Sonification of Health & Environmental Data

SoniHED
Conference 2022

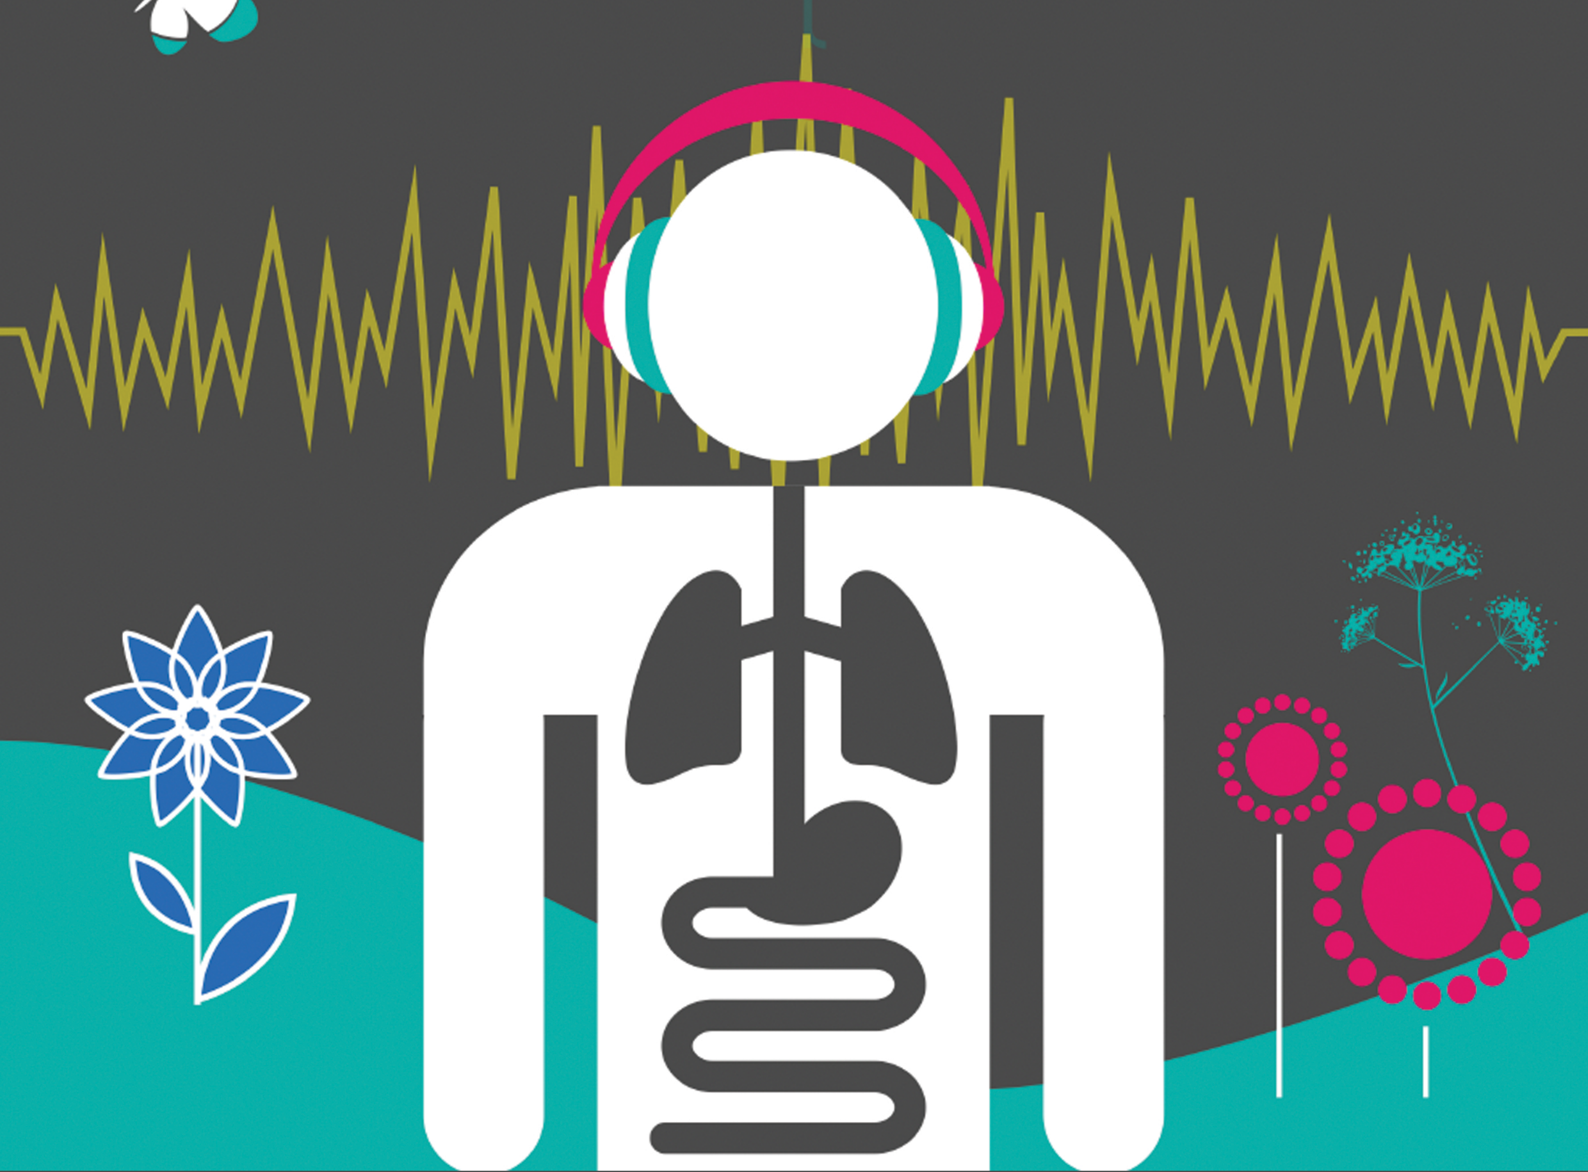

Sonification of Health and Environmental Data
27-28 October 2022, KTH, Stockholm, Sweden

Conference Proceedings

Energimyndigheten

KTH VETENSKAP OCH KONST

Sound for Energy

PaDS Participatory Sound

SoniHED 2022

Proceedings

of the

2nd Conference on the Sonification of Health and Environmental Data

27-28 October 2022
Stockholm, Sweden

Edited by Sandra Pauletto, Stefano Delle Monache, Rod Selfridge

ISBN: 978-91-8040-358-0

Table of Contents

ISBN: 978-91-8040-358-0

Introduction

Sonification, and more generally sound design and sonic interaction design, are concerned with displaying data and information in sonic form so that listeners (experts and/or non-experts) can perceive and engage with data structures, complex information and their meaning.

The domains of health and the environment are strongly related. In particular, the environment can have a direct and powerful effect on our health. In our increasingly digitalised world, we are gathering more and more data about health and the environment, both locally and personally, as well as on a wider level. Sound provides us with a complementary way to engage with this information opening up novel ways to access and reflect on these fields.

The Sound for Energy Project (https://soundforenergy.net) is delighted to host the Second Conference on Sonification of Health and Environmental Data (SoniHED 2022). The SoniHED Conference brings together experts in the fields of sonification, sound design, health and environmental science to evaluate and discuss novel sonic ways to engage with data from these fields.

This year we especially, but not exclusively, welcomed research addressing the theme: Sound and Energy. Energy is fundamental to life. Our bodies need energy to keep moving and be healthy. We can use our physical energy to replace environmentally damaging energy sources, for example, by cycling or walking to work instead of driving. We obviously consume a lot of energy at home and at work – using computers, electricity, heating, hot water and more – and often are unaware of how we can make our behaviours more efficient and sustainable. In this context, sound can provide new engaging ways to understand how we produce and use energy, helping us maintain good habits, while reducing unnecessary energy consumption.

The SoniHED Conference is a gathering of experts in sustainability, environmental science, energy, health, sonification, sound design and sonic interaction design that aims to discuss latest research in this interdisciplinary area.

The Conference is supported by the Sound for Energy Project which is funded by the Swedish Energy Agency and by KTH Royal Institute of Technology, and by the Marie Sklodowska-Curie project Participatory Sound (PaDS).

We are pleased to present a large variety of works both scientific and artistic during the conference that promises to be highly inspiring.

We hope that the work presented in these proceedings will inspire and stimulate new research in this area and more widely in the field of sustainable sound computing and design.

The SoniHED Team

Sandra Pauletto (KTH, Sweden), Rod Selfridge (Edinburgh Napier University, UK), Stefano Delle Monache (TU Delft, NL), Yann Seznec (KTH, Sweden), Vincenzo Madaghiele (KTH, Sweden)

Conference Programme

Thursday 27 October 2022

10:00-10:15 WELCOME with Sandra Pauletto

10:15-11:45 GUEST PANEL: COMPLEX SYSTEMS - Chair: Stefano Delle Monache

Linking Data to Action Cristian Bogdan (KTH)

The Sound of Mobility Nicolas Misdariis (IRCAM)

On Control, Sonic Art and Energy Sol Andersson (sound artist, Sweden)

11:45-12:15 Break

12:15-13:15 PAPER SESSION 1 - Chair: Katharina Groß-Vogt

Prithvi Ravi Kantan, *et al*. *Designing Sonified Feedback on Knee Kinematics in Hemiparetic Gait Based on Inertial Sensor Data*

Pascal Staudt, *et al. Automatic Classification of Interactive Gestures for Inter-Body Proximity Sonification*

Juan Ignacio Mendoza Garay, *et al. Musification of Accelerometry Data Towards Raising Awareness of Physical Activity*

13:15-14:00 Lunch break

14:00-15:30 GUEST PANEL: SOUND, LANDSCAPES, STORIES AND DATA - Chair: Sandra Pauletto

The Musical Landscape of the Human Body, Eva Bojner Horwitz (KMH)

Exploring Low Carbon Futures in the Past, Elina Eriksson (KTH)

Loud Numbers: Sound, Data and Stories, Duncan Geere (information designer, Sweden)

15:30-16:00 Break

16:00-17:30 Sonification and Music Composition Session - Chair: Yann Seznec

ISBN: 978-91-8040-358-0

Friday 28 October 2022

10:00-10:15 WELCOME with Sandra Pauletto

10:15-11:45 GUEST PANEL: ENVIRONMENTS - Chair: Rod Selfridge

Sound Ideas for Critical Care, Elif Özcan Vieira (TU Delft)

Control of the Indoor Temperature in Single Family Houses, Björn Palm (KTH)

Singing Windows, Robert Jarvis (sound artist, UK)

11:45-12:15 Break

12:15-13:15 PAPER SESSION 2 - Chair: Niklas Rönnberg

Vincenzo Madaghiele and Sandra Pauletto. *Experimenting techniques for sonic implicit interactions: a real time sonification of body-textile heat exchange with sound augmented fabrics*

Stephen Williams. *Can you hear the Earth breathing? Translation and disclosure in sound art data sonification*

Katharina Groß-Vogt and Iohannes Zmölnig. *Pushing sustainability topics in sound teaching*

13:15-14:00 Lunch break

14:00-15:00 PAPER SESSION 3 - Chair: Emma Frid

Yichun Zhao and George Tzanetakis. *Interactive Sonification for Health and Energy using ChucK and Unity*

Maria Mannone and Veronica Distefano. *Sonification and Clustering: a Multi-Sensory Perspective for Precision Medicine*

Brian Lindgren. *Composing 'Eaten Alive': Exploring the Intersection of Climate Change and Global Fish Production Through Sonification*

15:00-16:00 Round table, Annoucements and Conclusion - Chair: Sandra Pauletto

ISBN: 978-91-8040-358-0

Guest Talks and Biographies

Nicolas Misdariis, IRCAM, Paris

The sound of mobility

The subject of my talk will fall within the framework of the relationship between sound and mobility, by questioning the way in which sound accompanies human experiences of mobility, but also and especially, the way in which mobility objects – or artifacts – contribute to shape human sound environments. By considering the concept of mobility and observing the (old and new) forms it takes, we can make two postulates, among many others: firstly, the mobility artifact is inherently a source of sound, either intentional or consequential; secondly, mobility artifacts shape the environment and therefore have something to do with sound design and ecology. By bringing together these two assumptions, and based on different research studies – especially about electric vehicles –, we will then address, and try to support, the following issue: how can sound design improve the sound of mobility while taking into account environmental, ecological and sustainable dimensions, that are as essential as they are vital?

Biography

Nicolas Misdariis is an IRCAM Research Director, Head the Sound Perception & Design group/STMS Lab, and Deputy-Dead of the IRCAM's STMS Lab.Nicolas has a background on applied acoustics and synthesis/reproduction/perception of musical and environmental sounds. Working at IRCAM since 1995, he has developed many research works and industrial applications related to sound synthesis and reproduction, environmental sound and soundscape perception, auditory display, human-machine interfaces (HMI), interactive sonification and sound design. Since 2010, Nicolas is a regular lecturer in the Sound Design Master at the High School of Art and Design in Le Mans (ESAD TALM, Le Mans).

Elif Özcan Vieira, TU Delft

Sound ideas for critical care: How to leverage environmental intelligence for sound-driven digital solutions

Critical care is a data-driven professional environment. For the safety and survival of the patient and consequently the success of the treatment protocols not only patient vitals are observed but also the quality of the care environment. Within this context, collected 'healthcare' data offer ample design opportunities for sonification but also for visualization of sonic data. In this talk, I will showcase a couple of data- and sound-driven design examples that helped us, researchers at the Critical Alarms Lab of TU Delft, provide solutions for patient and clinician needs that concern sound. I will also propose a new way of monitoring the sonic environment and actively engage with the audience to help me deconstruct the issues that are inherent to sound observations such as automatic categorization of sounds, tension between speech privacy and data collection, sensing technology, and product ideas for easy adoption. Look forward to a session chockful of sound ideas!

ISBN: 978-91-8040-358-0

Biography
Elif Özcan conducts sound-driven design and research activities at the Faculty of Industrial Design Engineering of TU Delft and leads the Silent ICU project at the Department of Intensive Care of Erasmus Medical Centre. She is the director of Critical Alarms Lab at TU Delft where she mainly investigates sound-driven wellbeing of users and professionals in critical contexts (e.g., patients and clinicians). CAL aims to shape the future of product-user interactions in complex environments through audible and visual information design. The lab is a flexible consortium of individuals, institutes, and companies, and it offers multiple opportunities for student participation.

Eva Bojner Horwitz, KMH Royal College of Music, Stockholm

The musical landscape of the human body
Would it be possible to explore the musical landscape of the human body in the same way that the body explores musical soundscapes? Inside our bodies are internal structures, communication, and signals about the current state of health that could be investigated using various qualitative and quantitative methodologies. Different examples will be discussed.

Biography
Eva Bojner Horwitz, Professor of music and health at the Royal College of Music in Stockholm and researcher at the Department of Clinical Neuroscience Karolinska Institute (KI). She is associate professor and doctor in social medicine, cultural health researcher, specialized in psychosomatic medicine and in creative arts; co-founder of the Center for Social Sustainability (CSS), KI. She is anchored in interdisciplinary research, has doctoral students, authored scientific articles and books. She is known internationally for her implementation of music activities in health care and school systems and for her evaluation with video interpretation technique (VIT), combining quantitative (stress hormone analyses, heart rate variability) with qualitative research (micro phenomenology and phenomenological hermeneutics). Research focus: music and health (sing health in schools, performance evaluations, HeArtS: Health – Arts – Sustainability platform building); music and learning (knowledge concerts); music in end-of-life situations; music and social sustainability (inner transformation and creativity); arts and humanities; nature related soundscapes, resilience, aesthetics and flow.

Elina Eriksson, KTH Royal Institute of Technology, Stockholm

Exploring low carbon futures in the past
Elina Eriksson will present current research from the ongoing project “Beyond the event horizon: tools to explore local energy transformations”, in which the theoretical concepts of pastcasting and recasting are put into practice. The goal of the project is to develop a scenario-workshop methodology and to use it to carry out workshops that explore and support local energy transitions by helping local actors focus on, process, democratically anchor and expand "what is perceived to be possible". The methodology has

been cocreated together with the project partners from the Transition Network Sweden and been tested in both online and physical workshops. In discussion with the audience and SoniHED guests, we hope to explore what role sound could play in these novel methods for sustainability research.

Biography
Elina Eriksson is an Associate Professor in Human-Computer Interaction with a specialization in Sustainability at KTH Royal Institute of Technology. Her research is action-oriented and strives to support energy transitions. She also conducts research on how to integrate sustainability in computing education.

Björn Palm, KTH Royal Institute of Technology, Stockholm

Control of the indoor temperature in single family houses
In a single family house in Sweden the by far highest energy consumption is related to heating of the building. The indoor temperature makes a big difference for the overall consumption: Reducing the indoor temp by 1 deg C reduces the amount of energy consumed by about 5%. In Sweden it is customary to keep the temperature constant over day and night, even if no-one is home. By allowing the temperature to fluctuate, cooler at night, cooler when the house is empty, it could be possible to reduce the energy consumption. We have tested using thermostats which can be programmed beforehand, with some success. Some of the possibilities and potential problems with this way of controlling the indoor climate will be discussed in this brief presentation. At SoniHED, we hope to explore what role sound could play in this scenario.

Biography
Björn Palm is Professor in Energy Technology at the Department of Energy Technology, KTH. His research is related to energy use in buildings, and in particular to heat pumps, refrigeration machines, heat exchangers, natural refrigerants and the role of heat pumps in the energy system.

Cristian Bogdan, KTH Royal Institute of Technology Stockholm

Linking Data to Action: Designing for Amateur Energy Management
Design of eco-feedback has primarily aimed at persuading individuals to change behaviours into more environmentally sustainable ones. However, it has been questioned how effective such feedback is in supporting long-term change. As an alternative focus for energy feedback, we present a case study of amateur energy management work in apartment buildings owned by housing cooperatives, and the design of an app that aims to stimulate and support cooperatives in taking energy actions that significantly reduce the cooperative's collective energy use. By linking energy data to energy actions, the users can see how actions taken in their own and other cooperatives affected the energy use, learn from each other's experiences and become motivated as energy amateurs. Based on our

housing cooperative case, we reflect on design aspects to consider when designing for energy management in amateur settings. I look forward to discuss with the SoniHED researchers what role sound could have in such a scenario.

Biography

Cristian Bogdan is an Associate Professor at the Department of Media Technology and Interaction Design. In the area of IT support for energy systems, Cristian is interested in amateur energy activists' conceptual models of energy systems, from small, short-term energy use systems like electric vehicles, to large, longer term systems like buildings and housing associations. Cristian is especially interested in the ability to take action to improve the energy situation, beyond energy awareness.

Sol Andersson, sound artist, Stockholm

On control, sonic art and energy

The research draws upon a number of practices, such as sound art, guerilla art, music, and performance. The aim of the research is to explore aspects of freedom within the context of sound art, that might lead to a notion of freedom of objects. That is, objects free from certain biases, free from traditions. Therefore, making portable self-resonating sound objects driven by solar panels and batteries is a way to explore aspects of control.

Biography

Sol Andersson is a musician, sound artist and currently undertaking practice-based research in Sound Art at the department of Video, imaging and Sonic Arts, De Montfort University, Leicester, UK. The last years work includes performances, sound installations, happenings, commissioned pieces and record releases.

Duncan Geere, information designer, Helsingborg

Loud Numbers: Sound, Data and Stories

Everyone has a favourite song. It might be a pop song, it might be a classical piece, it might be rock’n’roll, metal or techno. But for some reason it’s our favourite - and that reason is almost always linked to the emotional connection that we have with it. In my SoniHED talk, I’ll explore the strong connection between sound, music and emotion, focusing on how sonification can be used to reach new audiences with data stories. Using examples from my own work and the work of others, I’ll discuss how this type of data storytelling with sound can be effective, the kinds of stories it’s best suited to, and why it’s important to use this powerful approach with care.

Biography

Duncan Geere is an information designer based in Helsingborg, Sweden. He’s the cofounder of the Loud Numbers data sonification podcast, and the Elevate Dataviz

ISBN: 978-91-8040-358-0

Learning Community. He's created sonifications for clients including The Museum of London, Walmart, and Radboud University.

Robert Jarvis, sound artist, UK

Singing windows

The surprising effect that a simple sonic intervention had across towns in the UK. Shop windows were converted into loudspeakers, in order to relay a multi-channel sound piece throughout a series of towns in the south east of England (Kent). The sonic intervention made use of manipulated recordings of birdsong, with the birds appearing to react and call out to each other. Hearing the sounds, passing pedestrians had their curiosity raised, and therefore related to the area in a different way.

Biography

Robert Jarvis' work as a sound artist lies somewhere between that of a composer and a creative researcher. He has composed musical works drawing from air quality data, plant genetics, bat echolocation, and insect recordings. His installations have involved multi-speaker soundscapes, interactive games, real-time astronomical sonifications, and covid19 transmission data. His work is concerned with encouraging people to rethink their environments and for them to question how they relate to their surroundings. His aim is to open up new worlds for those that come into contact with his work, posing new questions and enticing new appreciations of the sonic landscape. He also plays trombone.

SoniHED 2022 Conference Articles

Conference Scientific Committee

Avanzini Federico, University of Milan, Italy
Barahona-Ríos Adrián, University of York, UK
Delle Monache Stefano, Delft University of Technology, Netherlands
Di Donato Balandino, Edinburgh Napier University, UK
Eriksson Elina, KTH Royal Institute of Technology, Sweden
Fontana Federico, University of Udine, Italy
Groß-Vogt Katharina University of Music and Performing Arts, Graz, Austria
Holzapfel André, KTH Royal Institute of Technology, Sweden
McGowan John, Edinburgh Napier University, UK
McGregor Iain, Edinburgh Napier University, UK
Papetti Stefano, Zurich University of the Arts / ICST, Switzerland
Pauletto Sandra, KTH Royal Institute of Technology, Sweden
Selfridge Rod, Edinburgh Napier University, UK
Zmöling Johannes University of Music and Performing Arts, Graz, Austria

Conference Panels' Chairs

Delle Monache Stefano (TU Delft, Netherlands)
Frid Emma (KTH/Ircam, Sweden/France)
Groß-Vogt Katharina (University of Music and Performing Arts, Graz, Austria)
Pauletto Sandra (KTH, Sweden)
Rönnberg Niklas (Linköping University, Sweden)
Selfridge Rod (Edinburgh Napier University, UK)
Seznec Yann (KTH, Sweden)

ISBN: 978-91-8040-358-0

Designing Sonified Feedback on Knee Kinematics in Hemiparetic Gait Based on Inertial Sensor Data

Prithvi Ravi Kantan
Aalborg University, Denmark
prka@create.aau.dk

Sofia Dahl
Aalborg University, Denmark
sof@create.aau.dk

Helle Rovsing Møller Jørgensen
North Denmark Regional Hospital, Denmark
hrmj@rn.dk

Chetali Khadye
University of Dundee, Scotland
ckhadye001@dundee.ac.uk

Erika G. Spaich
Aalborg University, Denmark
espaich@hst.aau.dk

ABSTRACT

In recent years, interactive sonification based on data from wearable sensors has been explored as a feedback tool in movement rehabilitation. However, it is yet to be routinely adopted as part of physiotherapy protocols, partly due to challenges with designing solutions tailored to diverse patients. In this work, we propose a set of adaptable feedback paradigms on knee kinematics for hemiparetic stroke patients undergoing gait training. We first collected inertial data and video footage from 15 hemiparetic patients during overground walking. The video footage was then analyzed by a physiotherapist, who identified three main knee-related movement impairments - reduced range of motion, dysregulated extension, and hyperextension. Using a custom-built software architecture, we devised two music-based paradigms for providing tailored concurrent feedback on knee movement and the impairments identified by the physiotherapist based on inertial data. The paradigms will be clinically tested with patients as part of a future study, and we believe that their impairment-specificity and individual adjustability will make them an advancement of existing auditory feedback designs.

1. INTRODUCTION

The auditory medium has been increasingly explored as a means of providing interactive feedback to augment motor learning in rehabilitation [1] and sport [2]. Specifically, interactive sonification [3] has been applied as biomechanical biofeedback with the goal of enhancing self awareness by providing objective and accurate information about one's movements [4–6]. Sound has clear potential as a feedback route due to the excellent temporal resolution of the auditory system compared to vision [7], as well as the potential of sound to free up the visual route for other tasks [1, 8]. However, there are no established conventions or frameworks for how sonification should be designed and implemented in various motor learning scenarios, which likely limits the realization of this potential [1, 9]. In this work, we approach the design of sonification to aid motor learning in a specific but widespread clinical context - gait training in hemiparetic patients with a focus on knee-related impairments.

1.1 Post-Stroke Gait Disturbances

Walking is a crucial activity of daily life and is associated with longevity in older adults [10]. It is a highly complex movement whose kinematic properties are regulated by visual, auditory, tactile, and proprioceptive feedback [11, 12]. Walking patterns in hemiparetic stroke patients differ from normal walkers, particularly in that they are asymmetric in terms of spatiotemporal, kinematic, and kinetic parameters [13]. This is particularly seen in their *joint kinematics*, where stroke patients also exhibit great inter-individual variability [13].

The knee joint is an important contributor to the act of walking, and it has been shown that knee muscle strength on the most affected side is significantly correlated with lower gait speed and lower performance in multiple clinical gait tests in mild to moderate hemiparesis [14]. The muscles of the lower limb can also develop abnormal muscle firing patterns post stroke [15]. This results in the formation of muscle synergy patterns which disrupt the rhythmic co-ordination of movement as observed during a normal gait cycle [16]. Stroke also causes an alteration in the afferent sensory input from the muscle to the brain, resulting in improper muscle activation and less coordination during gait [15]. Common knee kinematic problems post-stroke during gait are reduced peak flexion (bending) when the foot is in the air and poorly regulated or excessive extension (straightening) as the foot makes contact with the ground [13, 17, 18]. These abnormal movements at the knee joint are compensated for by abnormal motion at the ankle and hip joints on both sides [19, 20]. Therefore, we propose that addressing knee issues can contribute to improving overall lower limb joint co-ordination in hemiparetic gait.

1.2 Enhancing Multisensory Integration Through Sonification

Given that the motor execution of gait relies on the integration of feedback from multiple sensory channels [11, 12], movement sonification has clear potential to augment reha-

ISBN: 978-91-8040-358-0

bilitation. Specifically, the auditory information it provides can be exploited by multisensory integrative brain areas and the transmodal mirror neuron system, which may lead to a more accurate representation of the movement [21,22] and benefit motor learning [23]. It is known that sound can impact bodily awareness, body movement, and body representations [24], and recent research supports the notion that when suitably designed, sonified kinematic information can have a significant impact on motor perception and estimation [22, 25, 26].

1.3 Designing Feedback for Motor Learning

Feedback design (irrespective of modality) has been found to be a key determinant of motor learning, specifically the perceptual information available to the learner while performing a task [9, 22, 23, 25–27]. Sigrist et al. [25] recommend a design approach wherein the feedback guides the learner toward desired movement patterns without diminishing the influence of intrinsic kinesthetic information (i.e. proprioception). Through the feedback, the linkage of key features of the movement to the existing kinesthetic information should be emphasized [23] - an example is feedback triggered only if a computed movement error metric exceeds a threshold [2]. Dyer et al. [9] phrase this as feedback that "draws attention to the aspects of movement that should be corrected, and can be perceived whilst maintaining attention on all intrinsic feedback sources". Such an approach can also economize cognitive workload, which is a key consideration in feedback design, particularly when cognitively impaired users are involved [4].

The *congruency* between perceptual streams is important in multisensory integration processes as well [22]. Specifically, congruent intermodal stimuli are more readily integrated, lead to more accurate internal representations of a movement, and enhance perceptuomotor processing [22, 26]. In the context of sonified feedback, it is clear that aside from temporal proximity [28], perceived congruency is tied to the transparency of the semantic *metaphor* [29, 30] that the action-sound mapping rationale is based on. For instance, we found in a previous study that an ascending melody was largely perceived by users to be congruent with the act of rising from a chair [31]. It has been shown that incongruency can increase motor performance errors [26], worsen judgments of small velocity changes [22], and adversely affect emotional experiences [32].

The goal of the present work was to apply these feedback design principles to the process of devising paradigms for knee kinematic feedback targeting hemiparetic gait. We first recorded motion data and video data from a sample of 15 hemiparetic patients. The videos were analyzed by a physiotherapist and prevalent knee impairments during walking in each patient were documented. We then identified characteristic graphical patterns corresponding to each of these impairments in the motion data and developed patient-adjustable sonification paradigms to provide congruent music-based feedback on them. The data collection, analysis, and sonification paradigms are detailed, demonstrated, and discussed in subsequent sections.

2. MOTION DATA COLLECTION

We obtained a gait dataset from hemiparetic patients during overground walking for kinematic analysis and feedback design.

2.1 Participants

A convenience sample of 15 hemiparetic stroke patients (9 men, 6 women) aged 65.53 ± 16.48 years volunteered to participate in the motion data collection. All of them had predominantly one-sided weakness of varied severity due to ischemic/hemorrhagic strokes, intracranial hemorrhages, or traumatic brain injury, and were admitted to Neuroenhed Nord, North Denmark Regional Hospital, Denmark. Each of them was briefed about the purpose and length of their participation beforehand, and informed that their data would be anonymized, and they could withdraw at any time. All procedures conformed to the ethics code of the Declaration of Helsinki. Informed consent was obtained prior to participation, and no sensitive or confidential information was collected from the participants.

2.2 Setup

The data collection was carried out at North Denmark regional hospital at the Neuroenhed Nord department Frederikshavn and Brønderslev. An approximately 10 m long stretch was designated as a walking track in a long corridor with no obstacles. For wireless inertial measurement, we used M5Stack Grey devices equipped with MPU 9250 9-axis inertial sensors [1]. These were fastened to the trunk (lumbar spine), thighs (lateral placement, just above knee), and shanks (lateral placement, just above ankle) of each patient using elastic straps bearing a silicone housing for the sensor. The sensors transmitted the data over a dedicated 2.4 GHz WiFi network (TP Link Archer C20 router) in the form of *Open Sound Control (OSC)* packets using the UDP protocol. A Dell Inspiron 15 7000 laptop was used to stream and log the inertial data using custom software. A data sampling rate of 100 Hz was used throughout. Additionally, a Motorola G8 Power mobile handset was used to capture video recordings of the patients' gait at a resolution of 720×1280 (portrait) and a frame rate of 30 Hz.

2.3 Procedure

The data collection was conducted on each patient in collaboration with a highly experienced physiotherapist. After obtaining informed consent and providing a short briefing, we calibrated the sensors for bias compensation purposes and mounted them on the patient's body. We then directed the patient to walk through the designated corridor segment at a comfortable pace using any walking aids they were accustomed to (physiotherapist support, rollator, training bench along one side). Inertial data from all five sensors were recorded during the entirety of this phase. Video footage was also captured at such an angle

[1] `https://shop.m5stack.com/products/grey-development-core`

that the head of the patient was omitted. The physiotherapist carefully supervised the walking activity and decided how many walking laps would be feasible depending on the abilities and endurance of the patient (minimum one lap). After completion, the sensors were dismounted from the patient's body, and all data were compressed and uploaded to a secure cloud server.

2.4 Data Analysis

To obtain a professional assessment of the gait impairments exhibited by each patient, the video recordings were sent for analysis to a physiotherapist not present at the data collection. As for the inertial data, we used an upgraded version of the technical framework in [34] to import the logs and reconstruct the gait kinematics of the patients. To study knee impairments, the following movement features were computed and subsequently logged:

- **Knee Angle:** The accelerometer and gyroscope readings were fused using the Madgwick gradient descent algorithm [35] to compute the inclinations of the thighs and shanks in the sagittal plane. The learning rate coefficient β was empirically set at 0.18 to balance the trade-off between excessive integration drift and overcorrection. The knee angle was then computed as the difference between shank and thigh inclination.
- **Knee Angular Velocity:** The bias-compensated gyroscope readings for each thigh and shank about the mediolateral axis were first smoothed using 2nd order Butterworth low pass filters (5 Hz cutoff). The angular velocity of each knee joint was then calculated as the difference between the shank and thigh reading.

The logged knee angles were then plotted in MATLAB 2018b to study whether the kinematic impairments reported by the physiotherapist for each patient were detectable in the plots (see Fig. 1).

3. RESULTS OF DATA ANALYSIS

On average, the patients traversed the path from end to end 4.13 times, and walked for a mean duration of 82.67 sec.

As summarized in Table 1, the physiotherapist identified three main knee-related impairments - *reduced range of motion* (11 patients), *hyperextension* (3 patients), and *dysregulated extension* (10 patients). We explain these below; note that the *stance phase* is when the foot is in contact with the ground, and *swing phase* is when the foot is in the air:

- **Reduced Range of Motion (RoM):** Reduced peak flexion was observed during swing and/or incomplete knee extension during stance, which led to the RoM being lower than in normal walkers [18].
- **Hyperextension During Stance:** Full straightening ($0°$) or overstraightening ($< 0°$) of the joint was observed during ground contact.
- **Dysregulated Joint Extension:** At the end of the swing phase (after peak knee flexion), some patients exhibited irregular extension of the knee joint [18], which was visible as a jerky or rapid straightening phase prior to ground contact.

The findings of our graphical analysis in context with the observations of the physiotherapist are shown in Table 1. We were able to detect the observed instances of reduced RoM and hyperextension in a reliable manner in the graphs, although this was true for 70% of the cases for dysregulated extension.

4. SONIFICATION DESIGN

4.1 Design Philosophy

Given the level of inter-individual variability prevalent in post-stroke gait [13], a foremost requirement was for the feedback to be *easily adjustable to suit the clinical profile of each patient*. This is in line with Stanton et al.'s suggestion [5] of a 'biofeedback toolbox' with options to tailor feedback parameters including feedback type, target, and method. We also decided to use *music as either the primary feedback medium or the underlying substrate* for manipulation. This was to leverage the known ability of music to motivate, monitor, and modify movement by mediating perception and action [36, 37] and induce feelings of self-efficacy [9]. As such, we devised two paradigms:

1. Continuous sonification of knee angle trajectories to generate a movement-congruent auditory representation of the joint movement. Salient (but not harsh) negative reinforcement is added to this in an on-off format to highlight dysregulated joint extension.
2. Intermittent feedback on hyperextension applied as an audio effect to user-selected music tracks.

The design rationale and technical implementation of each paradigm are detailed next.

4.2 Technical Setup

To design, develop, and adjust the sonification topologies for each feedback paradigm, we used the same technical framework [34] as mentioned earlier. This allowed us to stream the recorded raw inertial data from patients, compute and visualize the movement features, and transform them into intermediate mapping variables through a series of standard operations [38]. The mapping variables were transmitted in real-time as OSC messages to REAPER v6.23 [2], where they were mapped to relevant parameter controls of selected VST [3] synthesizers or effects for sound generation. An M-Audio M-Track Solo [4] device was used for audio I/O at a sampling rate of 48 kHz with a 256 sample software buffer. We chose this distributed software architecture to maximize freedom in terms of real-time mapping choices. Based on an assessment conducted on a system

[2] https://www.reaper.fm/
[3] https://www.steinberg.net/technology/
[4] https://www.m-audio.com/m-track-solo

Impairment	# Total Patients	# Detectable in Graphs	Detection %	Representative Characteristics in Graphs
Reduced RoM	11	11	100	Reduced peak-to-peak knee angle amplitude compared to normal walkers
Hyperextension	3	3	100	Periods with knee angle close to (or saturated at) 0 degrees
Dysregulated Extension	10	7	70	Sharp slope of descent from peak flexion (crest) Fluctuations following peak extension (trough)

Table 1. A summary of the knee impairments observed by the physiotherapist, followed by an overview of their detectability in visual graphs as well as graph characteristics that we judged to represent them.

