



## OPEN ACCESS

## EDITED BY

Chris Rhodes,  
University College London, United Kingdom

## REVIEWED BY

Paul Vickers,  
Northumbria University, United Kingdom  
Francesco De Pace,  
Polytechnic University of Turin, Italy

## \*CORRESPONDENCE

Paul Bremner,  
✉ paul2.bremner@uwe.ac.uk

RECEIVED 21 March 2024

ACCEPTED 04 June 2024

PUBLISHED 20 June 2024

## CITATION

Simmons J, Bremner P, Mitchell TJ, Bown A and McIntosh V (2024), The ballad of the bots: sonification using cognitive metaphor to support immersed teleoperation of robot teams.  
*Front. Virtual Real.* 5:1404865.  
doi: 10.3389/frvir.2024.1404865

## COPYRIGHT

© 2024 Simmons, Bremner, Mitchell, Bown and McIntosh. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# The ballad of the bots: sonification using cognitive metaphor to support immersed teleoperation of robot teams

Joe Simmons<sup>1</sup>, Paul Bremner<sup>2\*</sup>, Thomas J. Mitchell<sup>3,4</sup>, Alison Bown<sup>1</sup> and Verity McIntosh<sup>1,3</sup>

<sup>1</sup>Bristol VR Lab, University of the West of England, Bristol, United Kingdom, <sup>2</sup>Bristol Robotics Laboratory, University of the West of England, Bristol, United Kingdom, <sup>3</sup>Digital Cultures Research Centre, University of the West of England, Bristol, United Kingdom, <sup>4</sup>Creative Technologies Lab, University of the West of England, Bristol, United Kingdom Technologies Lab, University of the West of England, Bristol, United Kingdom

As an embodied and spatial medium, virtual reality is proving an attractive proposition for robot teleoperation in hazardous environments. This paper examines a nuclear decommissioning scenario in which a simulated team of semi-autonomous robots are used to characterise a chamber within a virtual nuclear facility. This study examines the potential utility and impact of sonification as a means of communicating salient operator data in such an environment. However, the question of what sound should be used and how it can be applied in different applications is far from resolved. This paper explores and compares two sonification design approaches. The first is inspired by the theory of cognitive metaphor to create sonifications that align with socially acquired contextual and ecological understanding of the application domain. The second adopts a computationalist approach using auditory mappings that are commonplace in the literature. The results suggest that the computationalist approach outperforms the cognitive metaphor approach in terms of predictability and mental workload. However, qualitative data analysis demonstrates that the cognitive metaphor approach resulted in sounds that were more intuitive, and were better implemented for spatialisation of data sources and data legibility when there was more than one sound source.

## KEYWORDS

virtual reality, sonification, robotics, nuclear decommissioning, teleoperation

## 1 Introduction

In an era of rapid technological convergence and advancement, many industries are investing in new strategies to reform high risk practices, and to reconsider the management of hazardous tasks in industrial settings. Settings in which direct human intervention might be considered unsafe or inappropriate. For example, in applications ranging from underwater surveillance to disaster relief, tools such as immersive technologies, artificial intelligence and autonomous robot systems are starting to be combined to give remote operators the sense of being “close to the action” whilst engaging at a safe physical distance.

In many cases virtual reality (VR) may be considered a desirable tool due to the psychological sensation of *presence*, which here we define as a sense of “being there” within remote environments (as defined in (Murphy and Skarbez, 2020)). A sense of presence can

elevate levels of trust (Salanitri et al., 2016), which can be vital in the operation of semi-autonomous robots where operators need to rely on correct behaviours. Perhaps more importantly the affordances associated with presence of embodiment and kinaesthesia have been shown to support users of teleoperation systems in task performance and decision making, seemingly through enhanced spatial awareness and sensory cognition in comparison with more traditional teleoperation interfaces such as 2D video and text (Whitney et al., 2019).

In this study researchers worked in partnership with Sellafield Ltd to explore a specific nuclear decommissioning scenario in which virtual reality (VR) is used to teleoperate a simulated team of semi-autonomous robots as they characterise and map a chamber with largely unknown contents within a virtual nuclear facility. In this scenario the operator is responsible for interpreting sensor data gathered by the robots, identifying invisible hazards so they can be labelled. Operators are also responsible for monitoring and protecting the robots from the negative effects of prolonged exposure to such hazards. The study scenario was co-designed with Sellafield workers to represent a plausible use case for near future working in which semi-autonomous robot teams are guided by a remote user to characterise potentially hazardous spaces using a VR digital twin environment built from robot sensor data (as opposed to the entirely simulated environment used here).

This study specifically examines the potential utility and impact of sound and sonification in the design of such systems. The human auditory system has a high temporal resolution, wide bandwidth and is able to localise and isolate concurrent audio streams within an audio scene (Carlile, 2011). Our approach leverages best practice from immersive systems design, using the 360° nature of sonic experience to reduce the “sensory overload” caused by an over-reliance on visual systems (Benyon, 2005), and to enable a user to engage in “attentional switching to best achieve aims and monitor for potentially salient distractions” (Kurta and Freeman, 2022). Building on preliminary studies (Bremner et al., 2022; Simmons et al., 2023), we leverage the capacity of VR as an embodied, multi-modal tool, spreading salient live data across visual and auditory channels, giving operators much of the information that they need to complete tasks in the form of carefully designed bespoke sonification sound sets without the need for overloading the visual channel with this information. Further, sound has previously been shown to enhance users’ sense of presence in virtual environments (Larsson et al., 2007), a desired condition as outlined above.

The bespoke sound sets used in this study have been created by professional sound designers to exemplify two distinct design approaches to sonification. The first sound set applies the principles of cognitive metaphor, engaging with theories of embodied image schema theory, narrative scaffolding and cognitive metaphor theory (Roddy and Bridges, 2020; Wirfs-Brock et al., 2021). This sound set has been developed in response to the literature (detailed in section 2) suggesting that naturalistic mappings may be experienced more intuitively by users, allowing them to apply socially acquired contextual knowledge that could make sounds automatically decodable and facilitate meaning making (Kantan et al., 2021; Wirfs-Brock et al., 2021).

The second sound set follows a more traditional approach to sonic information design, assigning auditory parameters such as pitch, loudness and signal brightness to task-critical features in the environment on the assumption that their sonic distinctiveness will

create a legible and functional sonic ecosystem (Dubus and Bresin, 2013).

Within this emergent industrial use case, and contrasting two distinct sonification design approaches, our work has significant novel contributions in the following areas of complementary interest:

- The capacity of a user of a VR teleoperation system to interpret and apply task data conveyed by sonification.
- The application of cognitive metaphor approaches (from sound design philosophy) to data sonification.
- The effect of the two sonification design strategies on task performance and user experience.
- Levels of trust, comfort, stress and habituation experienced by participants using a system with data sonification.

Our work demonstrates the utility of data sonification in VR teleoperation systems, and provides important insights into the efficacy of two different sonification design approaches.

In this paper we first present the background of what sonification is, current design approaches, as well as review related literature of sonification applications in VR and HRI. We then detail our hypothesis and the scenario used. Following this we describe the VR user interface for teleoperation of simulated robots in a virtual environment, consisting of visual and audio components; within which we describe the two compared sonification design approaches. The two sonification design approaches are evaluated in two user studies: the first evaluates user experience of them within the VR environment, the second is an online study of the objective evaluation of the efficacy of our sonifications in hazard identification. Finally we discuss our results and provide conclusions of our findings.

## 2 Background

### 2.1 Data sonification

Sonification is broadly understood as “the use of non-speech audio to convey information” (Kramer et al., 2010), and has been used effectively in a wide range of applications: sensory substitution (Meijer, 1992), accessibility (Holloway et al., 2022), medical diagnosis (Walker et al., 2019), science communication (Sawe et al., 2020), sound synthesis (Mitchell et al., 2020) and data-inspired music (Mitchell et al., 2019). However, the question of what form of “non-speech audio” should be used and how it can be applied in different applications is far from resolved. Roddy and Bridges (2020) introduce “the mapping problem, a foundational issue which arises when designing a sonification, as it applies to sonic information design” and refer to the aesthetic and structural challenges of parameter mapping sonification as a means to legibly communicate information. Dubus and Bresin’s meta-analysis highlighted that dominant forms of parameter mapping sonification have tended towards computational approaches by assigning discrete data sources to sonic dimensions. The dimensions tend to be those common to Western tonal music such as loudness, pitch, timbre, and tempo (Grond and Berger, 2011; Dubus and Bresin, 2013). Flowers argues that the

simultaneous plotting of numerous continuous variables, particularly using arbitrary sonic characteristics such as pitch, is unlikely to be meaningful to general listeners (Flowers, 2005).

It is important to extend sonification design research beyond how data is expressed as sound and consider how auditory displays are interpreted by listeners, considering factors relating to human perception and aesthetics. There is some debate on the role of aesthetics in sonification, with several authors stating that an auditory display becomes decreasingly useful as it becomes increasingly pleasant and musical (Gresham-Lancaster, 2012; Ferguson and Brewster, 2019). Vickers brought this dichotomy into question (Vickers, 2016), and, along with others (Kramer, 1994; Grond and Sound, 2014) argues that aesthetic design strategies have the potential to reduce fatigue while enhancing the usability and expressive qualities of a display. This observation has led to numerous calls for interdisciplinary collaboration to encourage the integration of artistry and craft into sonification research (Barra et al., 2002; Barrass, 2012).

Wirfs-Brock et al. (2021) note that one of the “reasons sonification has failed to reach broad audiences are the tension between cognitive *versus* ecological approaches to sound design”. Roddy and Bridges (2020) suggest that “the mapping problem can be addressed by adopting models of sound which draw from contemporary theories of embodied cognition to refine the more traditional perspectives of psychoacoustics and formalist/computationalist models of cognition”. In this paper we explore a model of sound which is based on the embodied cognition notion of cognitive metaphor and compare the approach with a more typical computationalist sonification approach.

