



AudioMiXR: Spatial Audio Object Manipulation with 6DoF for Sound Design in Augmented Reality

BRANDON WOODARD[†], Brown University, United States

MARGARITA GELETA[†], University of California at Berkeley, United States

JOSEPH J. LAVIOLA JR., University of Central Florida, United States

ANDREA FANELLI, Dolby Laboratories, United States

RHONDA WILSON, Dolby Laboratories, United States




Fig. 1. (Left) A schematic illustrating a user surrounded by virtual audio objects in 3D space. The arrows represent the virtual-audio-object panning initiated by the user's gesture; (Right) Real-world images of AudioMiXR users employing directed free-hand manipulation for music and cinematic sound design, respectively.

We present AudioMiXR, an augmented reality (AR) interface intended to assess how users manipulate virtual audio objects situated in their physical space using six degrees of freedom (6DoF) deployed on a head-mounted display (Apple Vision Pro) for 3D sound design. Existing tools for 3D sound design are typically constrained to desktop displays, which may limit spatial awareness of mixing within the execution environment. Utilizing an XR HMD to create soundscapes may provide a

[†]Both authors contributed equally to this research. Work primarily done at Dolby Laboratories.

Authors' Contact Information: [Brandon Woodard](#), Brown University, United States, brandon_woodard@brown.edu; [Margarita Geleta](#), University of California at Berkeley, United States, geleta@berkeley.edu; [Joseph J. LaViola Jr.](#), University of Central Florida, United States, jjl@cs.ucf.edu; [Andrea Fanelli](#), Dolby Laboratories, United States, andrea.fanelli@dolby.com; [Rhonda Wilson](#), Dolby Laboratories, United States, rhonda.wilson@dolby.com.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

ACM 2474-9567/2025/9-ART140

<https://doi.org/10.1145/3749478>

real-time test environment for 3D sound design, as modern HMDs can provide precise spatial localization assisted by cross-modal interactions. However, there is no research on design guidelines specific to sound design with 6DoF in XR. To provide a first step toward identifying design-related research directions in this space, we conducted an exploratory study where we recruited 27 participants, consisting of expert and non-expert sound designers. The goal was to assess design lessons that can be used to inform future research venues in 3D sound design. We ran a within-subjects study where users designed both a music and cinematic soundscapes. After thematically analyzing participant data, we constructed two design lessons: (1) Proprioception for AR Sound Design, and (2) Balancing Audio-Visual Modalities in AR GUIs. Additionally, we provide application domains that can benefit most from 6DoF sound design based on our results. To expand on these insights, we conducted a second within-subjects study comparing AudioMiXR to a 2D panner baseline. Results show that AudioMiXR significantly improved usability (SUS), reduced frustration and mental workload (NASA-TLX), and enhanced creativity across all subscales. These findings demonstrate that 6DoF AR interaction not only introduces novel experiential affordances, but also yields measurable gains in user experience and creative output, positioning AudioMiXR as a promising foundation for future AR-based sound design tools.

CCS Concepts: • **Human-centered computing** → **Mixed / augmented reality**; **Auditory feedback**; **Gestural input**.

Additional Key Words and Phrases: Augmented Reality, Spatial Audio, Sound Design, Interactive Technologies

ACM Reference Format:

Brandon Woodard, Margarita Geleta, Joseph J. LaViola Jr., Andrea Fanelli, and Rhonda Wilson. 2025. AudioMiXR: Spatial Audio Object Manipulation with 6DoF for Sound Design in Augmented Reality. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 3, Article 140 (September 2025), 41 pages. <https://doi.org/10.1145/3749478>

1 INTRODUCTION

Traditional approaches to spatial audio, such as channel-based surround sound systems, have been gradually supplanted by the more versatile and immersive object-based audio format [49, 66]. This format facilitates placement of each sound source as independent audio objects within a 3D space. Although existing tools for spatial audio panning within audio mixing frameworks provide more precise and flexible spatialization, they remain confined to 2D user interfaces [39, 47]. This limitation constrains the creative freedom needed for immersive mixes and spatialized audio editing, and the transition towards the object-based audio paradigm has introduced implementation complexities that remain unresolved. One such challenge lies in the sequential nature of 3D panning. Often, movement along the z-axis is constrained after the initial adjustments in the x-y plane, limiting the creative flexibility and the natural coherence of spatial audio compositions. Moreover, the complexity of user interfaces for spatial audio panners can result in a steep learning curve. Further, binaural rendering can complicate audio object localization within mixing environments due to the translation between 3D user interfaces (3DUI) rendered on desktop displays and real-world applications [12].

Currently, most commercial Digital Audio Workstations (DAWs) and sound design software still rely on planar controls to manipulate the placement and motion of 3D audio objects [47]. This setup forces audio engineers and producers to mentally map 3D sound onto 2D screens, often resulting in higher cognitive load [25]. Extended reality (XR) – comprising both augmented reality (AR) and virtual reality (VR) – offers an opportunity to overcome these limitations, especially when six degrees of freedom (6DoF) is employed. Within a 6DoF XR interface, users can move freely in real or virtual space (translation and rotation) rather than just relying on purely rotational or “point-and-click” interactions. Such physical freedom has been shown to facilitate more accurate cognitive map building [60, 62], and can enable direct manipulation of virtual audio objects situated within the user’s environment.

Motivated by the challenges identified through conversations with professional audio engineers – including the difficulty of creating immersive mixes, localizing binaurally rendered sound sources, and mastering steep learning curves for 3D panners – we developed AudioMiXR, an XR-based user interface for sound design and audio mixing, designed for XR-capable head-mounted displays (HMDs), offering free-hand gestural interactions

with visually rendered 3D audio objects. While it is focused on AR HMDs in its current implementation, it can be adapted to VR contexts. By leveraging how the human perceptual system understands spatial relationships, AudioMiXR aims to make immersive sound design and audio mixing more intuitive, and reduce the cognitive load of 3D panning.

This premise leads to our research questions, which seek to quantify and characterize the benefits of spatial audio interaction techniques in AR:

RQ1. How should we design an XR sound design interface that leverages 6DoF?

RQ2. What application areas will benefit most from an AR sound design interface with 6DoF?

2 RELATED WORK

We present an overview of relevant concepts and prior work in 3D sound design and rendering, spatial localization in XR and the utilization of expanded degrees of freedom, a brief outline of sonified experiences in XR, and XR tools for sound design and audio mixing.

2.1 3D Sound Design and Rendering

Modern audio production is frequently carried out within DAWs – software environments for music creation [47]. Popular DAWs such as Cubase, Nuendo, Logic Pro, Ableton, and Pro Tools provide visual interfaces for manipulating digital audio, often adopting a skeumorphic design to replicate the look and feel of legacy hardware (e.g., mixers, drum machines, or synthesizers). While this approach helps users transition from traditional studio setups by mimicking familiar controls and interaction paradigms, it can constrain the interfaces to 2D representations. Furthermore, most DAWs rely on a timeline-based structure which encourages the sequential arrangement of audio objects [47]. This temporal editing view does not take into account the “spatial” aspect of audio. Under the *object-based audio* paradigm, each sound source (or “object”) is treated as an independent entity accompanied by metadata describing its spatial attributes – namely, the location of *spatial audio* in a 3D space [8, 45]. For spatial editing, DAWs offer 3D panners, which are tools to control the position of each audio object in a 3D space [3, 49]. A prime example is the Dolby Atmos Panner available in Pro Tools.

There are several approaches to spatial audio rendering, and we can highlight two in the context of object-based audio workflows [66]. The first is *loudspeaker stereophony*, which consists in placing multiple speakers around the listener to create a surround sound experience commonly used in cinemas, home theaters, and gaming. A limitation of loudspeaker-based setups is that the listener must typically remain in an optimal “sweet spot” to fully appreciate the intended spatial effect, and phantom sound sources (i.e., perceptual illusions of sound sources) are restricted to positions between the speakers. The second approach, thanks to which many of audio-augmented XR applications are possible, is *binaural rendering*. This method reproduces 3D audio using only two channels, simulating how sound is perceived by the human auditory system [12, 50]. By convolving audio signals with head-related transfer functions (HRTFs), which model how sound interacts with anthropometric features, binaural rendering can simulate the localization, timbre, and externalization of sound sources, creating a realistic perception of spatial objects in a 3D environment. However, this technique can suffer from perceptual errors, such as front-back or up-down confusions and “in-head localization” effects [78].

Once a spatial mix is finalized, the sequence of objects is sent to the renderer – for example, the Dolby Atmos Renderer (potentially configured with personalized HRTFs). The renderer interprets the scene metadata to produce the final output for various configurations, whether loudspeaker arrays (e.g., 5.1.4, 7.1.4, 22.1) or headphone-based binaural playback.