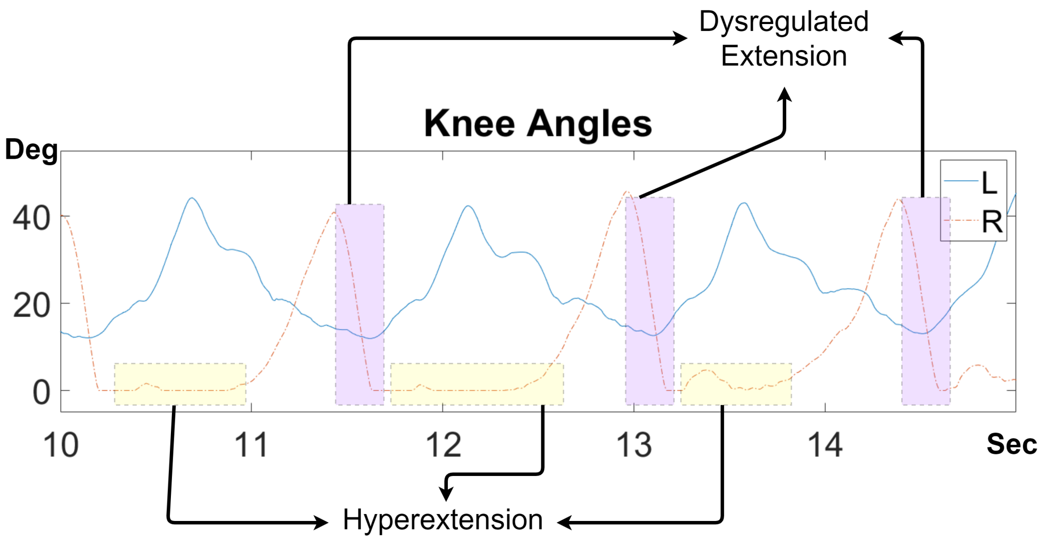

Figure 1. An illustration of the observed knee impairments summarized in Table 1. The dotted trace represents the paretic side, which exhibits both dysregulated extension and hyperextension. Both sides exhibit a reduced range of motion (max 40° compared to the normal range of about 65° [33].

with identical technical architecture, we estimate the mean round-trip latency to be 93 ms [39]. The technical framework is also capable of streaming inertial data from the sensors in real-time, making the setup viable for experimentation in clinical settings.

4.3 Paradigm 1: Knee RoM + Extension Feedback

The goal of this paradigm is to encourage knee patterns that exhibit (A) increased peak flexion during swing (compared to pre-training), (B) Left-Right symmetry, and (C) a regulated extension phase prior to foot contact.

Conceptually, the knee movement of each leg manipulates a single synthesis parameter of an ambient music piece, resulting in a continuous feedback signal whose spectral properties are directly correlated with the combination of the original knee angle trajectories. This is done in such a way that the timbre becomes brighter and richer as the flexion angle increases. This sound additionally serves as the substrate for providing feedback on dysregulated knee extension, wherein the otherwise continuous audio signal is subjected to interruptions if the knee is extended too fast (as deemed by the therapist). The block diagram of the paradigm is shown in Fig. 2.

In practice, the paradigm requires the therapist to first configure target ranges for knee flexion on either side as well as thresholds for acceptable knee extension angular velocity (see Fig. 2). As the patient walks, the mapping framework normalizes the measured knee angles within the target ranges of motion, and checks whether the angular velocity thresholds are exceeded. The normalized knee angles are sent to REAPER, where a major chord MIDI loop is synthesized using the PerFormant virtual instrument (VSTi) [5]. This instrument works by passing a set of sawtooth waves through a parallel bank of three 8^{th} order bandpass filters. The normalized knee angles are mapped to the center frequency of one of them, such that it varies from approximately 600 Hz - 4 kHz. Hence, the output sound becomes richer in high frequency spectral content as the knee angle increases, which is intended to generate an auditory reward for attaining the target flexion angle. The density of harmonics resulting from the chordal input coupled with the resonant properties of the filter ensure that the auditory changes are perceptually salient.

If the knee angular velocities during extension exceed the configured threshold, the overall audio output is attenuated to silence, leading to an audible interruption in the signal. As the phases of dysregulated knee extension tend to be rapid (see Fig. 1), the interruptions sound like brief but

[5] https://www.kvraudio.com/product/performant-by-elena-design

ISBN: 978-91-8040-358-0

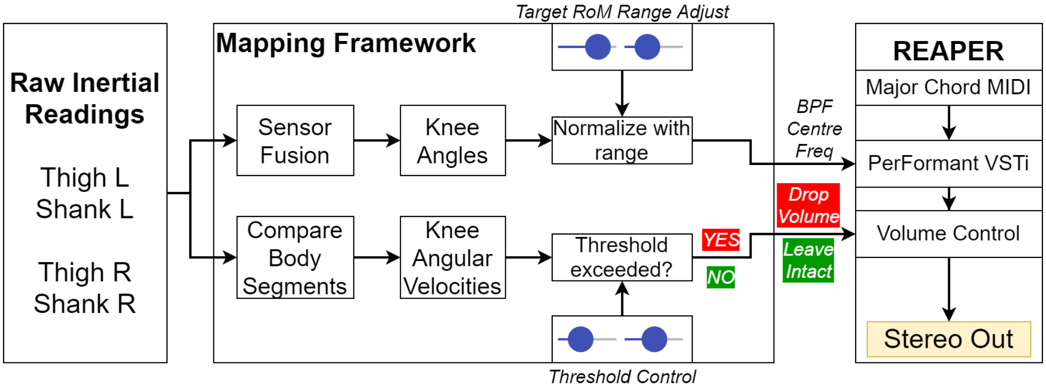

Figure 2. Block diagram of the first mapping paradigm, i.e. for reduced knee RoM and dysregulated extension feedback.

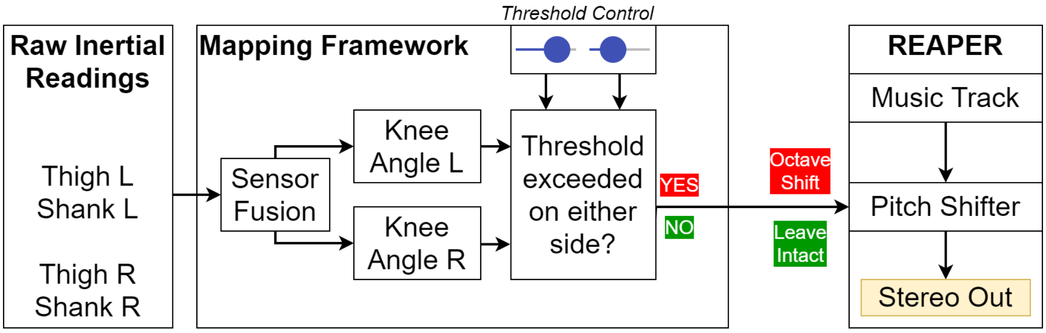

Figure 3. Block diagram of the second mapping paradigm, i.e. for knee hyperextension feedback.

clear audio dropouts. We applied the paradigm to inertial data recordings from several patients and synchronized them with the video footage, which yielded Videos 1 and 2 in the attached demo folder [6].

4.4 Paradigm 2: Hyperextension Feedback

The goal of the second paradigm is to provide clear on-off feedback that encourages the patient to avoid hyperextension by keeping the most affected knee slightly flexed during gait. The reason for the on-off choice is because it has previously been shown to be effective at reducing hyperextension [17], and would both be perceptually salient and straightforward to perceive and comprehend.

It is conceptually similar to [40] who successfully provided similar feedback but based on electrogoniometric readings as reviewed in [17]. Specifically, we designed it such that the substrate for feedback can be patient-selected music with vocals, and that the feedback is provided by applying a chipmunk-like effect to the music if (and for the entire duration that) the patient hyperextends their knee. The block diagram of the paradigm is shown in Fig. 3.

In practice, the therapist must first configure thresholds for hyperextension on either side (although the focus is typically on the most affected side). The patient-selected music track is played back in REAPER while the patient walks, and the mapping framework continuously checks whether the knee angles cross the respective hyperextension thresholds. The output is an on-off variable that is mapped to the octave shift control of a pitch shifter effect applied to the music track (REAPER's inbuilt ReaPitch in our case). If the threshold is exceeded, the music track is pitch-shifted up by an octave, or else it is left intact. Video 3 in the demo video folder illustrates this paradigm.

[6] Video Demos: https://drive.google.com/drive/folders/1eU-3LioaiszFDl8MGSPvbduwaxKB9ixh?usp=sharing

5. DISCUSSION

Based on inertial data and video footage from 15 patients, we developed two separate paradigms to provide real-time music-based feedback on knee kinematics in hemiparetic gait. We designed the feedback to be movement-congruent and based on transparent semantic metaphors.

Our primary design requirement for the paradigms was that they should be adjustable to suit the clinical profile of each patient. We believe that we were able to meet this through both the conceptual design and the technical implementation. Having a therapist configure the parameters of the feedback (and thereby the goals of the training session) fits well with proposed approaches in past literature [5]. As we implemented normalized feedback variables rather than absolute value mappings, it is possible to provide tailored, yet consistent feedback at a variety of impairment levels. Moreover, our mapping framework allows key parameters (i.e. target RoM, and angular velocity and hyperextension thresholds) to be adjusted on the fly, serving as a proof of concept for a future streamlined version. We also managed to integrate musical structures as core components of the feedback (our second requirement). With our distributed software architecture and the flexibility of modern music software at our disposal, it is possible for us to experiment with and iteratively hone the feedback paradigms [37] in a loudspeaker-based training paradigm to unlock the full potential of musical feedback in movement rehabilitation [36]. It would be conducive to widespread adoption if the final system can be made smartphone-based and utilize the same off-the-shelf sensors that we used here.

We argue that our paradigms align well with suggested design practices to enhance motor learning [9, 25, 26, 36]. The first paradigm features a direct continuous mapping between knee angle and audio spectral properties, which should readily allow the feedback to be integrated with knee proprioceptive information and help the patient form a more robust internal representation of their movements [22]. We aim to investigate the potential motor learning benefits of this in a future study. The impairment-specific feedback in both our paradigms is not provided as a continuous error signal but as an interruption or distortion to the underlying sound. This form of feedback, too, has been postulated as being potentially beneficial in the process of strengthening sensorimotor linkages [25]. Indeed, patients suffering from hyperextension were shown to benefit considerably from a threshold-based on-off feedback concept similar to our second paradigm [40], suggesting that our design can confer similar benefits but in a more pleasant and motivating format due to the use of a musical substrate.

We have not yet tested our paradigms with patients (part of future work), and therefore cannot comment on their perceived meaningfulness, pleasantness, perceptual salience, or intrusiveness, and further in the timeline, on their clinical effects. From the perspective of the patient's needs, it may not make sense to only provide feedback on hyperextension *or* dysregulated extension, so combining the paradigms may be necessary - but it is unclear whether the resulting cognitive load will be manageable for patients. It may also be necessary to design alternative sonic textures to cater to the needs of patients suffering from hearing loss in the frequency range where most of the information is conveyed by the described paradigms [41]. We aim to address these questions as part of a larger future study involving users (patients and therapists), and expect our implementation of the paradigms to function as intended in the clinical environment as they were designed based on data from patients. Lastly, we will also explore the design of feedback paradigms to address impairments in hip and ankle kinematics, as those are also relevant in restoring normal gait patterns [15].

6. CONCLUSIONS

Through this work, we developed music-based paradigms for sonified feedback on knee kinematics in hemiparetic gait based on inertial data collected from real patients. We demonstrated the sonification designs using recorded data from a group of patients with varied impairment levels, showing the feasibility of applying the paradigms as an informative feedback tool in gait rehabilitation. Future work includes testing the perceived meaningfulness, pleasantness, perceptual salience, and intrusiveness of the paradigms, as well as their clinical effects along with developing an integrated feedback system that therapists can use with ease. Overall, we believe that the impairment-specificity and individual adjustability of our paradigms can advance existing auditory feedback designs for hemiparetic gait rehabilitation.

Acknowledgments

We would like to thank the patients at Neuroenhed Nord (Brønderslev and Frederikshavn) who graciously agreed to take part in the study. This work is partially funded by by NordForsk's Nordic University Hub, Nordic Sound and Music Computing Network NordicSMC, project number 86892.

7. REFERENCES

[1] J. Guerra, L. Smith, D. Vicinanza, B. Stubbs, N. Veronese, and G. Williams, "The Use of Sonification for Physiotherapy in Human Movement Tasks: A Scoping Review," *Science & Sports*, vol. 35, no. 3, pp. 119–129, 2020.

[2] V. van Rheden, T. Grah, and A. Meschtscherjakov, "Sonification approaches in sports in the past decade: a literature review," in *Proceedings of the 15th International Conference on Audio Mostly*, 2020, pp. 199–205.

[3] A. Hunt and T. Hermann, "Interactive sonification," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Logos Verlag Berlin, 2011.

[4] A. Kos and A. Umek, *Biomechanical Biofeedback*. Springer, 09 2018.

[5] R. Stanton, L. Ada, C. M. Dean, and E. Preston, "Biofeedback improves performance in lower limb activities more than usual therapy in people following stroke: a systematic review," *Journal of physiotherapy*, vol. 63, no. 1, pp. 11–16, 2017.

[6] T. H. Nown, P. Upadhyay, A. Kerr, I. Andonovic, C. Tachtatzis, and M. A. Grealy, "A mapping review of real-time movement sonification systems for movement rehabilitation," *IEEE Reviews in Biomedical Engineering*, 2022.

[7] J. A. Grahn, "See what i hear? beat perception in auditory and visual rhythms," *Experimental brain research*, vol. 220, no. 1, pp. 51–61, 2012.

[8] E. Jovanov, K. Wegner, V. Radivojevic, D. Starcevic, M. S. Quinn, and D. B. Karron, "Tactical audio and acoustic rendering in biomedical applications," *IEEE Transactions on Information Technology in Biomedicine*, vol. 3, no. 2, pp. 109–118, 1999.

[9] J. F Dyer, P. Stapleton, and M. WM Rodger, "Sonification as Concurrent Augmented Feedback for Motor Skill Learning and the Importance of Mapping Design," *The Open Psychology Journal*, vol. 8, no. 1, 2015.

[10] S. Studenski, S. Perera, K. Patel, C. Rosano, K. Faulkner, M. Inzitari, J. Brach, J. Chandler, P. Cawthon, E. B. Connor *et al.*, "Gait speed and survival in older adults," *Jama*, vol. 305, no. 1, pp. 50–58, 2011.

[11] M. Alfuth and D. Rosenbaum, "Effects of changes in plantar sensory feedback on human gait characteristics: a systematic review," *Footwear Science*, vol. 4, no. 1, pp. 1–22, 2012.

[12] T. Lam and K. G. Pearson, "The role of proprioceptive feedback in the regulation and adaptation of locomotor activity," *Sensorimotor Control of Movement and Posture*, pp. 343–355, 2002.

[13] B. Balaban and F. Tok, "Gait disturbances in patients with stroke," *Physical Medicine and Rehabilitation*, vol. 6, no. 7, pp. 635–642, 2014.

[14] U.-B. Flansbjer, D. Downham, and J. Lexell, "Knee muscle strength, gait performance, and perceived participation after stroke," *Archives of physical medicine and rehabilitation*, vol. 87, no. 7, pp. 974–980, 2006.

[15] N. Arene and J. Hidler, "Understanding motor impairment in the paretic lower limb after a stroke: A review of the literature," *Topics in Stroke Rehabilitation*, vol. 16, no. 5, pp. 346–356, 2009.

[16] C. A. Oatis, *Kinesiology the Mechanics and Pathomechanics of Human Movement*.

[17] M. Geerars, N. Minnaar-van der Feen, and B. M. Huisstede, "Treatment of knee hyperextension in post-stroke gait. a systematic review," *Gait & Posture*, vol. 91, pp. 137–148, 2022.

[18] C. Beyaert, R. Vasa, and G. E. Frykberg, "Gait post-stroke: pathophysiology and rehabilitation strategies," *Neurophysiologie Clinique/Clinical Neurophysiology*, vol. 45, no. 4-5, pp. 335–355, 2015.

[19] D. A. Neumann, *Kinesiology of the Musculoskeletal System-e-book: Foundations for Rehabilitation*.

[20] S. B. O'Sullivan, T. J. Schmitz, and G. Fulk, *Physical Rehabilitation*.

[21] L. Shams and A. R. Seitz, "Benefits of multisensory learning," *Trends in cognitive sciences*, vol. 12, no. 11, pp. 411–417, 2008.

[22] G. Schmitz, B. Mohammadi, A. Hammer, M. Heldmann, A. Samii, T. F. Münte, and A. O. Effenberg, "Observation of sonified movements engages a basal ganglia frontocortical network," *BMC neuroscience*, vol. 14, no. 1, pp. 1–11, 2013.

[23] R. Sigrist, G. Rauter, R. Riener, and P. Wolf, "Augmented Visual, Auditory, Haptic, and Multimodal Feedback in Motor Learning: A Review," *Psychonomic Bulletin & Review*, vol. 20, 11 2012.

[24] E. Azañón, L. Tamè, A. Maravita, S. A. Linkenauger, E. R. Ferrè, A. Tajadura-Jiménez, and M. R. Longo, "Multimodal contributions to body representation," *Multisensory research*, vol. 29, no. 6-7, pp. 635–661, 2016.

[25] R. Sigrist, G. Rauter, L. Marchal-Crespo, R. Riener, and P. Wolf, "Sonification and haptic feedback in addition to visual feedback enhances complex motor task learning," *Experimental brain research*, vol. 233, no. 3, pp. 909–925, 2015.

[26] A. O. Effenberg and G. Schmitz, "Acceleration and deceleration at constant speed: systematic modulation of motion perception by kinematic sonification," *Annals of the New York Academy of Sciences*, vol. 1425, no. 1, pp. 52–69, 2018.

[27] D. Linnhoff, S. Alizadeh, N. Schaffert, and K. Mattes, "Use of acoustic feedback to change gait patterns: Implementation and transfer to motor learning theory—a scoping review," *Journal of Motor Learning and Development*, vol. 8, no. 3, pp. 598–618, 2020.

[28] P.-J. Maes, M. Leman, C. Palmer, and M. Wanderley, "Action-based Effects on Music Perception," *Frontiers in Psychology*, vol. 4, p. 1008, 2014.

[29] A. N. Antle, G. Corness, and M. Droumeva, "What the Body Knows: Exploring the Benefits of Embodied Metaphors in Hybrid Physical Digital Environments," *Interacting with Computers*, vol. 21, no. 1-2, pp. 66–75, 2009.

[30] S. Roddy and B. Bridges, "Addressing the Mapping Problem in Sonic Information Design through Embodied Image Schemata, Conceptual Metaphors, and Conceptual Blending," *Journal of Sonic Studies*, no. 17, 2018.

[31] P. Kantan, E. G. Spaich, and S. Dahl, "An embodied sonification model for sit-to-stand transfers," *Frontiers in psychology*, vol. 13, 2022.

[32] A. Tajadura-Jiménez, N. Bianchi-Berthouze, E. Furfaro, and F. Bevilacqua, "Sonification of surface tapping changes behavior, surface perception, and emotion," *IEEE MultiMedia*, vol. 22, no. 1, pp. 48–57, 2015.

[33] P. Rowe, C. Myles, C. Walker, and R. Nutton, "Knee joint kinematics in gait and other functional activities measured using flexible electrogoniometry: how much knee motion is sufficient for normal daily life?" *Gait & posture*, vol. 12, no. 2, pp. 143–155, 2000.

[34] P. R. Kantan, E. G. Spaich, and S. Dahl, "A Metaphor-Based Technical Framework for Musical Sonification in Movement Rehabilitation," in *The 26th International Conference on Auditory Display (ICAD 2021)*, 2021.

[35] S. O. Madgwick, A. J. Harrison, and R. Vaidyanathan, "Estimation of IMU and MARG Orientation Using a Gradient Descent Algorithm," in *2011 IEEE International Conference on Rehabilitation Robotics*. IEEE, 2011, pp. 1–7.

[36] P.-J. Maes, J. Buhmann, and M. Leman, "3MO: A Model for Music-Based Biofeedback," *Frontiers in Neuroscience*, vol. 1, 12 2016.

[37] M. Lesaffre, "Investigating Embodied Music Cognition for Health and Well-Being," *Springer Handbook of Systematic Musicology*, pp. 779–791, 01 2018.

[38] F. Grond and J. Berger, "Parameter mapping sonification," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Logos Verlag Berlin, 2011.

[39] P. Kantan, E. G. Spaich, and S. Dahl, "A technical framework for musical biofeedback in stroke rehabilitation," *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 2, pp. 220–231, 2022.

[40] M. E. Morris, T. A. Matyas, T. M. Bach, and P. A. Goldie, "Electrogoniometric feedback: its effect on genu recurvatum in stroke," *Archives of Physical Medicine and Rehabilitation*, vol. 73, no. 12, pp. 1147–1154, 1992.

[41] N. Koohi, D. A. Vickers, R. Lakshmanan, H. Chandrashekar, D. J. Werring, J. D. Warren, and D.-E. Bamiou, "Hearing characteristics of stroke patients: prevalence and characteristics of hearing impairment and auditory processing disorders in stroke patients," *Journal of the American Academy of Audiology*, vol. 28, no. 06, pp. 491–505, 2017.

Automatic Classification of Interactive Gestures for Inter-Body Proximity Sonification

Pascal Staudt, Emre Sarigöl, Marcello Lussana, Marta Rizzonelli, Jin Hyun Kim
Department of Musicology and Media Studies, Humboldt University, Berlin, Germany
pascal.staudt@hu-berlin.de
emresarigol@gmail.com
marcello.lussana@hu-berlin.de
rizzonma@hu-berlin.de
jin.hyun.kim@hu-berlin.de

ABSTRACT

The paper investigates the possibility of classifying interactive gestures from one-dimensional inter-body proximity. The sensor data for training and classification have been acquired within the interdisciplinary research project *Sentire*. The project combines artistic human-computer interaction, real-world research, and sonic interaction design and investigates how body movements mediated by musical sounds promote interactive and prosocial behaviours. The closed-loop interactive system 'Sentire' can sonify proximity and touch between two users in real time. We describe its sonic interaction design based on one-dimensional parameter input and our approach to learning interactive gestures from the continuous sensor data stream. Our model uses two convolutional neural networks (CNNs). One CNN processes the proximity signal and its speed calculation. The second learns frame-wise features from those two signals; an objective aligns the two CNNs through a sigmoid activation layer. Despite the small data set and the low dimensionality of our sensor data, we can predict interactive behaviours that intuitively affect the proximity signal. Especially gestures with arm movements perform significantly better than a baseline classifier based on linear regression.

1. INTRODUCTION

The Italian word sentire [senˈtiːre] means both to hear and to feel. Sentire is a digital system that mediates between bodily movements and musical sounds. Inter-body proximity and touch can be sensed and mapped to an algorithmic sound environment in real time, allowing for a closed-loop auditory interaction between two users. We map the continuous, one-dimensional output of our underlying sensor technology to a specific set of sound parameters, whereby inter-body touch is the only case for which we use a distinct parameter mapping from the continuous proximity sonification [1]. We hypothesise that adding specific interaction behaviours to the gestural input space can help to promote prosocial behaviours (prosociality). Therefore, we investigate whether a deep learning model can predict specific interaction behaviours from the one-dimensional sensor output (proximity). The data for training and classification has been acquired within the interdisciplinary research project Sentire, which combines artistic human-computer interaction, real-world research, and sonic interaction design. The coding system for the annotations was developed through structured observation (SO) from two studies with Sentire [1], [2]. SO is an empirical method for behavioural analysis and is particularly effective for analysing complex non-verbal behaviour. Although Rizzonelli et al. [1] did not develop the codes specifically for gesture recognition, a particular subset of the coding scheme is suitable as a basis for training and evaluating the classifier. Our approach to classifying interaction gestures is distinct from similar research since: (1) the data set was gathered from auditory interactions between two users (dyadic sonic interaction); (2) the output from the sensor used for training is one-dimensional; (3) our deep learning model architecture is suitable for real-time gesture sonification. First, we present related work and the theoretical background of our research. Then, we describe the data collection, the pre-processing of the sensor data, and the model architecture. Finally, we present the metrics for the trained models and interpret the results for different interactive gestures.

2. BACKGROUND AND RELATED WORK

Gesture and pattern recognition from motion sensor data is a topic researched across disciplines and with different aims; from the vast and debatable field of biometrics, which usually focuses on activity recognition (AR) or user recognition (UR) [3]–[9] to more human-centred applications like assisted learning [10]–[12], sonic interaction design [13]–[16] and art [17]–[19]. We distinguish the different approaches by sensor modality, dimensionality, computational method and real-time feasibility.

Vision-based systems for motion tracking usually produce a high dimensional sensor output and are linked to computational effort and real-time constraints. Whereas inertial measurement units (IMUs), which measure acceleration, orientation and angular velocity, are cheap and require less space and processing resources. Considering that each raw data modality has its limitations, hybrid solutions can

ISBN: 978-91-8040-358-0

increase recognition capabilities by combining different modalities and therefore increasing dimensionality [14]–[16], [20]. In contrast to these high-dimensional solutions, DeepEar [21] relies only on one-dimensional microphone input for real-time scene analysis, emotion recognition, stress detection and UR. DeepSense [22] is a classification framework that accommodates various applications and can combine different sensor modalities and dimensions. Visi and Tanaka [18] provide an overview of interactive machine learning techniques for analysing and designing musical gestures. Sonifying (musical) gestures does not necessarily include machine learning. It can also be based on raw sensor input or the computation of higher-level statistical features [10], [11], [23]–[25].

On the one hand, traditional machine learning approaches are still a good choice for gesture recognition, especially for real-time applications [8], [12], [14], [26]. On the other hand, deep learning is trending in current research [3]–[7], [9], [19], [22]. Recent image-based gesture recognition systems primarily build on generative adversarial networks (GANs, e.g. [19]). In contrast, systems with wearable sensors tend to use convolutional neural networks (CNNs) [3], [5], [6] or a combination of CNNs with recurrent neural networks (RNN) [7] respectively long short-term memory (LSTM) networks [4], [9]. Few research investigates gesture recognition in an interactive social context similar to Sentire [13], [27]. Other work does not implement machine learning or deep learning yet specifically focuses on interactive sonification strategies for motion [28], interaction design concerning embodiment [29], or artistic exploration of gesture sonification using a phenomenological approach [30]. To our best knowledge, our work is the first to evaluate a deep learning method for gesture recognition in the social context based only on one-dimensional sensor data (proximity) and suitable for real-time sonification.

3. SONIFYING DYADIC INTERACTION

In order to understand the sonification concepts behind sentire and our approach to sonifying dyadic interaction in general, it is essential to have an overview of the systems' architecture and some of its core characteristics.

Figure 1 shows the signal flow of our interactive system from sensor input to auditory feedback. To obtain a control signal in the digital domain, we use our software written in SuperCollider [31]. Converting the "raw" sensor signal (48kHz) into a proximity and touch signal is part of the sensor system (i.e. the gestural input). In contrast, despite using the same software framework, the sonification stage is part of the sound generator as defined in [32]. We map the sensor system's proximity and touch signal output to selected parameters of an algorithmic sound synthesis process to generate real-time feedback. Making this signal audible through a sound system enables closed-loop auditory interaction. We call the set of parameters that defines the tonal quality of the auditory feedback a "sound environment" (SE) [1]. A SE results from creative decisions regarding sound design, algorithmic composition and a specific parameter mapping that connects the gestural input (proximity and touch) to the sound synthesis processes

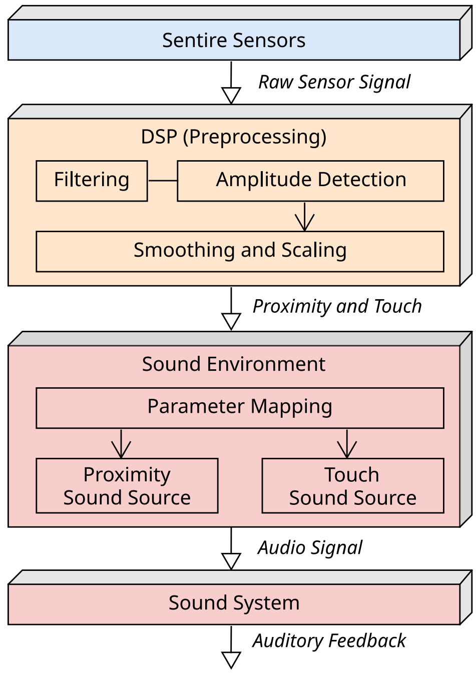

Figure 1. Schematic diagram of Sentire, including—from top to bottom—the sensor hardware (blue), the sensor data processing (orange) and sound generation (red).

("sound sources"). In the broader context of *auditory displays*, and following the taxonomy of [33], a SE constitutes a *parameter mapping sonification*. It combines *continuous* (proximity) with *event-based* (touch) sonification, and although the input data are *quantitative* (physical proximity), a SE also implicitly maps abstract, *qualitative* aspects of interpersonal distance, such as intimacy, to sound. Creating a SE and choosing a parameter mapping strategy resembles well-established paradigms of sonic interaction design. Still, the one-dimensional character of the gestural input and the nature of dyadic sonic interaction demands specific approaches. According to [33] our sonification techniques follow the function of *monitoring* or more specifically they "promote the awareness of task-related processes".

The task during a sentire interaction is sometimes not explicitly set or can be unknown to the interactant. However, when designing a SE we follow strategies to shape the interaction space in specific ways, and navigating this space can be seen as an implicit task during the sonic interaction. We generally aim to promote interactive behaviours and prosociality. For this purpose, we follow two key strategies: to increase *kinaesthetic perception* (see section 3.1) and to intensify *the intimate and personal space* (see section 3.3). Whereas the interaction designer and the performer or therapist are aware of these strategies, the regular Sentire user is not informed about them. Navigating both the gestural and the social space becomes an exploratory act that promotes the interaction's broader goal.

During years of artistic, performative practice with the system [34], we established the following sonic interaction design strategies: (a) mapping amplitude to proximity so that the interaction starts with silence and reaches its maximum loudness closely before the interactants touch; (b) mapping one "main" parameter that distinctly changes the tonal quality and thus is easy to track (usually frequency/pitch or duration/rhythm); (c) applying a non-linear mapping to the above parameters to make the perceived change more prominent in the intimate space; (d) mapping touch to a unique sound event; and (e) alternating the tonal quality of the touch sound based on the touch duration.

3.1 Kinaesthetic Perception

Based on the research of Kim et al. [35] we presume that social interaction is entrenched in the perception and selection of one's own and the other's behaviour and in distinguishing the self from the other. Accordingly, we ground social interaction on kinaesthetic perception and coordination [1]. Kinaesthetic perception—the perception of one's own movements [36] is one of the necessary conditions for coordination. We reflect these assumptions in our "one-to-many" parameter mappings, which map a one-dimensional gestural parameter (proximity) to multiple musical parameters simultaneously. With the above mapping strategies (b) and (d), we keep the action-perception cycle coherent and easy to understand, assuming this is crucial for kinaesthetic perception and explicit gestural agency.

3.2 Gestural Agency

Mendoza and Thompson define "gestural agency" as an agent's influence over other agents within a musical ecosystem [37]. The authors abstract the various signals connecting the network of agents as "gestural space". The low-dimensional gestural space provided by the nature of our sensor system, together with the design of the sound environments, allow for an experience of performing in real-time [37]–[39]. Although the dimensionality of the gestural space is theoretically unlimited, the capability of humans to navigate this space in real-time is not [40]. Doubtlessly, this depends on the individual skills of the human agents and the time they have already spent within the musical ecosystem. The current interaction design of Sentire aims to provide a gestural space accessible to novices. Future versions of the system should include options to challenge more trained users and may even automatically change the affordances based on how the interactants navigate the gestural space.

3.3 Intensifying the Intimate and Personal Space

The intimate (less than 0.5m) and personal (0.5m-1.2m) space, as classified by anthropologist Edward Hall describe culturally dependent zones of interpersonal distance influencing how people interact [41]. We conform with [42] that Hall's proxemics studies should be considered in Interaction Design. Furthermore, with (c-e), we apply the concept of digital proxemics [43], designing a sonic-aesthetic experience that aims to intensify and foster interaction in the intimate and personal space [1], [44]. Intensifying the experience does not necessarily mean using a higher pitch or louder volume in close range. For example, We can also achieve this by making the frequency distribution more harmonic, changing the amount and speed of automatic modulations or by choosing a specific tonal mode [44].

3.4 Promoting Prosociality

Sociality and social interaction are crucial aspects of human health. In line with [45], we believe that musical sound can be a valuable tool to facilitate prosocial behaviours [46], [47]. Studies show that synchronisation/entrainment related to movement and musical structures strengthens prosocial behaviour through empathy and affiliation [48], [49]. We think that sonifying specific interaction behaviours can help to further promote prosociality. Our approach seeks to detect specific movement patterns from the sensor input with the help of deep learning and add the resulting higher-level input cues to the gestural space (see figure 2).

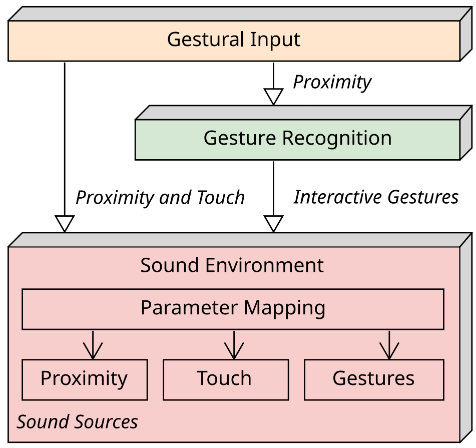

Figure 2. Schematic diagram of Sentire with integrated gesture recognition.

4. DATA AND MODEL DESCRIPTION

Our data set for gesture classification consists of recorded proximity sensor data from 40 sessions of dyadic interactions and annotations of gestures gathered from structured observation (SO). The process of collecting data from SO is cumbersome, and therefore we designed our data processing, training and inference pipelines to comply with a small corpus.

4.1 Data collection

During data collection, we needed to synchronise the sensor data recording with the video recordings used for labelling. Since no simple tool existed that could record multiple high-resolution camera streams and be synchronised via the network, we developed a python based application which we can connect to our software via *Open Sound Control* (OSC). We additionally use this tool to play back recorded interactions, simulating auditory feedback with SuperCollider during sonic interaction design. The software is open source and available as a python package [50].

Sentire's sensor system builds on capacitive coupling, and amplitude measurements through an audio codec [44]. Therefore, the raw sensor data is processed in the audio domain using the SuperCollider programming language (see figure 1). For practical reasons, our software stores the calculated proximity signal in the *Waveform Audio File*

Format (WAV). Yet, the input used for sonification and training does not encode an actual audio signal but movement data. In contrast to audio data, the proximity signal is unipolar (0–1, maximum distance–touch) and mainly has low-frequency content (f<1Hz). The sensor data's narrow bandwidth contributes significantly to our system's minimal design. It decreases the end-to-end training/inference times, where it takes only a few minutes to train, and inference ranges between 1 to 2 seconds, depending on the sample length.

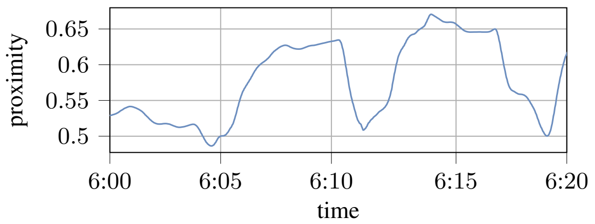

(a)

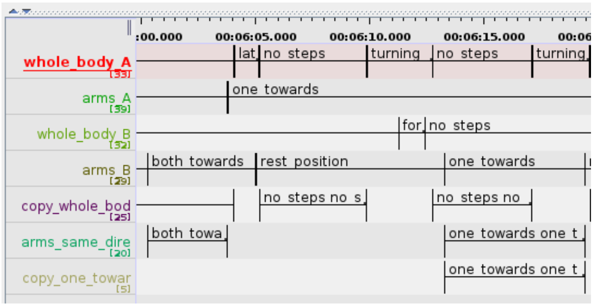

(b)

Figure 3. 20 seconds of recorded proximity data (a) and corresponding annotations in ELAN (b) gathered through structured observation (SO).