## 2.2 Sonification design and cognitive metaphor

Sound choice in parameter mapping sonification is of key importance for conveying information, the choice of sound should be meaningful and relate to the underlying data of a parameter mapping (Ferguson and Brewster, 2018) and help to reduce listener fatigue and annoyance (Vickers et al., 2014). The use of cognitive metaphor (or conceptual metaphor (Lakoff, 2006)) draws from the field of embodied cognition, a counter-Cartesian philosophical position that holds that “cognitive processes are deeply rooted in the body’s interactions with the world” and that sense-making is motivated by multi-sensory, embodied experience (Wilson, 2002).

In sonification, embodied cognition principles can be applied in the selection of more naturalistic or contextually linked mappings that are built on the principles of conceptual metaphors (Kantan et al., 2021). Such mappings take into account the embodied associations that humans learn over a lifetime of experience (Kantan et al., 2021), although Roddy and Bridges (2020) caution that these embodied associations (or embodied knowledge) can vary between different groups and cultures, recommending that “[s]uch factors must be taken into account during phases of design” and recommending a user-centred design approach in order to create systems that are tailored to the expected user groups. Indeed, closer associations between sounds and the data that they represent could

allow users to reduce cognitive load and memory burdens associated with sonification interpretation (Preziosi and Coane, 2017).

Wirfs-Brock et al. (2021) suggest that a core part of the sonification design process should be “to teach audiences to listen to data,” and that alongside cognitive metaphor, designers should consider what “narrative scaffolding” might be needed to support listeners as they learn to work with the complexities of a new sonic landscape. Roddy and Bridges (2020) propose the Embodied Sonification Listening Model (ESLM) as a means to track and understand not just the mapping of parameters to data sources, but the mapping process that listeners undertake when arriving at their own understanding of the information received. The model builds on ideas of the “circuit of communication” in Reception Theory in which the producer encodes their intentions within a media product, and invites the viewer to “decode” the work at distance, drawing on their own cultural frameworks and embodied knowledge to unlock layers of meaning (Hall, 1980).

Further to the use of cognitive metaphor, certain design strategies vaunted for interactive technologies and systems design have proved highly relevant to this research.

From the literature, and drawing on experience from early prototyping, we became aware that sonifications in this context could rapidly become congested and cacophonous, as multiple sound sources compete for attention. Similar to Hermann et al. (2011), Case (Case, 2016) reminds us of “the limited bandwidth of our attention,” recommending five core principles for the design of calm technology:

- Technology should require the smallest possible amount of attention.
- Technology should inform and create calm.
- Technology should make use of the periphery.
- Technology should amplify the best of technology and the best of humanity.
- Technology can communicate, but doesn’t need to speak.

Case calls for designers to adopt aesthetics that make sense in the environment in which they are encountered, and that call attention to themselves only when relevant and of use to the person who is expected to attend to that call.

Finally, in “Designing Interactive Systems,” Benyon (2005) discusses the importance of mapping the information architecture within interactive systems, and paying attention to the relationship between physical and conceptual objects when designing in 3D space. Simply put “a good mapping between conceptual and physical objects generally results in better interaction.” (Benyon, 2005)

## 2.3 Sonification in human robot interaction and virtual reality

In the context of robotics, Zahray et al. (2020) and Robinson et al. (2021) have used sonification to supplement visual information on the motion of a robot arm. Both studies showed that different sounds can alter the perception of the motion as well as the robot capabilities.

Hermann et al. use sound to convey information in a complex process monitoring system of a cognitive robot architecture

(Hermann et al., 2003). By dividing data features between the visual and auditory channels, operators were shown to establish a better understanding of system operation. This is conceptually similar to the aims of the work presented here.

Triantafyllidis et al., 2020 evaluate the performance benefits of stereoscopic vision, as well as haptic and audio data feedback on a robot teleoperation pick and place task; both haptic and audio feedback were collision alerts rather than for data representation as evaluated here. They found that stereoscopic vision had the largest benefit, with audio and haptics having a small impact. Similarly, Lokki et al. demonstrate that users can use sonified data to navigate a virtual environment (Lokki and Gröhn, 2005). By comparing audio, visual, and audiovisual presentation of cues, they found that audiovisual cues perform best, demonstrating multimodal cue integration.

It is clear that data sonification is an under explored area in HRI and telerobotics, with relatively few studies across a small number of application domains. Hence, the work we present here represents an important early step in understanding the utility of sonification in HRI. We are interested in understanding the utility of sonification in this novel problem domain, consequently we collaboratively design and compare two different sets of sonifications and systematically investigate their impact on task performance and workload of teleoperators. The two sets of sounds are cognitive metaphor based, i.e., more literal, calling on established representations of these parameters, and computationalist based, with no direct connection to established sound representations.

## 3 Methodology

### 3.1 Hypotheses

Drawing on the literature we will explore the impact of using a cognitive metaphor approach to sound design (cog) vs. a more traditional computationalist approach (comp), and hypothesise that sonifications designed in line with principles of cognitive metaphor will:

- H1** increase presence/immersion.
- H2** increase trust.
- H3** be easier to interpret and reduce workload.
- H4** reduce listener annoyance and fatigue.
- H5** meaningfully relate to the underlying data types.
- H6** will make data level understanding more difficult.


To test these hypotheses a between-subjects study was conducted in order to compare the effect of two sonification approaches that are intended to convey real-time measurements of radiation, temperature and gas. The first “cognitive metaphor” approach draws, where possible, from embodied and established associations that humans already have with the underlying data types. The second “computationalist” approach disregards any established experience and adopts the most prevalent parameter mappings found in the literature. We anticipate that by using a sonification design approach that is ecologically linked with the environment the sense of presence (Larsson et al., 2007; Kern and Ellerbeier, 2020) and consequently trust will be increased

(Salanitri et al., 2016) (H1 & H2). Sounds that more closely align with user expectations will be easier to interpret and will reduce the cognitive load (Schewe and Vollrath, 2020) (H3). Sonification mappings that take an ecological design approach inspired by cognitive metaphor will improve the user experience by reducing listener annoyance and fatigue when compared with mappings to common auditory parameters, which are designed to efficiently convey arbitrary information (H4). Aligning the sonification mappings with user expectations will make the data types easier to identify (Walker and Kramer, 2005; Ferguson and Brewster, 2018) (H5). We expect the data levels to be easier to understand with the utilitarian approach since there is a clearer mapping between sound and data with no consideration taken for listener comfort or environmental coherence (H6). However, our aim is to design both sound sets with understanding in mind, meaning that this difference should be minimal.

### 3.2 The scenario

A real world industrial use case developed with our partners at Sellafield Ltd. is that of characterisation and mapping of sites to be decommissioned. There exist many legacy sites for which the exact physical and environmental conditions are unknown, necessitating this characterisation process. An approach proposed to solving this challenge developed on the Robots for Nuclear Environments programme grant is the use of multi-robot teams (University of Manchester, 2024). In order to comply with the nuclear safety case human oversight is required in this process to ensure the environment is safely and correctly mapped (autonomous behaviours may have to be overridden to gather needed data), and decommissioning activities properly planned. Hence, the robots need to be teleoperated, and their gathered data interpreted by a human operator. The approach being developed is to use semi-autonomous robots to gather environment data in order to construct a digital-twin of the environment which can be safely navigated in VR, and used to allow the operator intuitive understanding of the environment and to issue commands to the robots. Within the environment there are hazardous conditions of which the operator needs to be aware in order to safely and usefully direct the robots and plan decommissioning activities. We have adapted this scenario to allow us to evaluate sonification approaches by sonifying hazard data so that an operator can be aware of hazard levels while still being able to visually focus on the physical environment and the activities of the robots.

In order to allow us to evaluate the sonification approaches while the robot teleoperation and environment reconstruction systems are being developed we have designed a simulated nuclear environment within which a characterisation and mapping task can be conducted. It simulates object reconstruction based on 3D scans of real objects, and sensor readings of environmental hazards (which are sonified). A simulated (and simplified) robot controller and teleoperation interface is used to direct the robots when needed. Thus, we have a facsimile of a digital twin even though no real environment or robots exist, we term this a simulated twin.




**FIGURE 1**  
The virtual reality environment. Voxel objects that have been revealed by the robots can be seen and one is tagged as radioactive. The laser pointer is in radiation tagging mode. A robot is selected in waypoint mode (green outlined robot), and two waypoints have been set. A hazard marker is visible by the barrels on the right.

## 4 Multi-modal interaction design for robot teleoperation

We have developed a virtual reality “simulated twin” environment for our use case. The virtual environment enables data observation and robot teleoperation within a simulated nuclear decommissioning (ND) scenario. In this scenario, a team of four robots move autonomously to “scan” and map an unknown chamber of a nuclear facility. The robots have virtual “sensors” that detect radiation, temperature, and flammable gas, environmental hazards detailed as important by experts from Sellafield Ltd. The virtual environment is comprised of a graphical user interface (GUI) for observing robot sensor data and operating the robots, and a data sonification system for relaying as sound data from simulated hazard sensors. Users must use the robot’s sensor data to characterise areas of high hazard levels (tag objects with an appropriate data label) for subsequent stages of the ND process. An additional goal for users of the system is to keep the robots as safe as possible to minimise damage to them, hence reducing the frequency of repair and replacement costs; further, inoperative robots that have become irradiated become additional waste that must be disposed of during decommissioning. In the following sections we outline the simulated environment, the user interface needed for the tasks, and the audio implementation. A more detailed description of the task is provided in [Section 5](#). A video of system operation can be found [here](#).