2.2 Spatial Localization with Expanded Degrees of Freedom in XR

XR spans both augmented reality (AR) and virtual reality (VR), combining the physical world with a digital twin world in an interactable environment [54]. A central factor in XR experiences is the notion of *degrees of freedom* (DoF), which refers to the number of independent ways users can move within these virtual or augmented environments, and consists of translational and rotational components. Per convention, 3DoF restricts head tracking to the three rotational axes (pitch, yaw, roll), allowing users to look around but not physically move in space. A more immersive XR experience comes from adding to the equation the translational component (right/left, up/down, forward/back), resulting in 6DoF movement, which enables a more immersive and naturalistic exploration and navigation. An essential component for a user's navigation in an XR environment is building the cognitive map – an internal representation of locations within a world-reference frame. Prior work has shown that 6DoF movement develops a more accurate cognitive map compared to 3DoF or purely joystick-based navigation [60, 62]. When users can physically walk around, rather than just pivot in place or point a controller, they gain a richer “movement fidelity” of real-world locomotion.

Beyond navigation, 6DoF also permits precise placement of virtual objects. Instead of pointing a device (e.g., controller) or hand gesture in 3DoF to interact with graphical user interfaces (GUIs) situated spatially around them, with 6DoF, users can physically move to the target location and place those virtual directly [41]. This approach magnifies the realism of XR soundscapes in 6DoF audio experiences for storytelling or entertainment. For instance, many 360 degree cinematic experiences are in 3DoF and allow the user to look around the scene as if they were also co-located with the actors. Video games involving 6DoF interactions facilitate realism due to the envelopment of a user's entire body as an input device to interact with the content, encouraging direct engagement with XR content [61].

Spatial audio is central to audio augmentation, such that virtual sounds are perceived as emanating from specific locations in 3D space [63]. Humans rely on multiple sensory modalities when they engage with their environment [41, 52, 77], and the auditory sense remains highly significant for localization even when visual cues is limited, sometimes replacing visual information altogether (e.g., “watching” television from another room) [18, 34]. Research indicates that the use of spatial audio encourages users to adopt a more active role in spatial navigation, leading to even more accurate cognitive maps [16], while simultaneously reinforcing the sense of presence in XR environments [42]. As reported in [77], audio interaction techniques should present directional and distance cues in contextually meaningful ways. Finally, an enhanced feeling of immersion often requires not just “looking and hearing around”, but also “moving around” in 6DoF to achieve the full benefits of spatial localization [25].

2.3 Sound Experiences in XR

Sound experiences in XR focus on interactions that leverage multimodal GUIs afforded by AR or VR head-mounted displays (HMDs), providing audio experiences that are more expressive and facilitate engagement for users. A notable example is *Spatial Orchestra*, an AR interface, presented by Kim et al. [43] which allowed users to walk into “bubbles” in a fixed position and co-located near each other emitting musical notes, that can only be heard once they were fully inside the bubble. This provided users with an audio experience where they could interact with music with their body without having to play a traditional instrument; although in order for the interface to work, users had to remain within a confined virtual space that was 3.3 m by 3.3 m, thus limiting their translational movements. This work is one example of audiovisual experiences that have taken advantage of the spatial cues to visualize and interact with audio using methods to render 3D visualizations of sound that are reported in the literature to form stronger connections with the audio content engaging multiple senses of users due to the multimodal feedback [2, 10, 43]. Other forms of sound experiences in XR include VR or AR music-based video games like *Beat Saber* where users wield a saber to hit oncoming blocks containing parts of

the song's beat [7]. In this game users are able to stand up and use their body's orientation in a fixed location to interact with the game. The visual-aural feedback provided in music-based video games provide an engaging experience where musical notes tend to take a physical form engaging users in a multisensory experience. These works consist of different combinations of haptic feedback, visual representations, 3DoF, or physical interactions. However, these approaches are experiential and do not allow for precise control of audio to compose soundscapes intended for 3D audio, for example, music, cinematic scenes, or video games. To illustrate, [57] and [66] discussed the use of spatial sound in VR for creating immersive virtual environments with a focus on computer games.

Research in audio augmented reality (AAR) has leveraged the physical environment to augment auditory feedback rendered to the user targeted *navigation and location-awareness*, as spatialized “audio beacons”, verbal instructions, or guides [44, 70], studies have underscored its potential for diverse domains, including *presentation and display*, particularly in museums or archaeological sites that use audio augmentations to deliver contextual information about cultural artifacts [33, 51], or increase engagement with visitors at art exhibitions through spatialized sonic artwork like *Sonic Sculpture* [48].

2.4 XR Tools for Sound Design and Audio Mixing

A subset of design tools in XR utilize the flexibility of a spatial interface to interact with floating GUIs comprising of drag-and-drop UIs that can be manipulated with controllers. A notable commercial example is *DearVR*, which uses virtual knobs and faders with a similar appearance of a music production studio [19]. Several research works in sound design [6, 40] have utilized a VR headset to recreate a virtual mixing environment, similar to the panning tools in DAWs, where users can control audio objects positioned at a distance with a controller and raycast to select the objects while remaining stationary. While all of these tools utilize some spatial aspects applied to sound design, they rely on a stationary user position and none of them exploit 6DoF direct manipulation, in which users can physically move around the environment to locate and place virtual audio objects. By contrast, AudioMiXR integrates 6DoF direct free-hand manipulation into the audio mixing workflow, allowing the user to traverse a real or virtual environment, directly grab virtual audio objects, and reposition them wherever they see fit – this way, enabling true *hands-on* editing of spatial sound as if the users were “inside the panner”, in ways conventional DAW panners or existing VR mixing tools cannot as seen in Figure 2.

3 DESIGN CONSIDERATIONS

To inform the early design of AudioMiXR, we conducted semi-structured formative interviews with four professional audio mixers (E1, E2, E3, E4), and followed sound design principles [25] and AR design guidelines [23] to ensure AudioMiXR aligns with both established spatial audio mixing conventions and ergonomic AR design. We identify experts as people who use DAW software at a professional capacity and recruited from an industry-leading company specializing in audio technology. The insights from the interviews are listed next.

3.1 Expert Formative Interview

For our expert interviews, we adapted the structure of the formative interview method by [35]. Each one revolved around a set of 12 questions designed to uncover established spatial audio workflows, technical challenges, and preferences. These were used to design AudioMiXR by maximizing the prospective benefits of an AR interface. Refer to the appendix for the complete list of questions.

3.1.1 Mixing Workflow. We began by asking experts to describe their general processes for mixing and panning spatial audio, as well as the software tools they typically use. When describing, experts emphasized how they use audio object placement in 3D to support the narrative of sonified experiences and how the visual-aural feedback is key to create a compelling experience. E2 noted that if a visual on-screen object moves, the sound

Interface	3DoF (Orientation)	6DoF (Orientation & Pos.)	Real-Time Spatialization	Real-Time Audio Control	VR	AR	Standalone	Unbounded Space
Bargum et al.					✓			
Spatializing music in VR			✓		✓		✓	
BoomRoom	✓		✓	✓				
SPARTA			✓	✓				
Dolby Atmos Panner	✓		✓	✓				
Avid Pro Tools Ultimate			✓	✓				
ASR – Augmented Sound Reality						✓		
Evaluating Polaris	✓		✓			✓		
DearVR Connect (panner)	✓		✓	✓	✓			
AudioMiXR	✓	✓	✓	✓	✓	✓	✓	✓

Fig. 2. Structured comparison of AudioMiXR versus traditional 2D audio tools across key dimensions. AudioMiXR combines 6DoF spatial manipulation, 6DoF binaural rendering with respect to the user’s position and orientation real-time audio control, and immersive unbounded environments. These features are absent or more limited in desktop-based workflows or previous AR/VR interfaces intended for mixing spatial audio. Interfaces (from top to bottom): [5, 6, 10, 19, 21, 22, 40, 56, 69]

should follow its motion to create a convincing illusion that they are “*in sync*” – an observation that highlights the importance of perceptual alignment between visual cues and audio sources. This notion reflects the AR design principle of *alignment between physical-virtual worlds* [23], ensuring that the listener’s perceptual system “*automatically accepts*” the audio-visual synchronization. Experts stressed that the overall focus of a spatial mix is to create a *compelling* environment – realistic or stylized – that meets the goal of “*telling a story or supporting*” a particular user experience. In other words, the spatial mix must reflect the intended function, whether it is cinematic storytelling, gaming, or sound meditation, mirroring the AR design principle of *fitting with the user environment and the task* [23].