Figure 3 shows an excerpt (20 seconds) of proximity data and the corresponding annotations made within the *ELAN* software using three high-resolution video recordings of the interaction from different angles. The available annotations result from a coding scheme developed within the project to investigate the relationship between sound and sociality [1].

4.2 Data Processing

We aimed to design a compact and rich training dataset without significant loss of proximity information and annotated observations. To that end, we applied the following data pre-processing steps:

Downsampling We downsampled the proximity sensor data by a factor of 1,000 from 48kHz to 48Hz, Though we process the initial sensor data in the audio domain, the slow nature of interactive movements limits the frequency content of the sensor signal. An comparison of the proximity signal before and after the downsampling showed that the reduced bandwidth does not change the signal characteristics. The lower sampling rate allowed us to design more compact and less resource-demanding models, which we ideally can use in a real-time prediction setting.

Speed vector calculation and feature extraction For each proximity sensor data sample, we computed the speed vector to be used as an additional input to the model. The speed vector is calculated through the first derivative of the proximity data, i.e. the amount of change in proximity over time. In addition, we extracted nine features at frame levels. For each frame, we computed the following features: IQR, centroid, signal magnitude area, energy and root-mean-square energy, mean, standard deviation, minimum, maximum and median absolute deviation values. We computed these nine features for the proximity signal and the speed vector, resulting in an additional 18-dimensional vector.

Input data batching Each batch has a size of 512 frames and contains data from a distinct session during training. We padded the start of each batch with a null vector to signal the start of the respective session. We also padded the end of each batch with null vectors to signal the end, where the number of null frames added up to 512 frames to ensure that each batch mapped to a particular session. Sessions with longer durations were split into multiple batches, where the padding at the end was applied only to the last batch.

Data augmentation During training, we simultaneously used batches of sessions constructed with different hop sizes, i.e. how much we advanced the time among consecutive frames. Each batch of data was introduced four times with hop sizes of 1.0, 0.75, 0.50, and 0.25. Smaller hop sizes generate more training samples. Therefore, we set the sample weight inversely proportional to the hop size. In addition to using different hop sizes, we also generated training batches using four different frame sizes of 1, 2, 3, and 4 seconds, whereby a frame of 1 second corresponds to 48 samples. We used combinations of these two sets of parameters to obtain 16 training rounds, assuming that the model can more reliably capture time dependencies at different granularities within a sequence.

4.3 Label Set Description

Out of all the annotated labels, we focused on those that might affect the proximity signal and therefore pruned labels describing no motion (e.g. 'rest position', 'gaze'). Table 1 provides an overview of the labels used for training. We use a total of eight labels, which are under two top-tier gesture classes, namely *whole body* and *arm* movements. A *tier* is a category for annotations that share the same set of characteristics. According to this definition, the *whole body* tier, for example, categorises all annotations for interactions where the whole body is involved. All the labels were annotated individually for the two participants, except the *whole body:copy* label. This label is positive only in cases where both participants performed a similar bodily gesture and corresponds to imitation.

We could map the timeline of annotated gestures to each second of processed data since the gestures were annotated

ISBN: 978-91-8040-358-0

Tier	Label
whole body	backward step/s
	lateral step/s
	forward step/s
	full circle
	turning steps
arms	copy (imitation)
	one towards
	both towards

Table 1. Labels used for training and testing the models.

with timestamps in milliseconds granularity, signalling the start and end of each gesture.

4.4 Model Description

We trained separate models for the aforementioned two top tiers (see Figure 4). Each model uses two convolu-

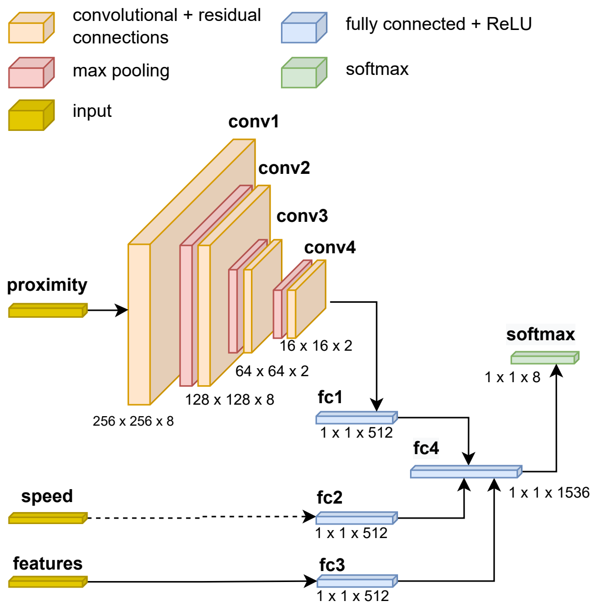

Figure 4. CNN model architecture. There are two identical CNNs, each consisting of four blocks with residual connections. The proximity and speed vector inputs are passed through the two CNNs separately. For simplicity, only one of the CNNs is plotted, where the speed input's stack is represented by a dashed arrow.

tional neural networks (CNNs) with residual connections that process sequences of proximity signals and their speed vector separately. The aforementioned frame-level features are processed by a separate dense layer, whose output is concatenated to the outputs of the two CNNs. We use a sigmoid objective to align these three outputs, and a weighted binary cross-entropy loss, suited for the multi-label classification task. Each label within the tiers is annotated for the two participants separately. However, at training time, we merged labels from both participants into sequences of multi-label target outputs. At inference time, the two models take as input the same input sensor and feature data to predict the respective labels within the tier. It would also be possible to train each tier-level model for the participants separately, but we think this would limit our capability to use the system in a setting with multiple participants.

The two tier-level models use four dilated convolutional blocks, each of which is constructed from residual stacks with weight normalisation. These convolutional blocks are constructed separately for the downsampled sensor and speed vectors. The outputs of the two blocks are concatenated and passed to the final multi-headed output layer with two heads, one for each participant.

5. MODEL EVALUATION

Out of the 40 recorded sessions, we used fixed sets of 30, 4 and 6 sessions for training, validation and evaluation of our models, respectively. We scored each model in the ensemble separately and aggregated the predictions. To compare performances across the labels, we fixed a precision level of 0.8 for each label and computed the f1-score at that precision level.

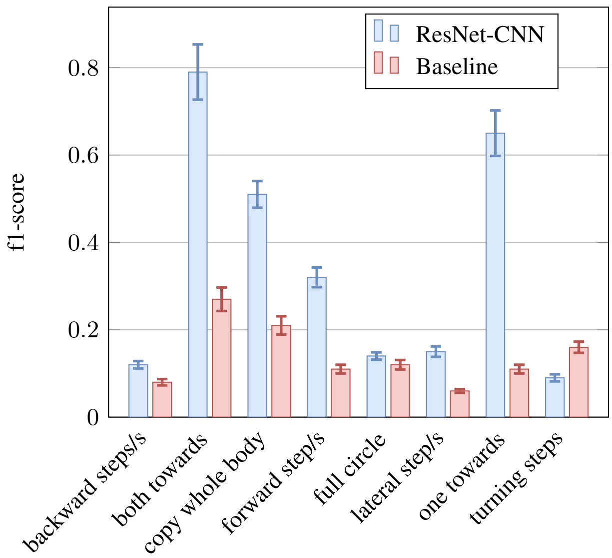

Figure 5. Average f1-scores and standard deviation per label for the ResNet-CNN model (blue) and the logistic regression baseline (red).

Figure 5 shows the resulting metrics, including a comparison with a baseline for each label based on a logistic regression model whose binary outcome is whether the gesture took place or not. We calculated scores by averaging the outcomes for each session. The resulting metrics vary significantly for each label. We observed considerable uplift for three gestures: *arms:one towards* (f1:0.11/0.65), *arms:both towards* (f1:0.27/0.79) and *whole body:copy* (f1:0.21/0.51). Gestures that are expected to produce a distinctive sensor signal perform better (e.g. arms movements involving one or both arms). In contrast, gestures that intuitively might produce sensor signals without a consistent pattern (e.g. *backward* and *turning steps*, *full circle*) are poorly detected. Gestures involving steps performed better than the baseline, though not with such a significant uplift than gestures including arm movements.

ISBN: 978-91-8040-358-0

As a proof of concept, we wanted to quantify and compare the level of misclassifications across and within tiers. Intuitively, if the model makes more mistakes within top tiers (e.g. mislabeling *one towards* with *both towards*), this would still show that the training did not fail significantly and so would help us conclude that there is space for improvement. To that end, we created inter-tier predictors by aggregating the predicted and true labels into their top tiers, where the occurrence of any of the labels within a tier is treated as a positive label, and computed *Cohen's Kappa Coefficient* [51] κ to empirically test the relationship between inter- and intra-tier predictions. We obtain κ for the intra-tier setting by aggregating the performance of the models for individual classes into a single average κ metric. Kappa Coefficient is defined in Eq. (1) as follows:

$$\kappa = \frac{Pr_a - Pr_e}{1 - Pr_e} \tag{1}$$

where Pr_a is the relative observed agreement between classification and test data, and Pr_e is the hypothetical probability of chance agreement, where the observed data are used to calculate probabilities of each rater to choose a label at random. In this case, $\kappa = 1$ if the raters are in complete agreement, and $\kappa <= 0$ if they are in agreement only to the extent of what is expected by chance (as defined by Pr_e). We chose Kappa coefficient since it captures a combination of precision and recall similar to the f1-score value, but includes the comparison to the baseline of a random classifier in its calculation. Using the f1-score for evaluating a random classifier would instead be dependent on the distribution of the classes, whereas Kappa Coefficient is independent of class size and thus is resilient to biases introduced by differing class sizes.

Tier	**Label**	κ
whole body	inter-tier	0.58
	intra-tier	0.13
arms	inter-tier	0.41
	intra-tier	0.14

Table 2. Kappa coefficients for inter- and intra-tier predictions.

Table 2 provides the intra- and inter-tier coefficients. There are significantly higher levels of agreement for the inter-tier setting for both tiers. Therefore we can safely hypothesise that prediction errors are related more to specific labels, and less to the question of the predictability of gestures at large.

6. CONCLUSIONS AND OUTLOOK

The described approach demonstrates the possibility of classifying interactive gestures from one-dimensional proximity data. Despite the ultra-low dimensionality of our sensor signal and a relatively small data set derived from structured observation, we can predict interactive behaviours that intuitively affect the proximity signal. Especially gestures with arm movements perform significantly better than a linear regression baseline. The approach is not restricted to proximity data and gestural input. In the context of sonification, the method can be abstracted to classifying time-based events in one-dimensional data streams. Granted that enough labelled data exist, the model architecture could be the ground for event-based sonification for a broad range of health and environmental data. Its real-time feasibility could be especially relevant for process monitoring tasks that require immediate sound feedback. Future work should include cross validation and investigate the effect of a larger corpus and more diverse augmentation/noising techniques. At the time of writing, the presented work serves as a proof of concept, and an actual implementation for real-time sonification remains. We think that adding interactive gestures to the gestural space of our interactive system could help to promote prosociality. Testing this hypothesis should be the subject of subsequent studies. Finally, we encourage the discourse on whether the effort spent to follow the deep learning trend is justified. A straightforward approach based on raw sensor input and higher-level statistic features, combined with well-designed parameter mappings, could achieve similar results while leaving more time to focus on the artistic facets of the interdisciplinary project.

Acknowledgments

We thank Karunya Somashekar, Marian Lepke and Anton Kogge for their student assistance. The Project was funded by the Federal Ministry of Education and Research (BMBF).

References

[1] M. Rizzonelli, J. H. Kim, P. Staudt, and M. Lussana, "Fostering Social Interaction through Sound Feedback: Sentire," *Organised Sound*, vol. 28, no. 01, pp. 317–326, 2022.

[2] B. Stahl, M. Lussana, M. Rizzonelli, P. Staudt, A. Milek, and J. H. Kim, "Sentire: Exploring the Suitability of Movement and Sound in Couple Therapy," presented at the Open Ground: Music Therapy in Collaboration & Exchange (London), 2021.

[3] M. Gadaleta and M. Rossi, "IDNet," *Pattern Recognition*, vol. 74, no. C, pp. 25–37, 2018.

[4] L. Chen, Y. Zhang, and L. Peng, "METIER: A Deep Multi-Task Learning Based Activity and User Recognition Model Using Wearable Sensors," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, no. 1, 5:1–5:18, 2020.

[5] M. Gil-Martín, R. San-Segundo, R. de Córdoba, and J. M. Pardo, "Robust Biometrics from Motion Wearable Sensors Using a D-vector Approach," *Neural Processing Letters*, vol. 52, no. 3, pp. 2109–2125, 2020.

[6] Y. Li, H. Hu, Z. Zhu, and G. Zhou, "SCANet: Sensor-based Continuous Authentication with Two-stream Convolutional Neural Networks," *ACM Transactions on Sensor Networks*, vol. 16, no. 3, 29:1–29:27, 2020.

[7] Q. Zou, Y. Wang, Q. Wang, Y. Zhao, and Q. Li, "Deep Learning-Based Gait Recognition Using Smartphones in the Wild," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3197–3212, 2020.

[8] J. Zhu, R. San-Segundo, and J. M. Pardo, "Feature extraction for robust physical activity recognition," *Human-centric Computing and Information Sciences*, vol. 7, no. 1, pp. 1–16, 2017.

[9] T. Zhu, L. Fu, Q. Liu, Z. Lin, Y. Chen, and T. Chen, "One Cycle Attack: Fool Sensor-Based Personal Gait Authentication With Clustering," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 553–568, 2021.

[10] T. Großhauser, B. Bläsing, C. Spieth, and T. Hermann, "Wearable Sensor-Based Real-Time Sonification of Motion and Foot Pressure in Dance Teaching and Training," *Journal of the Audio Engineering Society*, vol. 60, pp. 580–589, 2012.

[11] F. L. Cibrian, O. Peña, D. Ortega, and M. Tentori, "BendableSound: An elastic multisensory surface using touch-based interactions to assist children with severe autism during music therapy," *International Journal of Human-Computer Studies*, vol. 107, pp. 22–37, 2017.

[12] D. Dalmazzo and R. Ramirez, "Bow Gesture Classification to Identify Three Different Expertise Levels: A Machine Learning Approach," in *Machine Learning and Knowledge Discovery in Databases*, P. Cellier and K. Driessens, Eds., Springer, Cham, 2020, pp. 494–501.

[13] S. Mealla, A. Väljamäe, M. Bosi, and S. Jordà, "Sonification of Brain and Body Signals in Collaborative Tasks Using a Tabletop Musical Interface," in *Proceedings of the 17th International Conference on Auditory Display*, Budapest, Hungary, 2011, pp. 1–5.

[14] C. Volioti, S. Hadjidimitriou, S. Manitsaris, L. Hadjileontiadis, V. Charisis, and A. Manitsaris, "On mapping emotional states and implicit gestures to sonification output from the 'Intangible Musical Instrument'," in *Proceedings of the 3rd International Symposium on Movement and Computing*, New York, NY, USA: Association for Computing Machinery, 2016, pp. 1–5.

[15] A. Tajadura-Jiménez, F. Cuadrado, P. Rick, *et al.*, "Designing a gesture-sound wearable system to motivate physical activity by altering body perception," in *Proceedings of the 5th International Conference on Movement and Computing*, New York, NY, USA: Association for Computing Machinery, 2018, pp. 1–6.

[16] A. Tanaka, B. Di Donato, M. Zbyszynski, and G. Roks, "Designing Gestures for Continuous Sonic Interaction," in *The International Conference on New Interfaces for Musical Expression*, Porto Alegre, 2019.

[17] G. Burloiu, V. Mihai, and S. Damian, "Layered Motion and Gesture Sonification in an Interactive Installation," *Journal of the Audio Engineering Society*, vol. 66, no. 10, pp. 770–778, 2018.

[18] F. G. Visi and A. Tanaka, "Interactive Machine Learning of Musical Gesture," in *Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity*, E. R. Miranda, Ed., Cham: Springer International Publishing, 2021, pp. 771–798.

[19] J. Ren and J.-K. Park, "Dynamic Recognition and Analysis of Gait Contour of Dance Movements Based on Generative Adversarial Networks," *Computational Intelligence and Neuroscience*, vol. 2022, Article Id: 3276696, 2022.

[20] C. Chen, R. Jafari, and N. Kehtarnavaz, "A survey of depth and inertial sensor fusion for human action recognition," *Multimedia Tools and Applications*, vol. 76, no. 3, pp. 4405–4425, 2017.

[21] N. D. Lane, P. Georgiev, and L. Qendro, "DeepEar: Robust smartphone audio sensing in unconstrained acoustic environments using deep learning," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, New York, NY, USA: Association for Computing Machinery, 2015, pp. 283–294.

[22] S. Yao, S. Hu, Y. Zhao, A. Zhang, and T. Abdelzaher, "DeepSense: A Unified Deep Learning Framework for Time-Series Mobile Sensing Data Processing," in *Proceedings of the 26th International Conference on World Wide Web*, Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, 2017, pp. 351–360.

[23] B. Horsak, R. Dlapka, M. Iber, *et al.*, "SONIGait: A wireless instrumented insole device for real-time sonification of gait," *Journal on Multimodal User Interfaces*, vol. 10, no. 3, pp. 195–206, 2016.

[24] A. Giomi and J. Leonard, "Towards an Interactive Model-Based Sonification of Hand Gesture for Dance Performance," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Birmingham, United Kingdom, 2020, pp. 369–374.

[25] M. Raptis, D. Kirovski, and H. Hoppe, "Real-time classification of dance gestures from skeleton animation," in *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, New York, NY, USA: Association for Computing Machinery, 2011, pp. 147–156.

[26] B. Caramiaux, N. Montecchio, A. Tanaka, and F. Bevilacqua, "Adaptive Gesture Recognition with Variation Estimation for Interactive Systems," *ACM Transactions on Interactive Intelligent Systems*, vol. 4, no. 4, 18:1–18:34, 2014.

[27] S. Albawi and O. Ucan, “Social Touch Gesture Recognition Using Convolutional Neural Network,” *Computational Intelligence and Neuroscience*, vol. 2018, pp. 1–10, 2018.

[28] S. Landry and M. Jeon, “Interactive sonification strategies for the motion and emotion of dance performances,” *Journal on Multimodal User Interfaces*, vol. 14, no. 2, pp. 167–186, 2020.

[29] S. Klemmer, B. Hartmann, and L. Takayama, “How bodies matter: Five themes for interaction design,” in *Proceedings of the 6th Conference on Designing Interactive Systems*, vol. 2006, New York, NY, USA: Association for Computing Machinery, 2006, pp. 140–149.

[30] P. Polotti and M. Goina, “Gesture Sonification. An Interaction Design Approach to an Artistic Research Case,” *EAI Endorsed Transactions on Creative Technologies*, vol. 8, no. 27, e1, 27 2021.

[31] P. Staudt, B. Gola, and M. Lussana. “Sentire LE.” (2022), [Online]. Available: `https://audiophil-dev.github.io/sentire_le/` (visited on 09/02/2022).

[32] E. R. Miranda and M. M. Wanderley, *New Digital Musical Instruments: Control and Interaction Beyond the Keyboard*. Middleton, Wisconsin: A-R Editions, Inc., 2006.

[33] B. N. Walker and M. A. Nees, “Theory of sonification,” in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds., Berlin, Germany: Logos Publishing House, 2011, pp. 9–39.

[34] P. Staudt, O. Kozmanidze, and M. Lussana. “Sentire: A deep sensory experience.” (2019), [Online]. Available: `https://sentire.me/` (visited on 09/05/2022).

[35] J. H. Kim, A. Reifgerst, and M. Rizzonelli, “Musical Social Entrainment,” *Music & Science*, vol. 2, pp. 1–17, 2019.

[36] D. A. Rosenbaum, *Human Motor Control*, 2nd ed. Amsterdam [etc.]: Academic Press, 2009.

[37] J. I. Mendoza and M. R. Thompson, “Gestural Agency in Human–Machine Musical Interaction,” in *The Routledge Companion to Embodied Music Interaction*, M. Lesaffre, P.-J. Maes, and M. Leman, Eds., Routledge: Taylor & Francis Group, 2017, pp. 412–419.

[38] D. Arfib, J.-M. Couturier, L. Kessous, and V. Verfaille, “Strategies of mapping between gesture data and synthesis model parameters using perceptual spaces,” *Organised Sound*, vol. 7, pp. 127–144, 2003.

[39] V. Zappi and A. McPherson, “Dimensionality and Appropriation in Digital Musical Instrument Design,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Goldsmiths, UK, 2014, pp. 455–460.

[40] R. Tubb and S. Dixon, “An Evaluation of Multidimensional Controllers for Sound Design Tasks,” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, Seoul, Republic of Korea: Association for Computing Machinery, 2015, pp. 47–56.

[41] E. T. Hall, *The Hidden Dimension*. Garden City, N.Y.: Doubleday, 1966.

[42] S. Greenberg, N. Marquardt, T. Ballendat, R. Diaz-Marino, and M. Wang, “Proxemic Interactions: The New Ubicomp?” *Interactions*, vol. 18, no. 1, pp. 42–50, 2011.

[43] J. A. McArthur, *Digital Proxemics: How Technology Shapes the Ways We Move*, S. Jones, Ed. New York: Peter Lang Publishers, 2016, vol. 110.

[44] P. Staudt, A. Kogge, M. Lussana, M. Rizzonelli, B. Stahl, and J. H. Kim, “A New Sensor Technology for the Sonification Of Proximity and Touch in Closed-Loop Auditory Interaction,” in *Proceedings of the 19th Sound and Music Computing Conference*, Saint-Étienne, France: Zenodo, 2022, pp. 441–448.

[45] I. Cross, F. Laurence, and T.-C. Rabinowitch, “Empathy and Creativity in Group Musical Practices: Towards a Concept of Empathic Creativity,” in *The Oxford Handbook of Music Education*, G. E. McPherson and G. F. Welch, Eds., 1st ed., vol. 2, Oxford University Press, 2012, pp. 337–353.

[46] D. J. Rickson and W. G. Watkins, “Music therapy to promote prosocial behaviors in aggressive adolescent boys–a pilot study,” *Journal of Music Therapy*, vol. 40, no. 4, pp. 283–301, 2003.

[47] S. Kirschner and M. Tomasello, “Joint music making promotes prosocial behavior in 4-year-old children,” *Evolution and Human Behavior*, vol. 31, no. 5, pp. 354–364, 2010.

[48] T.-C. Rabinowitch, I. Cross, and P. Burnard, “Long-term musical group interaction has a positive influence on empathy in children,” *Psychology of Music*, vol. 41, no. 4, pp. 484–498, 2012.

[49] J. Stupacher, G. Wood, and M. Witte, “Synchrony and sympathy: Social entrainment with music compared to a metronome.,” *Psychomusicology: Music, Mind, and Brain*, vol. 27, pp. 158–166, 2017.

[50] P. Staudt and B. Gola. “Pyoscvideo: Synchronized multi-video recorder and player, controllable via OSC.” (2022), [Online]. Available: `https://github.com/Sentire-Dev/pyOscVideo` (visited on 08/13/2022).

[51] J. Cohen, “Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit,” *Psychological Bulletin*, vol. 70, no. 4, pp. 213–220, 1968.

Musification of Accelerometry Data Towards Raising Awareness of Physical Activity

Juan Ignacio Mendoza
University of Jyväskylä,
Department of Music, Art, and Culture Studies
juigmend@student.jyu.fi

Andrew Danso
University of Jyväskylä,
Department of Music, Art, and Culture Studies,
Centre of Excellence in Music, Mind, Body and Brain
andrew.a.dansoadu@jyu.fi

Geoff Luck
University of Jyväskylä,
Department of Music, Art, and Culture Studies,
Centre of Excellence in Music, Mind, Body and Brain
geoff.luck@jyu.fi

Timo Rantalainen
University of Jyväskylä,
Faculty of Sport and Health Sciences,
Gerontology Research Center
timo.rantalainen@jyu.fi

Lotta Palmberg
University of Jyväskylä,
Faculty of Sport and Health Sciences,
Gerontology Research Center
lotta.m.palmberg@jyu.fi

Sebastien Chastin
Glasgow Caledonian University,
School of Health and Life Sciences;
Ghent University,
Department of Movement and Sports Sciences
sebastien.chastin@gcu.ac.uk

ABSTRACT

Previous research has shown that the temporal dynamics of human activity recorded by accelerometers share a similar structure with music. This opens the possibility to use musical sonification of accelerometry data to raise awareness of daily physical activity. In this study a method was developed for quantifying the daily structure of human activity using multigranular temporal segmentation, and applying it to produce musical sonifications. Two accelerometry recordings of physical activity were selected from a dataset, such that one shows more physical activity than the other. These data were segmented in different time-scales so that segmentation boundaries at a given time-scale have a corresponding boundary at a finer time-scale, occurring at the same point in time. This produced a hierarchical structure of daily events embedded in larger events, which is akin to musical structure. The segmented daily data of each subject was mapped to musical sounds, resulting in two short musical pieces. A survey measured the extent to which people would identify the piece corresponding to the most active subject, resulting in a majority of correct answers. We propose that this method has potential to be a valuable and innovative technique for behavioural change towards reducing sedentary behaviour and increasing physical activity.

1. INTRODUCTION

Miniature sensors, wearable devices and mobile technologies can track daily activity of people, both in extent (i.e., amount of movement) and type (e.g., walking, sitting). This capability has been utilised as a behavioural change technique [1] in interventions to promote a healthier lifestyle, increase physical activity and reduce sedentary behaviour [1] [3, 4]. These technologies may be effective aids in interventions to increase physical activity and reduce sedentary behaviour [5], but only in the short-term. Long-term adherence is still a major challenge [6–10]. Recent reviews suggest that more engaging methods are needed to effectively produce a change in behaviour [11].

Sonification is a potential strategy to increase long term engagement and adherence, especially since it has been shown that the temporal dynamics of human motion and activity share similarity with that of music [12, 13]. Several studies have explored the use of real-time sonification of movement to aid sports performance and rehabilitation [14]. For example, Ley-Flores et al. [15] found that sonification of exercise with metaphorical sounds affect body perception, causing people to feel strong and thus increase their amount of physical activity. Other studies investigated presenting activity patterns as musical sound to raise awareness about behaviour. For example, Krasnoskulov [16] developed a system in which data measured by an accelerometer and optical heart-rate sensor were mapped to musical parameters such as pitch, timbre, tempo, space and loudness. This form of musical sonification is rather direct and may not result in a clear representation of events. Consequently, some

[1] A short article by O'Keeffe, Scheid and West [2] explains the differences and similarities between physical activity and sedentary behaviour.

ISBN: 978-91-8040-358-0

studies have considered segmentation of data, so that the resulting sonification is structured in blocks that preserve the temporal relations of events. Last and Usyskin [17] developed a sonification paradigm that segments data into a user-defined number of segments, which was successful to convey the desired information. Vickers and Höldrich [18] progressed this to produce segments using zero-crossing of a one-dimensional data-stream. Then the segments were mapped to sound. These studies show that sonification and musical sonification are feasible ways to convey activity data. Temporal segmentation may be a relevant part of the process, as it allows for mappings between data and sound that have a clear correspondence. However, the temporal segmentation methods used by the mentioned studies have important limitations, as they are based on threshold, zero-crossings or clustering. These methods require careful calibration of input parameters and do not generalise well when patterns in data are multidimensional.

The present study has focused on the development of a system to produce musical sonification (also referred to as *musification*) of daily activity data recorded by wearable devices. The method employs a novel approach to multi-granular temporal segmentation, that results in a clear correspondence between daily events and sound. Additionally, the system does not require the final user to do any fine-tuning of segmentation parameters. We propose this system as an aid in behavioural change, by raising awareness of people's own daily physical activity in an engaging way.

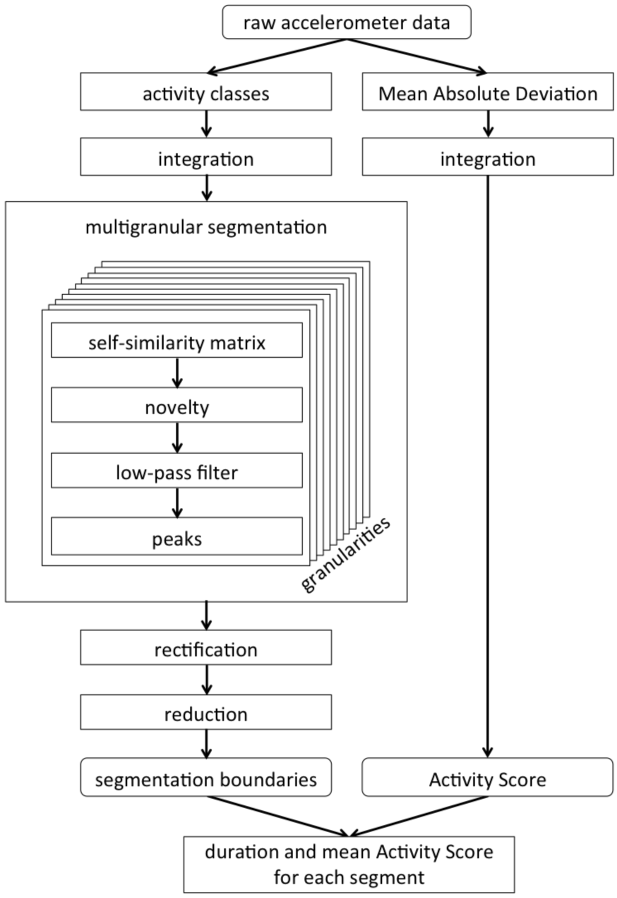

Figure 1. Multigranular segmentation of daily activity.

2. METHODS

2.1 Accelerometry Data

We used two multiple-day recordings of accelerometry from 75-year-old adults. These were chosen from the AGNES database [19, 20] so that one corresponds to a low-activity sedentary subject while the other corresponds to a high-activity non-sedentary subject. The data was obtained by two tri-axial accelerometers, one chest-worn and the other thigh-worn. These data were pre-processed to obtain features for successive non-overlapping epochs of 5 seconds. One feature is the Mean Absolute Deviation (MAD) of the square norm [21], from the thigh-worn accelerometer (Fig. 2a). The other features are the activities identified from the orientation of the accelerometers: lying, sitting, upright posture and walking [22] (Fig. 2b).

2.2 Segmentation

The segmentation procedure is shown in Fig. 1. After MAD is computed and activities are identified, numerosity is reduced by integrating in windows of 120 data points of 5 seconds each (10 minutes) with an overlap of half the length of the window. For MAD the integration is

$$A_i = \log_b \left(1 + \sum_{j=1}^{n} w_j \right), i = \{1 \ldots N\},$$

where vector A of length N is the Activity Score, N is the number of windows, the logarithmic base b is a free parameter to rescale A, and w is one window of length n. The logarithm preserves the data distribution, as the relation between time of inactivity and activity follows a power-law distribution [12]. For the examples reported in this article the logarithm has a base $b = 3$. Each activity is a binary vector, where an activity is represented by a one, otherwise a zero (Fig. 2b). The integration of each activity vector is the sum of the window, with the same length and overlap as for MAD. Additionally, integration acts as a low-pass filter removing unnecessary detail.

The next step is segmentation of the integrated data using the algorithm described by Foote [23]. That algorithm has been used for segmentation of musical audio and video. It can detect boundaries of segments at different *granularities* (i.e., time-scales). It has also been tested for segmentation of accelerometry data of dance [24] and daily activities [25].

The segmentation algorithm first computes a self-similarity matrix of the integrated activities (Fig. 3a). Then, a checkerboard kernel (i.e., a small matrix of four sections where the diagonal is negative units and the anti-diagonal is positive units) tapered by a normal distribution, is correlated along the diagonal of the self-similarity matrix. This was done several times, each with a checkerboard kernel of minimally different size. The size of the kernel corresponds to the granularity of the

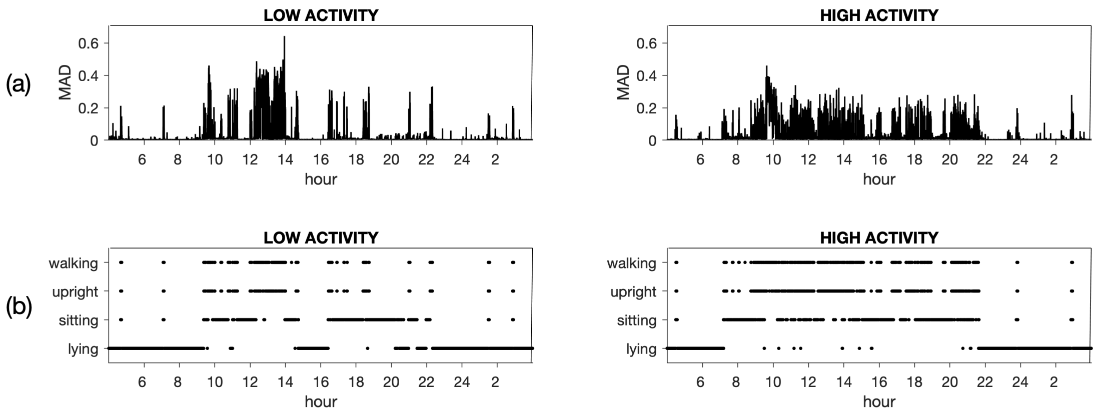

Figure 2. a) Mean Absolute Deviation every 5 seconds of accelerometer data; b) Classification of daily activities.

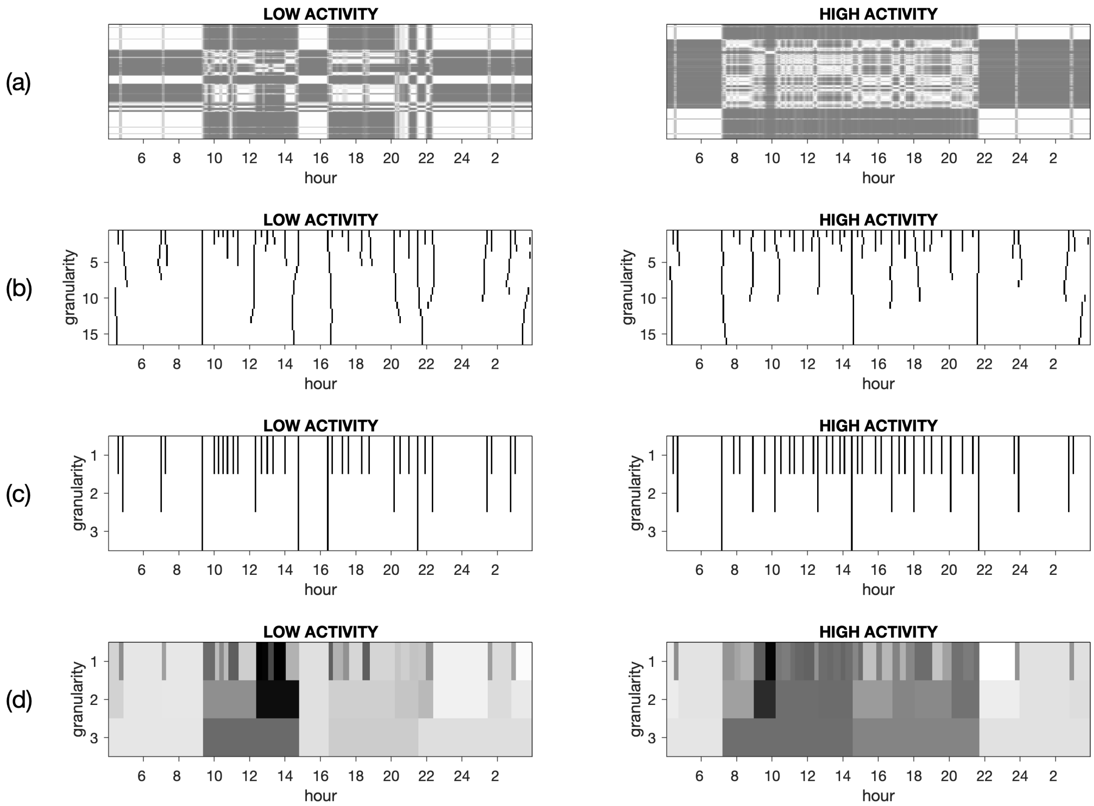

Figure 3. a) Self-similarity of activities; b) multigranular raw segmentation boundaries; c) rectified and reduced multigranular segmentation boundaries; d) segmented Activity Score, darker shades show greater average activity for a segment and vice versa.

segmentation. A smaller kernel detects finer granularity segments and vice versa. The size of the kernels was specified as the standard deviation σ_k of their normal distribution tapering. For the examples shown here, $\sigma_k = \{2, 4, \ldots, 32\}$ windows. This resulted in several novelty scores, one for each granularity, each of which was then smoothed with a normal-distribution (i.e., Gaussian) low-pass filter to remove irrelevant peaks. The peaks of each novelty score represent segmentation boundaries (Fig. 3b). The size of each filtering vector was set to each corresponding value in σ_k, while the standard deviation was set to 0.4 for all of them.