### 4.1 Graphical user interface

The graphical user interface (GUI) was built using Unity, utilising the SteamVR Unity package for additional assets and features chosen to resemble Sellafield nuclear facility. The GUI includes a simulated ND environment which is explored by a team of robots along with task performance interfaces needed for interacting with the environment and robots. An overview of the GUI is shown in [Figure 1](#).



**FIGURE 2**  
Layout of the simulated nuclear environment. Hazard spheres are made visible and are red for temperature, green for radiation, and blue for flammable gas. The robots and the user avatar start in the top right-hand corner.

#### 4.1.1 Environment design

The environment comprises voxel objects that are initially invisible, revealed when detected by the proximity sensors of passing robot(s). The objects have been produced by scanning real-world simulacrum of nuclear decommissioning objects using an RGB-D camera, the resulting pointclouds are then converted to voxel renderings. The floor and walls are always visible, representing information known *a priori*. This setup simulates a typical nuclear decommissioning scenario where the floor plan of the rooms is available from blueprints but the exact location of objects and hazards is unknown.

In order to simulate hazards in the VR environment that are not solid objects, the space is populated with invisible hazard spheres within which hazard levels increase linearly from the edge to the centre. Hazard spheres exist for each of the three types of hazard, with a variety of sizes and positioning, including overlapping spheres (which for matching hazard types are additive up to a maximum level), and are deployed to create hazardous areas in the environment for the user tasks. The hazard level at a given point in space drives the output of the sonification system described in [Section 4.2](#). A birds-eye view of the environment with hazard spheres made visible is shown in [Figure 2](#).

#### 4.1.2 User navigation

Following the conventions of contemporary, virtual reality experience design, users have six degrees of freedom (6DOF), and can freely walk in their immediate environment with a 1 to 1 mapping between their physical and virtual location. To support navigation over longer distances, two further movement modes (both commonly deployed in VR experiences) are also made available. The first is the teleportation system provided by SteamVR: the user holds a button on one of the VR controllers

and this projects an arc with a target point where the arc collides with the floor, when the target is placed as desired the user releases the button and teleports to that location. The teleportation arc is blocked by visible objects, i.e., when an area is unexplored by robots it appears to contain no objects so it can be teleported into freely. The second movement system utilises the D-pad on the VR controller. Clicking forward/backward jumps the user a short distance in that direction, allowing movement within the space without inducing nausea (as continuous controller based motion can), as well as navigation through obstacles to more easily reach specific locations. Clicking left/right jump rotates the user 45° in that direction, allowing the user to make large rotations without having to turn their head/body large amounts, increasing comfort and avoiding cable tangling issues.

### 4.1.3 Robots—autonomous behaviour

The robots use the Unity Navmesh navigation system to travel between a series of checkpoints which determine their exploration route through the environment. The Unity Navmesh system allows the navigable area to be specified, and Unity is then able to calculate the path between a Navmesh agent's current location and a specified point on the Navmesh. Initially the environment appears to be empty to the robots, so they navigate along the shortest path. As objects are revealed by the robots' sensors the Navmesh is updated to include non-traversable regions occupied by the revealed objects, and the path to the next checkpoint is updated.

In earlier testing (Bremner et al., 2022) we learned that users find it frustrating when robots appear to work against their own interests, dwelling in areas that have already been identified as hazardous at a risk to their own health. To allow the robots to avoid hazardous areas of the environment the hazard spheres are populated with Navmesh modifier tiles that can be used to increase the cost to travel through hazardous areas. When a robot detects a tile it occupies is hazardous, the Navmesh is updated to increase the cost to travel through that tile (proportionally by hazard level), and trigger path replanning. As the Unity Navmesh system uses cost to travel in its path finding, the robots will tend to avoid returning to hazardous areas unless directed to do so by the user.

### 4.1.4 Robots—control systems

Users are invited to engage in two main tasks: 1) utilise the robots' sensors to identify and tag hazardous areas in the environment, and 2) ensure the safety of the robots (Section 5). To support these tasks, users can issue instructions to the robots, overriding their autonomous behaviours.

### 4.1.5 Setting waypoints

Waypoint control utilises a laser pointer (attached to the controller in the users right hand) to place waypoints on the ground which the selected robot navigates between. Waypoints are placed by pointing and clicking at locations on the ground, the waypoints spawned at the targeted locations are numbered to indicate the order of navigation. Waypoints not yet visited can be removed with a laser pointer click. On entering waypoint control mode the robot stops moving and a "Go" button is visible directly over the robot, navigation to waypoints may be initiated by clicking of the "Go" button. As the robot arrives at each waypoint it is removed, waypoints can be removed by the user or by an object

being discovered at the waypoints location, and this results in the robot immediately re-planning its path. When a robot is deselected it will travel to its remaining waypoints and then resume its autonomous navigation.

### 4.1.6 Status indicators

Each robot has an outline that can be made visible and is coloured to indicate a particular status; to increase utility the outline visibility is not blocked by intervening obstacles. A red outline is used to give a visual indication of a priority alert (indicating the robot is in danger, see Section 4.2.4), an orange outline indicates real-time listening (RTL) mode is engaged (where hazard data is sonified for that robot alone, see Section 4.2.1), a green outline indicates the robot is in waypoint control mode (described above, RTL is also active in this mode). Real-time listening and waypoint control modes are selected for a particular robot using the laser pointer, clicking once on the robot enables RTL mode (it continues to navigate autonomously), clicking the robot again enables waypoint control mode, clicking the robot again deselects it. Only one robot may be in RTL mode at a time (to reduce audio clutter), so clicking a second robot deselects the previous robot (if one is selected).

### 4.1.7 Hazardous areas of the environment

To support the tagging of objects that are hazardous (see Section 5), two UI elements were designed: hazard markers, and object tagging. As robots move through the environment, they detect areas of high hazard levels and on doing so place a visible marker (provided there are no other markers within a threshold radius) allowing the operator to identify areas that require further investigation. The user can use the right controller to project a laser and tag objects as hazardous (i.e., allocating hazard types from those advised by Sellafeld). The laser is coloured according to the hazard tag that is to be assigned to an object: red for heat, green for radiation, blue for flammable gas, black for removing all tags. The tag colour may be toggled using the side button on the controller, and a text label appears over the controller to remind the user of the meaning of each laser colour. When the laser is pointed at a voxel chunk and a button is pressed the chunk changes colour to that of the laser. If an object is already tagged, and a new tag colour is assigned, the colours are added together using a light mixing paradigm, e.g., red and green tags make an object yellow; if the black laser is selected all tags are removed when an object is tagged.

## 4.2 Audio implementation

Whilst attending to the tasks, users experience *either* the sounds designed according to principles of cognitive metaphor (cog) *or* those designed in line with a computationalist approach (comp). All audio has been implemented using Audiokinetic's Wwise and integrated into the Unity environment using C# scripts. Unless stated otherwise, all sounds in our audio scene are spatialised, i.e., they are processed as if they emit from specified point sources within the environment. Artificial reverberation is also used to simulate the characteristics of a real environment. These global audio features not only enhance the realism of our "cyber-physical model" (Walker et al., 2022), but aid users in locating robots

and navigating the virtual environment. The following audio elements are present: Real-Time Listening, Self Real-Time Listening, Notifications, Priority Alerts, UI Feedback and Ambience, each described below.

#### 4.2.1 Real-time listening

The principal sonifications in our virtual environment comprise three parameter mappings communicating the radiation, temperature and flammable gas levels measured by each robot. This sonification is the only modality by which real-time hazard levels are conveyed (by mapping hazard levels to sound parameters) and referred to as “real-time listening” (RTL). RTL is the only information users receive to perform their task of labelling highly hazardous objects. For this reason, primary focus was given to the design of these sounds, with further sonification elements designed around them.

Upon trialling RTL within a VR test environment, and hearing concurrent spatialised sonification streams for a team of up to four robots, it was clear that the auditory scene could quickly become congested and stressful for users. Degradation in the perceptual clarity of the streams was especially pronounced when multiple robots were simultaneously sensing the same kind of data type, as the user would hear multiple audio streams of the same mapping from multiple point sources. On this basis, we chose to limit live hazard data sonification (RTL) to one point source at a time by requiring the user to select a single robot on which to activate RTL mode. Once selected, the user is then able to listen to the live, sonified sensor data as though the sound is output by the selected robot.

#### 4.2.2 Self real-time listening

If the user prefers to listen in to where they themselves are in the space, they may select themselves, and RTL is applied to their own avatar (i.e., they are the sound source), a practice we refer to as “Self-RTL.” When Self-RTL is enabled, any robot RTL is deactivated to ensure only a single RTL sound-source is active. To select this mode, the laser pointer is used to select the left hand of the user, which then acquires an orange outline to confirm this mode is active.

In order to preserve the integrity of the defined industrial use case (where information in the simulated space is generated only by the sensors of physical robots in the nuclear facility) when the user activates Self-RTL mode, they can only hear sounds in locations that have already been mapped by robots. As the user’s avatar has no physical analogue in the nuclear facility, they would be dependent on the digitally twinned robots to provide live, sonified data regarding hazard levels. To make this legible for users, when Self-RTL mode is enabled, the paths that robots have travelled through the environment are visualised. The navmesh modifier tiles within the hazard spheres (described in [Section 4.1.3](#)) only allow the sonifications to be audible if a robot has travelled through that tile.

#### 4.2.3 Notifications

When neither robots nor a users’ own avatar are selected, RTL is turned off. However, in order that users have a continued awareness of the state of the robots (even when they are not currently visible) we designed a set of notification sounds. Users are notified when a robot first encounters a hazard, a short notification sound, or “earcon” ([Dubus and Bresin, 2013](#)), - short snippets of RTL

counterparts are emitted when they encounter a hazard. These spatialised sounds are designed to communicate both hazard type and the respective robot’s relative location to the user, even if visibility is obscured. Case and Day list scenarios when it is appropriate to include notifications: “the message is short and simple,” “information is continually changing,” “the user’s eyes are focused elsewhere,” “the environment limits visibility” ([Case and Day, 2018](#)). Our notifications are designed in alignment with these criteria.