3.1.2 Spatialization. When asked about specific spatialization strategies, E1, E2, E3 agreed that it is ultimately “*done to taste*”, although certain heuristics commonly apply. For instance, low-frequency elements (like bass or kick drums) often remain front and center, while higher-frequency elements might be placed in overhead channels. These practicalities resonate with how the human perceptual system associates certain frequency bands with location [24, 29], such that the audio object arrangement feels coherent to the listener’s mental model of the environment. In other words, it should be familiar enough so that the listener’s expectations are met (i.e., “*the waves should be there, the birds over there*”).

3.1.3 Challenges of Spatial Audio. The experts identified localization issues while mixing – particularly front/back confusion in binaural rendering – as a frequent challenge. They noted that point-source audio objects can fail to convey an effective sense of immersion if the sounds remain narrowly localized or singular. Both E1 and E3 mentioned how changing the size or duplicating audio objects to occupy more of the soundscape could address this issue by providing multiple directional cues or a broader “*spread*”. Such references to *spread*, *duplication*, and *overlapping elements* yet again reflect the sound design principle of *layering* [25], where combining or stacking multiple sounds adds depth and complexity to the mix. It also connects with the AR principle of *form communicates function* [23], since manipulating the visual appearance of an audio object maps directly to certain auditory attribute changes, creating affordances.

3.1.4 Collaborative Design. Because spatial audio often involves interdisciplinary collaboration, we also inquired about the social and collaborative aspects of mixing. In current workflows, experts typically exchange binaural mixes remotely, receiving asynchronous feedback from listeners. An AR environment could potentially streamline this iterative process by allowing co-located or remote collaborators to view and hear each other’s manipulations in real-time. A scenario where collaborators can be in the same room with an AR mixing interface synced with loudspeakers in the physical space or while one ‘operator’ mixes was discussed.

3.1.5 Future Directions. We also asked participants about notable successes in their spatial audio practices, and future directions that might enhance their workflows. General directions pointed toward improving the UIs to be more user-friendly and intuitive and, interestingly, E3 hinted at how a more physical approach could make mixing easier and more accessible: “*certainly using body movement to create and automate panning. You could almost ‘dance the mix’. This also opens up accessibility*” – E3. This notion of using one’s body movement and gestures directly touches on the applicability of AR in an audio mixing task, encouraging gestures that feel natural rather than fatiguing, referencing the AR design heuristic of *fitting with user’s physical abilities* [23].

3.1.6 Evaluation of a Mix. Finally, experts reflected on how they typically evaluate the quality of a spatial audio mix, describing metrics for success. For instance, they consider whether it “*sounds more interesting than [in] stereo*” or whether it has “*a good sense of height, width, depth, externalization, and timbre*”. Although these mostly rely on subjective measures of audio mixer’s taste, E2 noted: “*the one thing that can be truly measured is the loudness*”. Generally measured in LUFS units for audio mixing applications, *loudness* serves as an objective metric for ensuring that the final mix meets the delivery standards across different platforms.

3.2 Key Takeaways

Experts in our formative interviews revealed several core considerations for spatial audio mixing in AR which we took into account for designing AudioMiXR. Foremost, we clarified why AudioMiXR could be better than traditional 2D audio panning interfaces. Its advantage lies in the natural conversion of virtual audio objects in a 3D space, which makes it easier for creators to mold compelling experiences for diverse applications. Experts stressed that ensuring audio objects are synchronized with their corresponding visuals is crucial for immersion. Moreover, experts highlighted the benefits of a more physical, gesture-driven interface. Following on their feedback, we made sure to design AudioMiXR with real-time aural-visual feedback and to provide free-hand spatial manipulation within our interface. However, at the same time, localization challenges remain an obstacle. We leverage the audiovisual sensory component of our AR interface since humans are better at that than in unisensory localization scenarios [53, 58].

Our experts also recognized the collaborative nature of spatial audio production. An AR approach could support more synchronous or alternative forms of collaboration – allowing multiple users to see and hear each

other's changes in real time. In all, the interviews point to the promise of AR-based mixing for bridging the mapping gap between the virtual spatial mix and the actual physical implementation and delivery. Taken together, these insights form the backbone of AudioMiXR. We present its detailed system design in the following section.

4 SYSTEM

This section describes the AudioMiXR system, an XR user interface for free-hand manipulation of spatial audio. We detail the system architecture, and provide an overview of user interaction mechanisms and the standard workflow.

4.1 System Architecture

To enable free-hand spatial audio manipulation, AudioMiXR relies on several hardware and software components. On the hardware side, the system runs on an XR-capable HMD which provides a built-in device camera and depth sensors to capture the user's hand gestures and movements, producing 3D coordinates for real-time gesture recognition. Additionally, the inertial measurement unit (IMU) within the HMD tracks the position and orientation relative to a defined origin. This ensures that the system consistently aligns the user's real world head movement with the superimposed virtual environment.

On the software side, AudioMiXR is comprised of a suite of modules. First, an *image tracking and gesture recognition* module processes the depth-camera feed and applies HMD's native computer vision algorithms to interpret how the user's hands move and which gestures they perform. Next, a *3D mapping* module translates these hand positions into virtual coordinates, accurately tracking and registering the hand positions within a 3D coordinate space. A *gesture I/O* module then assigns a class of a specific gesture to input and output functions in order to interact with the visual 3D audio objects. To render and update these virtual audio objects, AudioMiXR has a *3D simulation* module – whenever the system detects a relevant user gesture, the 3D simulation module handles how the the visual 3D objects should respond to the users input (e.g., the user moves their hand up and down along the vertical axis then the audio object vertically translates in accordance to the hand movement and positioned at the extent of their fingertips). The *communication module* updates the 3D hand coordinates and visual audio object positions for a set interval while the application is running (e.g. updates audio object position once every 30 frames via USB or wireless protocol). The HMD's spatial operating system, processor, and memory provide the necessary computational resources to execute computer vision algorithms, maintain stable 3D rendering, and track hand gestures while simultaneously updating audio objects.

4.2 Deployment Target

Although both AR and VR share common goals of augmenting user environments, they differ in key respects. VR often disconnects users visually from their surroundings, risking collisions or accidents while they are interacting with the digital twin world without having cues on their real environment [38, 73]. Even when using VR safety measures, such as boundaries (e.g., Meta Quest Guardian System) users may still break them, and safety comes at the expense of immersion [73]. VR immersion can also induce *cybersickness* – a visual-vestibular conflict from motion in a virtual world – which may intensify if, for instance, users rely on teleportation for locomotion [77] or omnidirectional treadmills for infinite-walking [14], which even after proper training do not provide firm locomotion, but rather an awkward gait [62]. Ensuring perceptual-motor fidelity between a user and their digital-twin avatar can mitigate these negative effects by “virtual body ownership” [4, 67].

Conversely, AR situates virtual content onto physical environments, allowing users to interact with virtual elements as though they were truly present in their surroundings [63, 68, 78], often avoiding some of VR's drawbacks, as users retain awareness of their physical environment [71, 73]. Given these considerations, our user study focused on the AR setting, although future extensions could address VR scenarios, since AudioMiXR

runs on headset capable of supporting both AR and VR modes. While the immersive 6DoF capabilities in AR and VR are comparable in principle, VR introduces additional aforementioned caveats that might require further design adaptations.

4.3 User Interaction and GUI

In AudioMiXR, we utilize rendered computer graphics to depict 3D audio object manipulations (e.g., position in space or distance from the listener). The visual representation of each sound objects is modified through a free-hand pinch gesture used for selecting and holding objects. AudioMiXR uses simple, direct gestures to give the sensation of physically touching and/or moving virtual audio objects in the air. As shown in Figure 3, users can “pinch” a virtual audio object to grab it, “push” it further away, or “pull” it closer – even if they remain standing at a distance.




Fig. 3. (Left-to-Right) AudioMiXR enables users to reposition audio objects within their reach envelope through direct gestural interactions, giving the impression that their hands are physically in contact with the objects. Users can push objects beyond their reach envelope by targeting the object with a pinch gesture and performing a pushing motion. If an object is out of reach and the user prefers not to physically walk towards it, they can align their pinch gesture with the object in their line of sight and “pull” it back into their reach envelope.

4.4 User Workflow

Once calibration is complete, the user can start the session by booting the XR interface, which initializes the *gesture recognition*, *hand tracking*, and *body tracking* modules. Once in the environment, the user can freely grab any virtual audio object by pinching it (Figure 4) and reposition it within or beyond their physical reach. If the target object lies out of reach, the user may “pull” it back by targeting it with the gaze, performing the pinch gesture, and making a dragging motion. Changes in position immediately translate to perceptible alterations in audio attributes.