The segmentation boundaries at different granularities are not perfectly aligned in time (Fig. 3b) because, as the checkerboard kernel gets larger, it incorporates more information causing the novelty peak to move slightly in

Table 1. Segmented Activity Score

fine		medium		coarse	
duration	mean Activity	duration	mean Activity	duration	mean Activity
5	0.45	8	0.45	38	0.48
3	1.11	-	-	-	-
30	0.48	30	0.48	-	-
8	1.03	21	0.92	88	1.48

First 4 lines of "high activity" segmented (corresponding to Fig. 3c and d). Duration is windows of integrated data. The headers of this table are not part of the actual list.

either direction. However, because the granularities were set with minimal difference ($\Delta\sigma_k = 2$ windows), it is safe to assume that they correspond to the same segment. Following this logic, every coarser granularity boundary has an origin in a finer granularity boundary, except for those at the borders. The temporal structure is hierarchical, where segments are embedded in larger segments. This reflects the structure of human daily activity. For example, a large portion of the day such as the morning, may contain activities like waking-up and getting ready, breakfast, commuting, and so forth. This hierarchical structure is also analogous to musical structure. For example, a song has sections like introduction, verse and chorus, each of which have sub-sections, such as melodic lines. However, in music the boundaries of each section exactly match in time, unlike the structure resulting from the procedure described above. If that multigranular structure were to be used as musical structure for sonification, it would result in a seemingly unnatural performance. For example, each granularity level may be assigned to a different musical instrument. If so, then instruments would begin and change sections of the song at different times.

Therefore, the segmentation boundaries were aligned to the finest-granularity boundary. Also the boundaries at the borders were removed. This resulted in sequences at different granularities being identical or slightly different. Thus, the finest and coarsest granularity sequences were kept, as well as the sequences that provide greatest variety in number of boundaries. For the examples given here, the reduction resulted in sequences at 3 levels of granularity: fine, medium and coarse (Fig. 3c). Finally, the median Activity Score was computed for each segment at each granularity level (Fig. 3d).

2.3 Musical Sonification

The result of the segmentation procedure is a list of paired columns, where the paired values are segment duration, in windows, and the mean Activity Score for the segment. If a segment's boundary doesn't have a corresponding boundary at a coarser granularity level, the values are omitted. The first line assumes a boundary at all levels of granularity. Table 1 shows an example.

The list was formatted as a CSV file and has a header line composed by the number of windows, the sum and grand mean of Activity Scores (from the matrix depicted by Fig. 3d), and the number of granularities. This file is the input to a separate sonification program consisting mainly of a sequencer and two synthesis modules (Fig. 4.). The sequencer loads the CSV file and immediately reads the header. The user specifies how long the performance will last and the program computes the duration of each window in real time units (e.g., milliseconds), using the first value in the header (number of windows). The second value of the header (sum of Activity Score) is used as the seed for all pseudo-random generators, to obtain a deterministic performance (i.e, the sonification of a CSV file will always be the same). This may help to perceive a strong connection between sonic material and actual daily activity information. The third value in the header (grand mean Activity Score) sets the tempo. The mean between the values of both subjects was mapped to 120 BPM (beats per minute) for crotchet notes (60 BPM for minim notes), as the typical healthy average heartbeat at rest is just over 60 BPM [26] and both preferred musical tempo and average walking steps have a period of about 120 BPM [27]. Hence, the sonification for the high-activity subject will have a slightly higher tempo than the sonification for the low-activity subject. The last element of the header (number of granularities) is used to compute the mean Activity for each combined segment. For example, for the first row in Table 1, all mean Activity values will be added and divided by 3. For the second row, the only value is for the finest granularity and will be divided by 3.

The user inputs a duration in seconds and clicks a button to start the performance. Then, the first line in the CSV file (i.e., the first row of Table 1) will be read and it will wait the duration given by the leftmost value multiplied by the duration of each window, then it will read the next line and so on. When each line is read, the values are sent to the synthesis modules as described below. This process continues until the final line is reached or until the user interrupts it by the click of a button.

Synthesis module 1 is composed by three synthesisers that produce bell-like sounds, whose pitches are pseudo-randomly produced according to a distribution that smoothly transitions from chromatic (i.e., all 12 tones allowed) to a user-selected scale. For this study a pentatonic scale was used. The transition is given by the mean Activity of all segments at the start of a finest-granularity segment. The higher this value is, the closer the distribution will be to the selected scale. For example, when the program begins playing the list in Table 1, it will compute the mean of the "mean Activity" values of the first row, which will determine the distribution of the pseudo-randomly produced notes. Given this distribution, each synthesiser produces a note at the start of each segment and the duration of the note is the duration of the segment. Each synthesiser has been set to play only at a distinct octave, with the synthesiser allocated to the coarsest granularity playing the lowest octave and vice versa.

The resulting sounds are somehow dissonant when activity is low and consonant within the user-defined scale, when the activity is moderately energetic. This defines

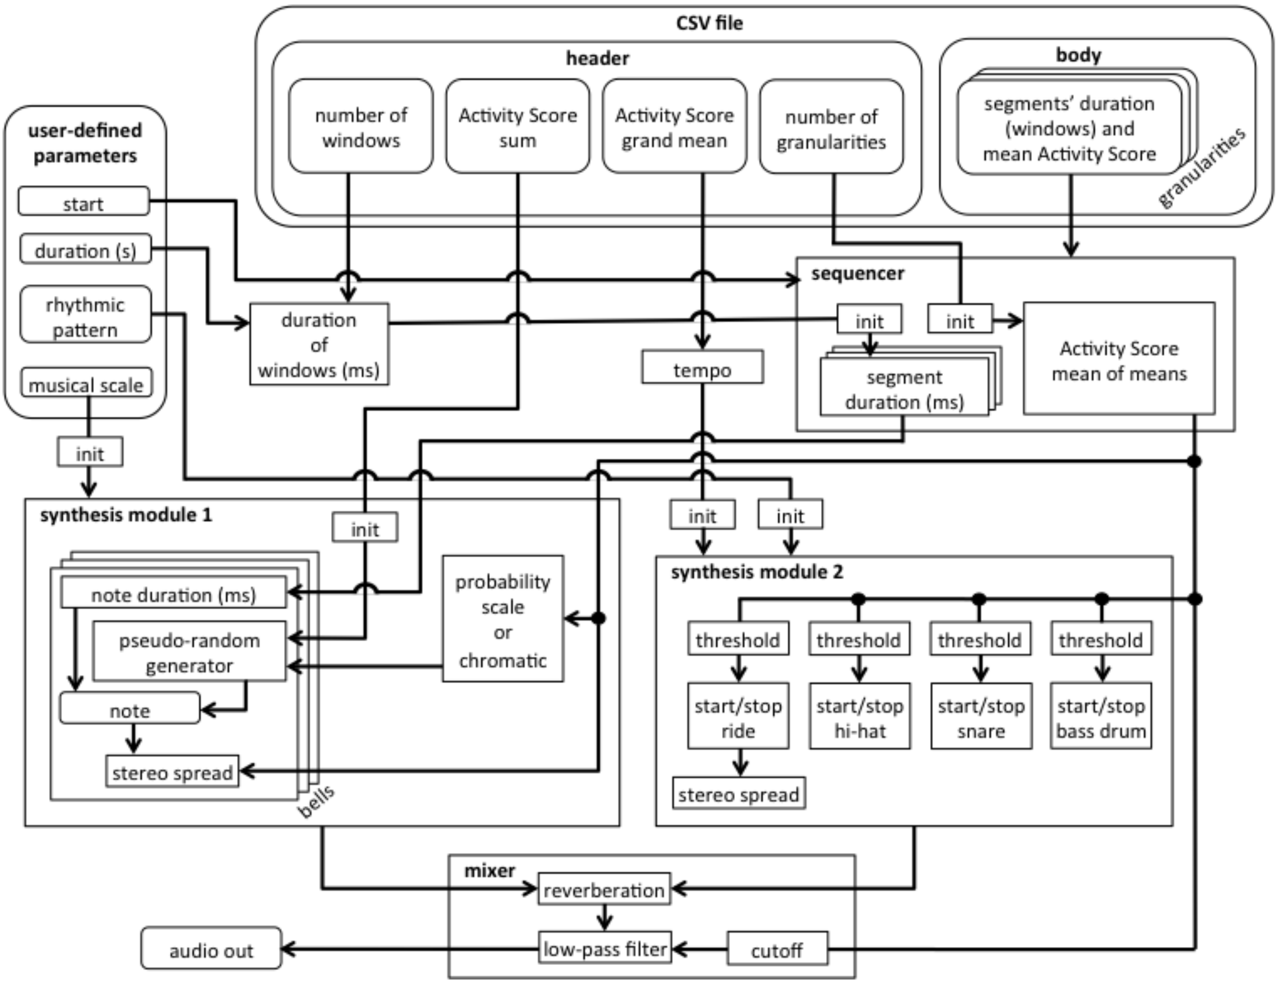

Figure 4. Sonification program.

high amount of activity as consonant and low amount of activity as dissonant. Also each synthesiser has a "stereo spread" capability that has been mapped so that the higher the mean Activity, the wider the stereo allocation of the notes, meaning the pseudo-random balance between their output to the two main output audio channels.

Synthesis module 2 is a drum-machine with 16 steps (quaver notes) and 5 voices produced by frequency modulation: ride cymbal, open hi-hat cymbal, closed hi-hat, snare drum and bass drum. The rhythmic pattern can be programmed by the user. The tempo is given by the grand mean Activity Score as explained previously and each instrument is activated when a level of mean Activity of the current segment exceeds a defined threshold. For this study the bass drum was set to be permanently active, while the ride cymbal was set to fade in when activity changes from very low to moderately low. Also the ride cymbal was set to permanently have a full stereo spread, resulting in a subtle and surrounding rhythmical noise. The open hi-hat was set to a medium threshold and the closed hi-hat was not used for this study. The snare drum was set to a moderately high threshold. The full drum set is active when activity is energetic.

In the examples (see Fig. 2 and Fig. 3), the full drum set and only notes within the defined scale play between about 12:30 and 15:00 for the low-activity subject and from about 9:00 to 10:00 for the high-activity subject. The output from the bell synthesisers and drum-machine is mixed and subtle reverberation is added to blend the sounds. Finally, a low-pass filter is applied to the final mix and its cutoff frequency is controlled directly by the mean activity of the current segment, so that the resulting sound is slightly brighter as there is more energetic activity and vice versa.

2.4 Perceptual Assessment

Two audio files were produced with the method described above, using excerpts from 6:00 to 23:00 of the data presented in the figures. These audio files were used as stimuli for a perceptual assessment. Data for this assessment were collected during 31 days by means of a short survey using QuestionPro, a service to make and publish questionnaires which can be answered with an internet browser, and Twitter. Participants were recruited via Twitter and Facebook using both free and paid adverts, the latter targeting Finland and major English-speaking countries, and via authors' direct contact within their acquaintance networks. In the survey, participants were asked to listen to each audio file, and indicate which of them represented the more active person. The order of presentation was randomised. The survey included the researchers' contact information, notified participants that no personal information would be collected, and that data collection complied with the General Data Protection Regulation of the European Union. The stimuli can be listened on Twitter: `https://twitter.com/listeningsurvey` and Facebook: `https://www.facebook.com/ListeningsurveyJYU`.

3. RESULTS AND DISCUSSION

The methods described in this report are firstly, a system that processes accelerometry data of daily activity, resulting in multigranular hierarchical segmentation akin to musical structure. The second method is a program devised as a proof of concept, to demonstrate a possible musical sonification of the daily activity data utilising the segmentation obtained. The resulting sonification has, by design, one main property, which is that there is a clear association between sonic events and daily activity. The perceptual assessment of two example sonifications produced with the system described measured the extent to which a person would correctly identify the sonification for high activity data, when presented along the sonification for low activity. A total of 1847 responses were collected by a survey on the internet, of which 1225 (66.3%) correctly identified the sonification corresponding to high activity. A one-proportion z-test was performed to evaluate the statistical significance of the results, yielding $z = 14.03$, with a p-value $< 1 \times 10^{-5}$. This may be sufficient to reject the null hypothesis, suggesting that the proportion of correct responses is significant.

The described musical sonification system may be useful in public health interventions towards increasing healthy physical activity or reducing sedentary behaviour, by making a person aware of their intraday activity in an engaging manner. In practice, the musical sonification system would be part of a portable system comprising hardware and software. Such a system would record daily activity, produce the musical sonification and possibly recommend actions to the user. The hardware may be composed of already existing technologies such as miniature accelerometers and mobile computing devices like a smartphone or smartwatch. Future research shall be carried out to implement the system and test it in ecologically valid conditions. Preliminary testing shall be carried out in order to explore the extent to which the musical sonification may work as an engagement strategy, and to identify the conditions in which it may be effective. These conditions may include personal characteristics of target users such as age, personality or income, as well as environmental conditions. Also it would be useful to compare the multigranular segmentation daily profiles of users with self-reports on their activities, to assess the extent of their correspondence.

While this report describes a method for multigranular segmentation and musical sonification of intraday activity of one subject, it is trivial to expand the method to work with different data. First, instead of using classified data for the segmentation, the Activity Score may be used alone. Also instead of using a single time period, like a day, an average of several days may be used, resulting in a representation of a typical day. Furthermore, instead of using data for a single subject, a group of subjects may be used. A population may be pre-clustered in groups with homogeneous characteristics, such as age, gender, and so on. The resulting multigranular temporal segmentation may be useful to examine the typical intraday behaviour of the group. Its musical sonification will represent the group and this may open new and interesting doors for community music making. For example, daily data of a person may be uploaded to a server, where it would be combined with data of other people in their social circle. This would enable them to produce music as a group, instead of individually. This way of collaborative music-making may be a relevant avenue for exploration in further research, as it has been observed that social support through collaboration was the primary motivator for adults to maintain activity tracker use [28].

4. CONCLUSION

This study has developed a system to produce musical sonification of daily activity data recorded by wearable devices. The sonification may be used as a tool for raising awareness and behaviour change by conveying daily activity information to users in a clear and engaging way. This capability may be used in interventions to increase physical activity (i.e., total amount of bodily motion) and reduce sedentary behaviour (i.e., proportion of time sitting or lying down) in hard to reach populations such as older adults, teenagers or people with visual or learning difficulties. A key property of the musical sonification is that it shows clearly not only the overall physical activity over a period of time, but of the temporal structure within, such as commuting to work, or taking a lunch break. This property would allow someone to identify, by listening to the sonification, the times of the day they were more or less active and spent more or less time sitting. That was achieved by devising a novel multigranular temporal segmentation procedure that preserves the time relations between events.

Acknowledgments

Petri Toiviainen suggested using probability to generate musical notes.

This work was supported in part by the Finnish Cultural Foundation (Suomen Kulttuurirahasto).

The AGNES study was financially supported by the Advanced Grant from the European Research Council (grant 310526) and the Academy of Finland (grant 693045), both to Taina Rantanen. The funders had no role in the design of the study and data collection, analysis, and interpretation of data, and in writing the manuscript. The content of this article does not reflect the official opinion of the European Union. Responsibility for the information and views expressed in the article lies entirely with the authors.

5. REFERENCES

[1] S. Michie, M. Richardson, M. Johnston, C. Abraham, J. Francis, W. Hardeman, M. P. Eccles, J. Cane, and C. E. Wood, "The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions," *Annals of behavioral medicine*, vol. 46, no. 1, pp. 81–95, 2013.

[2] N. O'Keeffe, J. L. Scheid, and S. L. West, "Sedentary behavior and the use of wearable technology: An editorial," *AInternational Journal of Environmental Research and Public Health* 17(12), p. 4181, 2020.

[3] R. Daryabeygi-Khotbehsara, S. M. Shariful Islam, D. Dunstan, J. McVicar, M. Abdelrazek, and R. Maddison, "Smartphone-based interventions to reduce sedentary behavior and promote physical activity using integrated dynamic models: Systematic review," *J Med Internet Res*, vol. 23, no. 9, p. e26315, Sep 2021. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/34515637

[4] F. Monteiro-Guerra, O. Rivera-Romero, L. Fernandez-Luque, and B. Caulfield, "Personalization in real-time physical activity coaching using mobile applications: A scoping review," *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 6, pp. 1738–1751, 2020.

[5] R. T. Larsen, V. Wagner, C. B. Korfitsen, C. Keller, C. B. Juhl, H. Langberg, and J. Christensen, "Effectiveness of physical activity monitors in adults: systematic review and meta-analysis," *BMJ*, vol. 376, 2022. [Online]. Available: https://www.bmj.com/content/376/bmj-2021-068047

[6] K.-J. Brickwood, G. Watson, J. O'Brien, and A. D. Williams, "Consumer-based wearable activity trackers increase physical activity participation: Systematic review and meta-analysis," *JMIR Mhealth Uhealth*, vol. 7, no. 4, p. e11819, Apr 2019. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/30977740

[7] S. A. Buckingham, A. J. Williams, K. Morrissey, L. Price, and J. Harrison, "Mobile health interventions to promote physical activity and reduce sedentary behaviour in the workplace: A systematic review," *Digital health*, vol. 5, 2019. [Online]. Available: https://doi.org/10.1177/2055207619839883

[8] M. I. Cajita, C. E. Kline, L. E. Burke, E. G. Bigini, and C. C. Imes, "Feasible but not yet efficacious: a scoping review of wearable activity monitors in interventions targeting physical activity, sedentary behavior, and sleep," *Current Epidemiology Reports*, vol. 7, no. 1, pp. 25–38, 2020. [Online]. Available: https://doi.org/10.1007/s40471-020-00229-2

[9] J. Y.-W. Liu, P. P.-K. Kor, C. P.-Y. Chan, R. Y.-C. Kwan, and D. S.-K. Cheung, "The effectiveness of a wearable activity tracker (wat)-based intervention to improve physical activity levels in sedentary older adults: A systematic review and meta-analysis," *Archives of Gerontology and Geriatrics*, vol. 91, p. 104211, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167494320302053

[10] A. V. Creaser, S. A. Clemes, S. Costa, J. Hall, N. D. Ridgers, S. E. Barber, and D. D. Bingham, "The acceptability, feasibility, and effectiveness of wearable activity trackers for increasing physical activity in children and adolescents: A systematic review," *International Journal of Environmental Research and Public Health*, vol. 18, no. 12, 2021. [Online]. Available: https://www.mdpi.com/1660-4601/18/12/6211

[11] W. Wang, J. Cheng, W. Song, and Y. Shen, "The effectiveness of wearable devices as physical activity interventions for preventing and treating obesity in children and adolescents: Systematic review and meta-analysis," *JMIR Mhealth Uhealth*, vol. 10, no. 4, p. e32435, Apr 2022. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/35394447

[12] S. Chastin and M. Granat, "Methods for objective measure, quantification and analysis of sedentary behaviour and inactivity," *Gait & Posture*, vol. 31, no. 1, pp. 82–86, 2010. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S096663620900602X

[13] D. J. Levitin, P. Chordia, and V. Menon, "Musical rhythm spectra from bach to joplin obey a 1/f power law," *Proceedings of the National Academy of Sciences*, vol. 109, no. 10, pp. 3716–3720, 2012. [Online]. Available: https://www.pnas.org/doi/abs/10.1073/pnas.1113828109

[14] N. Schaffert, T. B. Janzen, K. Mattes, and M. H. Thaut, "A review on the relationship between sound and movement in sports and rehabilitation," *Frontiers in psychology*, vol. 10, p. 244, 2019.

[15] J. Ley-Flores, L. T. Vidal, N. Berthouze, A. Singh, F. Bevilacqua, and A. Tajadura-Jiménez, "Soniband: Understanding the effects of metaphorical movement sonifications on body perception and physical activity," in *CHI '21 Conference, Yokohama, Japan, May 8-13, 2021*, Y. Kitamura, A. Quigley, K. Isbister, T. Igarashi, P. Bjørn, and S. M. Drucker, Eds. ACM, 2021, pp. 521:1–521:16. [Online]. Available: https://doi.org/10.1145/3411764.3445558

[16] A. Krasnoskulov, "Family album: How does your daily activity sound?" in *ICAD Conference, Northumbria University*, April 2019. [Online]. Available: https://www.researchgate.net/profile/Alex-Krasnoskulov/publication/341071964_FAMILY_ALBUM_HOW_DOES_YOUR_DAILY_ACTIVITY_SOUND

[17] M. Last and A. Usyskin (Gorelik), *Listen to the Sound of Data*. Cham: Springer International Publishing, 2015, pp. 419–446. [Online]. Available: https://doi.org/10.1007/978-3-319-14998-1_19

[18] P. Vickers and R. Höldrich, "Direct segmented sonification of characteristic features of the data domain," *arXiv*. [Online]. Available: https://arxiv.org/abs/1711.11368 2017.

[19] T. Rantanen, M. Saajanaho, L. Karavirta, S. Siltanen, M. Rantakokko, A. Viljanen, T. Rantalainen, K. Pynnönen, A. Karvonen, I. Lisko, L. Palmberg, J. Eronen, E.-M. Palonen, T. Hinrichs, M. Kauppinen, K. Kokko, and E. Portegijs, "Active aging –resilience and external support as modifiers of the disablement outcome: Agnes cohort study protocol," *BMC Public Health*, vol. 18, no. 1, p. 565, 2018. [Online]. Available: https://doi.org/10.1186/s12889-018-5487-5

[20] E. Portegijs, L. Karavirta, M. Saajanaho, T. Rantalainen, and T. Rantanen, "Assessing physical performance and physical activity in large population-based aging studies: home-based assessments or visits to the research center?" *BMC Public Health*, vol. 19, no. 1, p. 1570, 2019. [Online]. Available: https://doi.org/10.1186/s12889-019-7869-8

[21] H. Vähä-Ypyä, T. Vasankari, P. Husu, A. Mänttäri, T. Vuorimaa, J. Suni, and H. Sievänen, "Validation of cut-points for evaluating the intensity of physical activity with accelerometry-based mean amplitude deviation (mad)," *PloS one*, vol. 10, no. 8, p. e0134813, 2015.

[22] T. Rantalainen, K. Koivunen, E. Portegijs, T. Rantanen, L. Palmberg, L. Karavirta, and S. Chastin, "Is complexity of daily activity associated with physical function and life space mobility among older adults?" *Medicine and science in sports and exercise*, February 2022. [Online]. Available: https://doi.org/10.1249/MSS.0000000000002883

[23] J. Foote, "Automatic audio segmentation using a measure of audio novelty," in *2000 IEEE International Conference on Multimedia and Expo. ICME2000. Proceedings. Latest Advances in the Fast Changing World of Multimedia (Cat. No.00TH8532)*, vol. 1, 2000, pp. 452–455 vol.1.

[24] J. I. Mendoza and M. R. Thompson, "Modelling perceived segmentation of bodily gestures induced by music," in *ESCOM 2017 Conference*, E. Van Dyck, Ed. Ghent University, 2017, pp. 128–133.

[25] J. Rodrigues, P. Probst, and H. Gamboa, "Tssummarize: A visual strategy to summarize biosignals," in *2021 Seventh International conference on Bio Signals, Images, and Instrumentation (ICBSII)*, 2021, pp. 1–6.

[26] D. Nanchen, M. J. Leening, I. Locatelli, J. Cornuz, J. A. Kors, J. Heeringa, J. W. Deckers, A. Hofman, O. H. Franco, B. H. C. Stricker *et al.*, "Resting heart rate and the risk of heart failure in healthy adults: the rotterdam study," *Circulation: Heart Failure*, vol. 6, no. 3, pp. 403–410, 2013.

[27] B. Burger, M. R. Thompson, G. Luck, S. H. Saarikallio, and P. Toiviainen, "Hunting for the beat in the body: on period and phase locking in music-induced movement," *Frontiers in Human Neuroscience*, vol. 8, no. 903, 2014. [Online]. Available: https://doi.org/10.3389/fnhum.2014.00903

[28] A. Kononova, L. Li, K. Kamp, M. Bowen, R. Rikard, S. Cotten, and W. Peng, "The use of wearable activity trackers among older adults: Focus group study of tracker perceptions, motivators, and barriers in the maintenance stage of behavior change," *JMIR Mhealth Uhealth*, vol. 7, no. 4, p. e9832, Apr 2019. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/30950807

Experimenting techniques for sonic implicit interactions: a real time sonification of body-textile heat exchange with sound augmented fabrics

Vincenzo Madaghiele
KTH Royal Institute of Technology
vmad@kth.se

Sandra Pauletto
KTH Royal Institute of Technology
pauletto@kth.se

ABSTRACT

In this paper we present our prototype of a sound augmented blanket. With this artifact we aim to investigate the potential to achieve sonic implicit interactions through auditory augmentation of fabrics. We describe the development of a blanket that sonifies the approximate temperature exchange between the body and the fabric, using sound as a medium of interaction and a carrier of information. We propose different methods for auditory augmentation of fabrics through a piezoelectric contact microphone used for movement sensing. After describing the technical development of the prototype, we discuss our early findings from a qualitative standpoint, focusing on the process of sense-making of such an artifact from an evaluation based on free exploration. Our preliminary results suggest that different auditory augmentation models encourage different affordances, and are able to provide a simple creative and aesthetic experience. The ability of the chosen sonic interaction models to effectively communicate information should however be further investigated.

1. INTRODUCTION

Smart home automation systems, based on Internet of Things (IoT) technology, are growing in popularity in recent years also due to the benefits, in terms of energy efficiency, they can facilitate. However, research [1, 2] shows that the process of deployment of distributed computing in the household can, depending on how it is implemented and presented to users, have undesirable consequences, such as legitimization of unsustainable energy practices [2] and reinforcement of gender-related power dynamics [3]. The reasons for these shortcomings might be related to the ways these interfaces are designed, but also to the ideological positions they reflect. As argued by Gaver (2002) [4], digital technologies are associated to ideologies and values from workplaces - the contexts for which they were originally designed - and are usually devised to achieve a clear unique and specific function. This should be kept in mind when designing interfaces for the home, an environment in which priorities and intended functions of objects can be radically different. Several examples from HCI propose alternative designs that critically engage in this discussion. Some authors, for example, advocate for technologies that operate at the periphery of attention [5], encourage creativity [6], playfulness and wonder [4], designing for 'downtime' [7], 'slow time' [8] and prioritize simplicity [9].

Following these premises, we experimented with sound as an interaction medium and as a carrier of information in the household through auditory augmentation of textile materials, exploring the potential of sounds to deploy a creative, un-intrusive aesthetic experience. We developed the prototype of a sound-augmented blanket that sonifies the approximate energy exchange between the body and the fabric in terms of heat, reflecting on the discourses on energy efficiency and household digitalisation by raising awareness of the hidden or un-measurable energy savings in terms of heating that can be obtained with simple actions such as covering the body with a blanket. Specifically, in this paper we explore the potential of sonic implicit interactions, interactions that are not initiated by the user explicitly [10, 11], to communicate information at the periphery of attention while embedded in fabrics. In this contribution we describe the technical realization of our artifact, the employed sonic interaction strategies, and explore the initial approach and sense-making process of users to the auditory augmentation of a regular household object, in this case a blanket [1].

2. BACKGROUND AND RELATED WORK

2.1 HCI for energy in the household

Households are receiving increasing attention in research about energy efficiency, however developing effective ways to encourage less energy-intensive behaviour is not a trivial task. Different approaches have been attempted in the field of HCI, usually exploiting recent Smart Home technologies which can provide a large amount of sensor data about different aspects of energy consumption. These data are usually presented to users in the form of numeric, often visual, eco-feedback and used to drive home automation technologies in which users take on the role of managers. Despite substantial research, the degree of these systems' effectiveness in the long term is not clear [12]. Some authors argue that the assumptions behind this approach are too simplistic to suit the operations of a complex socio-

[1] All of the code used for this paper, including the interaction mappings and the sound models, can be found in open-source at https://github.com/vincenzomadaghiele/Sound-augmented-fabrics.

ISBN: 978-91-8040-358-0

technical system such as a household [1,6]. Strengers (2014) [13] argues that the reasons for the shortcomings might be found in the premises behind their design, which are deemed unrealistic and seem to overlook what people actually do in their home.

Some examples of such energy awareness artifacts include the Power-Aware Cord [14], the Energy-Aware Clock [15], lamps that change their color based on energy consumption [16], and audiovisual installations [17]. The *Static!* research group [18] built prototypes to promote critical reflection on energy usage. Design objects can also be related to unconventional ways of harvesting energy, such as micro-scavenging devices to capture energy from mechanical, thermal or light sources [19] or hand-powered electrical devices. In some research the focus has been shifted from individual users to communities, for example exploring interfaces for communication of demand response, load shifting [20] and energy forecasts in distributed settings [21]. Another way of introducing interactions in the home environment is to exploit the affordances present in the room itself. This feedback method has been proposed by Menon et al. (2021) [22] and promotes the idea of a shift in design focus from a user-centered perspective to a space-based approach, proposing the use of animistic principles to induce in users empathy for the space. Some relevant sound-based interventions for energy efficiency include the sonification of water consumption in the shower [23], water consumption regulation through singing [24], the sonification of energy data through interaction with a carpet [25], the real-time sonification of the electric power consumption of an institute's kitchen through auditory augmentation of the space [26], the *Powerchord* project (2016) [27] and *Bird-Wattching*, power consumption sonification in the kitchen using bird sounds [28]. Additionally, novel methods to explore connections between the household soundscape and concepts of energy efficiency [29] are being investigated.

2.2 Sonic implicit interactions

As our life is more and more surrounded by digital technology, research is starting to explore how we can make human-computer interactions less demanding, for example reducing the amount of attention interfaces require. Most of the currently employed digital interfaces involve some kind of explicit user input - pressing a button, entering a command and so on. These modes of communication are defined as **explicit interactions** [30]. These kinds of interfaces do not take into account the complexity of non-verbal communication that humans use to acquire knowledge about situations as well as to interact among each other and with the world. A way to obtain more seamless and user-friendly interfaces is by developing **implicit interactions**, which are defined as "an action performed by the user that is not primarily aimed to interact with a computerised system but which such a system understands as input" [10]. Often in implicit interactions the exchange of information occurs outside the attentional foreground of the user. When assessing implicit interactions, it is important to evaluate the amount of attentional demand they require, defined as "the degree of cognitive or perceptual focalization, concentration, and consciousness required of the user" [30]. A deeper evaluation of attentional demand involves looking at aspects such as spatiality, breadth and intensity. Sound is a particularly suited means for implicit interactions, as often it does not require the complete focus of the user's attention. In our everyday lives, sounds can be implicitly felt and understood, still providing feedback with minimal attentional demand.

Not much research has been dedicated to the pursuit of sonic implicit interactions [11], however the relationship between movements and sounds, often aspects of implicit interactions, has been object of substantial research. For example, the sonic affordances of gestures have been studied by Altavilla et al. (2013) [31], who described participants opinion of sound-gesture mappings. The sonic augmentation of everyday objects has been experimented by Franinovic et al. (2011) [32]. Sound-gesture mappings have been the object of research in the Sonic Interaction Design [33] and New Interfaces for Musical Expression (NIME) communities [34].

2.3 Sonic textile interfaces

The relationship between fabrics and sound is historically significant: the sound-absorption properties of textile materials [35–37] have been exploited for sound insulation for centuries. Moreover, the sounds themselves of fabrics of different materials have been shown to be linked to auditory comfort and emotional responses at a physiological level [38]. Pauletto (2022) [11] points out that the experience of Foley in cinema can be of great inspiration when dealing with sonic interactions, as bodily interactions with different textiles are certainly a large part of the Foley sonic palette. Elblaus et al. (2015) [39] developed an interactive sound-augmented garment to investigate the functional aestehtics that emerged with the introduction of sonic affordances, noting how the presence of sound changed the design priorities and the usage of the garment. Additionally, the works of Nabil et al. (2021) [40] and Prindl et al. (2020) [41], who are developing fabric speakers for application on clothes, are of great interest for this project.

3. MOTIVATIONS AND CONTEXT

In this project, we are developing sound-augmented fabrics which sonify real-time heat exchange between themselves and the body, with the aim to make users aware of energy exchanges and savings that are not usually measured and quantified in their daily life. More generally, these artifacts are tools to question the role of computing technologies in the household, proposing an alternative way to display data in this environment which is flexible, playful and leaves a degree of ambiguity. Specifically, in this paper we focus on the potential of sound to communicate information in the household through subtle, sporadic interactions, how can these interactions be embedded in textile materials through piezoelectric elements, and users' early approach to auditory augmentations of such materials.

We developed two alternative sonic augmentation models by using the normal sound of the blanket as input, sensed with a piezolectric contact microphone, such that the sounds would coordinate with the movements of the fabric. In our initial evaluation we started to explore the trade-off between clarity of information display and aesthetic acceptability of the sound model, with the aim of it to be integrated in daily practices and routines. The blanket proved to be an appropriate tool for this investigation, since it is extremely simple and flexible, it can cover a great range of functions and case uses.

4. SYSTEM DESIGN

We decided to utilise the natural sound of the body against the moving fabric as primary sonic material for our sound models. The data being sonified in real-time is the heat being exchanged between the body and the blanket. This objective presented a number of challenges on the technical side. For example, the flexibility and constant movements of the fabric requires to build a solid and testable prototype with relatively small components, in such a way that the electronics do not condition the usage too much. The most challenging aspect in the design phase - and a main subject of our evaluation - regards the complexity of the sound models to be developed: sound models that are too complex might be difficult to interpret; on the other hand sound models perceived to be too simple might be not engaging enough.

4.1 Physical prototyping

The base element of our prototype was a normal house blanket of polyester material of dimensions 114 cm by 152 cm. The electronic components were sewed to the fabric such that they were securely attached and resisted to the movement. As can be seen in Figure 1, all of the components were on the same side of the blanket, which was intended to be the side in contact with the body. We used a Bela Mini embedded platform [42] to process the input from a LM35 temperature sensor and a piezoelectric disk with the Pure Data programming language [43]. Since our aim was to augment the natural sound of the object, we decided to place the speaker directly onto the blanket. The sonic output from the Bela platform was therefore routed to a small 4 Ohm speaker, amplified through a Mono 2.5W Class D Audio Amplifier [44]. The Bela was powered by a small LiPo battery.

After some experimentation, we decided to place the Bela - the heaviest component - on one corner of the blanket, so that the body movements do not destabilize the electronic connections. The piezo disk was placed in the center of the blanket, slightly shifted to one side. The temperature sensor was also placed close to the center of the blanket, with the assumption that it would be the closest point to the body. Choosing the position of the sensor was challenging, since it is difficult to foresee how different users would wear the blanket. We placed the speaker facing the side of the electronics, closer to the longer edge of the blanket.

