-based environments to support creativity. *International Journal of Human-Computer Studies* 63, 4-5 (2005), 383–409.

[32] Josh Holinaty, Alec Jacobson, and Fanny Chevalier. 2021. Supporting Reference Imagery for Digital Drawing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2434–2442.

[33] Simon Holland, Andrew P McPherson, Wendy E Mackay, Marcelo M Wanderley, Michael D Gurevich, Tom W Mudd, Sile O'Modhrain, Katie L Wilkie, Joseph W Malloch, Jérémie Garcia, et al. 2016. Music and HCI. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 3339–3346.

[34] Andrew Horner, James Beauchamp, and Lippold Haken. 1993. Machine tongues XVI: Genetic algorithms and their application to FM matching synthesis. *Computer Music Journal* 17, 4 (1993), 17–29.

[35] Cheng-Zhi Anna Huang, David Duvenaud, and Krzysztof Z. Gajos. 2016. ChordRipple: Recommending Chords to Help Novice Composers Go Beyond the Ordinary. In *Proceedings of the 21st International Conference on Intelligent User Interfaces* (Sonoma, California, USA) *(IUI '16)*. Association for Computing Machinery, New York, NY, USA, 241–250. https://doi.org/10.1145/2856767.2856792

[36] Sicong Huang, Qiyang Li, Cem Anil, Xuchan Bao, Sageev Oore, and Roger B Grosse. 2018. Timbretron: A wavenet (cyclegan (cqt (audio))) pipeline for musical timbre transfer. *arXiv preprint arXiv:1811.09620* (2018).

[37] Andy Hunt and Marcelo M. Wanderley. 2002. Mapping performer parameters to synthesis engines. *Organised Sound* 7, 2 (2002), 97–108. https://doi.org/10.1017/