Fig. 4. A user interacting with virtual audio objects in AR using direct free-hand manipulation.

4.5 Implementation

We developed AudioMiXR using the Unity game engine (runtime version 2022.3.30f1) in C#, deploying it on an Apple Vision Pro (AVP) HMD, on its native spatial operating system visionOS 1.0. We used Unity for its robust developer community and comprehensive documentation and support for rendering and interaction. We used a Unity Professional license to gain access to the necessary AVP deployment modules and visionOS-specific libraries, and an Apple Developer license to sign, build, and deploy AudioMiXR to the AVP HMD.

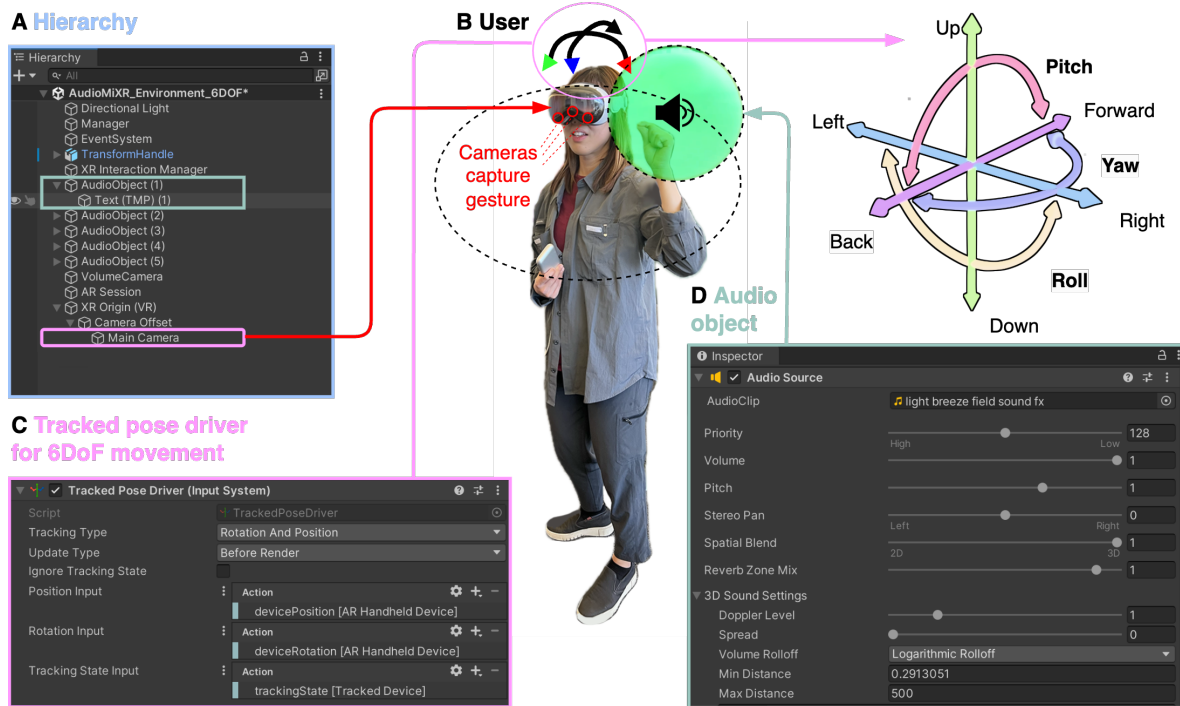


Fig. 5. (A) Hierarchy in the Unity3D engine showing all the game objects; (B) A user featuring the headset with a tracked pose driver enabling 6DoF movement and gesture-based interaction with a virtual audio object; (C) Configuration settings for Unity Engine's spatial tracking driver; (D) Settings of the audio source element attached to the virtual audio object.

4.5.1 Hardware and Tracking. The AVP provides 6DoF tracking through its built-in IMU sensors and LiDAR-based environment mapping, allowing real-time position (right/left, up/down, forward/back) and orientation (pitch, yaw, roll) updates. As illustrated in Figure 5.C, we leverage Unity’s *Tracked Pose Driver* component on the HMD camera to enable positional and rotational tracking, configured to update sample tracking input right before the rendering step for minimal latency. Our *XR origin* in Unity contains the audio listener – a component that captures the incoming sound – and references the HMD’s head pose using the XR Interaction toolkit, while an additional *AR session* object manages environment tracking (position and rotation) via ARKit, provided by the Unity AR Foundation library. To support free-hand spatial audio manipulation, we incorporate the AVP’s integrated 3D touch input using the PolySpatial engine, allowing users to “pinch” virtual objects and reposition them in 3D space when recognized by the AVP’s cameras (Figure 5.B). For audio playback, users can opt for AVP’s native near-ear headphones or Bluetooth-enabled headphones.

4.5.2 Audio Objects and Interaction. Each audio object in our scene (Figure 5.A) is represented as a spherical *GameObject* Unity’s fundamental object type, with the following added components: a Unity *audio source* component (Figure 5.D) configured for 3D spatialization with logarithmic rolloff for distance-based attenuation; a custom *opacity modulation script* that adjusts visual transparency relative to the distance from the *XR origin* (see left subfigure in Figure 6); a custom *pulse-transferring script* employing a Fast Fourier Transform (FFT) with a Blackman window to sample the audio signal in real time. We segment the FFT result into frequency bands and dynamically scale each sphere based on the amplitude of the selected band (e.g., high-frequency kicks cause the sphere to pulsate, providing immediate visual feedback). Our scenes feature multiple audio objects representing distinct sound samples. To accommodate large spatial layouts, we set Unity’s *volume camera* configuration to “unbounded”, giving users freedom to walk around and place objects anywhere in their physical environment. Unity’s PolySpatial engine also takes care of converting Unity’s materials and mesh renderers into AVP’s native rendering.

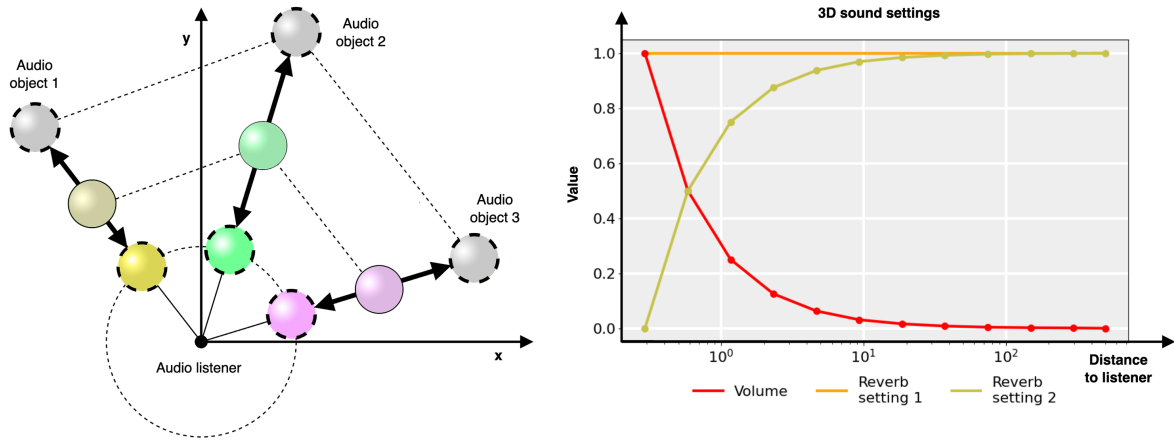


Fig. 6. (Left) Schematic representation of the AR panner. The spheres represent virtual audio objects, and the color saturation represents the level decay based on distance from the audio listener located at the HMD to the audio spheres. The *X-Y* axes quantify the amount of room response in the binaural rendering setup; (Right) The plot displays the 3D sound settings in the game engine. The virtual audio object level displays a logarithmic decay based on distance from the audio listener. For the reverb setting, we experimented with a uniform, full-strength reverberation (“Reverb setting 1”, our default setting for the user study) and a monotonically increasing reverb mode that adds more ambience as the distance from the listener grows (“Reverb setting 2”).