#### 4.2.4 Priority alerts

In order to allow users to perceive circumstances in which a robot has encountered environmental features that are hazardous to its health, we designed a set of priority alerts. Priority levels are useful as a safeguarding measure and combine a robot’s internal “health” and external risk factors, informing users when interventions are required to preserve robot safety. Each robot’s priority level calculations are normalised to a range of 0–1, with values above 0.9 triggering a high priority alert. High priority alerts are designed to take precedence over the rest of the auditory scene, but rather than making them louder, we dynamically applied gain reductions and filtering to all other audio streams to limit alarm fatigue. Even still, with three hazard priority sonifications and four robots, the user could hear up to twelve concurrent alerts, so all alerts are synchronised when simultaneously active, ensuring as much information be sonically communicated as possible without becoming cacophonous. When any robot’s priority level exceeds 0.95, an additional flanger effect is applied to the alert signal, communicating when a high priority scenario is becoming more serious.

To accompany high priority alerts, a medium priority alert system relays when a robot’s hazard priority level passes the 0.5 threshold. These short earcons, based on the same sounds as high priority alerts, differ depending on whether the priority level is rising or falling.

#### 4.2.5 UI feedback

The final sounds emitted by robots are earcons resembling tonal “grunts,” with the tonal melody dependent on the situation. These are used to acknowledge user interactions. They also occur before a high priority alert begins sounding, creating a “two-stage signal,” useful for conveying that a “complex piece of information is about to be delivered” ([Case and Day, 2018](#)). Feedback sounds for tagging, setting and removing waypoints, and confirming user interactions are stereo headlocked rather than spatialised.

#### 4.2.6 Ambience

This element represents a first order ambisonic sound bed responding to the user’s head movement. This aims to increase the sense of presence and helps reduce the detachment a user may feel when interacting with robots and objects within a cyber-physical model.

#### 4.2.7 Sonification design

The approach to sonification design proposed here follows the principles set out in [Section 2.2](#). One set of sounds (cog) utilise cognitive metaphor and ecological congruence in their design. A matching set of sonifications (comp) follow the archetypal,

computationalist approach to sonification design. Both sets of sonifications utilise a parameterised approach, whereby the value of a parameter is set according to the data level of the associated hazard. Fine details of the designs utilised are described in the following sections. We designed a comparative study to evaluate the two, contrasting approaches.

The Cognitive Metaphor (cog) approach seeks to metaphorically express the physical dimensions of data sources, as well as being ecologically congruous with the conditions in which our auditory scene exists. Furthermore, it is important that when multiple sonification streams occur concurrently, each stream of the resulting auditory scene can be differentiated easily by a listener (Neuhoff et al., 2002). With these considerations in mind, cognitive metaphor and gestalt principles played a key role in our sound design and implementation. Each sonification aims to:

- Clearly communicate the value of its underlying data feature by “performing a [...] simulation of underlying physical phenomena” (Dubus and Bresin, 2013).
- Cohere to existing sonic associations users are likely to possess about the hazard (Walker and Kramer, 2005; Cooper et al., 2018).
- Maintain a singular discernible audio stream that maintains “temporal coherence” (Shamma et al., 2011), e.g., a consistent timbre and pitch, with sounds in the stream occurring continuously or in rapid succession.
- Integrate with other sounds emanating from robots to provide a cohesive character.
- Maximise listener comfort and minimise annoyance and fatigue (particularly alarm fatigue (Sorkin, 1988)).

The RTL radiation sonification (example here) emulates the sound produced by a Geiger counter, consisting of short, high frequency clicks. Transient density varies in a rhythmically complex, yet subtle, way. Higher radiation levels are reflected by faster patterns of clicks. At very high levels, these transient clicks are interspersed with short chirps, designed to draw user attention.

The RTL sonification for flammable gas was designed to share the sonic characteristics with the sound of air inhalation or gas flowing in a pipe, the metallic quality of the sound corresponding with materials used in robotics. As can be heard here, higher gas density decreases the duration between sample onset, which results in the perceived sound of “more gas”. The sound was created using phased, reverse cymbals as source samples. These were passed through a phaser to create more variation over time. The samples were also filtered to remove low frequency content.

With the potential for three concurrent RTL sonifications, it was important to ensure the sonic qualities of each sonification remain unique and consistent, so gestalt principles took precedent in the design of the RTL temperature sonification. As heard here, it consists of a simple sine wave pattern at a fixed pitch of 220 Hz. Higher temperature increases LFO speed and FM amount. The rhythmic simplicity, smoothness of timbre, and fixed low-mid frequency of the sine wave contrasts with both the transient granularity of the radiation and atonality of the gas sonifications and their higher spectral content, facilitating auditory scene perception.

Priority alerts for each hazard are designed using unique source samples, inspired by, rather than derived from, their RTL

counterparts. This is to ensure they do not become confused with the RTL sonifications. Whilst there is correspondence—radiation consists of clicks, flammable gas of noise and temperature of sine tones—high priority alerts are more melodic in nature, consisting of an ascending arpeggiated sequence. As Case and Day posit, “a melody would be difficult to miss” reinforcing our “attention to such alarms rather than detracting” (Case and Day, 2018). However, they go on: “overt melody-making could run an additional risk of extreme annoyance from overuse.” For this reason arpeggiated notes are short and do not adhere to an explicit musical scale.

The Computationalist Approach (comp) has been developed in order to facilitate a fair comparison with the first sound set (cog). Hence, it is important that the comp sonifications avoid the inadvertent use of the novel approaches and considerations laid out for cog, whilst still adhering to an established and respected sonification framework. Fundamental in the design of comp is Dubus and Bresin’s systematic review of mapping strategies (Dubus and Bresin, 2013). We use this extensive review of 179 sonification publications as a framework for sound design due to the authors’ motivation to exploit and organise “the knowledge accumulated in previous experimental studies to build a foundation for future sonification works” (Dubus and Bresin, 2013). We consequently use the most common or successful auditory mappings based on the physical dimensions present in our system. It should be noted that in contrast to cog the sounds are designed with little attention paid to the application domain or user comfort.

As is the case for cog, “Loudness” and “Spatialisation” are reserved for “Distance” and “Location” respectively, to accurately convey the 3D environment of the cyber-physical model. Hence, we excluded these parameter mappings for use in our sonification design. The physical dimensions of the hazards present in our system are energy (radiation), density (flammable gas) and temperature; we utilised these physical properties in establishing appropriate mappings.

After Loudness, the next most common mapping for energy is pitch. Radiation is the primary hazard level we want to communicate and “pitch is by far the most used auditory dimension in sonification mappings,” “known to be the most salient attribute in a musical sound” (Dubus and Bresin, 2013). The RTL radiation sonification for comp consists of a pure sine tone with pitch mapped over two octaves, as heard here. Mapping to higher frequencies is avoided; as Kumar et al., 2012 have shown, “sounds with high unpleasantness have high spectral frequencies and low temporal modulation frequencies”.

The RTL sonification for flammable gas, heard here, utilises a white noise oscillator in accordance with Dubus and Bresin’s supported hypothesis that most sonification mappings follow the priorities of perception (Dubus and Bresin, 2013). After pitch, they found duration to be the most common auditory mapping for density, and that the range between 100 ms and 2 s is used in more than 50% of the reviewed studies. For this reason, the signal gain is attenuated using a sine LFO to create a beating effect, with frequency mapped between 0.5 Hz and 10 Hz. A sine LFO ensures the tone duration and inter-onset interval (IOI) are of equal length, no matter the LFO frequency, and that transitions between tone and IOI are gradual. This implementation method



ensures the sonification is not mistaken as a notification at low gas levels.

The RTL sonification for temperature is mapped to signal brightness for two reasons. Firstly, brightness can be thought of as perceptually orthogonal to both pitch and beating, improving clarity between audio streams (Ziemer and Schultheis, 2020). Secondly, after pitch and loudness—mappings reserved for radiation and distance—brightness is the next most successful auditory mapping for temperature in Dubus and Bresin’s sonification study review, being labelled to indicate that its “efficiency was found to be significantly better when tested in comparison to other mappings corresponding to ”temperature sonification (2013). Our temperature to auditory brightness mapping is achieved using two sawtooth waves with fixed base frequencies of 55 Hz and 110 Hz, reducing spectral interference with the radiation audio stream. PWM with a duty cycle of 80% is applied to the 110 Hz signal to introduce further harmonic content in the upper frequencies. A 100 Hz lowpass is applied at minimum temperature to filter out the higher signal and all harmonic content. The lowpass frequency is mapped to temperature level, reaching 20 kHz at highest temperature. Finally, the gain of each oscillator is also mapped to temperature level, with the higher frequency signal being more prominent at high temperatures and *vice versa*. To the listener, this results in a consistent volume, with the sound being duller when the environment is cooler and brighter when hotter. It can be heard here.

As in cog, the comp notifications, high priority and medium priority alerts for each hazard type are modelled on their respective RTL sonifications, with UI feedback SFX and ambience remaining unaltered.

## 5 User study

To evaluate our sonification approaches, a between subjects user study was conducted involving 50 participants (33 male, 16 female, 1 other, age =  $M28.5 \pm 5.26$ ) randomly split into two equally sized experimental groups, with sound set as the independent variable (cog and comp). The study used the HTC Vive Pro Eye head-mounted display, with the Vive controllers for robot control and navigation, and using the built-in headphones for audio. The software ran on a PC with an Intel Core i9-7900X CPU @ 3.30 GHz, with 32 GB of RAM, and an NVIDIA Titan V graphics card.