S1355771802002030
[38] Aaron Isaksen, Dan Gopstein, Julian Togelius, and Andy Nealen. 2015. Discovering unique game variants. In *Computational Creativity and Games Workshop at the 2015 International Conference on Computational Creativity*. Citeseer.
[39] Youngseung Jeon, Seungwan Jin, Patrick C Shih, and Kyungsik Han. 2021. FashionQ: an ai-driven creativity support tool for facilitating ideation in fashion design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–18.
[40] Colin G. Johnson and Alex Gounaropoulos. 2006. Timbre Interfaces Using Adjectives and Adverbs. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression* (Paris, France) *(NIME '06)*. IRCAM — Centre Pompidou, Paris, France, France, 101–102. http://dl.acm.org/citation.cfm?id=1142215.1142239
[41] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. 2007. The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (Baton Rouge, Louisiana) *(TEI '07)*. ACM, New York, NY, USA, 139–146. https://doi.org/10.1145/1226969.1226998
[42] Kevin Gonyop Kim, Richard Lee Davis, Alessia Eletta Coppi, Alberto Cattaneo, and Pierre Dillenbourg. 2022. Mixplorer: Scaffolding Design Space Exploration through Genetic Recombination of Multiple Peoples' Designs to Support Novices' Creativity. In *CHI Conference on Human Factors in Computing Systems*. 1–13.
[43] Shuhei Kodama, Pierre Poulin, Tomoaki Moriya, and Tokiichiro Takahashi. 2018. Creativity enhancement of painterly rendering using a suggestive interface. *Computers & Graphics* 71 (2018), 42–54.
[44] Joshua Ryan Lam and Charalampos Saitis. 2021. The Timbre Explorer: A Synthesizer Interface for Educational Purposes and Perceptual Studies. In *NIME 2021*. PubPub.
[45] Bruno Latour. 2007. *Reassembling the social: An introduction to actor-network-theory*. Oup Oxford.
[46] Yuyu Lin, Jiahao Guo, Yang Chen, Cheng Yao, and Fangtian Ying. 2020. It is your turn: collaborative ideation with a co-creative robot through sketch. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–14.
[47] Jörn Loviscach. 2008. Programming a Music Synthesizer through Data Mining. In *NIME*. 221–224.
[48] Matthieu Macret and Philippe Pasquier. 2014. Automatic Design of Sound Synthesizers As Pure Data Patches Using Coevolutionary Mixed-typed Cartesian Genetic Programming. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation* (Vancouver, BC, Canada) *(GECCO '14)*. ACM, New York, NY, USA, 309–316. https://doi.org/10.1145/2576768.2598303
[49] David Miller. 2016. AgentSmith: Exploring agentic systems. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 234–238.
[50] Thomas J. Mitchell. 2011. SoundGrasp: A gestural interface for the performance of live music.
[51] Mohammad Amin Mozaffari, Xinyuan Zhang, Jinghui Cheng, and Jin LC Guo. 2022. GANSpiration: Balancing Targeted and Serendipitous Inspiration in User Interface Design with Style-Based Generative Adversarial Network. In *CHI Conference on Human Factors in Computing Systems*. 1–15.
[52] Henry Newton-Dunn, Hiroaki Nakano, and James Gibson. 2003. Block Jam: A Tangible Interface for Interactive Music. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression* (Montreal, Quebec, Canada) *(NIME '03)*. National University of Singapore, Singapore, Singapore, 170–177. http://dl.acm.org/citation.cfm?id=1085714.1085755
[53] Peter O'Donovan, Jānis Lībeks, Aseem Agarwala, and Aaron Hertzmann. 2014. Exploratory font selection using crowdsourced attributes. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–9.
[54] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* (2016).
[55] Srishti Palani, David Ledo, George Fitzmaurice, and Fraser Anderson. 2022. "I don't want to feel like I'm working in a 1960s factory": The Practitioner Perspective on Creativity Support Tool Adoption. In *CHI Conference on Human Factors in Computing Systems*. 1–18.
[56] Chengzhi Peng. 1992. Survey of collaborative drawing support tools. *Computer Supported Cooperative Work (CSCW)* 1, 3 (1992), 197–228.
[57] Christian Remy, Lindsay MacDonald Vermeulen, Jonas Frich, Michael Mose Biskjaer, and Peter Dalsgaard. 2020. *Evaluating Creativity Support Tools in HCI Research*. Association for Computing Machinery, New York, NY, USA, 457–476. https://doi.org/10.1145/3357236.3395474
[58] Mitchel Resnick, Brad Myers, Kumiyo Nakakoji, Ben Shneiderman, Randy Pausch, Ted Selker, and Mike Eisenberg. 2005. Design principles for tools to support creative thinking. (2005).
[59] Matthew Rodger, Paul Stapleton, Maarten Van Walstijn, Miguel Ortiz, and Laurel Pardue. 2020. What makes a good musical instrument? A matter of processes, ecologies and specificities. In *Proceedings of the international conference on New Interfaces for Musical Expression*. Birmingham City University, Birmingham, UK, 405–410.
[60] Günther Schatter, Emanuel Züger, and Christian Nitschke. 2005. A Synaesthetic Approach for a Synthesizer Interface based on Genetic Algorithms and Fuzzy Sets. In *in Proc. International Computer Music Conference 2005 (ICMC 2005*. 664–667.
[61] Kevin Schlei. 2012. TC-11: A Programmable Multi-Touch Synthesizer for the iPad. In *New Interfaces For Musical Expression*. https://www.nime.org/proceedings/2012/nime2012_230.pdf
[62] Diemo Schwarz, Grégory Beller, Bruno Verbrugghe, and Sam Britton. 2006. Real-time corpus-based concatenative synthesis with catart. In *9th International Conference on Digital Audio Effects (DAFx)*. 279–282.
[63] Hugo Scurto and Frédéric Bevilacqua. 2018. Appropriating Music Computing Practices Through Human-AI Collaboration. In *Journées d'Informatique Musicale (JIM 2018)*. Amiens, France. https://hal.archives-ouvertes.fr/hal-01791504
[64] Hugo Scurto, Bavo Van Kerrebroeck, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Designing deep reinforcement learning for human parameter exploration. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 1 (2021), 1–35.
[65] Allan Seago, Simon Holland, and Paul Mulholland. 2004. A Critical Analysis of Synthesizer User Interfaces for Timbre. In *Proceedings of the XVIII British HCI Group Annual Conference HCI 2004*, Andy Dearden and Leon Watt (Eds.). Vol. 2. Research Press International, Bristol, UK, 105–108. http://oro.open.ac.uk/5688/
[66] Allan Seago, Simon Holland, and Paul Mulholland. 2010. A novel user interface for musical timbre design. In *Audio Engineering Society 128th Convention*. Audio Engineering Society.
[67] Jimmy Secretan, Nicholas Beato, David B D Ambrosio, Adelein Rodriguez, Adam Campbell, and Kenneth O Stanley. 2008. Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1759–1768.
[68] Reinhard Sefelin, Manfred Tscheligi, and Verena Giller. 2003. Paper prototyping-what is it good for? A comparison of paper-and computer-based low-fidelity prototyping. In *CHI'03 extended abstracts on Human factors in computing systems*. 778–779.
[69] Amir Semmo, Tobias Dürschmid, Matthias Trapp, Mandy Klingbeil, Jürgen Döllner, and Sebastian Pasewaldt. 2016. Interactive image filtering with multiple levels-of-control on mobile devices. In *SIGGRAPH ASIA 2016 Mobile Graphics and Interactive Applications*. 1–8.
[70] Yang Shi, Nan Cao, Xiaojuan Ma, Siji Chen, and Pei Liu. 2020. EmoG: supporting the sketching of emotional expressions for storyboarding. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
[71] Ben Shneiderman. 2009. Creativity Support Tools: A Grand Challenge for HCI Researchers. In *Engineering the User Interface: From Research to Practice*, Miguel Redondo, Crescencio Bravo, and Manuel Ortega (Eds.). Springer London, London, UK, 1–9. https://doi.org/10.1007/978-1-84800-136-7_1
[72] Anna Spagnolli, Diletta Mora, Matteo Fanchin, Valeria Orso, and Luciano Gamberini. 2020. Automation and Creativity: A Case Study of DJs' and VJs' Ambivalent Positions on Automated Visual Software. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
[73] Wei-Tse Sun, Ting-Hsuan Chao, Yin-Hsi Kuo, and Winston H Hsu. 2017. Photo filter recommendation by category-aware aesthetic learning. *IEEE Transactions on Multimedia* 19, 8 (2017), 1870–1880.
[74] Kıvanç Tatar, Matthieu Macret, and Philippe Pasquier. 2016. Automatic synthesizer preset generation with presetgen. *Journal of New Music Research* 45, 2 (2016), 124–144.
[75] Stuart Taylor and Jonathan and Hook. 2010. FerroSynth: A Ferromagnetic Music Interface. In *Proceeding of NIME 2010* (proceeding of nime 2010 ed.). Microsoft Research. https://www.microsoft.com/en-us/research/publication/ferrosynth-a-ferromagnetic-music-interface/
[76] Kristen Vaccaro, Sunaya Shivakumar, Ziqiao Ding, Karrie Karahalios, and Ranjitha Kumar. 2016. The elements of fashion style. In *Proceedings of the 29th annual symposium on user interface software and technology*. 777–785.
[77] Cyrus Vahidi, George Fazekas, Charalampos Saitis, and Alessandro Palladini. 2020. Timbre space representation of a subtractive synthesizer. *arXiv preprint arXiv:2009.11706* (2020).
[78] Peter-Paul Verbeek. 2015. Beyond Interaction: A short introduction to mediation theory. *Interactions* 22, 3 (2015), 26–31.
[79] Roel Vertegaal and Ernst Bonis. 1994. ISEE: an intuitive sound editing environment. *Computer Music Journal* 18, 2 (1994), 21–29.
[80] Shaun Wallace, Brendan Le, Luis A Leiva, Aman Haq, Ari Kintisch, Gabrielle Bufrem, Linda Chang, and Jeff Huang. 2020. Sketchy: Drawing inspiration from the crowd. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–27.
[81] Yunlong Wang, Priyadarshini Venkatesh, and Brian Y Lim. 2022. Interpretable Directed Diversity: Leveraging Model Explanations for Iterative Crowd Ideation. In *CHI Conference on Human Factors in Computing Systems*. 1–28.
[82] David L Wessel. 1979. Timbre space as a musical control structure. *Computer music journal* (1979), 45–52.