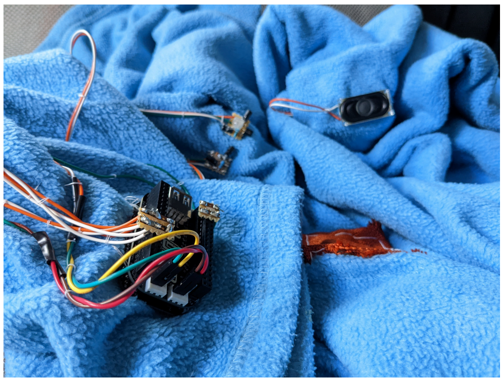

Figure 1. Sound-augmented blanket prototype. The piezo disk is inside the orange patch, which was tightly stitched to the blanket to maintain close contact.

4.2 Body heat and energy exchange

Heat exchange between the human body and the environment is a complex phenomenon which is a combination of multiple physical processes. Heat is generated in the human body by metabolic processes and it is exchanged with the environment through conduction, convection, evaporation and radiation. The amount of heat exchanged between the body and the environment can be approximated with Equation 1 [45, 46].

$$Q_i = h_i A_D(\overline{T_s} - T_o) \tag{1}$$

In Equation 1, Q_i is the amount of heat exchanged between the body and the environment at conditions of $\overline{T_s}$ average body surface temperature and T_o environmental temperature - in our case the supposed temperature of the blanket. In our software implementation, the average environmental temperature was approximated to $T_o = 30$ °C, while $\overline{T_s}$ is the temperature retrieved by the LM35 sensor, which we associate to the body temperature. h_i is the heat exchange coefficient, which is in the case of this model approximated to just the free convection value, which amount to $2.3W/m^2K$. A_D is in this case *Du Bois' body surface* approximation [47], as in Equation 2, in which m is the body mass in kilograms [kg] and h is the body height in meters [m].

$$A_D = 0.202m^{0.425}h^{0.725} \tag{2}$$

The values of m and h were approximated to the average weight and height in the Swedish population, respectively $1.74m$ and $74.3kg$ [2].

4.3 Sonic interactions and sonification design

We used the sound signal from a piezoelectric contact microphone to detect movements in the fabric, activating different sound models accordingly. The direct sound coming from the piezo disk was pre-processed as in Figure 2: to adjust the sensitivity of the piezo disk, the incoming sound was multiplied by a constant (the amount of sensitivity) and then processed through a hyperbolic tangent function,

[2] Information from Dagens Nyether - Svenskarna längre och tyngre.

ISBN: 978-91-8040-358-0

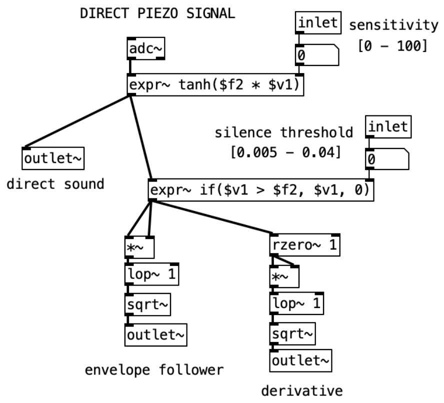

Figure 2. Digital pre-processing of the audio signal from the piezo disk.

which allowed to amplify the lower-volume components - corresponding to interactions that were further from the disk - without increasing the amplitude of the higher volume sounds. An adjustable threshold was applied to the signal to filter out background noise when the fabric is not moving. The envelope of the sound and its derivative were also calculated.

Our sonification models were inspired by recent research on shared semantic connotation of sound attributes [48, 49]. Specifically, we focused on the definition of *warm* sounds, as described by Rosi et al. (2020) [48]:

> A warm sound seems to be a low-pitched or mid-low-pitched sound. It gives a feeling of spectral richness in the mid-low frequencies. It has a rather soft attack and it is a fairly pleasant sound for the listener, giving a sensation of envelopment.

In our sonifications, a higher amount of heat exchange was expressed by a *warmer* sound. We experimented with the following techniques, which may be employed together or separately. The strategies to achieve the *warm* sonification metaphor are different for the two models, as is the way in which sound responds to interactions with the fabric.

1. **Real-time granular synthesis (Model 1):** the piezo audio signal is used as the input of a real-time granular synthesizer [3] whose parameters are mapped to the heat data. The *warm* sonification is in this case obtained by mapping increases in heat exchange to higher octave transposition of the sound and a shorter live-recording window. Higher heat was also mapped to lower frequency cutoff of a low pass filter at the end of the chain, to obtain a smoother sound with less high frequency components.

2. **Envelope-follower (Model 2):** the audio signal from the piezo element is used to drive an envelope follower which controls the activation of a synthesizer. To obtain an organic interaction, the derivative of the signal controls small pitch variations. The heat data influences the parameters of the synthesizer employed - in our case a bank of sawtooth oscillators with different amounts of detuning. The *warm* metaphor is obtained by increasing the amplitude of the oscillators that are detuned down as the heat increases, and decreasing the amplitude of higher harmonics to obtain a richer spectrum with more low-end components.

Since the envelope has been found to be an important perceptual component in this semantic sound characterization [49], in all of the aforementioned cases the data influences not only the spectral qualities of the sound, but also the reactivity of the sound envelope, with *warmer* sounds corresponding to a slower attack and decay. Finally, the direct audio signal coming from the piezo was also included in the mix, to emphasize the connection between the synthesized sound and the sound produced by the blanket movements.

5. EVALUATION METHOD

We performed an initial qualitative evaluation procedure with the aim of gathering an initial understanding how people relate to our prototype. The participants were six PhD researchers in our Division, three women and three men. The test lasted approximately 25 minutes, and was run in a quiet sonic environment, the MID Media Production studio at KTH Royal Institute of Technology.

Participants were told that the sounds on the blanket would respond to their movements and the temperature between the fabric and the body, and they were given freedom to explore the object and move freely. At this stage, we were interested in evaluating the experience and interaction as a whole, so we purposely did not put great emphasis on explaining that the sound would provide information about the warmth provided by the blanket. We did not want the participants to focus specifically on whether they ”heard” the increase in warmth, instead we wanted to observe what the object suggested to them and how they experienced the interaction with it. The two sound models were evaluated separately, starting from Model 1 and then changing to Model 2. Participants were asked to freely describe their experience while moving, and then to answer the following questions:

1. *How would you describe your experience with the blanket?*
2. *Which was your favorite aspect about it?*
3. *Is there something you would change? If yes, what and why?*
4. *Which one was your preferred sound model and why?*
5. *Could you imagine using an object like this in your daily life?*

[3] The real-time granular synthesizer employed is based on an implementation by Johannes Kreidler, released at `http://pd-tutorial.com/english/ch03s07.html`.

ISBN: 978-91-8040-358-0

We focused our evaluation on the interaction component, aiming to understand if and how the augmented sound component influenced the normal interactions with the blanket, and which interactions are encouraged by the two different sound models, which yielded two significantly different aesthetic experiences. Model 1 is more reactive, organic and complex. Model 2 has a slightly lower volume, it is milder and it has a lower internal complexity.

6. DISCUSSION

6.1 Initial approaches to sonic augmentation

As participants were given no instruction on how to interact with the blanket, diverse initial approaches to the exploration of the object emerged. A preliminary remark is that the exposed electronics limited the interactions, as most participants feared breaking parts of it.

After an initial stage of careful exploration of the sound-movement relationship, two different ways of approaching more complex interactions could be identified: some participants interacted with the object as they would with a normal blanket, trying to imagine how the sound component would influence the activity they normally do with a blanket; some other participants were not interested in finding real-life use cases and they focused on the sonic output, approaching the blanket as a sound-making object and talking about it as a musical instrument. This difference was reflected in the movements they produced. In the first case, the sound was a consequence of the movements, while, in the second, the movements were adjusted to explore the sonic response. These two ways of approaching the object generated different aesthetic experiences and specific use cases emerged.

6.2 Affordances of sonic interaction models

The potential for implicit interaction can be observed from the afforded movements generated by the sonic artifact. The two sound models in this experiment promoted different interaction strategies and produced different sonic content. Based on the interactions afforded by the two models, one could notice that the milder sonic quality of Model 2 stimulated more repetitive, less demanding interactions, while the higher perceived complexity of Model 1 prompted more ample gestures and focused exploration.

The affordances of the object could be recognized from the ways in which participants interacted with it. Some people were mainly seated, delicately exploring the surface of the fabric to investigate the sonic output with the blanket on their lap, while others preferred walking around the room and moving their body in contact with the blanket. A common way of engaging with it was to wear it as a cape and walk, even though this action was not described as a way they would use a blanket in real life. Another way to explore the interactions was to stand in one spot and move the body. For example, Participant 2 tried different yoga poses to with the blanket on, and Participant 5 tried wearing the blanket on their body in different ways, including putting it over their head.

Interactions with ample, complex movements were prompted by Model 1, as it was considered more organic and reactive, while Model 2 was perceived as being quieter and more stable by participants, encouraging fewer movements and subtle, repetitive actions. For example, most participants, when experiencing sound model 2, ended up with the blanket on their lap, slowly and repetitively stroking it in proximity of the piezo disc while answering the questions at the end of the test. Participant 1 decided to lay down on the floor with the blanket on their belly, slowly breathing as the warming body and the breathing movements generated a slow and repetitive sound. Meditation was a common theme that spontaneously emerged in multiple interactions, with different participants proposing different ways of meditating with it, through movement (Participant 2), breathing (Participant 1) or as a fidgeting device (Participant 5). Another proposed use cases including using it as a low-attention aesthetic experience alternative to watching television when tired.

6.3 Ambiguity and function

The connection between sound and temperature was not perceived clearly by any of the participants, some of which were however able to implicitly understand the connection between the position of the temperature sensor and the change in sound and using it consciously to produce changes in sounds in relation to movements, such as rubbing the blanket surface on the sensor. This is probably due to the position of the sensor, which made its measurement vulnerable to movements, and to the complexity of the sound models, which made it difficult to separate the effect of the temperature from the effect of movements on the sounds - especially since this aspect was not very clearly explained to participants before the test.

Despite these technical limitations, ambiguity was indeed, at this stage of prototyping, inherent to the design. Differently from most digital technologies, the sound-augmented blanket is purposely not focused on optimizing a specific action, but rather it is communicating, expressing an information - in this case temperature - about itself and the surface it is in contact with, which is usually a human body. Some participants found this ambiguity frustrating and were asking about its function, intended use cases and the information it was trying to communicate. For example, Participant 6 stated:

> A little disorienting. The association of a blanket with the sound is what I find not intuitive. I struggle a bit to see what, or how am I to interact with this or for what purpose. [...] It's important to have some sort of information about what I am supposed to do and be able to clearly understand how it reacts and how it responds.

Nevertheless, ambiguity brought up some interesting aspects, with some participants attributing a degree of agency to the artifact, saying that ”it's angry” or ”like a cat” (Participant 4), or saying: ”I find it interesting to communicate with a blanket. I think there's something intellectually stimulating about that idea." (Participant 6). These minimal attributions of agency could be related to the ambiguity of the artifact in terms of function, combined with the reactivity of its sonic response; this aspect however needs further investigation. Not all users found the ambiguity problematic; some approached the artifact as a musical instrument or a meditation device and were just curious about imagining or inventing what functions it could have.

7. CONCLUSIONS

In this paper we presented our initial prototype of a sound-augmented blanket. We argued for different priorities in design with data for the home, proposing an alternative approach to data sonification for the household environment. Our overarching aim is to expose hidden energy that is usually not accounted for by sensor-based digital energy measurements. We have experimented with different ways to implement this design approach to create sonic implicit interactions. Our sound design was informed by recent studies about shared semantic connotation of sounds. We have then run a simple evaluation with multiple users comparing different interaction strategies, to observe the inital approach to our artifact and the sense-making process related to the sonic augmentation of a common household object.

Overall, we observed that implementing sonic augmentation can radically change the perceived affordances of a common object. The sonic characteristics of the augmentation are crucial in shaping the afforded movements, as the two alternative sound models encouraged very different interactions, leading users to imagine different use cases for this object in their everyday life. The mapping of sound to movements in the two models was satisfactory in different ways for participants, and it was successfully able to adapt to different interactions. The association between a change in sound and a change in temperature did not clearly emerge from the experience, and we speculate that this could be due to peculiarities of the current sound models and the way in which participants interacted with the object. The sound-augmented blanket proved to be a very flexible object that was able to spark curiosity on a first approach and was indeed used creatively by participants to develop a range of simple and very personal aesthetic experiences, which were different depending on the users' imagination, personality and use cases.

Future work will include the development of multiple alternative prototypes, in which different materials will be tested, and the relationship between material, movement, sound and warmth will be further investigated, as well as a longitudinal study in which we would investigate if the acceptability of our sonic interactions is sustained over a long period of time.

Acknowledgments

This project was realized in the context of the Sounds for Energy project (Project No. 51645-1), funded by the Swedish Energy Agency.

8. REFERENCES

[1] T. Hargreaves, "Beyond energy feedback," *Building Research & Information*, vol. 46, no. 3, pp. 332–342, Apr. 2018. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/09613218.2017.1356140

[2] T. Hargreaves, M. Nye, and J. Burgess, "Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term," *Energy Policy*, vol. 52, pp. 126–134, Jan. 2013. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0301421512002327

[3] T. Hargreaves and L. Middlemiss, "The importance of social relations in shaping energy demand," *Nature Energy*, vol. 5, no. 3, pp. 195–201, Mar. 2020. [Online]. Available: http://www.nature.com/articles/s41560-020-0553-5

[4] W. Gaver, "Designing for homo ludens," *I3 Magazine*, vol. 12, 01 2002.

[5] M. Weiser and J. S. Brown, "Designing calm technology," *Power Grid Journal*, p. 5, 1995.

[6] Y. Strengers, *Smart Energy Technologies in Everyday Life*. London: Palgrave Macmillan UK, 2013. [Online]. Available: http://link.springer.com/10.1057/9781137267054

[7] G. Leshed and P. Sengers, ""i lie to myself that i have freedom in my own schedule": Productivity tools and experiences of busyness," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 905–914. [Online]. Available: https://doi.org/10.1145/1978942.1979077

[8] W. T. Odom, A. J. Sellen, R. Banks, D. S. Kirk, T. Regan, M. Selby, J. L. Forlizzi, and J. Zimmerman, "Designing for slowness, anticipation and re-visitation: a long term field study of the photobox," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2014, pp. 1961–1970. [Online]. Available: https://dl.acm.org/doi/10.1145/2556288.2557178

[9] P. Vannini and J. Taggart, *Off the Grid: Re-Assembling Domestic Life*. Imprint Routledge, New York, 2014.

[10] A. Schmidt, "Implicit Human Computer Interaction Through Context," *Personal Technologies*, vol. 4, Jul. 1999.

[11] S. Pauletto, *Foley performance and sonic implicit interactions*, 1st ed. Focal Press, 2022, pp. 265–278. [Online]. Available: https://www.taylorfrancis.com/books/9781003008217/chapters/10.4324/9781003008217-23

[12] Y. Strengers, M. Hazas, L. Nicholls, J. Kjeldskov, and M. B. Skov, "Pursuing pleasance: Interrogating energy-intensive visions for the smart home," *International Journal of Human-Computer Studies*, vol. 136, p. 102379, 2020. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1071581919301454

[13] Y. Strengers, "Smart energy in everyday life: are you designing for resource man?" *Interactions*, vol. 21, no. 4, pp. 24–31, Jul. 2014. [Online]. Available: https://dl.acm.org/doi/10.1145/2621931

[14] A. Gustafsson and M. Gyllenswärd, "The power-aware cord: energy awareness through ambient information display," in *CHI '05 Extended Abstracts on Human Factors in Computing Systems*. Portland OR USA: ACM, Apr. 2005, pp. 1423–1426. [Online]. Available: https://dl.acm.org/doi/10.1145/1056808.1056932

[15] L. Broms, C. Katzeff, M. Bång, A. Nyblom, S. I. Hjelm, and K. Ehrnberger, "Coffee maker patterns and the design of energy feedback artefacts," in *Proceedings of the 8th ACM Conference on Designing Interactive Systems - DIS '10*. Aarhus, Denmark: ACM Press, 2010, p. 93. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1858171.1858191

[16] J. Pierce and E. Paulos, "Materializing energy," in *Proceedings of the 8th ACM Conference on Designing Interactive Systems - DIS '10*. Aarhus, Denmark: ACM Press, 2010, p. 113. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1858171.1858193

[17] D. Lockton, G. Crawford, D. Singh, and S. Wu, "Electric Acoustic: Exploring Energy Through Sonic & Vibration Displays," in *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. Glasgow Scotland Uk: ACM, May 2019, pp. 1–6. [Online]. Available: https://dl.acm.org/doi/10.1145/3290607.3313014

[18] S. Backlund, M. Gyllenswärd, A. Gustafsson, S. Ilstedt, R. Mazé, and J. Redström, "Static! the aesthetics of energy in everyday things," in *DRS Wonderground conference*, 10 2006.

[19] J. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, no. 1, pp. 18–27, 2005.

[20] A. Woodruff, J. Hasbrouck, and S. Augustin, "A bright green perspective on sustainable choices," in *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*. Florence, Italy: ACM Press, 2008, p. 313. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1357054.1357109

[21] W. Simm, M. A. Ferrario, A. Friday, P. Newman, S. Forshaw, M. Hazas, and A. Dix, "Tiree Energy Pulse: Exploring Renewable Energy Forecasts on the Edge of the Grid," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Seoul Republic of Korea: ACM, Apr. 2015, pp. 1965–1974. [Online]. Available: https://dl.acm.org/doi/10.1145/2702123.2702285

[22] A. R. Menon, B. Hedin, and E. Eriksson, "Expanding Affective Computing Paradigms Through Animistic Design Principles," in *Human-Computer Interaction – INTERACT 2021*, C. Ardito, R. Lanzilotti, A. Malizia, H. Petrie, A. Piccinno, G. Desolda, and K. Inkpen, Eds. Cham: Springer International Publishing, 2021, pp. 115–135.

[23] J. Hammerschmidt, R. Tünnermann, and T. Hermann, "Infodrops: Sonification for enhanced awareness of resource consumption in the shower," in *ICAD*. International Conference on Auditory Display (ICAD), 2013.

[24] Y. Seznec and S. Pauletto, "The singing shower: A melody-sensitive interface for physical interaction and efficient energy consumption," in *Sound and Music Computing Conference, June 2022*, 2022.

[25] V. Madaghiele and S. Pauletto, "The sonic carpet: real-time feedback of energy consumption and emission data through sonic interaction design," in *27nd International Conference on Auditory Display (ICAD)*, Online, 06 2022.

[26] K. Groß-Vogt, M. Weger, R. Höldrich, T. Hermann, T. Bovermann, and S. Reichmann, "Augmentation of an institute's kitchen: An ambient auditory display of electric power consumption," *Proceedings of the 24th International Conference on Auditory Display - ICAD 2018*, 2018.

[27] D. Lockton, F. Bowden, and C. Matthews, "Powerchord: Exploring Ambient Audio Feedback on Energy Use," in *Living Labs*, D. V. Keyson, O. Guerra-Santin, and D. Lockton, Eds. Cham: Springer International Publishing, 2017, pp. 297–308. [Online]. Available: http://link.springer.com/10.1007/978-3-319-33527-8_23

[28] D. Lockton, F. Bowden, C. Brass, and R. Gheerawo, "Bird-wattching: Exploring sonification of home electricity use with birdsong," in *SoniHED: Conference on Sonification of Health and Environmental Data*, 09 2014.

[29] Y. Seznec and S. Pauletto, "Towards a workshop methodology for involving non-experts in the sonic interaction design process: Connecting household sounds and energy," in *International Conference on Auditory Display*, 2022.

[30] W. Ju and L. Leifer, "The Design of Implicit Interactions: Making Interactive Systems Less Obnoxious," *Design Issues*, vol. 24, no. 3, pp. 72–84, Jul. 2008. [Online]. Available: https://direct.mit.edu/desi/article/24/3/72-84/60190

[31] A. Altavilla, B. Caramiaux, and A. Tanaka, "Towards Gestural Sonic Affordances," *New Interfaces for Musical Expression (NIME)*, p. 5, 2013.

[32] K. Franinovic, "The flo)(ps: Negotiating between habitual and explorative gestures," *New Interfaces for Musical Expression (NIME)*, 10 2011.

[33] D. Rocchesso, S. Serafin, F. Behrendt, N. Bernardini, R. Bresin, G. Eckel, K. Franinovic, T. Hermann, S. Pauletto, P. Susini, and Y. Visell, "Sonic interaction design: sound, information and experience," *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, 2008.

[34] T. Magnusson, "Affordances and constraints in screen-based musical instruments," in *Proceedings of the 4th Nordic conference on Human-computer interaction changing roles - NordiCHI '06*. Oslo, Norway: ACM Press, 2006, pp. 441–444. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1182475.1182532

[35] K. Hoshiba, S. Hirata, and H. Hachiya, "Measurement of ultrasonic transmission attenuation characteristics of canvas fabric," *Acoustical Science and Technology*, vol. 36, no. 2, pp. 171–174, 2015. [Online]. Available: https://www.jstage.jst.go.jp/article/ast/36/2/36_E1451/_article

[36] T. S. Niroshan, J. W. A. Madushika, and S. N. Niles, "Analysis of the effect of fabric parameters on sound related properties of blended fabrics," *IESL Annual Conference 2020*, p. 10, 2020.

[37] F. Shahani, P. Soltani, and M. Zarrebini, "The analysis of acoustic characteristics and sound absorption coefficient of needle punched nonwoven fabrics," *Journal of Engineered Fibers and Fabrics*, vol. 9, no. 2, p. 155892501400900, 2014. [Online]. Available: http://journals.sagepub.com/doi/10.1177/155892501400900210

[38] J. Cho, E. Yi, and G. Cho, "Physiological responses evoked by fabric sounds and related mechanical and acoustical properties," *Textile Research Journal*, vol. 71, no. 12, pp. 1068–1073, 2001. [Online]. Available: http://journals.sagepub.com/doi/10.1177/004051750107101206

[39] L. Elblaus, V. Tsaknaki, V. Lewandowski, and R. Bresin, "Nebula: An Interactive Garment Designed for Functional Aesthetics," in *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. Seoul Republic of Korea: ACM, Apr. 2015, pp. 275–278. [Online]. Available: https://dl.acm.org/doi/10.1145/2702613.2725454

[40] S. Nabil, L. Jones, and A. Girouard, "Soft Speakers: Digital Embroidering of DIY Customizable Fabric Actuators," in *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. Salzburg Austria: ACM, Feb. 2021, pp. 1–12. [Online]. Available: https://dl.acm.org/doi/10.1145/3430524.3440630

[41] T. Preindl, C. Honnet, A. Pointner, R. Aigner, J. A. Paradiso, and M. Haller, "Sonoflex: Embroidered Speakers Without Permanent Magnets," in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. Virtual Event USA: ACM, Oct. 2020, pp. 675–685. [Online]. Available: https://dl.acm.org/doi/10.1145/3379337.3415888

[42] G. Moro, A. Bin, R. H. Jack, C. Heinrichs, and A. P. McPherson, "Making High-Performance Embedded Instruments with Bela and Pure Data," *International Conference of Live Interfaces*, p. 5, 2016.

[43] M. S. Puckette *et al.*, "Pure data," in *ICMC*, 1997.

[44] A. Industries. Adafruit mono 2.5w class d audio amplifier - pam8302. [Online]. Available: https://www.adafruit.com/product/2130

[45] A. Shitzer and R. C. Eberhart, "Heat transfer in medicine and biology," in *Heat Transfer in Medicine and Biology*, 2012.

[46] Thermopedia. Body (human) heat transfer. [Online]. Available: https://thermopedia.com/content/587

[47] B. Shuter and A. Aslani, "Body surface area: Du Bois and Du Bois revisited," *European Journal of Applied Physiology*, vol. 82, no. 3, pp. 250–254, Jun. 2000. [Online]. Available: http://link.springer.com/10.1007/s004210050679

[48] V. Rosi, O. Houix, N. Misdariis, and P. Susini, "Uncovering the Meaning of Four Semantic Attributes of Sound : Bright, Rough, Round and Warm," *e-Forum Acusticum 2020*, p. 8 pages, 2020, artwork Size: 8 pages Publisher: e-Forum Acusticum 2020. [Online]. Available: https://hal.archives-ouvertes.fr/FA2020/hal-03234057

[49] ——, "Investigating the shared meaning of metaphorical sound attributes: bright, warm, round, and rough," PsyArXiv, preprint, Mar. 2022. [Online]. Available: https://osf.io/tzk8r

CAN YOU HEAR THE EARTH BREATHING? TRANSLATION AND DISCLOSURE IN SOUND ART DATA SONIFICATION

Stephen Williams
Chalmers University of Technology
stephengarywilliams@
gmail.com

ABSTRACT

Lacey introduces the concept of disclosure in sound art meaning the sounds that are taken for granted and making them focus of attention [1]. The many technical and artistic choices made during data sonfication and performance are not usually made visible to the audience. This paper presents a reflection on the concept of disclosure in sonification sound art performance - 'Can you hear the earth breathing?' performed in Berlin in July 2022. The performance included both sonified data representing different elements of the CO2 system along with field recordings. Design and performance of the work included multiple methods of disclosing technical and artistic choices.

We reflect on our intentions with disclosure, how we put that into practice, and how we might further such practices in the future. We close by proposing new research directions including questions such as how might sonification choices impact the affect of a sound art performance? Does disclosure help or hinder public engagement with the work? And what techniques might we use – practically and aesthetically – to disclose our sonification choices

1. INTRODUCTION

Lacey introduces the concept of disclosure in sound art in terms of taking the sounds that are taken for granted and making them focus of attention. Lacey notes "disclosure is the approach that demonstrates the soundscape always exceeds perceptive capacity, and that beyond the dominant affective forces that shape everyday experience, there are hidden qualities waiting to be revealed" [1, p.166]. These qualities are also hidden inside soundscape recordings, sound art performances, and especially in sonification. Choices are made about how and when to record, how to edit into a piece, what processing and effects are applied, how "raw" data is turned into pitch, gate, velocity, timbre and which instruments are employed to play back the sounds; everything that happens, in Sterna's term, "before sound" [2]. However, these many technical and artistic choices are not usually made visible to the audience. This paper presents a reflection on an attempt to embed disclosure into the artistic performance of a sound art data sonification piece – our intentions, our thoughts on how well our intentions were met, and thoughts for future performances.

This paper builds on experiences in designing and performing a sustainability focused sound art performance 'Can you hear the earth breathing?' in Berlin in July 2022[1]. This work attempted to bring to life sustainability cycles and the natural (and un-natural) systems that lie beneath the production and absorption of carbon dioxide – key elements of energy production and consumption. The performance included both sonified data representing different elements of the CO2 system along with field recordings. Design and performance of the work included multiple methods of disclosing technical and artistic choices. First, the event program included a 'Cast of Characters' that presented the system elements, sources of data used for sonification, along with sound and field recording sources. In addition, the program included detail on sound processing, playback instruments, and the tools and processes used to harness/employ/highlight feedback processes between system elements and sound. Second, the sound performance was supplemented with video that mixes footage of CO2 sources and sinks, "technical" slides and screenshots that illustrated the scaffolding and "disclose" the behind the scenes data, sonification, and sonic choices made. These choices were illustrated with video, screenshots, and screen captures that show the process of data download and manipulation, exporting to MIDI, importing into Ableton, sound design, and processing. Each of these processes contain choices that influence the sound that is heard by the audience; challenging the conception what the audience really hearing.

2. APPROACHES TO SONIFICATION DISCLOSURE

Much has been written on aesthetics and preferences in sonification practices [3], but there is relatively little about the impact of revealing or hiding these choices from listeners. Coming from outside the sonification field, Demers notes that "many hip-hop DJs and artists view sampling as a means of insider communication with their fans, a way

[1] See https://www.drusnoise.com/can-you-hear-the-earth-breathing for video documentation of the performance

ISBN: 978-91-8040-358-0

of conveying messages to the select few informed enough to identify quotations [4, p.54]. Sterna takes a deep self-investigation into the role of technologies, spaces, and the artist in her work. Referring to the piece *Currents* installed at the Oslo Opera, she remarks that:

> whilst the work's contextual ambition and its relation to the environment were also fully articulated via textual information, this does not mean that the work acted as a representation of the scientific data or that the text was a representation of the work. In this sense, the information about the work that was accessible on site in the Oslo Opera should also be considered an affective body with a capacity to affect, to establish new affective connections, and thus found new transversal machines in complicity with the work's body and the multiplicity of bodies of visitors who together were establishing new affective relations [2, p.158].

It is these 'new affective relations' that are worthy of further exploration in sonification and sound art studies.

Sonification choices have been mapped along various spectra and frameworks such as Ars Musica by Vickers and Hogg [5] and indexality from Keefe et al. [6]. Rather than focusing on the relative merits of sonification choices, this paper explores how those choices can be made transparent to listeners and audiences. While there are several ways to engage with these choices, this work is grounded in the theoretical concept of disclosure - taking the sounds that are taken for granted and making them focus of attention. Lacey notes that this approach "reveals the complexity of the urban environment, which is typically hidden by the dominating sounds of the everyday" [1, p.166] And further that "disclosure is the approach that demonstrates the soundscape always exceeds perceptive capacity, and that beyond the dominant affective forces that shape everyday experience, there are hidden qualities waiting to be revealed" [ibid]. These qualities are also hidden inside soundscape recordings and especially in sonification. However, these many technical and artistic choices are not usually made visible to the audience.

Disclosure also applies to the methods used to plan, prepare, and perform this work such as sound processing (including choices made for sonification) and playback instruments, tools and processes used to highlight feedback loops between system elements and sound, and the process of data download and manipulation, exporting to MIDI, importing into Ableton, sound design, and processing. Each of these processes contain choices that influence the sound that is heard by the audience

3. EXPERIENCES FROM CAN YOU HEAR THE EARTH BREATHING?

3.1 Performance details

The breath is one of the most fundamental, most personal, parts of any living being. Mauss notes that "man's first and most natural technical object, and at the same time technical means, is his body" [9, p.13]. In the 1950s, Charles Keeling started measuring CO2 concentrations in the atmosphere at the Moana Loa Observatory. These measurements showed for the first time how the planet itself is breathing as forests in the Northern hemisphere grow leaves in the Spring absorbing CO2, then release it back into the atmosphere in the Fall when leaves drop off the trees. The data also shows the accelerating rise of CO2 in the atmosphere over time. The Scripps Observatory makes available this global CO2 concentration data. The data is available in time slices from one week to 800,000 years. Data back to 1958 is from the Moana Loa Observatory, further back the data is taken from analyzing ice cores.

This performance included both sonified data representing different elements of the CO2 system along with field recordings (see Figures 1 and 2). Sounds are performed with modular synthesizers, samplers, biofeedback pads connected to plants, FX pedals, and feedback. The sound performance is supplemented with video that mixes footage of CO2 sources and sinks, "technical" slides and screenshots that illustrate the scaffolding and "disclose" the behind the scenes data, sonification, and sonic choices made, and video that shows changes over time such as time lapse video of a forest changing over a year and CO2 data animations.

Figure 1: Live performance in Berlin July 2022. Photo: Lisa Knolle used with permission

Through sonification processes in collaboration with Data Artist Lisa Knolle, we translate each of the raw CO2 data sets into MIDI. The curves make for some strangely discordant arpeggios. And the MIDI data can also be used to modulate and change the sounds over time. But the measurement of CO2 is just the end product. CO2 is produced by (generally) manmade processes such as fossil fuel production and use, industry, transportation, and many forms of human activity. CO2 is absorbed by plants and oceans which are in turn affected by the sun in the form of solar radiation.

ISBN: 978-91-8040-358-0

Figure 2: Live performance in Berlin July 2022. Photo: Lisa Knolle used with permission

While described here as discrete elements, in reality, these are all part of interconnected systems and feedback loops [10,11], These systems are working at nested scales ranging from the microscopic, to the personal, the local, city, region, and global, ranging in time scales from the current moment to 800k years ago and into an uncertain future(s).

3.2 Sonification Process

3.2.1 Data sourcing and analysis

Data was sourced through the Keeling Curve website which collects multiple data record sets from different research sources and data collection methods. Data is available in several time series ranging from daily to the last 800,000 years [12]. Data from 1958 to present was measured at the Moana Loa observatory while pre-historic data is measured from ice core analysis. The different data sources and time scales also had different measurement periods. Recent data is available in regular, predictable 15 minute or daily increments. Older data includes measurements at unpredictable intervals ranging from 100 to over 500 years between data points.

3.2.2 Translation to MIDI

Once the data was downloaded, additional choices were made on MIDI translation. Decisions on upper and lower note bounds, note lengths and note velocity were standardized across time scale data. This meant that while the resulting MIDI file for 800k years was longer than the file for one year, it was not 800k times longer! Data points were from the different time scale data sets were standardized at 1/16 note length. Pitch was standardized with a range of MIDI note values 24 - 100 (C0 - C7) with the lowest recorded CO2 value at 24 and the highest at 100. Note velocity (loudness) was scaled between a range of 0 – 127 using the same values as pitch. The data was then run through a conversion process using Python to generate a set of 10 MIDI files representing the different time scales of data available.

3.2.3 Playback and Recording

MIDI files were then playing back from Ableton Live through various modular synthesizers to generate different sounds with a range of timbres as loops. In addition, the two and 100k year MIDI files were slowed down so the audio recordings extended over 20 minutes. Recordings were made of "clean" audio along with versions with effects and modulations coming from houseplants connected via TENS Sensor Cable to the modular system. Pads of this cable, placed at two separate areas on the plant, measure the changes in surface conductance to generate control voltage and gate signals that are used for modulating sounds (see Figure 3).

Figure 3: TENS Sensor pads on plant leaves. Photo by author

3.2.4 Composition and Arrangement

An early decision was made to create a composition with multiple themes in roughly historical order. The themes were Water, Plants (both that absorb CO2), Animals, Industry, and Traffic (producing CO2). Field recordings representing the themes were arranged on a timeline reflecting the roughly 30 minutes allocated to the performance. Recorded loops were then placed on the arrangement timeline. The initial concept was to arrange the time-scale recordings in chronological order. However, after first experimenting with the recordings, we noted this would result in approximately 29 minutes of drone or slowly changing notes followed by a rapidly rising arpeggio in the final seconds of the performance. The choice was made to instead layer multiple time scale recordings rather than a strict chronological order, in effect changing the indexality balance of the sonification. The final arrangement was exported as 7 x 30-minute tracks that were loaded into the Octatrack sampler for live playback.

3.3 Live performance

The live performance instruments included a Soma Lyra-8 synthesizer, modular synthesizer with voices, effects, MIDI file playback, and clock dividers, houseplant connected to the Instruo Scion module via biofeedback pads, FX pedals (looper, delay, and reverb), mixer, and

ISBN: 978-91-8040-358-0

MIDI/CV controller (see Figure 4). The performance started by playing a sine wave at 419Hz representing the CO2 concentration level of 419.1 ppm measured that day. The pre-recorded arrangement of loops and field recordings played from the Octatrack. Throughout the performance, different MIDI files were played back triggering different modular synthesizer voices and (using the clock dividers) at different playback rates. Pre-recorded and improvised sounds were processed through FX pedals.

Table 1 provides a complete list of sound sources and playback methods. Note that this information was provided to audiences as described in the following section on Disclosure Methods.