Participants were tasked with identifying and labelling objects in the VR environment that they deem to be highly hazardous. To do this they must use RTL to determine hazard levels for objects within the environment and tag objects that are highly hazardous with the correct hazard type. Participants will need to redirect the robots to gather complete sensor data of some areas of the environment as the robot navigation route does not completely cover the area and hazard avoidance behaviour may also contribute to an incomplete data capture. Additionally participants are asked to keep the robots as safe as possible. This will involve setting waypoints around areas they perceive to be hazardous on the robot’s exploration route, and minimising time inside hazardous areas that need additional sensor data capture. The scenario is kept constant across all participants.

## 5.1 Study protocol

The study was undertaken at Bristol Robotics Laboratory based at UWE Frenchay campus, Bristol United Kingdom, and was divided into the following stages:

1. Instruction and consent: Introducing the study purpose and itinerary and obtaining consent.
2. On-boarding: Tasks, system operation and assigned sonification set were introduced using a video and 10 short tutorials in VR.
3. Main Test Scenario: Participants completed the tasks using the assigned sound set for a maximum duration of 20 min.
4. Survey: Gathering quantitative data about participants’ experiences of the Main Test Scenario.

While the system is designed to be as self-explanatory as possible, there are a number of features that need to be understood for system operation. To streamline the learning of these features we designed a series of on-boarding (tutorial) scenes, conducted in VR, to enable participants to learn about system operation and data presentation. Each scene introduces a subset of system features, and allows the user to experiment with the system features presented within that tutorial.

## 5.2 Dependent measures

In order to test our hypotheses on the efficacy of different approaches to data sonification we have used a variety of subjective dependent measures. In order to evaluate workload, trust in the system, and presence in the virtual environment we have used questionnaires that are common in the literature for assessing these facets of system performance: the NASA task load index (NTLX) (Hart and Staveland, 1988), trust in automation (TinA) (Körber, 2019), and the iGroup presence questionnaire (IPQ) (Schubert et al., 2001) respectively. In order to evaluate experiential factors directly related to the sonifications and their utility we have designed a set of questions presented in Table 1. Finally we gathered qualitative feedback on the system as a whole to hopefully illuminate our findings from qualitative analysis.

TABLE 1 Sound evaluation questions (SEQ). Each question is scored on a 5-point Likert Scale.

|      |  |
|------|--|
| SEQ1 | I found it easy to identify the gas sounds   |
| SEQ2 | I found it easy to identify the radiation sounds   |
| SEQ3 | I found it easy to identify the temperature sounds   |
| SEQ4 | I was able to determine different hazard levels  |
| SEQ5 | I found the alarms annoying  |
| SEQ6 | I found the real-time listening sounds annoying  |
| SEQ7 | I would be happy to have these sounds present all day if this were my job                      |
| SEQ8 | After a while I ignored the alarms   |
| SEQ9 | After a while the sounds helped me understand the scene without needing to concentrate on them |

Initially we had intended to log task performance metrics such as tagging accuracy and task performance time, however, during study pilots it immediately became clear that individual differences would result in noisy and useless data. Firstly, task interpretation was highly variable with some participants tagging everything slightly hazardous, some only the object at the centre of the hazard, others doing as requested and tagging highly hazardous objects. Despite improving the phrasing for the main study participants still interpreted the instructions quite differently. Secondly, there were large differentials in system operation capability which resulted in tagging errors that were not the result of sonification utility. Consequently, we conducted a secondary study to objectively evaluate the performance of the sonifications for data understanding (see Section 6), while the main study allows us to evaluate subjective user experience of the sonifications in the VR environment they were designed for.

### 5.3 Results

Where data were normally distributed, independent *t*-tests were used for the comparison (*t* and *p* values are reported); however, in the majority of cases data were not normally distributed (Shapiro-Wilk tests with  $p < 0.05$ ) so Mann-Whitney *U* tests were used (*U* and *p* values are reported).

For the NASA TLX data the mean (non-weighted) workload value was calculated for all contributory factors and compared using an independent measures *t*-test, no significant difference was found between the conditions  $t = 0.557$ ,  $p = 0.580$ . Decomposing the NTLX into separate factors, mental workload was found to be significantly higher for cog ( $M68.75 \pm 18.80$ ) than comp ( $M53.85 \pm 19.61$ ),  $U = 189.5$ ,  $p = 0.015$ . No other factors were found to be significant.

For the IPQ data the data was processed according to Schubert et al. (2001), and thus analysed as presence (P), spatial presence (SP), involvement (I), experienced realism (ER). No factors of the IPQ were found to be significant.

For the Trust in Automation (TinA) data the data was processed according to Körber (2019), and thus analysed as reliability/competence (RC), understanding/predictability (UP), familiarity (F), intent of developers (IofD), propensity to trust (PtoT), and trust in system (T). UP was found to be significantly greater for comp ( $M3.97 \pm 0.82$ ) than cog ( $M3.65 \pm 0.64$ ),  $U = 418$ ,  $p = 0.038$ . No other factors were found to be significant.

For the sound evaluation data each question was compared individually. The gas sounds were found to be easier to identify in comp ( $M4.54 \pm 0.71$ ) than in cog ( $M3.13 \pm 1.33$ ),  $U = 502.5$ ,  $p = 9.1e - 5$ . The radiation sounds were found to be easier to identify in cog ( $M4.66 \pm 0.70$ ) than in comp ( $M3.88 \pm 1.07$ ),  $U = 156.5$ ,  $p = 0.0009$ . The alarms were found to be more annoying in cog ( $M3.83 \pm 1.09$ ) than in comp ( $M2.763 \pm 1.28$ ),  $U = 163.5$ ,  $p = 0.003$ . No other results were found to be significant.

At the conclusion of the survey participants were invited to add “any additional comments.” Responses to this question provide additional insights, particularly in relation to elements that users found challenging. Of the 50 study participants, 30 provided a free text response to this question, 15 of whom had experienced the cog condition and 15 who had experienced

the comp condition. Twenty participants neglected to give a response to this question.

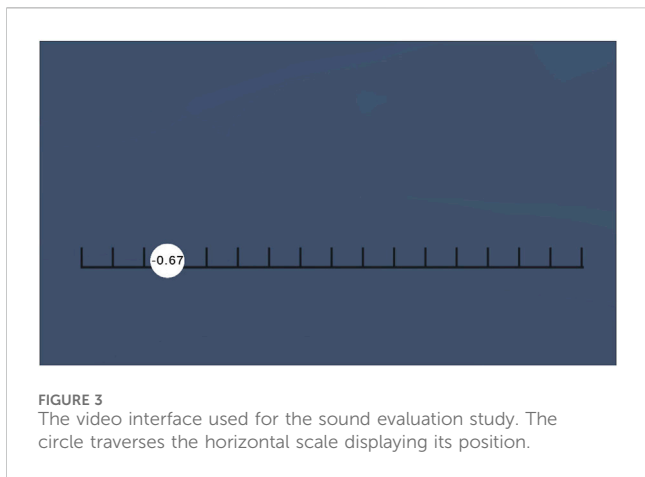
Although a relatively small sample set, there do appear to be some pronounced differences in the self-report of those experiencing cog and comp conditions. For our analysis we primarily used Linguistic Inquiry and Word Count (LIWC).

cog—User feedback is more focused on technical aspects, relating to areas that users felt could be practically improved (LIWC—tech cog: 1.93, comp: 0.2) The language used also suggests that greater attention was paid to the spatial qualities for users of cog (LIWC—space cog: 7.49, comp: 4.01) possibly inferring higher levels of situational awareness. cog participants also use more emotional language, both positive and negative in the way in which they report their experience (LIWC—emotion cog: 1.93, comp 1.0). Language such as “overwhelmed,” “disorientating” and “confusing” is commonly used in negative feedback, “fantastic” and “instinctive” are used in positive feedback.

Comp—Feedback from those who experienced comp give multiple reports of sounds being difficult to tell apart from one another. Six of 15 participants who responded to this question having experienced comp reported difficulty in distinguishing between some combination of the radiation, temperature and flammable gas sonifications. None of the 15 respondents to this question who experienced cog expressed such concerns, however two recounted difficulties distinguishing between the more granular expressions of each sound, e.g., high, medium and low levels of hazard. Negative language most commonly used to describe comp includes “distracting,” “mixed up” and “difficult.” Positive language included “good” and “entertaining”.