4.5.3 Binaural Rendering and Reverb. Unity’s native audio engine allows binaural rendering by applying 3D sound settings (Figure 5.D) that simulate how audio signals reach each ear from different directions and distances. While additional proprietary plugins exist for HRTF modeling, AudioMiXR relies on Unity’s built-in audio spatializer. For that reason, we manually deactivated AVP’s native spatializer. We also experimented with two reverb schemes for the audio objects. The first, *constant reverb*, applies a uniform reverberation effect at full strength across all distances between the virtual audio object and the audio listener, creating the sensation of a large, uniformly reverberant area. This setting served as the default option during our user studies. As an alternative, we used the *distance-based reverb* mode in our supplementary demos, which starts with no reverb at close range and increases logarithmically with distance. This approach provides an effect of ambiance and space as objects move farther away from the listener. The right subfigure in Figure 6 illustrates both the logarithmic distance-based attenuation and the corresponding reverberation curves.

5 METHOD

To evaluate AudioMiXR, we conducted two formal user studies designed to uncover design lessons, usability issues, usability, and creative output. First we outline the methods for the exploratory study intended to identify user values and interaction patterns to inform design lessons; the second, A/B experiment (Study 2), comparing AudioMiXR to a 2D spatial audio panner, is described in the next subsection.

5.1 Study 1 - Exploratory Design Lessons Experiment

5.1.1 Participants. We recruited 27 participants via snowball sampling from Dolby Laboratories, which is typically sufficient to reach saturation of recurring themes in participant responses [30, 36]. The sample included 17 male and 10 female participants, ranging in age from 22 to 59 ($\mu = 31.22, \sigma = 8.70$). Participants have diverse audio mixing and AR backgrounds, which we used to categorize them into DAW *experts* or *non-experts*. Specifically, those who self-reported as *unexperienced*, *beginners*, or *intermediate level* DAW users were considered “non-experts”, while those who self-identified as *advanced* and *expert* were considered “experts”, resulting in 18 non-experts and 9 experts. This study aimed to quantitatively measure workload, observe user interaction patterns, and collect qualitative feedback on AudioMiXR’s usability and potential professional applications. We used the AVP without optical inserts; consequently, we collected participants’ glasses prescriptions and diopters, if available (for control). Of the 27 participants, 10 had perfect vision and 17 had different mild eye conditions (e.g., mild nearsightedness or astigmatism) but felt comfortable taking part in the experiment without vision correction.

5.1.2 Setup. The study took place in a room measuring 7.62 m \times 6.09 m (width \times depth), as illustrated in Figure 7 (left). Participants wore an AVP HMD and stood at the origin point marked with a cross on the floor. We positioned an additional camera near one wall of the room to capture each participant’s movement, while the AVP simultaneously recorded the first-person view of virtual audio object manipulation.

5.1.3 Procedure. We ran a within-subjects study where each user attempted the instructed directed tasks in 6DoF. Each session began with a brief *calibration phase* (~2 minutes), in which participants aligned hand-tracking and eye-tracking for the AVP, followed by a *training phase* (~3 minutes) during which the participants were able to practice essential gestures – pinching to grab and reposition virtual audio objects, walking around them (clockwise/counter-clockwise), and moving objects closer or farther away to observe changes in audio attributes. After the training session, participants were instructed to complete tasks in two distinct *scenes*, each containing five unique audio objects. We choose five as a lower bound of *Miller’s rule* [55] – which defines 7 ± 2 as the number of objects an average human can hold in short-term memory – to leverage the attention span of our participants [25]. In the *music mixing scene* (Figure 8), users placed and arranged objects labeled “Synth”, “Bass”,




Fig. 7. (Left) Schematic representation of the room where the user study was run, the dimensions of the room were 7.62 m in width and 6.09 m in depth. The origin was marked right at the center of the room with a white cross on the floor; (Right); Snapshots of users interacting with AudioMiXR during the user study—a variety of interesting user behaviors were capturing, including but not limited to: crouching, attempt to manipulate audio objects with two hands at once, placing audio objects over the head, etc.; (Bottom) Panoramic view of the room where the user study was run.

“Kick”, “Filtered Key”, and “Drum” to construct what they feel is the *most* immersive listening experience. We described “immersive” to participants based on the definition presented by Agrawal et al.: “*Immersion is a phenomenon experienced by an individual when they are in a state of deep mental involvement in which their cognitive processes (with or without sensory stimulation) cause a shift in their attentional state such that one may experience disassociation from the awareness of the physical world*” [2].

In the *cinematic scene* (Figure 9), they positioned “Bird” chirps, “Frog” croaks, “River” stream, “Breeze”, and “Campfire” sounds to simulate an outdoor camping environment. We randomized the order of these two scenes for each participant, and each scene had a 5-minute limit (or ended earlier if participants were satisfied of their mixing output). Throughout the sessions, participants were free to navigate the room, placing virtual objects anywhere in the physical room (from floor-level to near the ceiling) to get the sonic effect they envisioned. We

recorded the positions of virtual audio objects and participant's head movements. External and internal HMD cameras captured a full account of each session for subsequent analysis.




Fig. 8. Users interacting with virtual audio objects in the music mixing scene. For illustration purposes, note that the users are using Meta Quest Pro. Since AudioMiXR has been developed in Unity3D, it can be deployed on both platforms, AVP and Meta Quest Pro.

5.1.4 Measures. Upon completing each scene, participants filled out the NASA-TLX [31] to report on mental, physical, and temporal demands, as well as perceived effort and frustration (refer to the full list of questions in appendix, 9). We used a 7-point Likert scale, where 1-3 corresponded to lower workload, 4 was neutral, and 5-7 indicated higher workload. Participants also completed a *demographic questionnaire* and a *short-answer survey* (see appendix, 9) describing any difficulties they found, suggestions for improvements, what they enjoyed the most, and – if they were experienced DAW professionals – potential applications of AudioMiXR in their own workflows. We also collected placement and movement logs (trajectories), identifying any patterns in how they arranged and organized virtual audio objects.




Fig. 9. A user building a cinematic experience in the outdoor camping scene. Observe how the users crouches to place the “River” and “Frog” audio objects on the floor, and positions the “Bird” audio object close to the ceiling.

5.2 Study 1 - Data Analysis Approach

Two authors conducted a thematic analysis to identify themes in participant responses in order inform design lessons for mixing with 6DoF. Open coding was done independently and followed by axial coding to form larger

code groups. Lastly, the two authors formed overarching themes specific to the interactive, aural, and visual elements with 6DoF.

5.3 Study 2 – 2D Panner Comparison

To contextualize the usability, creative outcomes, and design insights from the initial study, we conducted a comparative experiment between AudioMiXR and a custom-built 2D panner interface developed in-house for desktop interaction.




Fig. 10. **The 2D spatial audio panner.** *From left to right:* *Left:* Facilitates audio object panning in the vertical and horizontal axes and to the, *Right:* audio objects can be panned in the horizontal and forward axes. This functionality parallels 2D desktop spatial audio panners present in Ableton Live, Reaper, FL Studio, and SPARTA DAWs. [1],[17],[37],[69]

5.3.1 Setup. The experiment was conducted in a university lab environment without any physical obstructions that may interfere with the experiment. The 2D panner interface was displayed on a 27-inch ASUS ProArt monitor, with participants interacting via a standard wired mouse. Audio objects were represented as draggable spheres, analogous to those present in the AudioMiXR interface. The graphical user interface (GUI) of the 2D panner consisted of two circular planes that allowed users to manipulate the positions of audio objects along the x - y and x - z dimensions (see Figure 10), following conventions commonly seen in 2D panners within widely used digital audio workstations (DAWs). Spatial audio objects are displaced in the same relative locations as the AudioMiXR interface with respect to the fixed origin. Spatial audio is rendered to the user depending on the directionality and distance of audio objects about the head. The same push-pull parameter control used in AudioMiXR is applied to the 2D panner, but instead uses a drag-and-drop function with a mouse. In AudioMiXR, users must push objects out to 4 m (i.e., 4 world units in Unity). We mirror this functionality in the 2D panner by implementing logarithmic distance-based attenuation, such that objects approaching the edge of the rings become quieter until they become silent.

5.3.2 Participants. We recruited 18 participants via snowball sampling from Brown University. The participant pool ($N = 18$; 9M, 9F) was composed of non-experts in audio production (mean age = 23.3 years, range = 18–37). 10 participants reported perfect vision, while 8 indicated they were nearsighted. This group was chosen to evaluate interface usability and creative potential in novice users without bias from prior tool familiarity.

5.3.3 Procedure. Each participant received 3 minutes of training on each interface (order randomized), followed by 4 trials each 4 minutes long—composing both music and cinematic soundscapes in both AR and 2D (trial

order randomized). The same audio stimuli were used in both interfaces. All sessions were conducted using AVP headsets, and spatialization was controlled via identical Unity spatial audio settings. Importantly, while the AR interface supports 6DoF interaction, we restricted it to 3DoF (rotation + orientation only) to match with the directional placement capabilities of the 2D interface—since the 2D interface does not allow for translation we designed the experiment to control confounding variables that may arise if the AR user can translate. Participants were instructed to create a soundscape that is most immersive to them based on the definition of immersion described by Agrawal et al. [2].