System element	Data source	Sound source	Processing	Playback Instrument
CO2 Concentration levels (today)	Scripps Observatory		CO2 ppm converted to Hz	Sine wave through modular synthesizer and Soma Lyra 8
Breath		Field recordings	Time stretch and looped	Octatrack
CO2 Concentration levels (historical)	Scripps Observatory		Convert data to MIDI, stretch/normalize time scales. Send MIDI to modular synthesizers	Octatrack playing recorded loops Live modular synthesizers
Water		Field recordings		Octatrack
Plants		Field recordings	Time stretch and looped	Octatrack
Solar radiation	Earth/local magnetic field			Soma Synths Ether
Plants	Plant on table		Convert electricity resistance to pitch, gate, modulation with Instruo Scion	Modular synthesizers
Industrial activity		Field recordings		Octatrack
Transportation		Field recordings		Octatrack
Interconnected systems			Feedback loops	Mixer, Soma Cosmos looping pedal, Erica Synths Zen Delay

Table 1: Cast of characters (in order of appearance). From exhibition text provided to attendees 22 & 23 July 2022

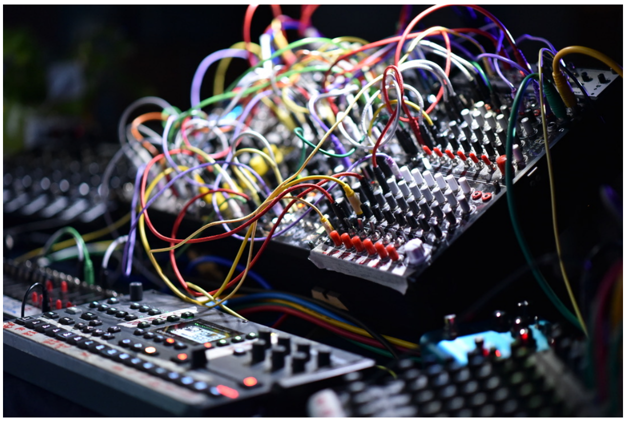

Figure 4: Live performance gear setup. Photo: Lisa Knolle

3.4 Disclosure process

We took two paths towards disclosure for the performance. First was utilizing photo and video documentation of the sonification process (see Figures 5 and 6). The video included photos of early design sketches, video capture of data manipulation and translation processes, audio recording, and arrangement. While the video clips were not long or detailed enough to show the audience every aspect of our choices, we attempted to communicate that *some* choices had been made even though we were not clear about the impacts of those choices. The video also included footage representing the themes (e.g. ocean and waves for the Water theme, city streetscapes and in-car video for the Traffic) theme.

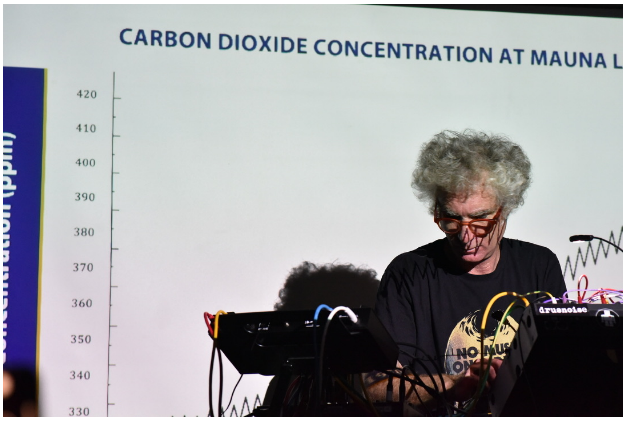

Figure 5: Background video showing CO2 concentration data from Mauna Loa. Photo: Lisa Knolle

ISBN: 978-91-8040-358-0

Figure 6: Example of 'thematic' video in background. Photo: Lisa Knolle

In addition, we created a two-page information sheet given to audience members at the venue. This information sheet described (in brief detail) the aims of the composition, the sonification process, and the technical and aesthetic choices made (see Figure 7 in Appendix 1). We also included a table that listed the 'cast of characters' in the composition including the data or field recording source, processing applied, and instruments used to play back the data or field recording (see Table 1 above)

3.5 Audience experience

As this was the first performance of the piece, we were not able to capture audience experience in a robust and statistically valid format. Anecdotal feedback was mostly positive with a few notable exceptions. Two attendees commented that the sounds (they had not paid attention to the video) made them feel very anxious. Another attendee commented that the video was 'competing' for attention and found it difficult to both listen to the sound art performance and watch the video at the same time.

4. DISCUSSION AND REFLECTION

The premiere performance of 'Can you hear the Earth breathing?' and this paper raise more questions than answers. We attempted through this piece to both provide a different way of engaging with climate change data and to disclose the sonification process to the audience.

Rather than present quantitative data, here we reflect on the disclosure choices made. The video documentation and overall presentation was well received by audiences and formed an integral part of the artistic experience. However, the video was attempting to both convey data sonification process details and convey the underlying sustainability message through footage of forests, oceans, traffic, etc. Instead of conveying both messages, the video in fact may have done neither. The disclosure process was presented 'free floating' in that the video content was not labelled or directly tied to the sound art performance. This was a deliberate choice to avoid becoming a lecture video but may have undercut our goals. In addition, the content of the video was visually stimulating and engaging – a good thing! – but also drew audience attention away from the sound. Audiences questioned whether this was a video presentation with accompanying audio or a sound art performance with supporting visuals. Artistically this is not a binary choice but the reminder of the split attention of audiences was valuable. In future performances, we will attempt to more closely link the visual and sound so that audiences will (or might?) be able to more clearly see and hear the relationship between what is happening on screen to what is being performed. Another approach is to leverage the power of live coding. In this approach, the data manipulation and translation would be performed live and projected on screen before the data is played back through synthesizers. This would bring the data sonification *process* itself into the performance and make the data-sound link clear to the audience.

The text handout was a more direct explanation of the sonification and performance process. Audiences attending sound art performances are accustomed to reading an artist statement or program description, so the handout fit within that context. We question how much the technical description contributed to the disclosure process or served more as a technical guide. Future options for the text disclosure might be to format more closely to, for example, an orchestral program which is organized into movements that highlight performers/vocalists for each, often including text of what the vocalists are singing. This would provide a more clear and direct explanation of the data-sound connection. However, as we discuss further below, asking audiences to read text during a sound art and video performance may distance the audience even further.

4.1 Open questions

The experience of the performance and our reflections lead us to a series of open questions that were not resolved in our exploration. Open questions can be clustered into three broad categories. First is around knowledge and understanding – did our disclosure processes help audiences understand the context of the piece and the sonification process itself? The underlying assumptions of this question (i.e. that audiences should understand the context and meaning of an artistic work) is of course one that has been questioned for many years. In a sustainability context, we may also question the value and efficacy of art as simply another instrumental communication tool [13,14].

Second is the effect of such disclosures on the aesthetic experience. To what extent is disclosure part of the artistic experience? Does disclosure impact the wonder and experience of an audience? Audience members and installation attendees arrive with wildly different understandings and expectations of the artistic work. An analogy might be an orchestral performance where some attendees may only know the title of work performed, others may read every page of the program, while others may bring along their own copy of the score to follow along. We wish to explore the balance between disclosure and maintaining engagement with the artistic work.

ISBN: 978-91-8040-358-0

Third is returning to Sterna's suggestion that disclosure may serve a role in "establishing new affective relations" between the audience, the work, and the space. However, how and why this may happen is the biggest open question that we will continue to explore through future projects and accompanying research.

4.2 Areas of further research

Further research will be valuable for the open questions discussed above. In addition, it would be fruitful to explore differing methods of disclosure. These can range from the typical wall mounted information found in galleries and installation spaces all over the world, to more detailed information sheets, artist talks, project web sites and documentation videos. Of particular interest is how disclosure might become part of the performance or installation itself (through audio, video, text or integration of all of the above) with the potential to take on not just a communicative role but an aesthetically engaging one. It would also be of interest to study the impact of when disclosures are made to the audience. For example, disclosing the context and sonification process before (for example by an artist talk or lecture), during (through the practices we have described), or after (e.g. with a post-show artist talk, audience Q&A, or dialogue session) a performance may have very different impacts on an audience.

5. CONCLUSIONS

'Can you hear the Earth breathing?' was an experiment in sonification of climate data. More importantly perhaps, the piece was an experiment in disclosure of data sonification processes. In challenging the seeming simple question of "what does climate change sound like?", we hoped to spark a different affective relationship with the audience that experienced the performance.

As we note, the performance – and our reflection – has raised more questions that answers. However, we hope that this paper will spark further discussion within the sound art and data sonification communities. We also hope to encourage further exploration of the challenges and opportunities of disclosure in sound art data sonification.

6. ACKNOWLEDGMENTS

First in acknowledgement must be the work of Data Artist Lisa Knolle who took on the task of sourcing, analyzing and translating the climate data into MIDI and also captured documentary photos of the performance. Steve Burnside from Alive.BLN provided video production and documentation. We are grateful to faculty and students at the Berlin School of Sound 'Practice-Based Sound Studies for Installation and Performance' course for feedback on early development of the sound art performance and review of an early draft of this paper. We are also grateful to ACUD Macht Neu in Berlin for providing space for the premiere performance of 'Can you hear the Earth breathing?' and to everyone who attended.

7. REFERENCES

[1] J. Lacey, Sonic Rupture. Bloomsbury Academic, 2017.

[2] Å, Stjerna, Before Sound: Transversal Processes in Site-Specific Sonic Practice, University of Gothenburg. Faculty of Fine, Applied and Performing Arts, 2018.

[3] S. Barrass and P.Vickers, Sonification Design and Aesthetics in T. Hermann, A. Hunt, and J.G. Neuhoff, The Sonification Handbook, COST Office and Logos Verlag Berlin GmbH, 2011, pp. 145-172.

[4] J. Demers, Listening through the noise: The Aesthetics of Experimental Electronic Music, Oxford University Press, 2010.

[5] P. Vickers and B. Hogg. Sonification abstraite/sonification concrète: An 'æsthetic perspective space' for classifying auditory displays in the ars musica domain. In T. Stockman, et al. editors, ICAD 2006 - The 12th Meeting of the International Conference on Auditory Display, pages 210–216, London, UK, 20–23 June 2006.

[6] D. F. Keefe, D. B. Karelitz, E. L. Vote, and D. H. Laidlaw. Artistic collaboration in designing VR visualizations. IEEE Computer Graphics and Applications, 25(2):18–23, March–April 2005.

[7] A. Engström & Å Stjerna. Sound Art or Klangkunst? "A reading of the German and English literature on sound art" in Organized Sound, 2009, 14(1), 11–18. http://doi.org/10.1017/S135577180900003X

[8] C. Andersson, "Complexity science and sustainability transitions" in Environmental Innovation and Societal Transitions, 2014, 11, 50–53. http://doi.org/10.1016/j.eist.2014.03.001

[9] F. Westley, P. Olsson, C. Folke, T. Homer-Dixon, H. Vredenburg, D. Loorbach, et al., Tipping Toward Sustainability: Emerging Pathways of Transformation in Ambio, 2011, 40(7), 762–780. http://doi.org/10.1007/s13280-011-0186-9

[10] C. D. Keeling, S. C. Piper, R. B. Bacastow, M. Wahlen, T. P. Whorf, M. Heimann, and H. A. Meijer, "Exchanges of atmospheric CO2 and 13CO2 with the terrestrial biosphere and oceans from 1978 to 2000". I. Global aspects, SIO Reference Series, No. 01-06, Scripps Institution of Oceanography, San Diego, 88 pages, 2001.

[11] D. Maggs & J. Robinson, Sustainability in an Imaginary World: Art and the Question of Agency. Routledge, 2021.

[12] R. Bendor, D. Maggs, R. Peake, J. Robinson & S. Williams, "The imaginary worlds of sustainability: observations from an interactive art installation", Ecology and Society, Vol 22, Issue 2, 2021.

8. APPENDIX 1

Breaths / LP / Deaths

Can you hear the Earth breathing?

Performance (30 minutes) at ACUD Macht Neu 22 & 23.07.22 20.0-23.00
by Steve Williams (drusnoise)

'Can you hear the Earth breathing?' brings to life sustainability cycles and the natural (and un-natural) systems that lie beneath the production and absorption of carbon dioxide. In the 1950s, Charles Keeling started measuring CO_2 concentrations in the atmosphere at the Moana Loa Observatory. These measurements showed for the first time how the planet itself is breathing as forests in the Northern hemisphere grow leaves in the Spring absorbing CO_2, then release it back into the atmosphere in the Fall when leaves drop off the trees. This performance includes both sonified data representing different elements of the CO_2 system along with field recordings. Sounds are performed with modular synthesizers, samplers, biofeedback pads connected to plants, FX pedals, and looping feedback. The sound performance is supplemented with video that mixes footage of CO_2 sources and sinks, "technical" slides and screenshots that illustrate the scaffolding and "disclose" the behind the scenes data, sonification, and sonic choices made, and video that shows changes over time such as time lapse video of a forest changing over a year and CO_2 data animations.

The core of this work is work is grounded in the concept of *disclosure* - taking the sounds that are taken for granted and making them focus of attention. Lacey notes that this approach "reveals the complexity of the urban environment, which is typically hidden by the dominating sounds of the everyday" (2017, p. 166). And further that "disclosure is the approach that demonstrates the soundscape always exceeds perceptive capacity, and that beyond the dominant affective forces that shape everyday experience, there are hidden qualities waiting to be revealed" (ibid). These qualities are also hidden inside soundscape recordings and especially in sonification. Choices are made about how and when to record, how to edit into a piece, what processing and effects are applied, how "raw" data is turned into pitch, gate, velocity, timbre and which instruments are employed to play back the sounds.

Disclosure also applies to the methods used to plan, prepare, and perform this work. Here I present the system elements (i.e. primary components of natural systems that produce and absorb CO_2), sources of data used for sonification, sound and field recording sources. In addition, I detail sound processing (including choices made for sonification) and playback instruments. In addition, I describe tools and processes used to harness/employ/highlight feedback processes between system elements and sound. Note that there are feedback loops and implications within recording, processing, and playback too. These choices are be illustrated with video, screenshots, and screen captures that show the process of data download and manipulation, exporting to MIDI, importing into Ableton, sound design, and processing. Each of these processes contain choices that influence the sound that is heard by the audience.

Conception, recording, production, performance: Steve Williams/drusnoise
Data artist: Lisa Knolle (@lets.learn.machines)
Data source: C. D. Keeling, S. C. Piper, R. B. Bacastow, M. Wahlen, T. P. Whorf, M. Heimann, and H. A. Meijer, Exchanges of atmospheric CO_2 and $^{13}CO_2$ with the terrestrial biosphere and oceans from 1978 to 2000. I. Global aspects, SIO Reference Series, No. 01-06, Scripps Institution of Oceanography, San Diego, 88 pages, 2001.

Berlin School of Sound

ACUD MACHT NEU

Figure 7: Text provided to audience for performances in Berlin 22 & 23 July 2022

ISBN: 978-91-8040-358-0

PUSHING SUSTAINABILITY TOPICS IN SOUND TEACHING

Katharina Groß-Vogt and IOhannes m zmölnig
Institute of Electronic Music and Acoustics, University of Music and Performing Arts Graz, Austria
vogt@iem.at, zmoelnig@iem.at

ABSTRACT

In teaching in the field of sound engineering and sound design, the urgent debate of sustainability issues is still underrepresented. Some students bring along their personal commitment, while many are uninformed — amazingly, within tertiary education. We present the website of our SIDlab, i.e. Sonic Interaction Design lab [1], where we highlight projects of sonic interaction design and sonification that have a focus on sustainability. Mainly, these projects have been developed within our courses that bring the students in contact with topics originally foreign to their discipline. The main strategy of the courses is limiting the choice of topics for their sound projects to an issue related to sustainability. In the first place, the website presents the SIDlab of the Institute of Electronic Music and Acoustics, fertilizing new project ideas within the university. Furthermore, the website and this paper are intended both as a model for other educational institutions and as an evolving project of communicating environmental facts in new ways. At the moment, two members of the institute curate the collection, but a further community driven approach is envisaged.

1. MOTIVATION

When teaching university students, the knowledge of the climate crisis and the need for a fundamental system change may be assumed. Or may it not? For instance, a basic understanding of the Earth's climate system, including the anthropogene influence on it, supposedly should be part of primary and secondary curricula, as education for sustainable development [1]. Still, when we included such topics at the periphery of sound related courses, as will be discussed in this paper, it turned out that a considerable number of the attending Master's students were quite ignorant of this topic. This is the initial motivation to write this paper. It is also supported by literature, e.g. Mann et al. [2] performed measures of ecological worldview and ethical sophistication with freshman computer students, and found that these students had poor ethical understanding and were only partly pro-ecological – the authors state that a substantial shift in their worldview would be needed for them to become "sustainable practitioners". In a different perspective, Easterbrook [3] claims the need for a shift from computational thinking, where each problem has a technological solution, to systems thinking, of a more fundamental and interdisciplinary nature.

[1] https://sidlab.iem.sh

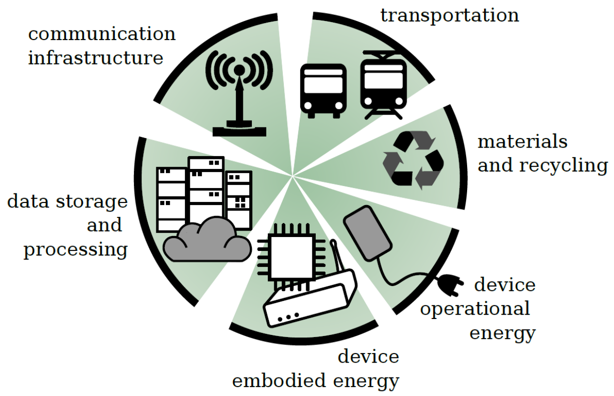

Figure 1. Key areas that contribute to the ecological footprint of the Internet of Sounds. (The proportions are not representative.) Source: [4]

Recently, Gabrielli & Turchet [4] discussed the sustainability of the internet of sounds — or, more generally, of the internet itself with a focus on sound related technologies. The authors found key areas to monitor the impact of the Internet of Sounds, as shown in Fig. 1: materials and recycling, device operational energy, device embodied energy, data storage and processing, and communication infrastructure. All of these areas contribute to the ecological footprint of the internet and related devices. In a life cycle analysis, the greenhouse gases (GHG) can be attributed to all phases of the life cycle of a product, from the extraction of raw materials, the manufacturing, and the usage to its disposal. These factors can be assessed for all areas of Fig. 1. We refer to a handful of prototypical facts resulting from such and similar analyses in the following. The internet is responsible for about 2% of worldwide CO_2 emissions — which is roughly equivalent to the carbon footprint of the aviation industry. [2] More detailed information can be found in [5], who estimate the carbon footprint of the ICT (Information and communication technology) sector with 1.4% of the total carbon footprint. Out of this, multimedia streaming has become the dominant causer with around 87% of share [6]. With respect to the focus of SoniHED, we may also refer to the energy use of the internet;

[2] https://www.ovoenergy.com/blog/green/the-carbon-footprint-of-the-internet

ISBN: 978-91-8040-358-0

for instance, the ICT sector is responsible for 3.6% of the global energy footprint [5].

In the same time, the internet may help to reduce emissions, for instance, teleconferences have a positive impact under certain preconditions, and can reduce the carbon footprint by up to 44% in academia in a study from 2013 [7] and even up to 94% in a recent study [8]. In modelling alternative scenarios for the showcase of one international research project, the total of emitted CO_2 equivalents could have been reduced to 53% of the emissions of this project by applying standard measures or to only 2.4% in a strict scenario [9]. Still, it has to be considered that, while the internet may be part of the solution, the general trend may overcome the positive effects, a rebound effect known in Economics as Jevons paradox [10].

To come closer towards a more sustainable internet, the authors of [4] certainly have no ultimate solution. Instead, they suggest simple strategies for the individual internet user and producer of sound: reusing — on the hardware side reusing equipment with similar functionality [11], on the software side creating software with backwards compatibility to older versions; consuming less in general avoiding to follow every trend; and enhancing video meetings if they substitute traveling.

General trends of Sustainability in Human Computer Interaction are explored in the field of SHCI for which a recent review is a good starting point [12].

On a meta level, the ubiquity of "the internet" can be used to share ideas and concepts of (education for) sustainable development, by giving examples on how to raise awareness for the "Grand Challenges". This ranges from stereotypical "life hacks" to full-fledged platforms to share and disseminate learning material and didactic methodologies such as the "Sustainicum Collection" [3] of the University of Natural Resources and Life Sciences, Vienna [13].

We are worried about climate change and we teach sound related courses at the Institute of Electronic Music and Acoustics (IEM), University of Music and Performing Arts Graz in Austria. Therefore, in the last years, we have explored formats to make sustainability part of this teaching. In this paper, we present a collection of mainly student (co-) works on the SIDlab website, see Fig. 2. They are part of our SIDlab, the Sonic Interaction Design lab at our institute, for which we have established and intend to further stimulate a focus on the topic of sustainability (SID4future).

2. TEACHING APPROACHES

For teaching topics related to sustainability within the disciplines of ICT, Eriksson and Pargmann [14] discuss how to actually reach students in this field. They argue that future engineers and professionals of ICT have to be sensibilized to the topic, reflecting the outcome of their future work. Furthermore, students can be engaged on a personal level. The course described in their paper included didactical elements such as the use of challenging, open questions where no simple answer can be given; a learning game where discussions in small groups can take place to explore the students' own values without ideolocially bias from the educators; and use energy as an intermediary concept to make environmental effects more graspable than abstract CO_2 units. While different authors have suggested for a more integrative way of introducing sustainability into a whole curriculum (see [14]), the minimal approach is a modular one, which can be implemented more easily. In Sweden, dedicating one course of any curriculum to sustainability issues is even required by law, but the authors of [14] state that this is not enough in the long run.

Teaching in the area of sound design, sonification, and creative coding, our advantage is the project-oriented nature of these courses. Students learn in self-chosen projects where they prepare, present and discuss sonic prototypes or algorithms. In all our lectures, the idea is to use sound as a medium to communicate *something* — a data set, an interaction feedback, an algorave. The *something* can be circumscribed to our chosen topic, i.e. sustainability. For instance, the United Nation's Sustainable Development Goals provide an extensive set of topics. [4] For sonification, the choice of data with a background of sustainability is straightforward. Sonification allows itself to be used as an objective perceptualization of scientific data. At the same time, sonification creates interest for the data, as people are still not used to "hearing data".

Students of the courses are typically at Master's level (with a few exceptions of elective subjects in Bachelor's degree) and the sizes of the classes are rather small, with about 5-15 students. They have slightly varying backgrounds as they study either electrical engineering-sound engineering, sound design, computer music, or musicology. Therefore the classes already have an interdisciplinary nature.

On the part of teaching sonic interaction design, the chapter of Rocchesso et al. in [15, ch. 4] provides a great resource. Rocchesso describes five methods for teaching sonic interaction design that can actually be adapted to other sound related studies. All of these can also be put in a specific thematic spotlight and are thus useful for conveying side-information on sustainability issues:

1. When performing sound walks outside, students can reflect on the relation of natural and human-made sounds.
2. Creating audio dramas, i.e. sound-based but "speechless" stories, can already be a project on its own, where, for instance, nature sounds can be applied.
3. In sketching and prototyping sonic interactions, basic mockups can be created to prepare for the real implementation of prototypes and test ideas; in this stage, the topic of the chosen project is already central.
4. With their fourth suggested method, Problem-Based Learning (PBL), we come to the core of the didactic strategy, that we use at the core of the SID classes. In PBL, students are regarded as active learners who collaboratively solve problems. This is especially

[3] `https://sustainicum.at/`

[4] `https://sdgs.un.org/goals`

ISBN: 978-91-8040-358-0

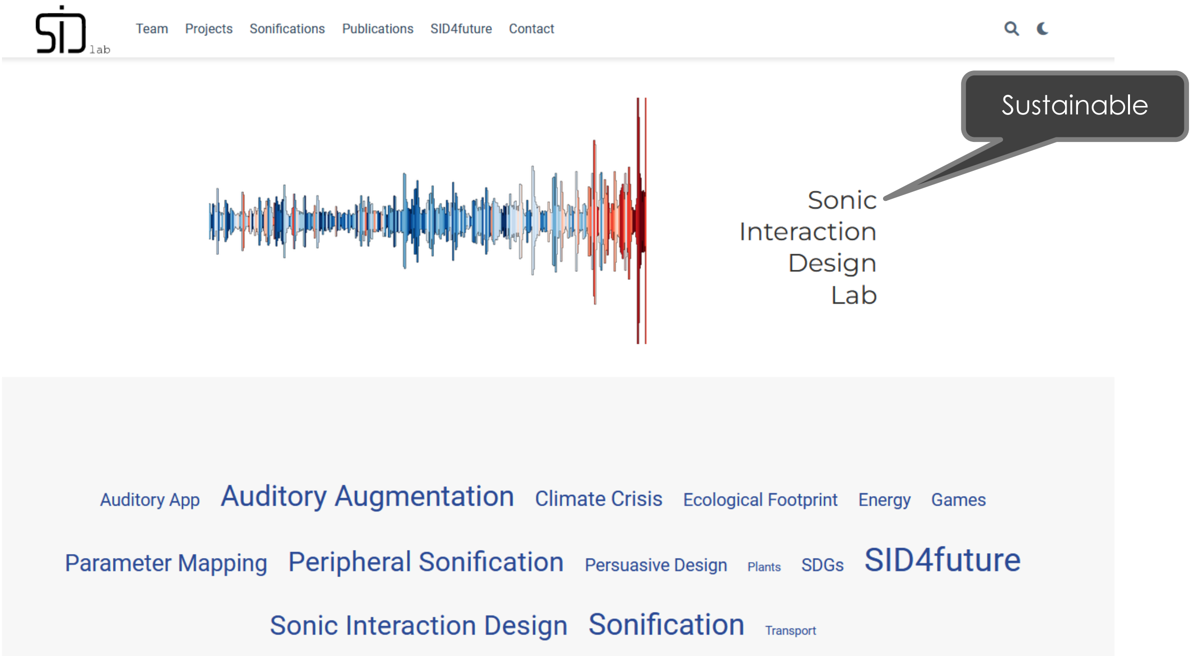

Figure 2. Screenshot of the `SIDlab` starting page. The word "sonic" is altered (blinking) with "sustainable". The SIDlab logo is based on the *Warming Stripes*, reminding of a digital sound wave. The tag cloud gives an overview over the currently added projects.

interesting when a variety of disciplines need to be incorporated to address a problem — certainly true for all systemic problems of humankind. Problems should be chosen by the students themselves, but we suggest that they make a choice within a limited field of options. Rocchesso cite [16] and [17] where PBL has proven to be particularly suitable for education dealing with design of interactive systems and multidisciplinary settings.

5. The final stage is the physical prototyping with interactive sound — which can be seen as the realization of the mockup prototypes and the PBL based preparations. Videos and sounds of some chosen prototypes are presented on the website.

Creative brainstorming methods in the area of sound design have been explored by Franinovic et al. [18]. Especially the speed dating approach proved to be fun and to open new perspectives in some of our courses. In that way, students come up with "crazy ideas" for the PBL approach. The speed dating is a series of short brainstormings in alternating pairs of students, on a pair of –equally alternating – topics. The topics' pairs are defined by a matrix of, for instance, SDGs, on one dimension, and a location on the other (the location can be anything, e.g. the tramway, the kitchen, the office desk). Students then have 2 minutes time to come up with a sound design idea in the context of the location with a focus on the respective SDG. In the following, students elect the best ideas and choose what they want to work on.

3. THE WEBSITE AND EXEMPLARY PROJECTS

On occasion of a new setup of our university's website we decided to present the results of our lab with a focus on sustainability projects. In this way, we do want to (1) fuel new projects of other student of our curricula who might become interested, (2) promote the many smaller and larger ideas of students for sounds for sustainability, and (3) inspire others in any scientific disciplines to set up similar projects in their teaching and research.

The website is a simple collection of our projects, searchable by tags, at `https://sidlab.iem.sh`. A screenshot is shown in Fig. 2. For reasons of completeness, there are also projects that do not have a focus on sustainability. In case of very prominent or larger projects, they have been added to the website. The majority, especially smaller student projects, are curated out of all projects for their focus on sustainability issues, and are tagged by "SID4future". The website is hosted by ourselves in order to have the freedom to adapt to new demands or formats of usage depending on how this collection develops in future.

In the following, a handful of exemplary projects at the lab are shortly presented; they show the range of activity from single-student works to externally funded research projects. The latter ones have their own ways of being disseminated, while the first ones, small student works, account for the critical mass of the website.

3.1 Individual student's works

These projects and the next ones (student's group works) shall constitute the core part of the website and will be

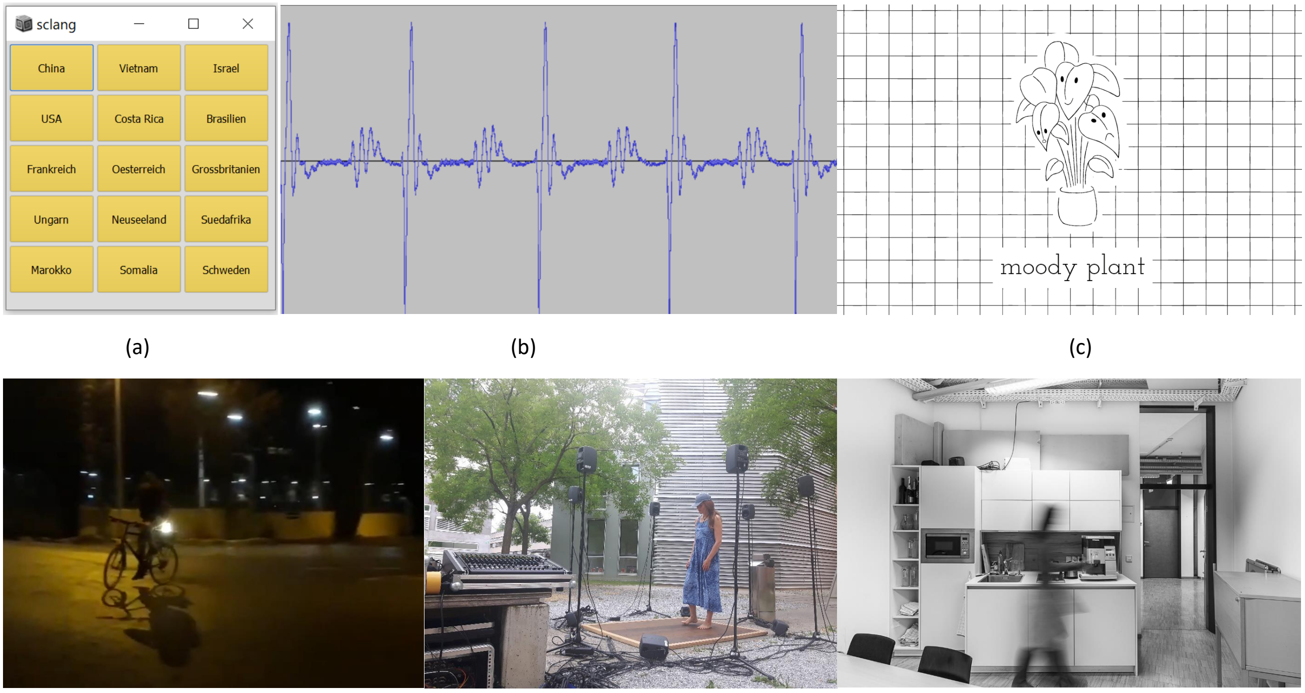

Figure 3. An exemplary selection of works presented at the SIDlab website.

added up in future. They are normally not presented elsewhere except on the SIDlab website. In the following, we present a few sonifications very shortly:

In a straightforward approach of parameter mapping, Robert Hofer mapped the amplitude of different sound samples (fire burning, water, alarm sound, wind) to the ration of types of energy sources (fossil fuels, water power, nuclear power, and wind energy). By interactively choosing different countries (see GUI in Fig. 3(a)), the listener may experience our dependence on fossil fuels.

Within a more metaphoric approach, Simon Windtner mapped the data on the worldwide temperature rise to the regularity of a human heartbeat and used a breathing sound to map precipitation data. Over the time span of the data (1880–2020), the physiological sounds appear to become more and more "unhealthy" as they speed up and become more intense (see Fig. 3(b)).

3.2 Student's group works

The only difference to the above is a more extensive nature of the projects, as several students are involved in their creation (typically, 2–5). We give two short examples:

The *Moody Plant* (see Fig. 3(c)) was developed by Hanna Brühwiler, Elisabeth Hacker, and Michael Reiter in 2021. By equipping a house plant with simple sensors for conductivity, they could measure reactions when the plant's leaves where touched. The result was an anthropomorphous design to make people more aware of, otherwise silent, plants, as co-creatures that need to be cared about. The system has been awarded a price for start-up ideas and is considered to be further developed.

The *Sonic Bicycle* (see Fig. 3(d)) was developed by Felix Wagner and Thomas Alpers in 2021. The prototype works with a motion detection sensor and a Bluetooth loudspeaker, mounted in the back of an ordinary bicycle. The faster the bike is going, the faster a beatbox rhythm is played. The prototype is a playful motivation to use bikes more (and, faster).

3.3 Student's works that have been extended beyond the class

At a few occasions, student's works have been continued over the end of the course, for instance as initial research projects where the students had the opportunity to present their research in conferences or as artworks. An example of the latter is the *Anthropocene Maze*, an interactive listening experience of our ecological footprint on selected habitats, see Fig. 3(e) and [19]. This interactive sound installation consists of a wooden floor equipped with Pyzoflex® technology and an Ambisonics audio environment. Footstep sounds are modified in the ambiance of different habitats' soundscapes that suffer from human influence. The story is following a narrative, where a future computer travel guide leads a tour into the past — which really is our present. The installation has been exhibited and evaluated at the university's campus, and was subsequently published and presented as an installation within the *DACA 2022 — data art for climate action.* [5]

[5] `https://dataclimate.org/`

3.4 The lab's research projects

The research of the staff of the institute has less focus on sustainability, as we do more basic research concerning sound synthesis and perception. Still, we try to add this focus more and more. An example is the *PilotKitchen*, depicted in Fig. 3(f), [20]. In this project we built a pilot system for sonifying the electric power consumption of our institute's kitchen — or, any room in general. The reverberation of the kitchen is changed depending on the actual consumption and its difference to a weekly baseline. If the actual consumption is low, it is mapped to a plausible kitchen reverberation that is similar to the real reverb of the small room. If it is high — as compared to the baseline — the reverberation becomes unnatural, in extreme cases the kitchen then sounded like a big cathedral. Evaluating the system gave insights on perceptibility and acceptance of auditory augmentation in a semi-home context. As outlook, such systems could be used as peripheral displays in smart homes in order to help their inhabitants to save electricity. This prototype has been followed by a more perception-oriented study to understand how many degrees of reverb can be differentiated at the periphery of our perception, i.e. while the subjects were performing a primary task without being distracted. See [21] for details.

3.5 Externally funded projects

In the externally funded research project *SysSon* [6] we explored a systematic procedure to develop sonifications, working with climate data. Within this project, a sonification platform [7] has been developed that can be used to create sonification on climate model's data to allow mainly for exploratory analyses. The project ran from 2013 to 2015 and was conducted in cooperation with the Wegener Center for Climate and Global Change, a major institute for climate research in Austria. The established research cooperation still proved very useful for a variety of smaller projects that have been realized since then.

4. CONCLUSIONS AND OUTLOOK

We motivated our paper by facts on the footprint of the internet and sound related activities to show the necessity to become active even if working within disciplines that are not directly related to climate change. From this, our major handle lies in public outreach and teaching. Both fields are especially apt for sonic displays and sonic interaction design, as new media create awareness. We discussed approaches in teaching sound, and suggest to use them in combination with topics of sustainability. In that way, both students gain a better understanding for the current crises, and ideas can be taken up to create more public awareness. For both goals we set up a website for our lab, the sonic/sustainable interaction design lab at the IEM. The website is a collection of projects of various sizes, ranging from individual student's sonifications to large research projects. The presented works are curated in a way to highlight the ones that are related to SID4future. We hope that these works inspire students to new ideas and other research institutions to similar activities.

[6] `https://sysson.kug.ac.at/index.php?id=14007`
[7] `https://sysson.iem.at/`

Acknowledgments

Many thanks to all parties involved, especially the engaged and creative students of the highlighted projects.