## 6 Sonification evaluation study

To evaluate the efficacy of the sonifications for data representation, we have conducted an online study where participants had to identify the underlying data from its sonification. We used a within subjects design conducted on Prolific Academic, 50 participants were recruited (31 female, 17 male, 2 removed for giving random answers, age =  $M43.85 \pm 14.32$ ), and compensated £4.50 for their time, recruitment was restricted to those without a hearing impairment, and they were instructed to wear headphones for the study. For each sonification participants were tasked with observing a video in which data values changed according to a hidden function, and had to identify at which point the sonifications indicated a maximum level for the data. For this task a set of videos were created in which a circular marker moved across the frame from left to right and back again, while displaying the numerical value of its position (Figure 3), a hidden Gaussian function converts that location into a data value that was then sonified using the RTL approach described in Section 4.2 (there was no spatialisation so the audio was equivalent to self-RTL); the centre for the Gaussian in each video was randomly generated in the range  $-1$  to  $+1$ . Each question block relates to one of the 6 sonifications, and consists of a set of 3 videos with randomised maximum values, displayed in a random order. To answer each question participants were permitted to re-watch the video as required, and then had to pause it at the location for which the data was at a maximum and enter the position value displayed. At



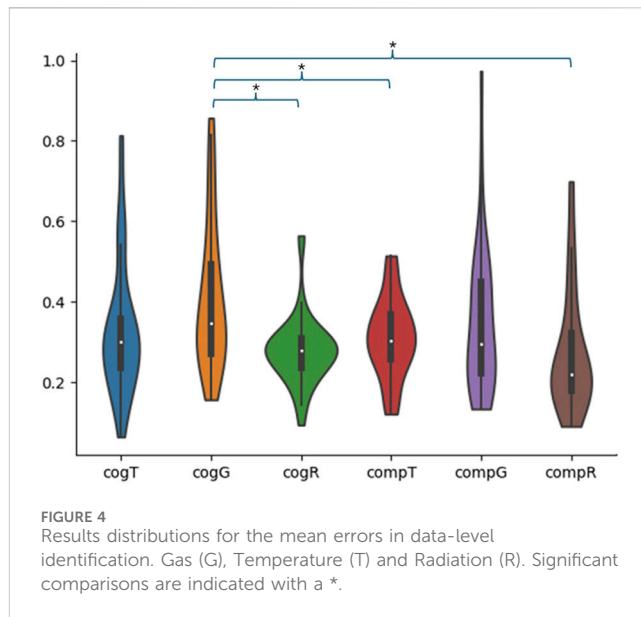
**FIGURE 3**  
The video interface used for the sound evaluation study. The circle traverses the horizontal scale displaying its position.

the end of each block participants had to identify which hazard they believed the sound represented.

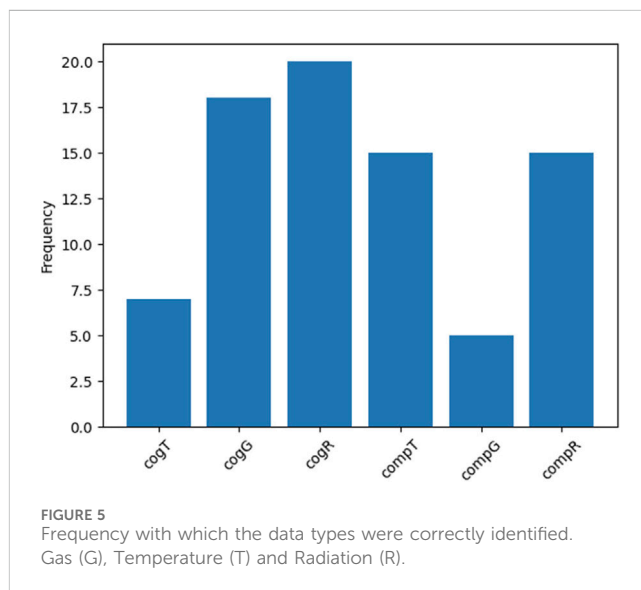
### 6.1 Results

To analyse the data for the hazard level identification accuracy, participant responses for each question were subtracted from the hazard maximum, and the absolute values for these calculated distances were then averaged for the three trials for each sound for each participant. On observation of the results distribution it was apparent that there were a number of outliers that would impact subsequent statistical analysis. Hence, outliers more than 3 standard deviations from the mean were filtered out prior to statistical analysis, this reduced the participant sample to 40 participants. To analyse the data we used a 2-way repeated measures ANOVA. However, the data violates the normality assumption (Shapiro-Wilk tests show that for 5 out of 6 data groups  $p < 0.05$ ), with excessive skew ( $ABS(skew) > 2$  for some conditions). To compensate for the excessive skew the data is transformed using a square root transform. The results following this data processing are shown in Figure 4. For higher order ANOVA ranked approaches have challenges for *post hoc* comparisons, and other non-parametric alternatives (e.g., ART-ANOVA) are not fully investigated. As such we have relied on the robustness of ANOVA to normality violations, and the fact that our transformed data is within reasonable bounds for skew ( $-2 < skew < 2$  for all conditions) and kurtosis ( $kurtosis < 5$  for all conditions); hence, our analysis and statistical conclusions are still valid.

Maulchy’s test indicates the data adheres to the assumption of sphericity  $\chi^2(2) = 1.812, p = 0.404$ . A 2-way repeated measures ANOVA shows that there is a main effect of hazard type ( $F(0.41, 2) = 8.23, p = 0.0006$ ), but no main effect of sound-set ( $F(0.06, 1) = 2.66, p = 0.11$ ) or interaction effect ( $F(0.04, 2) = 1.16, p = 0.31$ ). As there is no effect of sound set or interaction effect, pairwise comparisons are conducted between all 6 sounds (grouping the data by hazard type and ignoring sound set would not facilitate useful evaluation). Due to the large number of comparisons we have used False-Discovery Rate



**FIGURE 4**  
Results distributions for the mean errors in data-level identification. Gas (G), Temperature (T) and Radiation (R). Significant comparisons are indicated with a \*.



**FIGURE 5**  
Frequency with which the data types were correctly identified. Gas (G), Temperature (T) and Radiation (R).

(FDR) with  $\alpha = 0.05$  to compensate for the multiple comparisons (Storey, 2011). The results of the analysis (Figure 4) show that the gas sonification from cog ( $M0.41 \pm 0.18$ ) performs significantly worse than cog radiation ( $M0.18 \pm 0.09$ ),  $t(39) = 4.30, p = 0.0016$ ; comp temperature ( $M0.31 \pm 0.10$ ),  $t(39) = 3.457, p = 0.0066$ ; comp radiation ( $M0.27 \pm 0.15$ ),  $t(39) = 3.49, p = 0.0066$ . All other comparisons were not significant.

In order to analyse the hazard identification data, we have used a Cochran’s Q omnibus test to investigate if there are differences in the number of hazards identified correctly for the different sound designs. To do so the data has been binomially coded (i.e., as wrong or right). The test shows that there is a significant difference between sounds  $\chi^2(5) = 19.42, p = 0.002$ , hence, we conducted pairwise Cochran’s Q-Tests (with FDR correction,  $\alpha = 0.05$ ) between all of the sounds. We found

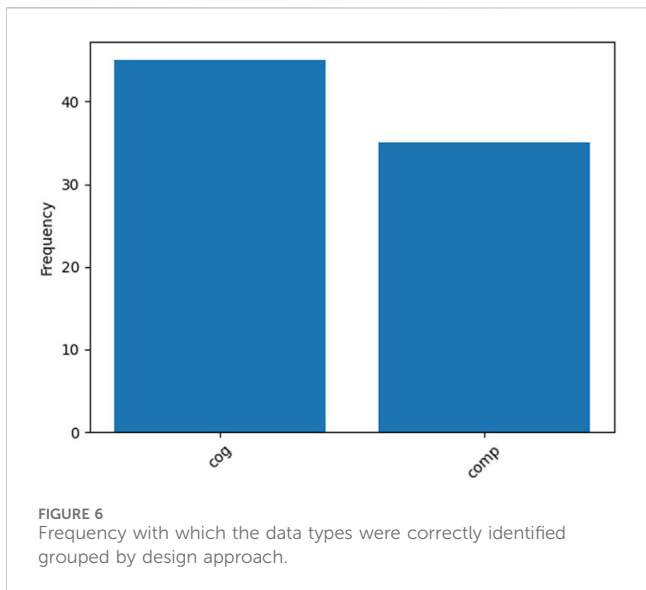


FIGURE 6 Frequency with which the data types were correctly identified grouped by design approach.

that sounds cogT and compG were harder to identify than the other sounds. The results are shown in Figure 5. Grouping the sounds by design approach we found there was no significant difference in hazard identification between the approaches. The results are shown in Figure 6.

## 7 Discussion

The central finding from the results suggest that when comparing cognitive metaphor (cog) and computationalist (comp) sound design approaches there was little impact on subjective experience or utility for data understanding. Where differences were found the results tended to favour the more conventional approach to sonification design. However,

particularly in light of this being contrary to our expectations, it is instructive to analyse our findings, how they relate to our hypotheses, and by reasoning on them suggest avenues for further research.

H1 suggests that cog sounds would give users an increased sense of presence, and was refuted. We posit that engaging the auditory channel regardless of sound design was sufficient for an increased sense of presence, as all factors of the IPQ were positively skewed for both sound design approaches. This finding corroborates (Larsson et al., 2007), and it seems reasonable to suggest, that while our design intention was for the sounds to be better situated in the environment, they could not be made sufficiently more congruent, while maintaining data representation, to have a positive effect on presence. Indeed, they may not be sufficiently aligned with the participants associations and expectations for the hazards types (Kantan et al., 2021). However, this finding does suggest that using either sonification approach would lead to task performance increases connected to an increased sense of presence (Bowman and McMahan, 2007), while not resulting in complete immersion that may not be desirable in a hazardous environment (Malbos et al., 2012).

The absence of any differences in measures relating to trust between the design approaches, refuting H2, are directly linked with the absence of differences in presence (Salanitri et al., 2016). Further, as the scenario was entirely simulated, beyond risks to the participants' task performance (of keeping the robots safe), participants may have seen it as a low risk scenario reducing their need to trust in the system. In future work we intend to have a real robot fleet which mirrors the actions of the simulated robots, and investigate what implications there might be for trust in a system with real-world consequences.