5.3.4 Measures. We assessed subjective workload and usability using two established instruments. Task-related workload was evaluated using the NASA-TLX as described in section, which includes five dimensions: mental demand, physical demand, temporal demand, effort, and perceived performance [31]. Participants rated each NASA-TLX dimension on a 7-point Likert scale (1 = very low workload, 4 = neutral, 7 = very high workload), with higher values indicating greater task demand or lower perceived performance. The System Usability Scale (SUS) included ten items assessing perceived usability, also rated on a 7-point Likert scale from 1 (strongly disagree) to 7 (strongly agree) [13]. Additionally, we assessed creative output by asking 7-point Likert scale questions informed by Study 1's results, probing participants' perceived creativity, their inclination to explore spatial arrangements, and the extent to which visual representations of sound objects supported creative insight. Lastly, we asked a general user satisfaction question for both the AudioMiXR and 2D interface on another 7-point Likert scale (1 = very low, 4 = neutral, 7 = very high).

Data Analysis Approach. For analysis, TLX, SUS, and creative output scores were rescaled to a standardized 0–100 range. Objective performance metrics included task completion time and the number of object edits. Because the study employed a within-subjects design, all comparisons were analyzed using paired non-parametric tests, with Holm corrections applied to control the family-wise error rate.

6 RESULTS

This section begins by presenting the quantitative and qualitative results of our exploratory AR study (Study 1), including NASA-TLX workload measures, statistical tests, and spatial visualizations of virtual audio object placement, which collectively offer objective insights about AR sound design patterns. We then present the qualitative findings from participant feedback, summarizing sentiments, codes, and themes on using virtual audio objects for sound design in AR across cinematic and music mixing tasks. To contextualize and validate these findings, we conducted a follow-up comparison study (Study 2) that quantitatively evaluates user experience across both the XR interface and a 2D desktop panner. Results from this complementary study provide a controlled contrast that highlights the specific usability and creativity benefits of immersive spatial interfaces.

6.1 Study 1: Cognitive Demands

First we examined the normality assumptions with the Shapiro-Wilkinson test, which indicated that only the dimensions of Frustration, Efficiency, Satisfaction for DAW expert users followed a normal distribution; but the remaining dimensions and groups violated the normality assumption (all $p < 0.05$). Consequently, we opted for the non-parametric Mann-Whitney U to assess the statistical significance of trends. A comparison of NASA-TLX scores between DAW expert and non-expert users, as depicted in Figure 11, revealed no statistically significant differences across all workload dimensions based on the Mann-Whitney U tests (all p -values > 0.05), suggesting that AudioMiXR imposes a similar perceived workload on users irrespective of their prior audio mixing experience and is thus suitable even for novice users. Overall, participants reported low levels of mental demand ($\mu = 2.30, \sigma = 1.20$), physical demand ($\mu = 2.22, \sigma = 1.25$), and temporal demand ($\mu = 1.85, \sigma = 1.13$) when manipulating virtual audio objects in AR, along with low frustration levels ($\mu = 2.04, \sigma = 1.22$). These factors

help explaining the high satisfaction scores ($\mu = 5.89, \sigma = 0.93$) with the system usability. Participants perceived a moderate level of effort ($\mu = 2.48, \sigma = 1.28$) required to complete the tasks, and the performance dimension received neutral ratings ($\mu = 2.70, \sigma = 1.44$), reflecting on the success of positioning and manipulating virtual audio objects accurately without excessive difficulty. Additionally, the ratings on the efficiency of completing the tasks are consistently high ($\mu = 5.11, \sigma = 1.45$), highlighting that participants were able to complete tasks without overhead.




Fig. 11. Comparison of mean Likert scores on NASA-TLX assessment profiled by DAW expertise level (with *unexperienced*, *beginners*, and *intermediate level* users labeled as “non-experts”, and *advanced* and *expert level* users classified as “experts”).

6.2 Study 1: Coordinates Visualization

We analyzed the HMD logs from each user trial to aggregate the 3D coordinates of all virtual audio objects within each audio scene. This aggregation involved combining all logged timestamps across users and averaging the x, y, z coordinates of each virtual audio object, thereby obtaining a mean 3D position for each object in each scene. Next, we binned the positions of each virtual audio object into two 2D histograms: one for the X-Z plane (top-floor view) and another for the X-Y plane (side view), and applied a Gaussian filter to the counts to produce smoother heat map visualizations of the average positions of each audio object class (Figures 12 and 13). The color intensity within each heat map indicates the regions where a virtual audio object is more likely to be placed on average, across our trials. We observed significantly different spatial patterns between the two scenes, the cinematic and the music mixing scene, which are described below. We also computed the distances between pairs of virtual audio objects within each scene and estimated the densities of histograms of distances using a Gaussian kernel (Figure 14). These densities allowed us to compare the proximity relationships between pairs of virtual audio objects and correlate that with their semantics.

In examining the heat maps for the cinematic scene, as illustrated in Figure 12, the “Breeze” sound is consistently positioned at the origin in X-Z plane, corresponding to the location where the user initiates the application, and approximately at the headset height in the X-Y plane. Because “Breeze” is an ambient sound, participants naturally tended to place it near their headset to emulate its pervasive presence. In contrast, the “Bird” sound, characterized as a point-source sound, is generally positioned above the headset in the X-Y plane.

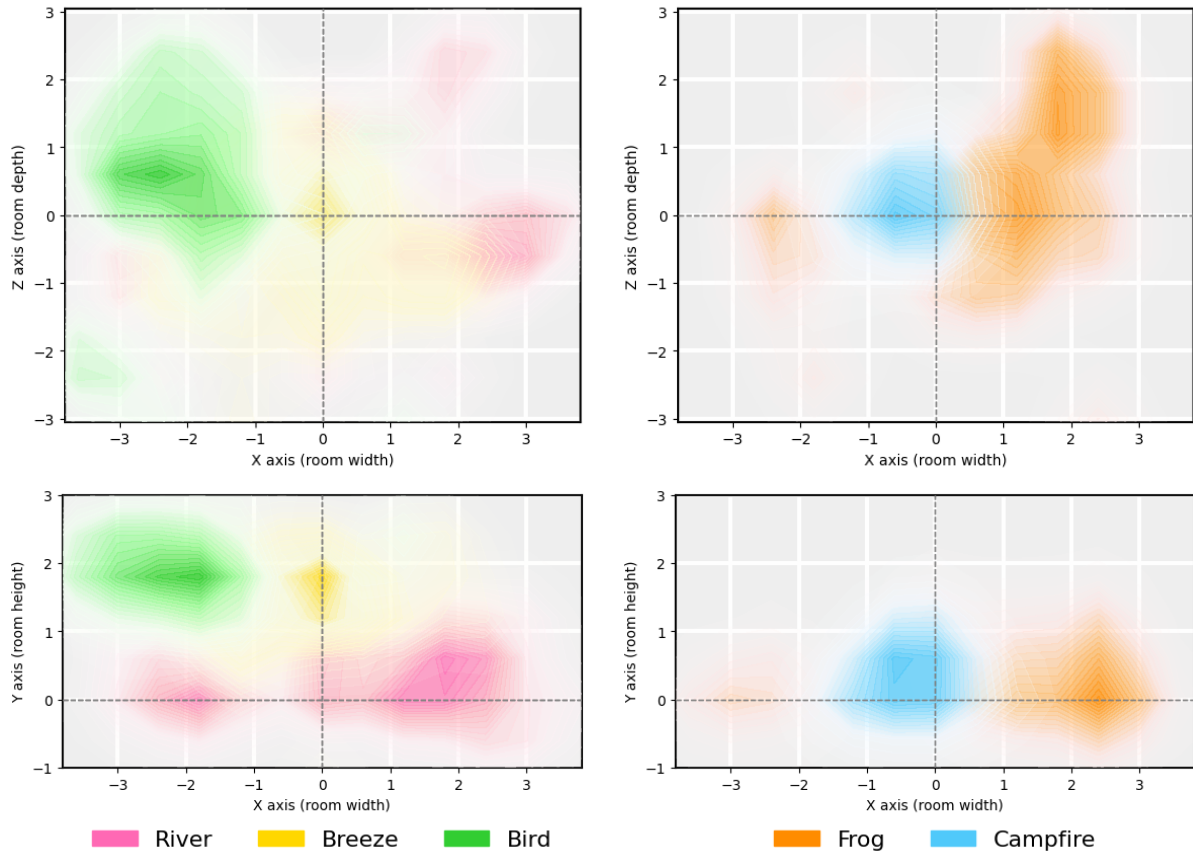


Fig. 12. Heat maps of the mean audio object 3D positions in the environment scene (first row: X-Z plane, second row: X-Y plane).