[83] Jarke J van Wijk and Cornelius WAM van Overveld. 2003. Preset based interaction with high dimensional parameter spaces. In *data visualization*. Springer, 391–406.
[84] Matt Wright. 2002. OpenSoundControl Specification 1.0. https://opensoundcontrol.stanford.edu/spec-1_0.html. Accessed: 2023-05-07.
[85] Matthew Yee-King and Martin Roth. 2008. Synthbot: An unsupervised software synthesizer programmer. In *ICMC*.
[86] Matthew John Yee-King, Leon Fedden, and Mark d'Inverno. 2018. Automatic programming of VST sound synthesizers using deep networks and other techniques. *IEEE Transactions on Emerging Topics in Computational Intelligence* 2, 2 (2018), 150–159.
[87] Linping Yuan, Ziqi Zhou, Jian Zhao, Yiqiu Guo, Fan Du, and Huamin Qu. 2021. Infocolorizer: Interactive recommendation of color palettes for infographics. *IEEE Transactions on Visualization and Computer Graphics* (2021).
[88] Chao Zhang, Cheng Yao, Jiayi Wu, Weijia Lin, Lijuan Liu, Ge Yan, and Fangtian Ying. 2022. StoryDrawer: A Child–AI Collaborative Drawing System to Support Children's Creative Visual Storytelling. In *CHI Conference on Human Factors in Computing Systems*. 1–15.
[89] Jamie Zigelbaum, Amon Millner, Bella Desai, and Hiroshi Ishii. 2006. BodyBeats: Whole-body, Musical Interfaces for Children. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (Montreal, Quebec, Canada) *(CHI EA '06)*. ACM, New York, NY, USA, 1595–1600. https://doi.org/10.1145/1125451.1125742