As with H1 and H2, H3 was similarly refuted with a lack of evidence that designing sounds which allow the application of acquired contextual knowledge are easier to interpret and utilise. Indeed, the opposite was found for the mental workload component of the NTLX. We think this opposing finding is

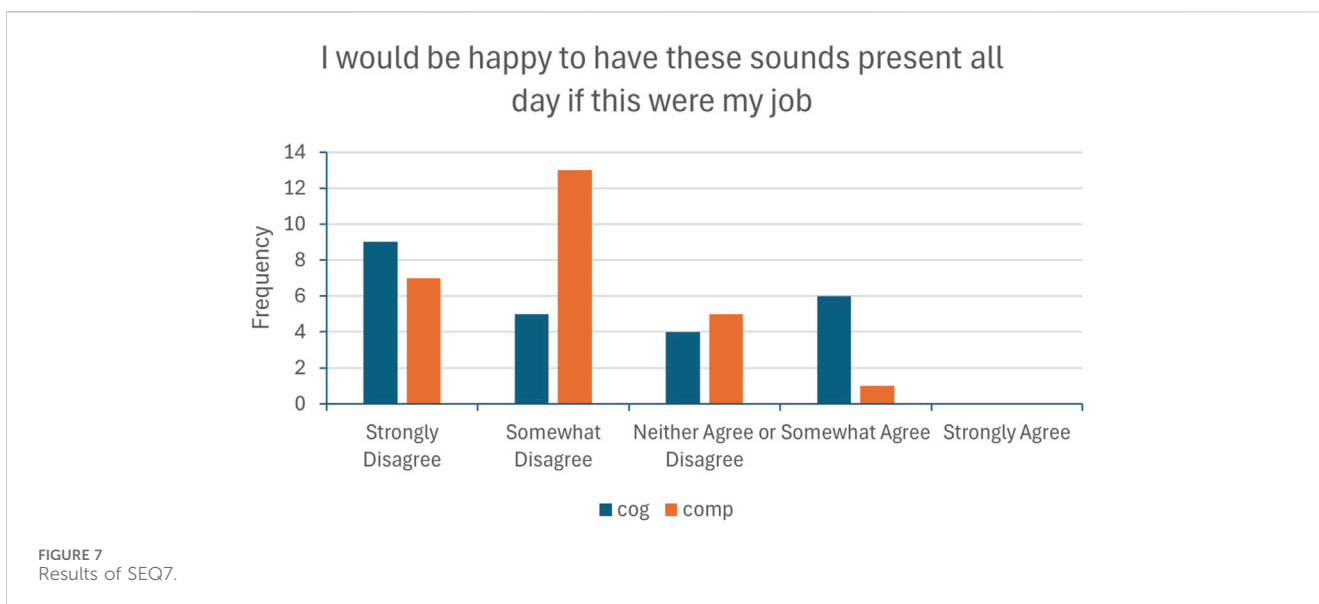


FIGURE 7 Results of SEQ7.

linked to the reduced predictability of the cog sounds relative to the comp sounds, as participants had to pay more attention to the cog sounds as they were less sure about how they would change. However, the findings from the qualitative data suggest that the cog sounds were more intuitive and had better spacialisation so could be localised more easily. To reconcile these conflicting pieces of evidence, the qualitative data suggests that the novel sonic landscape, and complex system usage may have obfuscated any difference in the impact on cognitive load of the relative intuitiveness of the different sound design approaches. Consequently, for future studies we intend to evaluate the systems over successive interactions to allow users to become naturalised to the sounds and system operation. The change over time of task performance, and final cognitive load of the system will be instructive as to the utility of different design approaches.

We found that there was no difference in subjective listening experience between the two RTL sound-sets, and the notification sounds for the cog sounds were found to be more annoying. These findings refute H4. However, looking at a graph of the results for SEQ7 (Figure 7), there is a clear trend for participants to prefer using the cog sounds as part of their job: 77% disagreed they would be happy to listen to the comp sounds all day and only 4% agreed, whereas 58% disagreed for the cog sounds and 25% agreed. Two factors suggest this observation might be of use: firstly with more participants this difference may become significant, secondly, and more importantly, in future studies with longer exposure times to the sounds these results are likely to shift, becoming more pronounced. It is also noteworthy that modifications may need to be made to sonification design for users to be happy with workplace usage of data sonification, though this is caveated with the fact that users were exposed to the sounds for only a short time period.

In the follow-up study we observed that how easy it was to connect sounds to data type was dependent on individual sound design choices rather than the design approach, refuting H5. By following the suggestions in the literature for physical property to sound mappings the hazard types were as easy to identify as those designed with cognitive metaphor. It is apparent that participants were as able to connect the mappings of sound parameters to physical properties as suggested in (Dubus and Bresin, 2013) for the Comp sounds, as they were able to associate the contextual appropriateness and utilise their embodied knowledge (Roddy and Bridges, 2020) for the cog sounds.

The results of the follow-up study also demonstrated that due to adhering too closely to cognitive metaphor, data levels of the gas sound from that sound-set were harder to determine than for any other sound, partially confirming H6. However, the qualitative data from the main study demonstrated that participants experiencing the comp sounds found it harder to discern data levels when multiple sounds played simultaneously. This provides some validation to the cog approach where sounds are carefully designed to be used simultaneously. This finding suggests that should the comp design approach be more appropriate for a particular setting, the sounds could be improved by designing the sounds to better suit simultaneous

play. To do this the sounds may well need to deviate from the proposed mappings, i.e., a blend of the two approaches to sonification design.

## 8 Conclusion

In this paper we have detailed a novel approach to data sonification based on cognitive metaphor and ecological congruency (cog), applied to a virtual reality robot teleoperation system for use in nuclear decommissioning. While prior work has applied those principles to sound design generally, to the best of our knowledge prior work has not applied them to sonification design. Further, we evaluate this approach through experimental comparisons with a computationalist approach to sonification design (comp) based on recommended sound mappings in the literature. Our main finding was that the computationalist approach performed slightly better than the cognitive metaphor approach on a small subset of metrics, principally on predictability and mental workload. However, qualitative data analysis demonstrated that the cog approach resulted in sounds that were more intuitive, and were better implemented for spatialisation of data sources and data legibility when there was more than one sound source.

Our results also highlighted the need for more prolonged testing periods so that users could become more naturalised to system operation and the soundscape. Such longer form studies would allow formation of a better picture of the impact of design decisions on task performance and subjective experience without the obfuscation of system complexity. Further, the impact on system learning time and prolonged exposure effects on noise fatigue would both be valuable factors to study. It is important to note that a limitation of our findings is that the objective evaluation of the sonifications was done outside of the context of VR so our conclusions on their effectiveness in this context less definitive. Our proposed extended user testing protocols would allow us to effectively evaluate them in a VR context without the obfuscation of individual user differences.

## Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

## Ethics statement

The studies involving humans were approved by the Faculty Research Ethics Committee for the Faculty of environment and technology, The University of the West of England. The studies were conducted in accordance with the local legislation and institutional requirements. The participants provided their written informed consent to participate in this study.

## Author contributions

JS: Investigation, Methodology, Software, Writing—original draft, Writing—review and editing. PB: Conceptualization, Data curation, Formal Analysis, Funding acquisition, Investigation, Methodology, Software, Supervision, Writing—original draft, Writing—review and editing. TM: Conceptualization, Funding acquisition, Methodology, Supervision, Writing—original draft, Writing—review and editing. AB: Conceptualization, Software, Writing—original draft, Writing—review and editing. VM: Conceptualization, Formal Analysis, Funding acquisition, Investigation, Project administration, Supervision, Writing—original draft, Writing—review and editing.

## Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This research has been made possible thanks to funding from the UKRI Trustworthy Autonomous Systems Hub (EP/V00784X/1).

## References

- Barra, M., Cillo, T., De Santis, A., Petrillo, U. F., Negro, A., and Scarano, V. (2002). Multimodal monitoring of web servers. *IEEE Multimed.* 9, 32–41. doi:10.1109/MMUL.2002.1022857
- Barras, S. (2012). Sonifications for concert and live performance. *AI Soc.* 27, 281–283. doi:10.1007/S00146-011-0348-0
- Benyon, D. (2005). *Designing interactive systems: a comprehensive guide to HCI and interaction design*. 2 edn. Pearson Education Ltd.
- Bowman, D. A., and McMahan, R. P. (2007). Virtual reality: how much immersion is enough? *Computer* 40, 36–43. doi:10.1109/mc.2007.257
- Bremner, P., Mitchell, T. J., and McIntosh, V. (2022). The impact of data sonification in virtual reality robot teleoperation. *Front. Virtual Real.* 3. doi:10.3389/frvir.2022.904720
- Carlile, S. (2011). “Psychoacoustics,” in *The sonification handbook* (Berlin, Germany: Logos Publishing House), 41–61.
- Case, A. (2016). *Calm technology*. O'Reilly Media.
- Case, A., and Day, A. (2018). *Designing with sound: principles and patterns for mixed environments*. O'Reilly Media, Incorporated.
- Cooper, N., Milella, F., Pinto, C., Cant, I., White, M., and Meyer, G. (2018). The effects of substitute multisensory feedback on task performance and the sense of presence in a virtual reality environment. *PLoS one* 13, e0191846. doi:10.1371/journal.pone.0191846
- Dubus, G., and Bresin, R. (2013). A systematic review of mapping strategies for the sonification of physical quantities. *PLoS ONE* 8, e82491. doi:10.1371/journal.pone.0082491
- Ferguson, J., and Brewster, S. (2019). “Modulating personal audio to convey information,” in Conference on Human Factors in Computing Systems - Proceedings (New York, NY, USA: Association for Computing Machinery), 1–6.
- Ferguson, J., and Brewster, S. A. (2018). “Investigating perceptual congruence between data and display dimensions in sonification,” in Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (New York, NY, USA: Association for Computing Machinery), 1–9.
- Flowers, J. H. (2005). “Thirteen years of reflection on auditory graphing: promises, pitfalls, and potential new directions,” in International Conference on Auditory Display (ICAD).
- Gresham-Lancaster, S. (2012). Relationships of sonification to music and sound art. *AI Soc.* 27, 207–212. doi:10.1007/S00146-011-0337-3
- Grond, F., and Berger, J. (2011). *Parameter mapping sonification*. Berlin, Germany: Logos Publishing House.
- Grond, F., and Sound, T. H. (2014). Interactive Sonification for Data Exploration: how listening modes and display purposes define design guidelines. *Organised Sound.* 19, 41–51. doi:10.1017/S1355771813000393
- Hall, S. (1980). *Encoding and decoding in the television discourse*. 1 edn, 2. Routledge.