5. REFERENCES

[1] K. D. Madsen, "Unfolding education for sustainable development as didactic thinking and practice," *Sustainability*, vol. 5, no. 9, pp. 3771–3782, 2013. [Online]. Available: https://www.mdpi.com/2071-1050/5/9/3771

[2] S. Mann, M. Lopez, D. Lopez, and N. Smith, "Educating for ict4s: Unpacking sustainability and ethics of ict student intakes," in *Proceedings of EnviroInfo and ICT for Sustainability 2015*. Atlantis Press, 2015/09, pp. 229–241. [Online]. Available: https://doi.org/10.2991/ict4s-env-15.2015.27

[3] S. Easterbrook, "From computational thinking to systems thinking: A conceptual toolkit for sustainability computing," in *Proceedings of the 2014 conference ICT for Sustainability*. Atlantis Press, 2014/08, pp. 235–244. [Online]. Available: https://doi.org/10.2991/ict4s-14.2014.28

[4] L. Gabrielli and L. Turchet, "Towards a sustainable internet of sounds," in *Proc. of the Audio Mostly '22*, 2022.

[5] J. Malmodin and D. Lundén, "The energy and carbon footprint of the global ICT and e&m sectors 2010-2015," *Sustainability*, vol. 10, no. 9, 2018. [Online]. Available: https://www.mdpi.com/2071-1050/10/9/3027

[6] T. Tabata and T. Y. Wang, "Life cycle assessment of co2 emissions of online music and videos streaming in japan," *Applied Sciences*, vol. 11, no. 9, p. 3992, 2021.

[7] W. M. Achten, J. Almeida, and B. Muys, "Carbon footprint of science: More than flying," *Ecological indicators*, vol. 34, pp. 352–355, 2013.

[8] Y. Tao, D. Steckel, J. J. Klemeš, and F. You, "Trend towards virtual and hybrid conferences may be an effective climate change mitigation strategy." *Nat Commun*, vol. 12, no. 7324, 2021.

[9] V. Reyes-García, L. Graf, A. B. Junqueira, and C. Madrid, "Decarbonizing the academic sector: Lessons from an international research project," *Journal of Cleaner Production*, vol. 368, p. 133174, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0959652622027627

[10] L. M. Hilty, A. Köhler, F. Von Schéele, R. Zah, and T. Ruddy, "Rebound effects of progress in information technology," *Poiesis & Praxis*, vol. 4, no. 1, pp. 19–38, 2006.

[11] L. Costalonga, D. Hora, M. Pimenta, and M. Wanderley, “The Ragpicking DMI Design: The Case for Green Computer Music,” in *10th International Conference on Digital and Interactive Arts*, 2021, pp. 1–10.

[12] L. r. E. J. Hansson, T. Cerratto Pargman, and D. S. Pargman, “A decade of sustainable hci: Connecting SHCI to the sustainable development goals,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: https://doi.org/10.1145/3411764.3445069

[13] R. Kromp, “Sustainicum Collection – Lehrmaterialien für Bildung für nachhaltige Entwicklung,” *GW unterricht. Vol. 141. S. 64*, vol. 68, 2016.

[14] E. Eriksson and D. Pargman, “ICT4S reaching out: Making sustainability relevant in higher education,” in *Proceedings of the 2014 conference ICT for Sustainability*. Atlantis Press, 2014/08, pp. 40–47. [Online]. Available: https://doi.org/10.2991/ict4s-14.2014.5

[15] K. Franinovic and S. Serafin, *Sonic interaction design*. Mit Press, 2013.

[16] N. Schultz and H. P. Christensen, “Seven-step problem-based learning in an interaction design course,” *European Journal of Engineering Education*, vol. 29, no. 4, pp. 533–541, 2004.

[17] J. V. Kimmons and P. R. Spruiell, “Using problem-based learning in a multidisciplinary setting,” *Clothing and Textiles Research Journal*, vol. 23, no. 4, pp. 385–395, 2005.

[18] K. Franinovic, L. Gaye, and F. Behrendt, “Exploring sonic interaction with artifacts in everyday contexts,” in *Proceedings of the 14th International Conference on Auditory Display*, vol. 40. Citeseer, 2008, pp. 1–4.

[19] A. Kobzar, B. Brands, V. Drack, J. Leitner, K. Groß-Vogt, and C. Amon, “The anthropocene maze – an interactive listening experience of our ecological footprint on selected habitats,” in *Proc. of the DACA - data art for climate action*, 2022. [Online]. Available: http://dataclimate.org/daca-2022-proceedingscatalogue/

[20] K. Groß-Vogt, M. Weger, R. Höldrich, T. Hermann, T. Bovermann, and S. Reichmann, “Augmentation of an institute’s kitchen: An ambient auditory display of electric power consumption,” in *Proc. of the ICAD*. Georgia Institute of Technology, 2018.

[21] K. Groß-Vogt, M. Weger, M. Frank, and R. Höldrich, “Peripheral sonification by means of virtual room acoustics,” *Computer Music Journal*, vol. 44, no. 1, pp. 71–88, 2020.

ISBN: 978-91-8040-358-0

Interactive Sonification for Health and Energy using ChucK and Unity

Yichun Zhao
University of Victoria
yichunzhao@uvic.ca

George Tzanetakis
University of Victoria
gtzan@ieee.org

ABSTRACT

Sonification can provide valuable insights about data but most existing approaches are not designed to be controlled by the user in an interactive fashion. Interactions enable the designer of the sonification to more rapidly experiment with sound design and allow the sonification to be modified in real-time by interacting with various control parameters. In this paper, we describe two case studies of interactive sonification that utilize publicly available datasets that have been described recently in the International Conference on Auditory Display (ICAD). They are from the health and energy domains: electroencephalogram (EEG) alpha wave data and air pollutant data consisting of nitrogen dioxide, sulfur dioxide, carbon monoxide, and ozone. We show how these sonfications can be re-created to support interaction utilizing a general interactive sonification framework built using ChucK, Unity, and Chunity. In addition to supporting typical sonification methods that are common in existing sonification toolkits, our framework introduces novel methods such as supporting discrete events, interleaved playback of multiple data streams for comparison, and using frequency modulation (FM) synthesis in terms of one data attribute modulating another. We also describe how these new functionalities can be used to improve the sonification experience of the two datasets we have investigated.

1. INTRODUCTION

The importance of interactivity in sonification is generally accepted, but many existing data sonifications are still limited in terms of their interactivity and when they are interactive they tend to be singular prototype systems created specifically for a particular dataset. In fact, most of the recent sonifications of health or energy data published in the International Conference on Auditory Display (ICAD) are limited in this sense, as summarized in Section 2.2.

In this paper, we introduce an interactive sonification framework which builds upon existing toolkits and utilizes new functionalities to improve the sonification experience so that it is easier for domain experts with no programming experience to adopt and is reusable, avoiding reinventing the wheel. Moreover, we focus on sonifications of data in the health and energy domains in recent years as case studies to demonstrate the framework.

2. RELATED WORKS

In this section, we summarize works related to sonification toolkits with graphical user interfaces (GUI) and sonifications of energy and health data from recent years, and organize them in chronological order. We then discuss specifically our contributions in relation to previous works.

2.1 Sonification Toolkits with GUI

The Sonification Sandbox was developed by Walker and Cothran [1], allowing mappings between data attributes and timbre, pitch, volume and pan. Percussive context could also be added to the sonification. This was later upgraded by Davison and Walker [2] to rebuild the system to allow integration into other systems and saving descriptions of the sonification representation in various formats. Pauletto and Hunt [3] described an interactive sonification toolkit where sonification methods could be changed rapidly, and the current position in the dataset sonified is controllable. The sonification research group at Berne University of the Arts (BUA) [4] developed project Sonifier which supports frequency modulation (FM) and notably sonified electroencephalogram (EEG), magnetic resonance imaging (MRI), and seismological data.

More recently, the sonification workstation was developed by Phillips and Cabrera [5] to increase the accessibility and ease of use of a general sonification toolkit especially for users who might not be familiar to sonification. Similarly, the WebAudioXML Sonification Toolkit (WAST) built by Lindertorp and Falkenberg [6] aims for the same goal but in the web browser. WAST was also evaluated with user studies. The paper suggested that involving instruments and features from Digital Audio Workstations (DAW) would create a more creative environment which is something that has also informed our research. The company Sonify has also developed a web tool, TwoTone [7] that allows users to map data to musical pitches with different instruments.

2.2 Sonifications of Health and Energy Data

Works sonifying data related to health from ICAD include CardioSounds which is a portable system developed by Blanco et al. [8] to sonify electrocardiogram data and enhance rhythmic details in real-time to diagnose and monitor cardiac pathologies. Winters et al. [9] sonified the internal workings of artificial neural networks for melanoma

ISBN: 978-91-8040-358-0

diagnosis such as the activation of nodes, seriousness of the conditions of lesions classified, and cluster centers after dimensionality reduction. A sound art piece *The Alchemy of Chaos* was created by Mitchell et al. [10] to discretely sonify episodes from Touette's syndrome and captured both the contextual information of the episodes and the aesthetics of sound design. Kantan et al. [11] developed a domain-specific prototype system that involves interactive "real-time mapping control and physical modelling-based musical sonification" to support biofeedback for movement rehabilitation. Frid et al. created sonification of the interactions between a COVID-19 patient and a medical team by using electrocardiogram (ECG) data and mapping between properties of heart signals and harmonic tension and dissonance.

In terms of data related to energy, Cowden and Dosiek [12] presented auditory representations of power grid voltage data with easy-to-use audification and sonification algorithms. Groß-Vogt et al. [13] sonified the electric power consumption of an institute's kitchen by altering the reverberation of the kitchen based on consumption data.

As for the two sonifications that are used as case studies later in this paper, one of them is by Steffert et al. [14] who used amplitude and frequency modulation to sonify EEG data and proposed an evaluation method. The other is by St Pierre and Droumeva [15] who used FM and granular synthesis to sonify air pollutant data with the notion of "harmonic identities". They are further described in Section 4.

2.3 Our Work

Our interactive sonification framework generally builds upon the work on accessible sonification toolkits [5] [6] by utilizing modern technology for both the audio generation and the interactive interface to support portability across platforms and software maintenance. It is created using ChucK [16] (a programming language creating concurrent strongly-timed procedural audio), Unity [17] (a game engine creating the user interface, and spatializing and visualizing the sonification in this project), and Chunity [18] (which combines both ChucK and Unity allowing integrated audiovisual programming). In comparison to a domain-specific system [11], our work can be used for any dataset in a tabular form. In addition to supporting most functionalities of existing sonification toolkits, we also further support de Campo's sonification design space [19] to distinguish between continuous and discrete sonification. Furthermore, we introduce novel methods which existing toolkits do not support, such as interleavingly playing sonifications of multiple data streams for comparison, using FM synthesis in terms of one data attribute modulating another, and supporting synchronized interactive 2D and 3D visualizations.

3. INTERACTIVE SONIFICATION

The general design of our work is inspired by a typical DAW where the top is general controls of data loading (selecting datasets and normalization type), audio playing (play rate, starting, stopping, and resetting), and mode selection (switching to audio visualization). After a dataset is loaded, each data attribute or stream populates a track view, and similar to how a typical DAW looks like, the tracks are horizontal rows stacked together. The top half of a track consists of sound synthesis controls and parameter mappings, and the bottom half is a visual plot of the data stream as a line graph. Each track is randomly and differently coloured not only for aesthetic reasons but also for easy identification of data attributes in the audio visualization view. When the sonification is played, there is a cursor for each track on the visual plot to signify the location of the current data point being sonified. Lastly, when multiple tracks are populated exceeding the current field of view, the tracks could be scrolled up and down to view on the desired track(s). Figure 1 shows a sample layout.

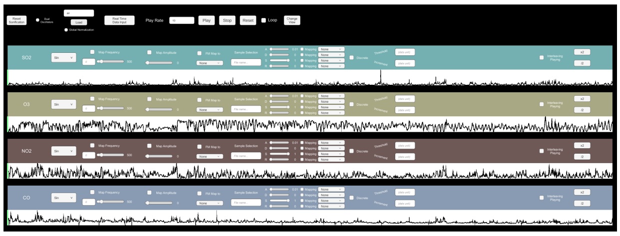

Figure 1. The main track layout of the sonification framework when a dataset consisting of 4 data attributes is first loaded with the default configuration.

4. CASE STUDIES

In this section, we describe the two datasets that we sonify as case studies: EEG and air quality data. They are chosen because they are the most recent publicly available data from ICAD that are related to the health and energy domains. In each subsection of them, we describe the replication of sonification from the original paper, and discuss how our sonification framework achieves more with interactivity and new functionalities that are appropriate to the context of the dataset.

4.1 EEG Data Sonification

The EEG data was recorded under the two conditions of having the subject's eyes being open and eyes being close. The alpha waves were separated into low alpha and high alpha frequency, resulting in 4 sets of data for eyes-closed and eyes-open in the high and low alpha frequency conditions [14].

4.1.1 Replication of Sonification

After the data is loaded, a minimum frequency can be set in the interactive interface and it is 261.6 or 523.2 Hz from the original paper. The play rate is set to 2 milliseconds per data point to replicate the duration of 1 minute for the sonification. For frequency modulation, the data is configured to be mapped to the frequency of the sine oscillator in the interface; The range of frequencies is set to 600 Hz using a slider as seen in Figure 2. For amplitude modulation, the

Figure 2. Frequency modulation by EEG data with a minimum frequency of 261.6 Hz and a frequency range of 600 Hz.

data is mapped to the amplitude (the data is automatically scaled to be between 0 and 1 inclusively); The frequency stays at 261.6 or 523.2 Hz.

4.1.2 Achieving More

The original paper used Pure Data [20] to create the sonification and needs to load each set of data individually. In our work, all 4 sets could be loaded together and muted or unmuted separately, eliminating the need to reload data when the user wishes to focus on another dataset. Moreover, the sound design parameters and the play rate can be modified in real-time. Although this was partially achieved in the original Pure Data patch, certain controls are easier in our work such as directly setting the frequency range instead of constructing calculations. Changing the oscillator type in Pure Data would require the removal of the oscillator object and re-patching to a different one, but here we could just select another sound option in the drop-down menu for the sound source.

The original sonification is only continuous, but our work also supports discrete events. Two parameters "threshold" and "increment" could be set. Threshold refers to a certain value for the data to reach in order to produce a sound, as demonstrated in this video [1] ; Increment refers to a value that could be added or subtracted to the threshold every time the threshold is reached. An example is shown in Figure 3 where the threshold is set to 1 and the increment is 2 in this case; A discrete sound is played when the EEG data reaches 1, 3, 5, ..., and so on. (This is especially useful when the data has an increasing or decreasing trend.) Conditional discrete events allow more ways of analyzing EEG data through sonification to only hear a sound when certain conditions are met, instead of hearing the whole audio. Furthermore, discrete events not only support the use of oscillators but also user-defined audio samples by selecting the "Sample" option in the sound source drop-down menu and specifying the file name of the sample. The audio samples can be further manipulated by changing their speed and amplitude and thus affecting their pitch and volume. This enables any kind of audio to be played as the user desires.

Envelope controls are useful here to help create more pleasantly discrete-sounding sonification and the controls of attack, decay, sustain, and release (ADSR) is partially automated to configure their values depending on the nature of sonification being discrete or continuous. The attack is set to 0.01 by default to prevent potential clipping at the start of a sound. If continuous, then the sustain is set to 1; If discrete, the sustain is set to 0 with the decay being 0.2, also shown in Figure 3. These values could be further modified by the user interactively with sliders. If the user changes the default values, the values are saved when the user toggles between discrete and continuous.

[1] `https://youtu.be/EHmMB9ddKAU`

Figure 3. The configuration of discrete sonification of the EEG data with envelope parameters modifying the shape of the sound and threshold parameters to specify the conditions for discrete sounds to be played.

4.2 Air Quality Data Sonification

The air quality data is collected from the Canadian provincial websites for British Columbia, Alberta, and Ontario for the year of 2014. Sulphur dioxide (SO_2), ozone (O_3), nitrogen dioxide (NO_2), carbon monoxide (CO), and particulate matter ($PM_{2.5}$) were used as the metrics of air quality, and the data was sampled hourly.

4.2.1 Replication of Sonification

Like the original paper, each data attribute is configured to have 2 sine oscillators being the modulator and carrier for FM synthesis. The data is mapped to the modulation index (or the amplitude of the modulator), and also the amplitude (of the carrier) [15] [21]. The panning of the audio for each data attribute is configured in the audio visualization view by dragging the speaker to the left or right of the middle listener. The colours of the speakers are matched with those of the tracks.

Since the exact configurations of FM synthesis used for SO_2, O_3, NO_2, CO were not provided, we follow the general descriptions closely resulting in the following configurations. As shown in Figure 4, for SO_2 the fundamental frequency is 750 Hz set for the carrier, and the carrier:modulator frequency ratio is 1:4, resulting in a frequency of 3000 Hz for the modulator; The panning is on the far left and configured in Figure 5 by dragging the speaker of the corresponding colour of the data attribute and moving its location to the left or right for different panning. For O_3, the fundamental frequency is 300 Hz and the modulator frequency is 1500 Hz with a ratio of 1:5; The panning is on the far right. NO_2 has a fundamental frequency of 1500 Hz and modulator frequency of 4500 Hz with a ratio of 1:3, and its panning is on the slight right. Lastly, CO has a fundamental frequency of 50 Hz and modulator frequency of 100 Hz with a ratio of 1:2, and its panning is on the slight left. The play rate is the same where each data point gets sonified for 0.2 seconds. Another attribute $PM_{2.5}$ is not sonified because granular synthesis was used but is not currently supported by our sonification framework.

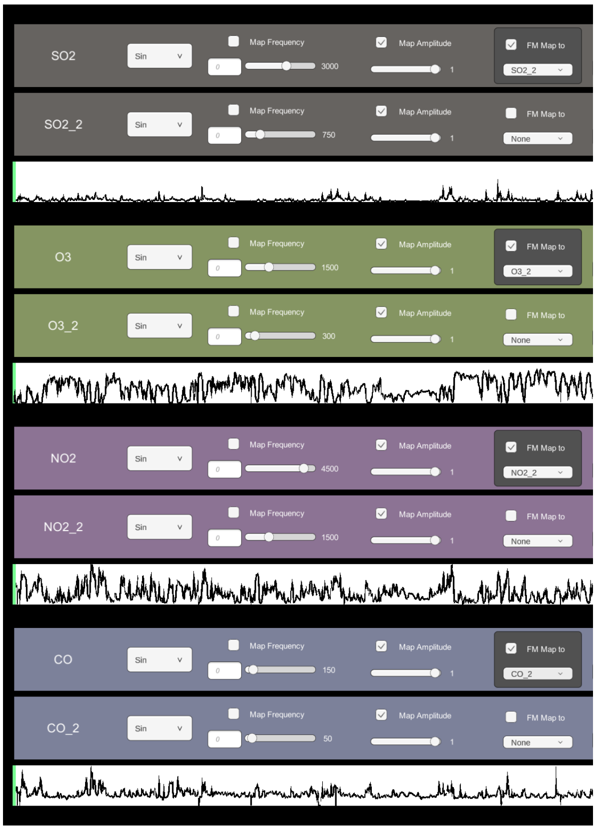

Figure 4. The configuration to replicate the sound design of the sonification of air quality data, where the first row of each data attribute serves as the modulator, and the second is the carrier for FM synthesis.

4.2.2 *Achieving More*

The FM synthesis in the original paper was used to synthesize sounds specific to each data attribute. In our work, we could link the oscillators (being the modulator or the carrier) together without any restrictions and control them interactively in real-time, by mapping any oscillator to another target oscillator. This allows the data attaching the modulator to influence the sound texture of the carrier which could be attached to another data attribute. This is especially helpful to hear how one data stream affects another. For example, if we want to investigate the relationship between CO and NO_2 data, and if data attribute CO has a modulator (where the data is mapped to the amplitude) connected to a carrier of data attribute NO_2 (where the data is mapped to the frequency), the sonification sounds more metallic in its texture when the value of CO data gets larger (and vice versa), and sounds higher in its pitch when the value of the NO_2 data gets larger (and vice versa) to hear the potential effect of changes in CO levels on NO_2 levels. This is demonstrated in this video [2]. More oscillators can be connected forming an algorithm to produce more complex sound involving multiple data attributes. (Sample selection is not supported because FM synthesis only applies to oscillators.) Moreover, the type of oscillators could also be chosen dynamically, instead of being restricted to only sine waves.

[2] https://youtu.be/yjWBIGvqRbU

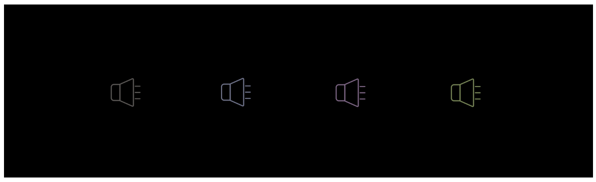

Figure 5. In the audio visualization view, the panning of the sonification of air quality data could be replicated and configured. The center is the location of the listener. The differently coloured speakers refer to the data attributes with the corresponding colours, and their locations indicate the panning; For example, the green speaker refers to O_3 in Figure 4 and its panning location is on the far right.

Interleaved playback is also supported in our framework allowing easier comparison between data attributes. If this mode is turned on for the selected data attributes, the attributes are sonified one by one at each position in the dataset for the user to hear each data individually in an interleaving way. For example in this case if the data is mapped to the frequencies of the sound source, and the 4 pollutants are to be sonified using interleaved playback, we could hear 4 sounds played sequentially coming from the 4 attributes; The highest-pitched sound signifies the corresponding attribute has the highest value at this time. This then repeats for the next position in the dataset, as demonstrated in this video [3]. Interleaved playback allows more exact and clearer comparisons, instead of hearing everything at once.

In the audio visualization view of our framework, the user can drag and move the locations of the speakers to achieve and experiment with different audio pannings in real-time. Moreover, as seen in Figure 6 and this video [4], the audio is visualized with the particle system in Unity where the particles populated get more clustered when the audio gets louder or higher-pitched depending on the parameter mapping. As opposed to the visual plots shown in the track view, this aims to provide a more artistic and momentary visualization of the current audio instead of the data with the rendering power of Unity.

5. CONCLUSION AND FUTURE WORK

To conclude, this research project succeeds in developing a prototype that achieves various methods of sonification and visualization, improves the sonification experience of two datasets related to health and energy, and supports more functionalities. Future work includes support for more parameter controls and more synthesis methods allowing for more expressiveness of the sound; support for real time data input and more communication protocols such as Open Sound Control (OSC) [22], Serial and Musical Instrument Digital Interface (MIDI); being able to select the time scaling between data time and play time; being able to select regions of data to be played; 3D audio spatialization;

[3] https://youtu.be/QceUDfS8d4s
[4] https://youtu.be/VVkg1i4y_ss

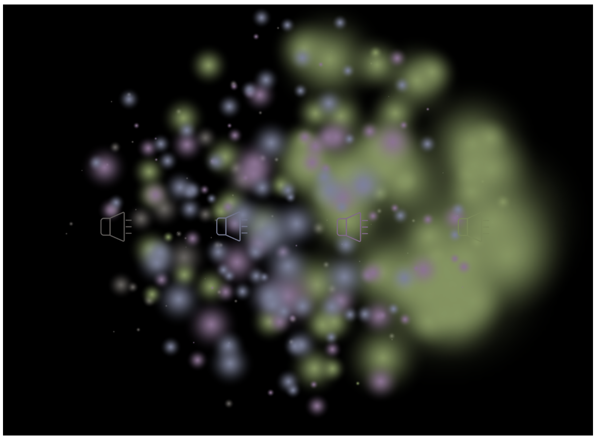

Figure 6. This is the audio visualization view during the playing of the sonification of air quality data. At this time instance, green is the predominant colour meaning that O_3 is the most air pollutant.

and more comprehensive visualization. Moreover, it would be interesting to design interactive mechanisms within the framework to conduct certain tasks to evaluate the sonifications users create. Lastly, the framework itself should also be evaluated. This could be achieved by using case studies of various datasets sonified using the framework, or involving a user study asking users to perform tasks to evaluate the effectiveness of the framework.

Acknowledgments

This research is supported by the Visual and Automated Disease Analytics (VADA) graduate program.

6. REFERENCES

[1] B. N. Walker and J. T. Cothran, "Sonification sandbox: A graphical toolkit for auditory graphs," in *Proceedings of the Ninth International Conference on Auditory Display (ICAD 2003)*, 2003, p. 3.

[2] B. K. Davison and B. N. Walker, "Sonification sandbox reconstruction: Software standard for auditory graphs," in *Proceedings of the Thirteenth International Conference on Auditory Display (ICAD 2007)*, 2007, p. 4.

[3] S. Pauletto and A. Hunt, "A toolkit for interactive sonification," in *Proceedings of the Tenth International Conference on Auditory Display (ICAD 2004)*, 2004, p. 7.

[4] Sonification research group at Berne University of the Arts (BUA), "Sonifyer." [Online]. Available: http://www.sonifyer.org/softhardware/?lang=en

[5] S. Phillips and A. Cabrera, "Sonification workstation," in *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*. Newcastle upon Tyne: Department of Computer and Information Sciences, Northumbria University, Jun 2019, p. 184–190. [Online]. Available: http://hdl.handle.net/1853/61529

[6] H. Lindetorp and K. Falkenberg, "Sonification for everyone everywhere - evaluating the webaudioxml sonification toolkit for browsers," in *Proceedings of the 26th International Conference on Auditory Display (ICAD 2021)*. Virtual Conference: International Community for Auditory Display, Jun 2021, p. 15–21. [Online]. Available: http://hdl.handle.net/1853/66351

[7] Sonify, "Twotone: An open-source data sonification and data-driven music web app," Jun 2022. [Online]. Available: https://github.com/sonifydata/twotone

[8] A. L. A. Blanco, S. Grautoff, and T. Hermann, "Cardiosounds: A portable system to sonify ecg rhythm disturbances in real-time," in *Proceedings of the 24th International Conference on Auditory Display - ICAD 2018*. Houghton, Michigan: The International Community for Auditory Display, Jun 2018, p. 20–27. [Online]. Available: http://hdl.handle.net/1853/60089

[9] R. M. Winters, A. Kalra, and B. N. Walker, "Hearing artificial intelligence: Sonification guidelines results from a case-study in melanoma diagnosis," in *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*. Newcastle upon Tyne: Department of Computer and Information Sciences, Northumbria University, Jun 2019, p. 262–267. [Online]. Available: http://hdl.handle.net/1853/61501

[10] T. J. Mitchell, J. Thom, M. Pountney, and J. Hyde, "The alchemy of chaos: A sound art sonification of a year of tourette's episodes," in *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*. Newcastle upon Tyne: Department of Computer and Information Sciences, Northumbria University, Jun 2019, p. 163–168. [Online]. Available: http://hdl.handle.net/1853/61514

[11] P. R. Kantan, E. G. Spaich, and S. Dahl, "A metaphor-based technical framework for musical sonification in movement rehabilitation," in *Proceedings of the 26th International Conference on Auditory Display (ICAD 2021)*. Virtual Conference: International Community for Auditory Display, Jun 2021, p. 176–183. [Online]. Available: http://hdl.handle.net/1853/66350

[12] P. Cowden and L. Dosiek, "Auditory displays of electric power grids," in *Proceedings of the 24th International Conference on Auditory Display - ICAD 2018*. Houghton, Michigan: The International Community for Auditory Display, Jun 2018, p. 113–120. [Online]. Available: http://hdl.handle.net/1853/60076

[13] K. Groß-Vogt, M. Weger, R. Höldrich, T. Hermann, T. Bovermann, and S. Reichmann, "Augmentation of an institute's kitchen: An ambient auditory display of electric power consumption," in *Proceedings of the 24th International Conference on Auditory Display - ICAD 2018*. Houghton, Michigan: The International Community for Auditory Display, Jun 2018, p. 105–112. [Online]. Available: http://hdl.handle.net/1853/60088

[14] T. Steffert, S. Holland, P. Mulholland, S. Dalton, and A. Väljamäe, *Prototyping a method for the assessment of real-time EEG sonifications*. Georgia Institute of Technology, Jul 2015, accepted: 2015-11-06T14:37:34Z. [Online]. Available: https://smartech.gatech.edu/handle/1853/54142

[15] M. St Pierre and M. Droumeva, "Sonifying for public engagement: A context-based model for sonifying air pollution data," in *Proceedings of the 22nd International Conference on Auditory Display (ICAD 2016)*. Canberra, Australia: The International Community for Auditory Display, Jul 2016, p. 60–64. [Online]. Available: http://hdl.handle.net/1853/56580

[16] G. Wang and P. R. Cook, "Chuck: A concurrent, on-the-fly, audio programming language," in *Proceedings of the 2003 International Computer Music Conference*, 2003, p. 8.

[17] Unity Technologies, "Unity real-time development platform." [Online]. Available: https://unity.com/

[18] J. Atherton and G. Wang, "Chunity: Integrated audiovisual programming in unity," in *Proceedings of the 2018 conference on new interfaces for musical expression*, p. 6.

[19] A. de Campo, "Toward a data sonification design space map," in *Proceedings of the 13th International Conference on Auditory Displa (ICAD 2007)*, 2007, p. 6.

[20] M. Puckette, "Pure data," in *Proceedings of the International Computer Music Conference*, Jan 1996, p. 41.

[21] M. St Pierre and M. Droumeva, "Sonification of air quality data in four canadian cities." [Online]. Available: https://soundcloud.com/marcstpierre/sets/sonification-of-air-quality

[22] M. Wright and A. Freed, "Open soundcontrol: A new protocol for communicating with sound synthesizers," in *Proceedings of the International Computer Music Conference*, p. 4.

ISBN: 978-91-8040-358-0

Sonification and Clustering: a Multi-Sensory Perspective for Precision Medicine

Maria Mannone
Department of Engineering,
University of Palermo, Italy;
DAIS, Ca' Foscari University of Venice, Italy
mariacaterina.mannone@unipa.it

Veronica Distefano
European Centre for Living Technology (ECLT);
DAIS, Ca' Foscari University of Venice, Italy
veronica.distefano@unive.it

ABSTRACT

Melodies are sequences of pitches distributed over time. In medicine, clinical variables, measured at regular intervals, can also be treated as quantities distributed over time. Thus, they can be mapped to sounds, creating melodies. These sonified trajectories can provide a multi-sensory perspective on disease progress for a sample of patients. Musical similarity becomes a tool to group patients with similar disease progress. Here, we focus on diabetic kidney disease (DKD) patients, proposing a sonification of their filtering efficiency information (through eGFR). We obtain Hertz for pitches and spectral centroids. Then, we use this information to build clusters of patients. We adopt a statistics method based on curve-parameter similarity (Traj), and then we propose a technique based on Fourier coefficients for symbolic-music analysis. We compare clusters built upon original values and Hertz, gaining hints on information found via proposed sonification. Possible future developments point to deeper connections between music theory, computer science, statistics, and precision medicine.

Keywords: diabetic kidney disease, eGFR, sonification, Fourier coefficients

1. INTRODUCTION

Sonification [1, 2] is the art (and science) of mapping data to sound. According to [1], "Sonification is the use of non-speech audio to convey information or perceptualize data." It allows for a multi-sensory perspective on different phenomena, potentially catching hidden or hard-to-find information. Recent examples of its use in the medical domain involve sonification of the Covid genome [3]; enhanced information detection from medical imagery [4]; 3-dimensional representation of the human brain [5]. Sonification can also enter the domain of precision medicine. In this branch of medicine, the heterogeneity of patients is specifically taken into account, aiming to find the best individualized therapeutic strategies according to personal characteristics. It is the overturning of the "one-model-fits-all" paradigm.

Several diseases require the observation of patients across a certain period. Longitudinal datasets contain the information of patients that are observed at different data points. Analyzing the variation of some information through time means that patients are characterized by longitudinal trajectories of a variable (or more variables). It is possible to cluster patients according to their path similarity. This approach is particularly relevant to finding patterns of behavior, and finding how subgroups of similar patients responded to the treatment. Such an approach aiming to foster individualized therapeutic strategies is crucial for a disease characterized by a high degree of heterogeneity, such as type-2 diabetic complications.

Let us focus on diabetic kidney disease (DKD). Patients with DKD are mainly monitored according to their estimated glomerular filtration rate (eGFR) over time, a measure of kidney filtering efficiency. The information represented by eGFR trajectories can be mapped into sound. Thus, the temporal information of each patient can become a melody.

In this research, we sonify trajectories of kidney-filtering efficiency, using this information to build clusters of similarly-behaving patients. There are existing approaches to sound-information clustering [6]; however, here we focus on symbolic content first derived as MIDI information and then mapped to Hertz values.

We consider a dataset of 256 DKD patients. We cluster them according to the *Traj* method, based on curve parameters comparison [7]. Traj has already been used to cluster diabetic patients according to their time evolution of glycated hemoglobin (HbA1c) [8]. We compare the patients' grouping obtained with original data and with the sonified ones. We also express comments according to their age, mapped to spectral centroids of sound. In fact, timbre can help differentiate amongst people with similar eGFR path but with different age ranges. Finally, we sketch a new method of clustering, based on Fourier coefficients and derived from mathematical music theory [9], to be used in future research.

Our article is structured as follows. Section 2 summarizes the proposed sonification strategy, the Traj clustering technique, and the proposed Fourier-based strategy. Section 3 presents the first sonification outputs and the clusters' comparisons. Finally, section 4 ends the article with discussion and conclusions.

2. DATASET AND METHODS

2.1 The Dataset

In a preliminary study, we sonified the eGFR trajectories of 10 DKD patients, comparing clustering results based on original data and on derived ones [10]. Here, we focus on 256 patients, allowing the application of different clustering techniques (section 2.3). The longitudinal dataset contains information on demographic and clinical variables, characterizing disease evolution over time, diabetes-related pathologies, and family history of disease for each patient. We focus on eGFR and age.

The dataset has been accessed during a project funded by the Horizon 2020 research and innovation program, Action RIA Research and Innovation action Call: H2020-SC1-BHC-2018-2020; Topic: SC1-BHC-02-2019.

2.2 The Sonification Strategy

In this section, we describe the considered sonification technique. We mainly focus on the pitch. To enrich our analysis, we also consider timbre to distinguish between people of different ages: The higher the age, the lower the overall harmonic component. For each patient, we can thus consider the spectral centroid in addition to melodies. The sonified dataset will thus contain 256 arrays (melodies) with a sequence of pitches as a MIDI number and Hertz, jointly with an indication of spectral centroid to characterize timbres. The idea is to join symbolic and sound information.

The technique to obtain notes is the following.

We first start with the array of values of eGFR. We consider four measurements, taken yearly. Thus, the time distance between two consecutive measurements is the same, and in our musical rendition, there are no rhythmical variations. The one-year difference, in our mapping, becomes a quarter-note rhythmic distance in a $4/4$ tempo. For each patient, we obtain a four-note melody at regular time intervals. We develop a 3-step procedure:

- trajectory rescaling between 0 and 1;
- values rescaling in a chosen MIDI interval;
- final mapping to Hertz.

Let us describe each step in detail.

Values of eGFR are measured in mL/min/1.73m^2. The maximum possible value is 100, while the threshold value for sufficient kidney efficiency is 60. Measurements of eGFR comprised between 15 and 29 denote a severe loss of kidney function, and, below 15, kidney failure [1]. In the considered dataset, there are no patients with kidney failure. Thus, in our analysis, we first choose 100 as the max, and 20 as the minimum, rescaling eGFR values to numbers strictly comprised between 0 and 1:

$$x = (eGFR - 20)/(100 - 20) \tag{1}$$

The second step involves a rescaling to obtain MIDI pitch note numbers in a range approximately between 49 (C#3) and 72 (C5), that is, in two octaves playable by several musical instruments and the human voice. This choice constitutes a good mapping because the resulting notes are within our central hearing range.

$$y = f(x) = x * 23 + 49. \tag{2}$$

The results are rounded to natural numbers. Finally, the third and last passage is a mapping from MIDI to frequencies in Hertz:

$$z = g(y) = 2^{\left(\frac{y-69}{12}\right)}(440Hz) \tag{3}$$

In formula 3, we added 'Hertz' because we transform an adimensional number, the MIDI height information, into a physical quantity (number of oscillations per second). Conceptually, we are overall mapping a filtering rate into an oscillation rate. The rounding step is a choice allowing us to work with Western-tuned musical notes. However, we could have been running the whole analysis without it, also considering that other musical systems use different tunings, and contemporary music often uses microtones. From the point of view of sonification, the rounding step may emphasize data variations.