A STUDY PARTICIPANTS

Table 2 lists all the participants who were a part of this research. Because some participants participated in multiple stages of the project, the rightmost column indicates the studies in which each participant had a part. *Interview* refers to the interview study described in Section 3, *Design* refers to the participatory design process described in Section 4, and *Evaluation* refers to the user study described in Section 6.

Table 2: The details of our participants.

ID	Age	Gender	Musical Role	Years of Experience	Experience Level	Participation in Our Study
P1	38	M	Producer	10+	Advanced	Interview, Design
P2	37	M	Composer, Producer, DJ	10+	Advanced	Interview, Design
P3	27	M	Keyboardist	0-5	Beginner	Interview
P4	23	M	Composer	0-5	Beginner	Interview
P5	45	M	Producer	5-10	Beginner	Interview
P6	26	M	Producer	5-10	Intermediate	Interview
P7	32	F	Composer, Vocalist	0-5	Beginner	Interview, Design, Evaluation
P8	26	M	Producer, Sound Designer	5-10	Intermediate	Interview
P9	39	M	Composer, DJ	10+	Intermediate	Interview, Design, Evaluation
P10	26	M	Composer	5-10	Beginner	Interview
P11	37	M	Composer, DJ, Keyboardist	0-5	Beginner	Evaluation
P12	29	M	Sound Engineer	5-10	Intermediate	Evaluation
P13	38	M	Synthesist	5-10	Intermediate	Evaluation
P14	28	F	Composer	5-10	Intermediate	Evaluation
P15	39	F	Synthesist	5-10	Intermediate	Evaluation
P16	33	M	Producer	10+	Intermediate	Evaluation
P17	DNA	F	Composer	0-5	Beginner	Evaluation
P18	31	F	Composer, Synthesist, DJ	0-5	Beginner	Evaluation

B INTERVIEW STUDY SEMI-STRUCTURED INTERVIEW QUESTIONS

Table 3: Formative interview study questions

How long have you been making music?
What is your level of musical education? How did you learn to make music?
What are your primary instruments?
Can you tell me a bit about a recent musical project where you used synthesizers?
What instruments did you use?
How did you choose the timbres you used for the project?
Did you use presets? Did you program your own timbres? Tweak presets?
Were you satisfied with the sounds you used?
Did you tweak the presets? To what extent?
What was the experience with synthesizer X like?
Did you feel like you understood how synthesizer X was working?
What strategies did you use for programming synthesizer X?
Did you have particular timbres in mind that you tried to program?
Did you have any problems or difficulties while working on this project?
If you did, how did you overcome them?
Did you feel confident with the synthesizers you were using?
What are some of your most used synthesizers?
Why do you use them?