## Acknowledgments

We would like to thank Prof. Paul White for providing help with the statistical analysis and Philip Tew for supporting the development of the virtual environment.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Hart, S. G., and Staveland, L. E. (1988). Development of NASA-TLX (task load index): results of empirical and theoretical research. *Adv. Psychol.* 52, 139–183. doi:10.1016/S0166-4115(08)62386-9

Hermann, T., Hunt, A., and Neuhoff, J. G. (2011). *The sonification handbook*. Berlin: Logos Publishing House, 586.

Hermann, T., Niehus, C., and Ritter, H. (2003). “Interactive visualization and sonification for monitoring complex processes,” in Proceedings of the 2003 International Conference on Auditory Display, 247–250.

Holloway, L. M., Goncu, C., Ilisar, A., Butler, M., and Marriott, K. (2022). “Infosonics: accessible infographics for people who are blind using sonification and voice,” in Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (New York, NY, USA: Association for Computing Machinery), 1–13.

Kantan, P. R., Spaich, E. G., and Dahl, S. (2021). “A metaphor-based technical framework for musical sonification in movement rehabilitation,” in The 26th International Conference on Auditory Display (ICAD 2021).

Kern, A. C., and Ellermeier, W. (2020). Audio in vr: effects of a soundscape and movement-triggered step sounds on presence. *Front. Robotics AI* 7, 20. doi:10.3389/frbot.2020.00020

Körber, M. (2019). “Theoretical considerations and development of a questionnaire to measure trust in automation,” in Proceedings of the 20th Congress of the International Ergonomics Association (IEA 2018) Volume VI: Transport Ergonomics and Human Factors (TEHF), Aerospace Human Factors and Ergonomics 20 (Springer), 13–30.

Kramer, G. (1994). “Some organizing principles for representing data with sound,” in *Auditory display: sonification, audification and auditory interfaces* (Reading, MA: Addison-Wesley), 285–221.

Kramer, G., Walker, B., Bonnebright, T., Cook, P., and Flowers, J. H. (2010). *Sonification report: status of the field and research agenda*. International Community for Auditory Display.

Kumar, S., von Kriegstein, K., Friston, K., and Griffiths, T. D. (2012). Features versus feelings: dissociable representations of the acoustic features and valence of aversive sounds. *J. Neurosci.* 32, 14184–14192. doi:10.1523/jneurosci.1759-12.2012

Kurta, L., and Freeman, J. (2022). “Targeting impact: a new psychological model of user experience,” in *International conference on human-computer interaction* (Springer), 196–212.

Lakoff, G. (2006). “Conceptual metaphor,” in *Cognitive linguistics: basic readings*, 34, 185.

Larsson, P., Västfjäll, D., Olsson, P., and Kleiner, M. (2007). “When what you hear is what you see: presence and auditory-visual integration in virtual environments,” in Proceedings of the 10th annual international workshop on presence, 11–18.

Lokki, T., and Gröhn, M. (2005). Navigation with auditory cues in a virtual environment. *IEEE Multimed.* 12, 80–86. doi:10.1109/MMUL.2005.33

- Malbos, E., Rapee, R. M., and Kavakli, M. (2012). Behavioral presence test in threatening virtual environments. *Presence Teleoperators Virtual Environ.* 21, 268–280. doi:10.1162/pres\_a\_00112
- Meijer, P. B. (1992). An experimental system for auditory image representations. *IEEE Trans. Biomed. Eng.* 39, 112–121. doi:10.1109/10.121642
- Mitchell, T. J., Jones, A. J., O'Connor, M. B., Wonnacott, M. D., Glowacki, D. R., and Hyde, J. (2020). "Towards molecular musical instruments: interactive sonifications of 17-alanine, graphene and carbon nanotubes," in Proceedings of the 15th International Audio Mostly Conference (New York, NY, USA: Association for Computing Machinery), 214–221.
- Mitchell, T. J., Thom, J., Pountney, M., and Hyde, J. (2019). "The alchemy of chaos: a sound art sonification of a year of tourette's episodes," in 25th International Conference on Auditory Display (ICAD 2019) (International Community for Auditory Display), 163–168.
- Murphy, D., and Skarbez, R. (2020). What do we mean when we say "presence". *Presence Virtual Augmented Real.* 29, 171–190. doi:10.1162/pres\_a\_00360
- Neuhoff, J. G., Wayand, J., and Kramer, G. (2002). Pitch and loudness interact in auditory displays: can the data get lost in the map? *J. Exp. Psychol.* 8, 17–25. doi:10.1037/1076-898X.8.1.17
- Preziosi, M. A., and Coane, J. H. (2017). Remembering that big things sound big: sound symbolism and associative memory. *Cognitive Res. Princ. Implic.* 2, 10–21. doi:10.1186/s41235-016-0047-y
- Robinson, F. A., Velonaki, M., and Bown, O. (2021). "Smooth operator: tuning robot perception through artificial movement sound," in ACM/IEEE International Conference on Human-Robot Interaction, Boulder, CO, USA, 09–11 March 2021 (IEEE), 53–62.
- Roddy, S., and Bridges, B. (2020). Mapping for meaning: the embodied sonification listening model and its implications for the mapping problem in sonic information design. *J. Multimodal User Interfaces* 14, 143–151. doi:10.1007/s12193-020-00318-y
- Salanitri, D., Lawson, G., and Waterfield, B. (2016). "The relationship between presence and trust in virtual reality," in Proceedings of the European Conference on Cognitive Ergonomics (New York, NY, USA: Association for Computing Machinery).
- Sawe, N., Chafe, C., and Treviño, J. (2020). Using data sonification to overcome science literacy, numeracy, and visualization barriers in science communication. *Front. Commun.* 5. doi:10.3389/fcomm.2020.00046
- Schewe, F., and Vollrath, M. (2020). Ecological interface design effectively reduces cognitive workload—the example of hmis for speed control. *Transp. Res. part F traffic Psychol. Behav.* 72, 155–170. doi:10.1016/j.trf.2020.05.009
- Schubert, T., Friedmann, F., and Regenbrecht, H. (2001). The experience of presence: factor analytic insights. *Presence Teleoperators Virtual Environ.* 10, 266–281. doi:10.1162/105474601300343603
- Shamma, S. A., Elhilali, M., and Micheyl, C. (2011). Temporal coherence and attention in auditory scene analysis. *Trends Neurosci.* 34, 114–123. doi:10.1016/j.tins.2010.11.002
- Simmons, J., Bown, A., Bremner, P., McIntosh, V., and Mitchell, T. J. (2023). "Hear here: sonification as a design strategy for robot teleoperation using virtual reality," in VAM-HRI Virtual, Augmented, and Mixed-Reality for Human-Robot Interactions at HRI, Stockholm, April 27 2023 (Stockholm, Sweden: Association for Computing Machinery), 1–9.
- Sorkin, R. D. (1988). Why are people turning off our alarms? *J. Acoust. Soc. Am.* 84, 1107–1108. doi:10.1121/1.397232
- Storey, J. D. (2011). "False discovery rate," in *International encyclopedia of statistical science* (Springer), 504–508.
- Triantafyllidis, E., Mcgreavy, C., Gu, J., and Li, Z. (2020). Study of multimodal interfaces and the improvements on teleoperation. *IEEE Access* 8, 78213–78227. doi:10.1109/access.2020.2990080
- University of Manchester (2024). *RNE: robotics for nuclear environments*.
- Vickers, P. (2016). "Sonification and music, music and sonification," in *The routledge companion* (London: Routledge), 135–144.
- Vickers, P., Laing, C., Debashi, M., and Fairfax, T. (2014). Sonification aesthetics and listening for network situational awareness. arXiv preprint arXiv:1409.5282.
- Walker, B., Rehg, J., Kalra, A., and Winters, R. (2019). *Dermoscopy diagnosis of cancerous lesions utilizing dual deep learning algorithms via visual and audio (sonification) outputs: Laboratory and prospective*. Elsevier.
- Walker, B. N., and Kramer, G. (2005). Mappings and metaphors in auditory displays: an experimental assessment. *ACM Trans. Appl. Percept. (TAP)* 2, 407–412. doi:10.1145/1101530.1101534
- Walker, M., Phung, T., Chakraborti, T., Williams, T., and Szafir, D. (2022). Virtual, augmented, and mixed reality for human-robot interaction: a survey and virtual design element taxonomy. arXiv preprint arXiv:2202.11249.
- Whitney, D., Rosen, E., Phillips, E., Konidaris, G., and Tellex, S. (2019). "Comparing robot grasping teleoperation across desktop and virtual reality with ros reality," in *Robotics research: the 18th international symposium ISRR* (Springer), 335–350.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bull. Rev.* 9, 625–636. doi:10.3758/bf03196322
- Wirfs-Brock, J., Fam, A., Devendorf, L., and Keegan, B. (2021). Examining narrative sonification: using first-person retrospection methods to translate radio production to interaction design. *ACM Trans. Computer-Human Interact.* 28, 1–34. doi:10.1145/3461762
- Zahray, L., Savery, R., Syrkett, L., and Weinberg, G. (2020). "Robot gesture sonification to enhance awareness of robot status and enjoyment of interaction," in 29th IEEE International Conference on Robot and Human Interactive Communication, RO-MAN 2020, Naples, Italy, 31 August 2020–04 September 2020 (IEEE), 978–985.
- Ziemer, T., and Schultheis, H. (2020). Linearity, orthogonality, and resolution of psychoacoustic sonification for multidimensional data. *J. Acoust. Soc. Am.* 148, 2786. doi:10.1121/1.5147752