While we build up clusters on eGFR and, separately, on Hertz to compare them, we also examine additional information to characterize our sample. From the age of patients, we computed the spectral centroids. In this way, timbre can be balanced according to the patient's age, creating a perceptual-equivalent effect. To compute the centroids, we converted age to Hertz with an inverse formula: $2^{-((y-69)/12)} * 440Hz$, to have an inverse relationship between age and luminosity of the sound (high-frequency component). That is, the younger the person, the brighter the sound. While dealing with timbre, we are more interested at mean values, thus we did not make specific hypothesis on tuning, and we skipped the rounding step. In next research, this method will be refined. In fact, in a few cases, since we are sonifying different data, we obtain values in Hertz of the spectral centroid (obtained from the age values) that are numerically lower than the mean Hertz value of the melody (obtained from the eGFR value), that is, of the fundamentals: this is not possible. This would require lower harmonics. A possible fix is the multiplication of the resulting number by a constant factor, to rise up all the values. A method to modify timbre independently by the pitch of a sound is discussed in [11].

2.3 The Clustering Techniques

In a former study [10], we had adopted a shape-similarity criterion [12], based on Fréchet distance [13]. Here, the greater dimension of the sample allows for a different trajectory-clustering technique, based on curve-parameter similarity, called *Traj* [7]. In this section, we summarize its main idea. Then, we sketch a new method for future research, specifically based on mathematical music theory [9].

2.3.1 The Traj method

The Traj method, developed by [7], allows one to classify trajectories in a longitudinal dataset. The method quanti-

[1] https://kdigo.org/

fies how heterogeneous is a cohort concerning patterns of change, explaining how to use these patterns to group patients. The Traj method generates 24 measures for each given trajectory. The measures that appear as the most relevant ones for the considered problem are selected using factorial analysis. They constitute the starting point to compute trajectory clusters. The quantified information includes monotonicity degree, eventual abrupt changes and their direction, and the overall trajectory trend. For our case study, monotonicity, trend, and presence of sudden variations are important, corresponding to the response to the therapeutic treatment. Measures can be grouped into four main classes: measures of change, non-linearity, information of non-monotonicity and sudden changes, and early-versus-later change information. The detailed formulas are provided in the Appendix of [7].

2.3.2 The Fourier analysis

Fourier coefficients have been used to retrieve patterns of variation in genetics [14]. They have also been exploited for trajectory analysis, in the domain of Aeronautics [15]. As the main advantage, Fourier coefficients allow for a dimensional reduction of the examined problem. The development of Fourier coefficients for sound-signal processing is well-known. A bit less known is probably their application to characterize chords and scales in the realm of mathematical music theory [9].

Our idea is to merge all these aspects, sketching a new clustering method where:

- trajectories are first mapped to sound creating melodies;
- for each melody, Fourier coefficients are computed;
- melodies (and thus trajectories) are classified according to the position of their highest-valued coefficients.

In [9, section 5.2] (page 140 and *sequitur*), it is provided the formula to compute coefficients, a less common version of discrete Fourier transform (DFT):

$$c(t) = \sum_{k=1}^{n} e^{\frac{1}{6} i\pi \left(a_k - \frac{12kt}{n}\right)}, \tag{4}$$

where n indicates the rank of the considered object; here, we have melodies of 4 notes: $n = 1, ..., 4$, a_k is the k-th element of the melody, and t is the coefficient index $t = 1, ..., 4$.

To apply this formula to our sonified trajectories, an additional step is needed. The formula considers musical notes *modulo 12*, that is, up to an octave. The information on precise pitches is thus forgotten, in favor of equivalence classes. To obtain pitch classes, we have to consider the MIDI values and subtract from each note the number '12' until a number comprised between 0 and 11 is obtained. Then, we can apply formula 4.

3. RESULTS

3.1 Sonification output

As an example of the output, we list, as data strings, the Hertz values and note names of the melodies obtained by sonifying the first three patients of the considered sample. See Table 1 and Figure 1. The listed examples can be heard clicking **here**. They have been created with software Sibelius. Clicking **here**, it is possible to listen to the three sequences with electronically-modified spectral centroids. The effect has been obtained with software Melodyne. Figure 2 presents the superpositions of three melodies, and Figure 3 shows a single melody whose timbre has been adjusted (starting on a cello sound) to approximate the spectral centroid based on age.

patient	melody			
1	311.13	277.18	277.18	293.66
	E♭4	C#4	C#4	D4
2	246.94	293.66	277.18	329.63
	B3	D4	C#4	E4
3	293.66	311.13	349.23	415.30
	D4	E♭4	F4	A♭4

Table 1. Hertz values obtained for the first three patients (One note for each eGFR yearly measurement).

Figure 1. Melodies obtained from the first three patients. Click **here** to listen with an arbitrary flute rendition (centroid: approximately 1600 Hz), and **here** with an electronically-modified rendition (centroid: circa 800 Hz).

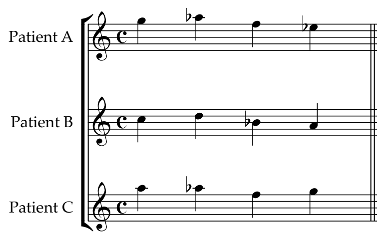

Figure 2. An example of superposition of melodies, with sonification of the eGFR trajectories of three patients. The first two patients belong to the same cluster (cl. 1), and present a parallel motion. The third one belongs to cluster 3. Click **here** for the audio file, with flute sounds.

3.2 Clustering results

3.2.1 Clustering of 256 patients with Traj

Figures 4 and 5, obtained applying Traj on eGFR and Hertz, respectively, indicate mean trajectories in each cluster. Clusters are numbered according to their numerosity.

Figure 3. The melody corresponding to a single eGFR trajectory. Clicking **here** it is possible to hear the audio, whose timbre has been electronically adjusted to approximate the age-based mean spectral centroid of 267.73 Hertz.

eGFR-based clusters have the following number of patients: 103, 64, and 89, respectively. Hertz-based clusters present the numerosity of 77, 170, and 9, respectively.

The eGFR-based mean trajectories (Figure 4) present a switch between patients of the first and second clusters, probably due to the change of drug after the baseline. The eGFR variation provides information on the status of controlled / uncontrolled DKD. We notice a cluster with improving disease (cl. 2), of stable disease (cl. 3) stuck on below-threshold filtering values, and a cluster with slightly-decreasing values, denoting a poorly-controlled DKD (cl. 1).

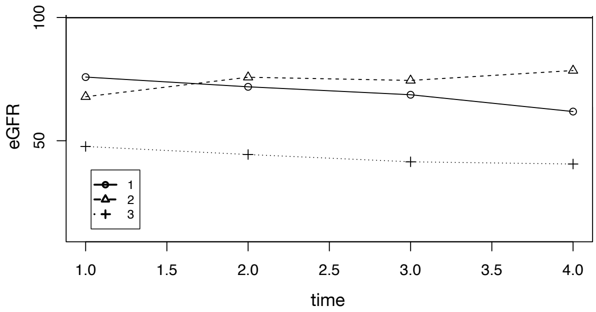

Figure 4. Clusters obtained applying Traj to original data.

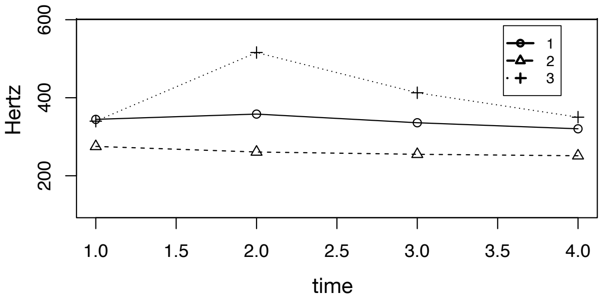

Figure 5. Clusters obtained applying Traj to sonified data.

At a first glance, Hertz-based mean trajectories (Figure 5) seem to loose information on patients' switch. However, we notice an unexpected behavior of cluster 3 (the one with higher values), rising and then decreasing. The change happens in correspondence to the first follow-up, when patients have been tested after the change of drug. Thus, it looks like that the sonification procedure emphasized the switch occurred during measurements at the first follow-up. Cluster 2 in Figure 5, approximately corresponding to cluster 3 in Figure 4 (stable on bad filtering values), is not affected by any abrupt change. However, the sonified information returns a slightly younger subset of patients, thus, with a more brilliant sound (higher spectral centroid).

Variations of disease progress can be evident in other variables before becoming manifest in the eGFR. It is the case, for example, of UACR (creatinine to albumine ratio) [16]. Here, let us consider age to gain further information (Table 2), with a closer eye on clusters 1-2 in Figure 4 and clusters 1-3 in Figure 5. Patients with an overall better disease progress are the younger ones. This information is confirmed by both techniques. However, the age gap amongst 1-2 in the first analysis and 1-3 in the second one is different. Patient distribution found with Hertz-based clustering presents a larger age gap. For a future analysis, it could be interesting to investigate the characteristics of DKD patients around 59 years old under therapy change and other variables. We will also test the effect of tuning choices on fine-graining cluster distribution.

method	cluster	mean age	mean centroid	mean eGFR	mean Hertz
original	1	64	587.33	76	353.89
	2	63	622.25	68	314.32
	3	71	392.00	48	222.62
sonified	1	64	587.33	74	344.58
	2	68	466.16	59	275.23
	3	59	783.99	73	339.95

Table 2. Overall comparison between mean values of clusters. Values for the 256 patients of the sample. Method 1: we compute clusters starting from the original values and then we mapped to sound the mean values in each cluster. Method 2: we first sonify the trajectories, and then we compute the clusters.

3.2.2 Clustering of 10 patients with Fourier

Applying the considered technique to the first ten patients of our dataset, we obtain notes in terms of pitch classes and the label of their higher two Fourier coefficients (Table 3). The cluster attribution is based on the label of the higher coefficient. In our example, we have a prevalence of c_2 (cl. 3), then c_4 (cl. 1), and then the others (cl. 2). To perform these computations, we coded a Jupyter program, which is available upon request. Table 3 contains a comparison with Traj-based clusters: interestingly, in 6 out of 10 patients, the classification is the same. In next research, such a comparison will involve all patients.

4. DISCUSSION AND CONCLUSIONS

In this article, we proposed a sonification strategy to analyze medical longitudinal data. We applied our method to a dataset of 256 DKD patients, sonifying eGFR trajectories. We built clusters on original data and on sonified ones, discussing which characteristics are emphasized with the proposed strategy. Finally, we sketched a clustering approach derived by Fourier coefficients.

ISBN: 978-91-8040-358-0

melody as pitch classes	higher coefficients	Fourier clusters	Traj clusters
3, 1, 1, 2	c_1, $\mathbf{c_4}$	1	1
11, 1, 2, 4	$\mathbf{c_3}$, c_4	2	2
2, 3, 5, 8	$\mathbf{c_1}$, c_4	2	2
9, 8, 5, 7	$\mathbf{c_2}$, c_3	3	3
6, 7, 5, 5	$\mathbf{c_2}$, c_3	3	2
7, 8, 5, 3	$\mathbf{c_2}$, c_3	3	1
0, 9, 9, 10	$\mathbf{c_1}$, $\mathbf{c_4}$	1 & 2	3
8, 8, 9, 11	c_1, c_1, $\mathbf{c_2}$	3	3
8, 8, 8, 2	$\mathbf{c_2}$, c_1, c_3	3	2
3, 1, 1, 2	c_1, $\mathbf{c_4}$	1	1

Table 3. Results of Fourier-based method on ten patients, and comparison with Traj-based clusters.

A possible application of this method could lead to a new interface, where auditory clusters are automatically updated when data of a new patient are added, provided that they respect the time distribution of four different observations, from 1 to 4. Auditory detection of cluster variation would exploit human ear sensitivity in retrieving patterns, corresponding to patterns of behavior.

The interdisciplinary endeavor of precision medicine can profit from knowledge developed in deeply different field such as sonification and mathematical music theory. Such a multi-sensory perspective might help physicians retrieve hidden information, fostering new strategies toward individualized approaches to disease. Such an interdisciplinary challenge stresses the importance of collaboration between different research fields. And, ultimately, we hope that aesthetics, beauty, and music may help improve people's lives.

Acknowledgments

The authors thank Prof. Emmanuel Amiot for the insightful discussions on Fourier coefficients for musical analysis.

5. REFERENCES

[1] G. Kramer, *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Santa Fe, Texas: Santa Fe Institute Studies in the Sciences of Complexit, 1994, vol. XVIII.

[2] T. Hermann, "Taxonomy and definitions for sonification and auditory display," in *Proceedings of the 14th International Conference on Auditory Display, Paris, France June 24–27, 2008*. [Online]. Available: https://smartech.gatech.edu/handle/1853/49960

[3] M. Temple, "Real-time audio and visual display of the Coronavirus genome," *BMC Bioinformatics*, vol. 21, p. 431, 2020.

[4] V. Chiroiu, L. Munteanu, R. Ioan, C. Dragne, and L. Majercsik, "Using the Sonification for Hardly Detectable Details in Medical Images," *Nature Scientific Reports*, vol. 9, no. 431, p. 17711, 2019.

[5] A. Rogińska, K. Friedman, and H. Mohanraj, "Exploring sonification for augmenting Brain scan data," in *International Conference on Auditory Display*, 2013.

[6] S. Krey, U. Ligges, and F. Leisch, "Music and timbre segmentation by recursive constrained K-means clustering," *Computational Statistics*, vol. 29, pp. 37–50, 2014.

[7] K. Leffondré, M. Abrahamowicz, A. Regeasse, G. A. Hawker, E. M. Badley, J. McCusker, and E. Belzile, "Statistical measures were proposed for identifying longitudinal patterns of change in quantitative health indicators," *Journal of Clinical Epidemiology*, vol. 57, pp. 1049–1062, 2004.

[8] T. Karpati, M. Leventer-Roberts, B. Feldman, C. Cohen-Stavi, R. I., and R. Balicer, "Patient clusters based on HbA1c trajectories: A step toward individualized medicine in type 2 diabetes," *PLoS ONE*, vol. 13, no. 11, p. e0207096, 2018.

[9] E. Amiot, *Music through Fourier Space*. Heiderlberg: Springer, 2016.

[10] M. Mannone and V. Distefano, "Trajectory-based and sound-based medical data clustering," in *BCB '22: 13th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*. Northbrook, Illinois: Association for Computing Machinery, 2022. [Online]. Available: https://doi.org/10.1145/3535508.3545102

[11] J. H. McDermott, A. J. Lehr, and A. J. Oxenham, "Is relative pitch specific to pitch," *Psychological Science*, vol. 19, no. 12, pp. 1263–1271, 2008.

[12] C. Genolini, R. Ecochard, M. Benghezal, D. Driss, S. Andrieu, and F. Subtil, "kmlShape: An Efficient Method to Cluster Longitudinal Data (Time-Series) According to Their Shapes," *PlosOne*, 2016.

[13] M. Fréchet, "Sur quelques points du calcul fonctionnel," *Rendiconti del Circolo Mathematico di Palermo*, vol. 2, pp. 1–74, 1906.

[14] J. Kim and H. Kim, "Clustering of change patterns using Fourier coefficients," *Bioinformatics*, vol. 24, no. 2, pp. 184–191, 2008.

[15] R. Annoni and C. H. Q. Forster, "Analysis of Aircraft Trajectories Using Fourier Descriptors and Kernel Density Estimation," in *15th International IEEE Conference on Intelligent Transportation Systems*. Anchorage, Alaska, USA: IEEE, 2012.

[16] B. L. Neuen, W. Weldegiorgis, W. G. Herrington, T. Ohkuma, M. Smith, and M. Woodward, "Changes in GFR and Albuminuria in Routine Clinical Practice and the Risk of Kidney Disease Progression," *American Journal of Kidney Diseases*, vol. 78, no. 3, pp. 350–360, 2021.

COMPOSING 'EATEN ALIVE': EXPLORING THE INTERSECTION OF CLIMATE CHANGE AND GLOBAL FISH PRODUCTION THROUGH SONIFICATION

Brian Lindgren
Composition & Computer Technologies
Department of Music
University of Virginia
wdt7wr@virginia.edu

ABSTRACT

The climate crisis has become one of the most pressing existential issues facing humankind. In a past era, the connection between ocean health and life on Earth might have been less clear. Today, however, the evidence pointing to the ocean's vital role in sustaining planetary life is undeniable. The mass extraction of fish since the burgeoning of the fishing industry in the twentieth century has had a devastating impact on the ocean's biodiversity and ecosystem. The composition 'Eaten Alive for String Quintet, Tape, and Electronics Processing' sonifies the intersection of climate change and global fish production. The process of composing the piece unfolded though: (1) gathering existing research describing the relationship between these two phenomena, (2) developing the compositional concept and sourcing the relevant datasets, (3) programming a sonification tool that would both satisfy the requirements of the piece's framework and support the expression of musical aesthetics, and (4) composing 'Eaten Alive.' The composition represents the author's attempt at bringing awareness to the relationship between climate change and global fish production through musical sonification. The process and considerations behind the data sonification and composition of the work are detailed in this paper.

1. INTRODUCTION

The climate crisis alarm has been sounding for decades. In 1965, U.S. President Lyndon B. Johnson's advisory board wrote, "pollutants have altered on a global scale the carbon dioxide content of the air" leading to effects that "could be deleterious from the point of view of human beings." [1] Despite nonstop warnings for more than 50 years and having failed to sufficiently alter our present trajectory, the planet has begun to experience unprecedented numbers of life-threatening weather conditions. The causes and potential solutions to this catastrophe are many – yet why does the ocean, one of the Earth's major carbon sinks [2], seldom find itself at the center of mainstream discussion? According to the United Nations Environment Programme: "Oceans, and all marine life that lives under and above the water, play a central role in stabilizing the Earth's climate. They provide a vital source of food to a vast number of land and water species and regulate the amount of CO2 (carbon dioxide) that stays in the atmosphere by absorbing 30 percent of global emissions." [3] The ocean's ecosystem and biodiversity are of vital importance in its ability to absorb carbon dioxide and produce oxygen. Sylvia Earle (former Chief Scientist of the U.S. National Oceanic and Atmospheric Administration) and John Bridgeland (former Director of the U.S. Domestic Policy Council) write, "Since the 1950s, half of the coral reefs, much of the phytoplankton, and 90 percent of many kinds of fish in the sea [...] have been extracted as commodities. While fishing to feed populations and sustain economies is critical, technologies have been applied to finding, catching and marketing ocean wildlife on a scale that is both unprecedented and unsustainable. [...] Climate policies take into account the release of carbon dioxide when a forest is cut or burned, and the losses caused by destruction of terrestrial systems that naturally sequester carbon. But where on the balance sheet is an accounting of losses incurred by industrial fishing that clear-cut ocean systems [...]?" [4] While the science seems to acknowledge the correlation between climate change and global fish production, fish capture has doubled globally since 1990 which may indicate a lack of public awareness regarding the issue. [5]

For a composer, data sonification techniques can be useful in bringing awareness to social issues through music. For example, in her work 'Heat and the Heartbeat of the City' Andrea Polli uses temperature data from Central Park to illustrate the current and predicted impacts of climate change on the residents of New York City. [6] Over the course of the work, the intensity of the sonified data increases, reflecting the number of consecutive days of extreme heat. Matthew Burtner, a native of Alaska, draws awareness to glacial melt in his composition 'Elegy (Muir Glacier 1889-2009).' [7] Recordings of melting glaciers are intertwined throughout the work which elucidates another dimension of the data's origin. Chris Chafe and Greg Niemeyer's 'Smog Music' sonifies smog data from approximately two dozen cities around world. [8] In contrast to the Polli and Burtner, this piece was composed for instrumentalists which, perhaps, creates a different experience

ISBN: 978-91-8040-358-0

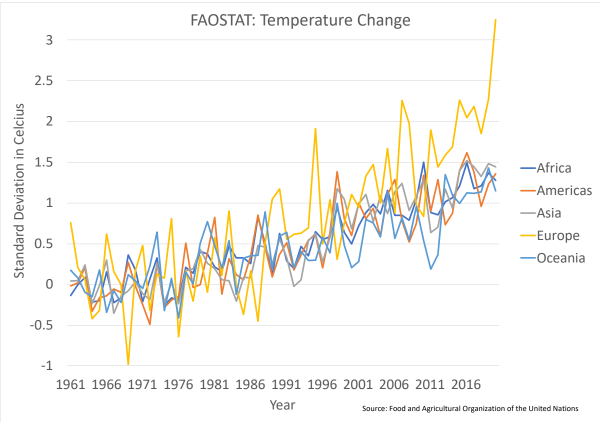

Figure 1. Global temperature change 1961-2020.

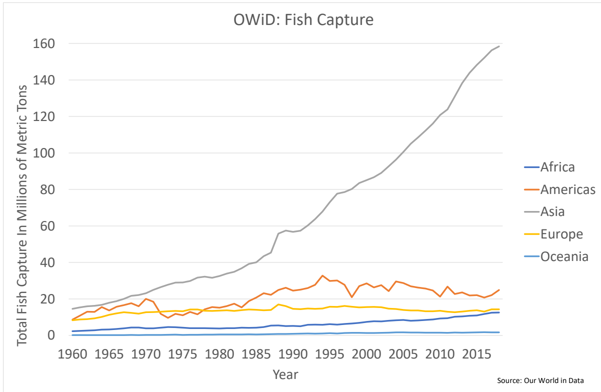

Figure 2. Global fish capture 1960-2018.

than a digitally conjured sonification. Each of these sonification pieces utilize techniques (sonic intensity for expressive impact, inclusion of related audio recordings, and performance by instrumentalists) that have been implemented within 'Eaten Alive for String Quintet, Tape, and Electronics Processing.' [1] This composition, similar to the aforementioned works, attempts to bring awareness to planetary health, specifically the relationship between global fish production and climate change. 'Eaten Alive' is conceived as a fixed-media video performance, common for music created during the COVID-19 pandemic lockdowns. It is scheduled to be premiered at the Technosonics festival at the University of Virginia in February 2023.

2. STRUCTURING THE COMPOSITION

2.1 Representing the Two Datasets

As increased musical pitch can often signify suspense, tension, or drama within a composition, it was determined that the climate data would be represented in the frequency (or pitch) domain, with lower temperatures being represented by lower pitches and higher temperatures by higher pitches. Because this work is intended to depict the harm that global fish extraction enacts on the oceans (and by extension the planet), fish production was chosen to be represented by a spectral filter and delay effects processing, which, when applied to the recorded string quintet, would degrade the fidelity of the recordings. That is to say, the greater the amount of fishing, the greater the degradation of the string quintet sound.

2.2 Data Sourcing

The climate dataset was sourced from The Food and Agriculture Organization of the United Nations [9] website and was available in monthly increments from 1961 through 2020 (Fig. 1). The fishing data was sourced from the Our World In Data [5] website and was available from 1960 to 2018 in yearly increments (Fig. 2). With the mid-20th century signaling a dramatic rise in both climate change and global fish production, this date range was sufficient. As each dataset was subdivided by country, the decision was made to consolidate the data roughly by continent in order to accommodate a string quintet: Oceania and Australia would be represented by Violin 1, Europe by Violin 2, Asia by Viola 1, the Americas by Viola 2, and Africa by the Cello. Next, the climate data was scaled to the range of its respective instrument (represented by MIDI values). The fishing data was normalized to a 0-1 scale for controlling the effects processing amount.

2.3 Compositional Considerations

Four considerations were identified at the outset of the composing process:

(1) One of the challenges when sonifying data that is sampled at periodic intervals is that, when musically represented, its rhythm sounds like a barrage of quarter or eighth notes. In order to create rhythmic variation, a Pure Data (PD) patch (detailed in Section 3) was created that would average data samples into longer notes, thereby creating the possibility for sustained notes of varying lengths. Besides being able to turn this feature on and off, the number of notes being averaged, or window size, would be dynamic.

(2) It would be important for the PD patch to have the capability to not only create longer (averaged) notes but also be able to hear the composition as the note lengths were being shaped. This way, compositional decisions could be auditioned on the fly, allowing for a flexible and creative workspace. This would allow the composition to be shaped, thereby reflecting the intensity of our climate trajectory. For example, Earth Day was established in 1970 to bolster public awareness of environmental pollution, thus the rhythmic activity of the score after 1970 increases.

(3) It was decided to allow for a variable-speed data playback rate. In effect, some months and years would pass by slower or faster than others. A film plot might span many years; however, the director may choose to spend more time focusing on certain dramatic points than others. To realize this concept within 'Eaten Alive,' a variable-speed data playback rate was incorporated into the patch.

(4) In order to illustrate the transformation of climate awareness over time, 'Eaten Alive' would also include au-

[1] 'Eaten Alive' video performance: https://youtu.be/diqBjdQohv8 and score: https://www.scribd.com/document/591685665/Eaten-Alive

ISBN: 978-91-8040-358-0

dio clips of news reports and scientists over the past 60 years that narrate this history. The dates of these clips are synchronized to the sonification playback, thereby creating a connection with their time period. The variable-speed playback allows the composition to focus in on some of these periods of dynamic change and zoom out during other time spans.

3. REALIZING DATABUMP – A CUSTOM-BUILT SONIFICATION PATCH IN PD

Two overarching goals were identified when designing DatabuMp [2] : (1) create a simple, yet effective, GUI interface to allow the composition process to unfold as naturally as possible and, (2) enable the software to encapsulate the entire scope of the composition process by both sonifying the datasets into a note-based score and applying the effects processing to the recorded stems. Figure 3 depicts the composition process shaped by DatabuMp.

3.1 Automation-based Approach

Drawing inspiration from traditional digital audio workstations (DAWs) such as Ableton Live or Pro Tools, it was decided to model DatabuMp on a track-based automation system (Fig. 4). The variable-speed data playback rate is set by a main automation lane (named "Metro"). Each instrument is assigned two automation lanes: (1) the zoom amount (i.e. the window size, or the number of data points that are averaged into one note) and (2) the mode of operation, of which there are three options: Normal, Moving Average, and Quantized Average. "Normal" indicates there is no averaging and the data is sonified directly from the scaled dataset. With "Moving Average" the average is recalculated at each new datapoint and effectively smooths out the stream of notes. "Quantized Average" calculates the average once at the beginning of each new window – this creates longer notes. Automation is integral to this compositional approach; averaging modes and window sizes can be continually changing, thereby creating a dynamic composition while faithfully representing the data. In addition, automation can be recorded on the fly, giving the composer the ability to react to and re-shape the composition while listening to playback.

In many cases, the composition process can take days, weeks or even months. The DatabuMp GUI provides access to features that have been implemented to streamline this process, allowing the composer to focus on shaping the data into a musical composition. These features allow for the following capabilities: save automation data in order to continue composing at a later time, change the point of playback (i.e. move the 'playback head' to the end or middle of the composition), quantize note lengths to a specific BPM (beats per minute) so notes will align to the desired grid within notation software, and calculate the length of the composition (which can change dramatically depending on the rate of data playback). Additional features are configurable by editing a text file within the

[2] DatabuMp GitHub: https://github.com/brianlindgren/DatabuMp

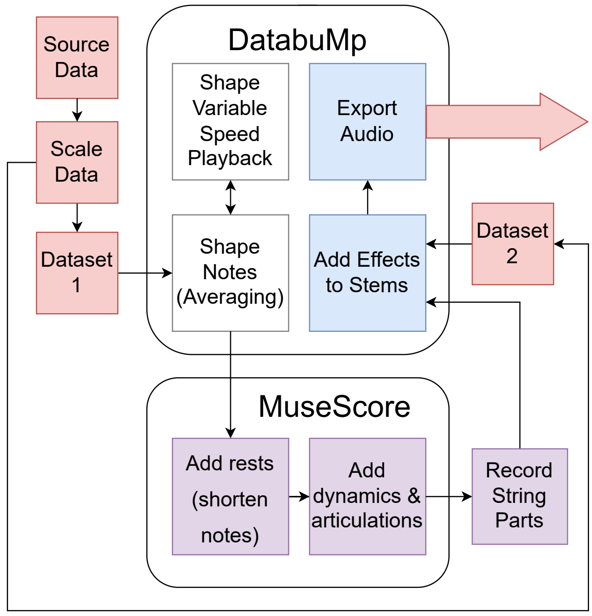

Figure 3. Composition process from beginning to end.

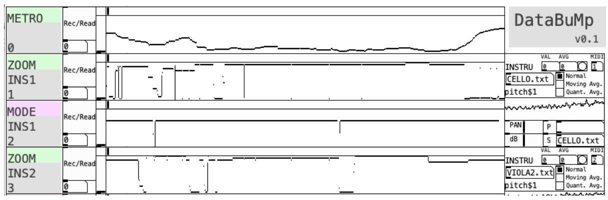

Figure 4. DatabuMp's track-based automation.

DatabuMp folder. These include the ability to: set the automation lane minimum and maximum values; change the ratio of note data playback to effects processing data playback (in the case of 'Eaten Alive' the climate change data was sampled monthly while the fishing data was sampled yearly, thus requiring a 12:1 playback ratio); set the intensity of the effects processing; and specify the types of note lengths available to the quantization engine (e.g. quarter note triplets, thirty-second notes, etc.).

3.2 Modular System

With the intention to allow for any number of instruments, DatabuMp is a modular system. When adding or subtracting an instrument, the automation, playback, and effects processing modules can be dynamically created or removed. The software is designed to be adaptable for future compositions.

4. COMPOSING 'EATEN ALIVE'

After sourcing the data, outlining the compositional concept, and building the software, the next step was to begin composing 'Eaten Alive.' While the preparation phase was conducted over the course of several weeks, the composing process took a few days. The software performed its primary task well: DatabuMp allowed for a focused and

ISBN: 978-91-8040-358-0

streamlined composition process. Working with the automation lanes felt ‘natural’ and the composer was able to ‘get in the zone.’

4.1 Working Within DatabuMp

As with any artistic process, there is a back-and-forth response between the artist and the medium. At the outset of the composition process, both the sound of the raw data and the conceptual material (the sonified data and the historic audio clips) evoked certain feelings. To illustrate, the sonified data in the early 1960s created consonant harmonies at some moments and dissonant harmonies at others, evoking emotions teetering between suspense and resolution. DatabuMp was used to contour longer notes that accentuated these harmonic qualities.

As the composing processs unfolded, the raw data was shaped using automation to create longer (averaged) notes. This allowed for both sparsely spaced notes, which depicted feelings of uncertainty reflected in the awareness of climate change at the time, and also longer phrases that created a gradual dramatic unfolding as the composition developed. As this material was played back (using either raw sine tones or MIDI strings) it inspired further tweaking and adjusting. For example, the growing public awareness of climate change in the 1980s is depicted by a dissonant melody played by the second violin and an unyielding *pizzacato* duel between the violas. Seasonal cycles and other weather phenomena such as erratic temperature fluctuations created natural sensations of tension and release within the work.

As the audio clips were added into playback (DatabuMp playback was synced with Ableton Live via MIDI), new ideas were conjured and the notes were further shaped. For instance, over the course of the piece, the spoken content of the audio clips becomes progressively urgent due to the increasing severity of climate change. Because global temperatures are rising, the sonified data rises higher and higher in pitch as the composition unfolds, creating an increasing intensity. Toward the end, it was decided to slow the rate of playback significantly, in effect, depicting the perception of time slowing down at a moment of calamity or crisis. ‘Eaten Alive’ culminates in a *fortissimo* within the upper range of each instrument’s register, accompanied by an audio clip of reporters ardently describing the disastrous California wildfires of 2019.

4.2 Moving into the Notation Environment

Next, the notes were recorded via the IAC driver (Inter-Application Communication on macOS) into Live first and then exported to MuseScore. Because each month had a corresponding temperature value (no month had a blank value) there were no musical rests in the composition. Miles Davis once allegedly said, “In music, silence is more important than sound.” The artistic decision was made to shorten many of the notes (without shifting the subsequent note ahead), thereby creating all-important silence (rests) as well as notes of varying lengths. This was approached with the intention of shaping expressive instrumental reactions to the unfolding events of this decades-long climate drama. Sometimes these choices created recurring thematic ideas (Fig. 5) and other times not. Articulations and dynamics were added to support both the thematic material and expressive musical gestures. Once the score was completed, it was recorded with each instrument individually miked.

Figure 5. Example of thematic material in ’Eaten Alive’.

4.3 Adding Effects Processing With DatabuMp

Lastly, the audio stems were imported into DatabuMp for effects processing. In the processing signal chain, the audio is first run through a spectral filter and then through a spectral delay. In the spectral filter, each frequency bin is boosted or attenuated by a random amplitude, but maximally limited at the value of fish production that year. Similarly, each frequency bin in the spectral delay is delayed by a random time value, but limited to the fishing value of that year. The result is that as fish production increases, the audio fidelity is increasingly degraded, sounding unstable and askew. The intensity of the effects processing was tweaked until satisfaction was achieved via the editable configuration file.

5. CONCLUSION

This paper details the process of sonifying climate change and global fish production data into the composition ‘Eaten Alive for String Quintet, Tape, and Electronics Processing.’ Sonification can be a powerful tool for transforming data into artistic material and thereby helping to decipher meaning behind otherwise mysterious numbers. By the author’s assessment this work represents a provocative musical depiction of two consequential datasets. The sonification, incorporated with historic audio clips, narrates a winding yet unmistakable trajectory brought about by disregard for the health of our ocean systems. While the rising pitch creates an expected increase in musical intensity, the addition of effects processing has a surprisingly persuasive effect in underscoring the intersection of climate change and global fish production. However, the premiere next year will function as a meaningful test in determining the work’s impact. It is the hope of the author that ‘Eaten Alive’ will make a contribution, no matter how small, in understanding and rectifying our planet’s current climate catastrophe.

ISBN: 978-91-8040-358-0

6. REFERENCES

[1] U. S. P. S. A. C. E. P. Panel and J. Tukey, *Restoring the Quality of Our Environment: Report*, ser. Report of the Environmental Pollution Panel President's Science Advisory Committee. White House, 1965. [Online]. Available: https://books.google.com/books?id=mdXanDtshX8C

[2] G. Mariani, W. W. L. Cheung, A. Lyet, E. Sala, J. Mayorga, L. Velez, S. D. Gaines, T. Dejean, M. Troussellier, and D. Mouillot, "Let more big fish sink: Fisheries prevent blue carbon sequestration – half in unprofitable areas," *Science Advances*, vol. 6, no. 44, p. eabb4848, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/sciadv.abb4848

[3] S. Lutz, "Business unusual: How 'fish carbon' stabilizes our climate," Oct 2018. [Online]. Available: https://www.unep.org/news-and-stories/story/business-unusual-how-fish-carbon-stabilizes-our-climate

[4] S. Earle and J. M. Bridgeland, "The big blue elephant in the room," Nov 2014. [Online]. Available: https://www.huffpost.com/entry/the-big-blue-elephant-in-_b_5908844

[5] H. Ritchie and M. Roser, "Biodiversity," Oct 2021. [Online]. Available: https://ourworldindata.org/biodiversity

[6] A. Polli, "Heat and the Heartbeat of the City: Sonifying Data Describing Climate Change," *Leonardo Music Journal*, vol. 16, pp. 44–45, Dec 2006. [Online]. Available: https://doi.org/10.1162/lmj.2006.16.44

[7] M. Burtner, "Elegy (muir glacier 1889-2009)." [Online]. Available: https://www.matthewburtner.com/product-page/elegy-muir-glacier-1889-2009

[8] H. Hart, "Smog musicians turn pollution data into jagged melodies," Oct 2010. [Online]. Available: https://www.wired.com/2010/10/smog-music/

[9] Food and A. O. of the United Nations, "Faostat temperature change." Feb 2021, license: CC BY-NC-SA 3.0 IGO. [Online]. Available: https://www.fao.org/faostat/en/#data/ET

ISBN: 978-91-8040-358-0