For instance, during the experiment a participant expressed they felt *the birds should be placed up high like there are flying or in a tree*. The opposite occurs with the “River” audio object, expected to emanate from the ground and, thus, heard from below. Consequently, in the X-Y plane, it is typically situated on the floor. The “Frog” audio object exhibits greater variance across both planes, likely because, as an animal, it could virtually be placed anywhere on the X-Z and still be coherent with the scene. Nonetheless, it exhibited a considerable overlap with the “River” heat map, indicating a tendency for participants to place the frog close to the river, reflecting nature-related biases. This is a possible manifestation of the Gestalt principle of similarity [25, 65], where the strongest proximity between subjects relates to those that have overlapping semantic associations, and led our participants to arrange those virtual audio objects together. This is also observed in the density estimates in Figure 14, where the closest virtual audio object to “River” is “Frog”, and vice versa. Additionally, the “Frog” audio object is generally placed below the headset height, similar to the “Campfire” audio object. The “Campfire” is typically positioned at the origin, close to the user. This pattern may stem from participants viewing the campfire as the central point of the camping setup, thus leaving it at the origin. Alternatively, it might be perceived as an ambient sound with a larger spread, justifying its proximity to the user.

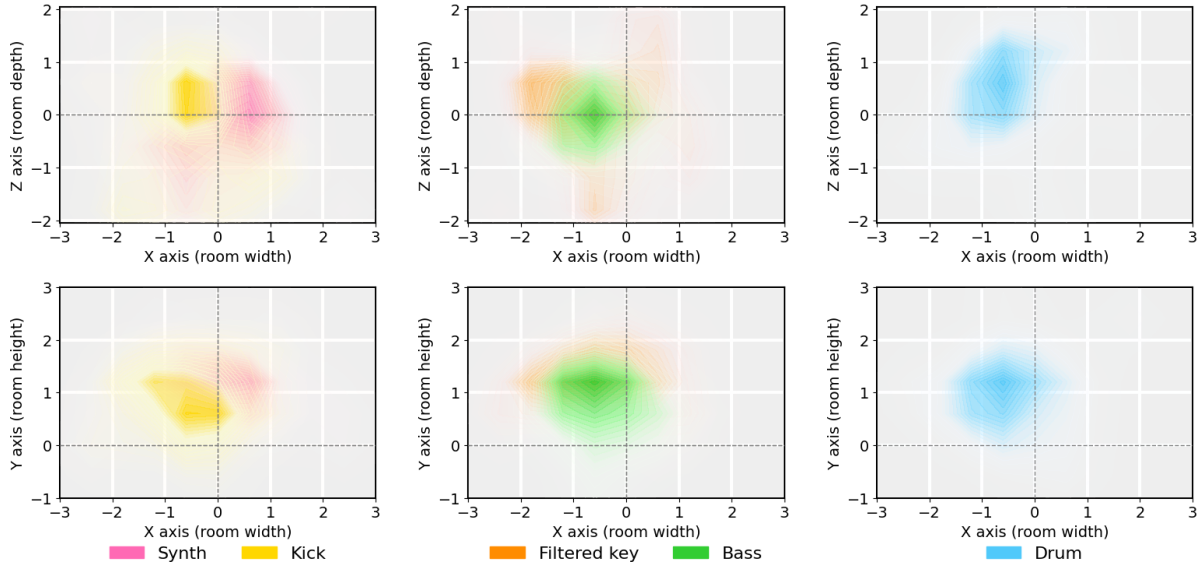


Fig. 13. Heat maps of the mean audio object 3D positions in the audio mixing scene (first row: X-Z plane, second row: X-Y plane).

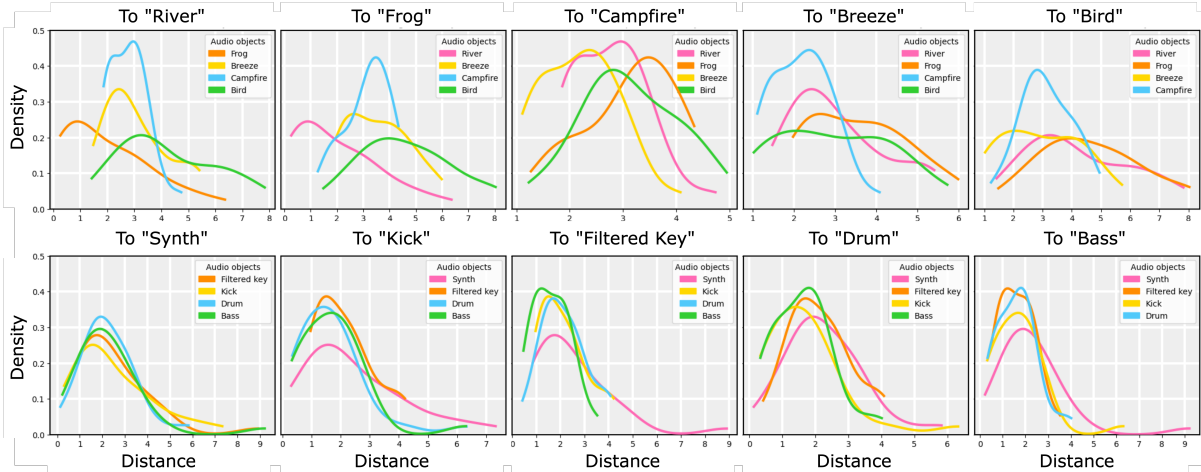


Fig. 14. Density estimates using a Gaussian kernel over histograms of distances between pairs of virtual audio objects.

In the music mixing scene (Figure 13), the heat maps present considerably more overlap and are primarily more concentrated around the origin (notice the range of the X-Z axes), extending roughly to the user's reach envelope. In standard spatial audio panning the loudness of spatial audio objects is kept constant, only the room response is modified. In those scenarios, expert audio mixers [20, 29] have noted that a common practice is to place high frequencies (called “treble” bands in sound design) higher and low frequencies (also known as “bass” bands) lower in the spatial field, reinforced by psychophysics research on congruence between pitch and visual

vertical positioning [9, 24, 53]. However, we have not observed such a pattern with audio objects with stronger high-frequency content (e.g. the “Drum”) versus audio objects with richer low-frequency content (such as the “Filtered Key”), likely because loudness was not kept constant. In a 6DoF application, raising or lowering a virtual audio object spatially substantially increases its Euclidean distance from the headset, effectively causing larger loudness attenuation than if the object was kept at the HMD level. That might explain why we observe that most of the music mix audio objects are placed approximately at the HMD level, producing notable overlap among different objects (Figure 14).

6.3 Study 1: Variance Analysis

To assess the homoscedasticity of final audio object placement between music and cinematic mixes, we employed Levene’s test. Specifically, we examined the distributions of x, y, z spatial components of the final audio object placement, as well as the overall distances of audio objects from the origin (see Figure 15). Our analysis revealed p -values below the significance level of $\alpha = 0.05$ for all tested components (x, y, z , and distance-based: $F_X = 36.94$, $p < 0.05$; $F_Y = 30.52$, $p < 0.05$; $F_Z = 4.93$, $p < 0.05$; $F_d = 13.90$, $p < 0.05$), indicating that the variances between the music and cinematic mixes are significantly different.

To visualize and further understand these differences, we represented the experimental room setup with the final placements of virtual audio objects across all users and scenes in Figure 15. We fitted a multivariate mixture-of-Gaussian (MoG) model using the expectation-maximization algorithm for density estimation, choosing one component per mix to capture the distinct spatial distributions of music and cinematic mixes. Each component was assigned its own covariance matrix. The resulting 3D Gaussians ellipsoids (see Figure 15) illustrate the spatial dispersion of final audio object placements. The estimated means for the cinematic scenes ($\hat{\mu}_X = 0.071$, $\hat{\mu}_Y = 1.247$, $\hat{\mu}_Z = 0.242$) and mix scenes ($\hat{\mu}_X = -0.064$, $\hat{\mu}_Y = 1.324$, $\hat{\mu}_Z = 0.209$) indicate slight positional shifts between the two mix types, while the variances ($\hat{\sigma}_X^2 = 2.654$, $\hat{\sigma}_Y^2 = 0.766$, $\hat{\sigma}_Z^2 = 1.610$ for cinematic and $\hat{\sigma}_X^2 = 1.016$, $\hat{\sigma}_Y^2 = 0.384$, $\hat{\sigma}_Z^2 = 1.156$ for music) reveal that audio object placement in cinematic mixes exhibit greater variability, particularly along the X and Z axes (width and depth dimensions of the room), corroborating the results from Levene’s test.