C ALGORITHM PSEUDO CODE

C.1 Evolution

Algorithm 1 Evolutionary Mode: repeated to generate each offspring timbre

Input:
$timbre[]$: vector containing the parameter values of the resulting timbre,
$timbre_{max}[]$: vector containing the maximum possible values of each parameter,
$mutation\_rate$: between 0.00 and 1.00 and set with the X knob,
$mutation\_range$: between 0 and 100 and set with the Y knob,
not_to_mutate: array of parameter indices to exclude from mutation,
$total\_params$: total number of synthesis parameters

Initialize:
$num\_to\_mutate \leftarrow total\_params * mutation\_rate$
$num\_mutated \leftarrow 0$
$matched \leftarrow false$

while $num\_mutated < num\_to\_mutate$ **do**
  $i \leftarrow$ random.uniform(0, $(total\_params - 1)$)
  if i exists in not_to_mutate **then**
    $matched \leftarrow true$ # i.e. do not mutate this parameter.
  end if

  if $matched$ is $false$ **then**
    $a \leftarrow (timbre[i] - mutation\_range)$.clip(min = 0)
    $b \leftarrow (timbre[i] + mutation\_range)$.clip(max = $timbre_{max}[i]$)
    $timbre[i] \leftarrow$ random.uniform(a, b)
    $num\_mutated \leftarrow num\_mutated + 1$
  else
    $matched \leftarrow false$
  end if
end while

C.2 Morphing

Algorithm 2 Morphing Mode: to calculate the resulting timbre for the current morph location

Input:
$timbre[]$: vector containing the parameter values of the resulting timbre,
$\{timbre_a[], timbre_b[], timbre_c[], timbre_d[]\}$: vectors containing the parameter values of the preset timbres,
p_M: current morph location,
$\{p_a, p_b, p_c, p_d\}$: locations of the four preset timbres,
$total\_params$: total number of synthesis parameters

Initialize:
$dist_{min} \leftarrow \sqrt{50^2 + 50^2} \approx 70.71$ # shortest distance between two preset timbre markers
$weights \leftarrow []$

for p_n in $\{p_a, p_b, p_c, p_d\}$ **do**
$dist_n \leftarrow Dist(p_n, p_M)$ # in screen space
$weights.append((1 - (dist_n / dist_{min})).clip(0, 1))$
end for

for $weight_n$ in $weights[]$ **do**
$weight_n \leftarrow weight_n / sum(weights)$
end for

for i in range(0, $(total\_params - 1)$) **do**
$timbre[i] \leftarrow weight_a * timbre_a[i] + weight_b * timbre_b[i] + weight_c * timbre_c[i] + weight_d * timbre_d[i]$
end for

D CONSIDERATION ASSESSMENT RATINGS

Table 4: Considerations and corresponding assessment statements. Note: questions 8 and 9 were reverse-scored.

Consideration	Statement
Musicians' desire for new sounds should not be equated with a desire to do more timbre programming.	I was able to create new sounds in a way that did not feel tedious or frustrating.
Explorations can be initiated with presets rather than by tweaking individual parameters.	Exploring sounds with the system was more satisfying than the methods I normally use.
Presets should be viewed as a starting point for timbre exploration, rather than as a solution for it.	The sounds I was able to create were pleasant and/or useful.
Timbre exploration designs should be conscious of and support the use of external effects processing.	The system integrated well with my existing effects and/or processing techniques.
The exploration process should be fast, fluid, and immediate.	The sound creation process was fluid and immediate.
A few key parameters or controls should always be readily available.	I felt like I had access to the parameters I needed.
A tactile, physical interface should be considered for the exploration process.	The physical interface was beneficial to my experience and/or sound creation.
Direct access to all parameters should not be equated with ease of exploration.	I felt overwhelmed or intimidated by the number of parameters the interface provided.
The timbre exploration process should be guided by meaningful limitations.	I felt overwhelmed or otherwise hindered by how many possibilities the system provided.

E USER STUDY SEMI-STRUCTURED INTERVIEW QUESTIONS

Table 5: Semi-structured interview questions

Do you feel like the interface allows you to make sounds more easily than before? What specifically about it is helpful?
Do you feel that this system can integrate with your existing workflow well? What about it works/doesn't?
What kind of sounds have you been able to make? Are they the types of sounds you could see yourself using?
How do you see this system working with your other equipment / effects? Would having this device change your process at all?
How did you feel about the speed of the exploration / sound creation? Could you proceed at a speed that was satisfying? What do you think specifically influenced this?
How did you feel about the controls the interface provided? Were they adequate? Were there too many? Too few?
How did you feel about the scope of the device? Did it give you enough control? Was it overwhelming? Specifically, why?