6.4 Study 1: Qualitative Analysis

This section outlines the findings from our thematic analysis, where we organized the codes into three themes: Embodied Mixing, Mapping Sound Attributes to AR GUIs, and Prototyping with Spatial Interface. These themes underscore the overarching insights gathered from participant responses to provide context for our analysis and design lessons.

6.4.1 Theme 1: Embodied Mixing. We first present codes describing participant responses ($N = 20$) related to their bodies’ influence on their mixing experience. Overall, respondents expressed that the combination of their body movements and the feedback they received contributed to their experiential design process and supported their mix.

6DoF Audio Presence. About half of the participants ($N=14$; 8 Non-Expert, 6 Expert) shared positive sentiments highlighting the effects of 6DoF, which adapted to their position and orientation-induced feelings of *presence* when designing their soundscape. We adopt the two-dimensional conceptualization of *presence* suggested by Weber et al. and define it as, “as the allocation of attentional resources to the mediated world and the sensation of perceptually being surrounded by the virtual environment or the user’s judgment about the degree of realism” [74]. Participants found the directional audio feedback to feel “*real-time*” in response to their position and orientation. Additionally, some participants described enveloping themselves in their mix by navigating freely through the responsive soundscape or surrounding themselves with their audio objects:

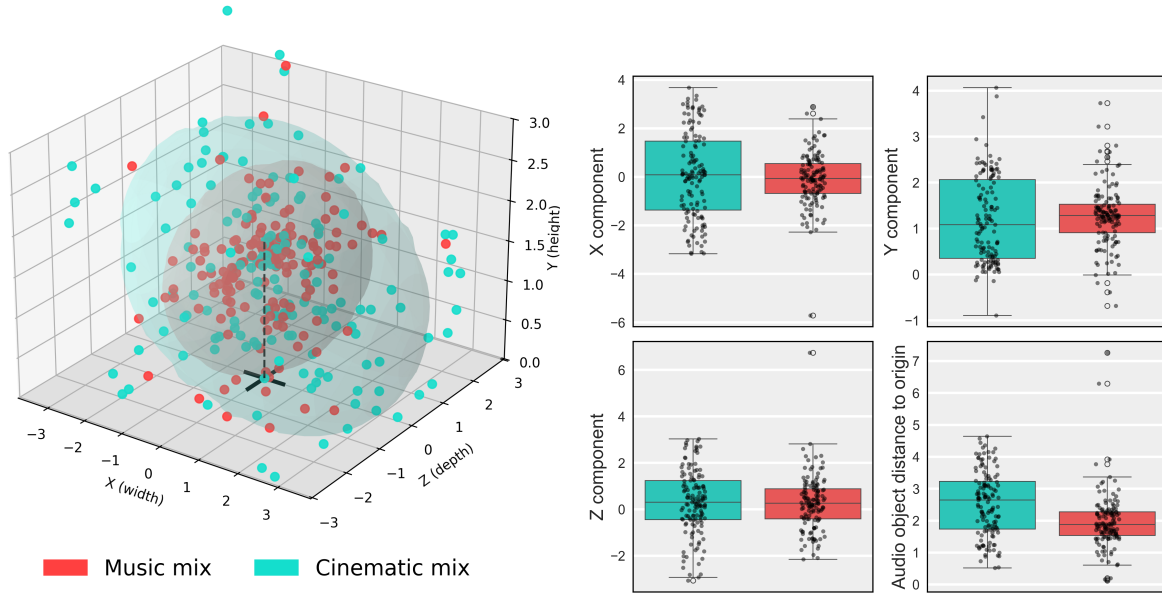


Fig. 15. (Left) A room representation showing the final placement of audio objects across all experiments and users, color-coded by mix. A Gaussian mixture with two components (one for each mix) is fitted to the data. The cross at the origin represents the initial position, while the vertical dashed line represents the average user height; (Right) Boxplots illustrating the distribution of the final placement of audio objects along the X, Y, Z axes, as well as the distribution of distances from the final audio object positions to the origin.

“Being able to place the different parts of the environment where you wanted was super cool. In the camping environment, I closed my eyes and it felt like I was actually camping. Being able to put specific tracks/channels above and behind me in the music mixing environment felt like I was truly immersed in the song.” – P21 (expert)

Participants noted that when they placed objects above or behind themselves, the 6DoF spatial audio accurately translated those physical positions into perceivable directional cues resulting in an “immersive” experience. This aligns with research linking mediated experiences involving body movement in physical spaces to an enhanced sense of presence [45, 64, 76] – specifically, when individuals engage more of their physical selves, the experience often feels more real. Some respondents delved deeper into why they were particularly intrigued by the sound’s directionality, noting how it interacted with and changed based on their physical movements:

“I liked being able to hear the sound coming from different directions as I moved it around, it is cool that instead of just hearing the sound around you can interact with the sound.” – P19 (Non-Expert)

The interface may have added a level of interactivity for experts and non-experts, which appeared to be unique from their everyday audio experiences, possibly contributing to feelings of envelopment. The dynamic audio changing due to their movements can be attributed to how the spatial audio responds to their body with 6DoF, which allows for more precise directionality and includes translation (e.g., audio attributes changing as a user walks away from origin) versus 3DoF interfaces which, in the context of audio mixing applications, generally support only rotational components.

Table 1. Final Codebook of Identified Themes, Corresponding Codes, and Exemplar Quotes

Theme	Codes	Exemplar Quote
1. Embodied Mixing (N = 20, 74%)	6DoF Audio Presence (N = 14, 52%)	P19: “I liked being able to hear the sound coming from different directions as I moved it around, it is cool that instead of just hearing the sound around you you can interact with the sound ”
	6DoF Audio Perceptibility (N = 8, 30%)	P6: “I chucked my head into a sphere, so I could hear just what was going on in there”
	Visual-Aural Responsiveness (N = 7, 26%)	P22: “There is a natural intuition on how to manipulate the objects. The immediate acoustic feedback made it easy to converge on manipulations that were most effective”
2. Mapping Sound Attributes to AR GUIs (N = 18, 67%)	Audio Attribute Level Control (N = 15, 56%)	P24: “For example close distances always made the object much louder than I wanted. I’d like to specify volume across different distances more precisely”
	Object Appearance (N = 8, 30%)	P23: “I think the size of audio object should be associated with spread”
	Modularity (N = 4, 15%)	P25: “I could see this integrated as a plugin in my DAW while looping / experimenting with multitrack placement.”
3. Prototyping with Spatial Placement (N = 12, 44%)	Vertical Positioning (N = 7, 26%)	P17: “Make it easier to move objects far away. Make it clear to look up if user is trying to move an object overhead”
	Spatial Exploration (N = 6, 22%)	P25: “Also Experimented with dragging objects while moving – Moving sound around while walking throughout the space, also finding different mix levels within a room to compare by walking through it.”
	Shaping soundscape (N = 5, 19%)	P6: “The wind should be everywhere, so how would I do I turn a wind to an object that isn’t spatial, it’s like a skybox in a computer game, so I put it above myself”

6DoF Audio Perceptibility. Some participants (N = 8; 5 Non-Expert, 3 Expert) shared specific strategies leveraging spatial audio with 6DoF, free-hand manipulation, and head orientation enabling strategies of perceiving and isolating sound emitted from audio objects. Generally, increased users’ perceptibility was due to the physical actions coupled with binaural rendering, amplifying interaural level differences (ILD). ILD is a phenomenon where the head occludes acoustics depending on the direction, altering the amplitudes’ distribution to either ear [11, 25]. For instance, during the experiments, people often moved the relative position between their ears and the audio objects to hear different perspectives of their mix:

“I really enjoyed the head tracking aspect. I felt that being able to bring an object close and then tilt my head from side to side made it easier to get a sense of the timbre of the sound. Also, being able to isolate objects in space was helpful.” – P22 (Non-Expert)

Head orientation plays a key role in sound localization, and actions like turning the head toward or away from the source assist in determining its physical location [11, 25]. Therefore, head movement may have felt advantageous to users when assessing the quality of their mix. Another technique worth noting was expressed by an expert during the experiment where they state,

“I tried putting my head inside the audio sphere to isolate the sound.” – P21 (Expert)

P21 defined their strategy as inserting their head inside a virtual audio object, thereby attenuating all other sound sources within the scene. The sound isolation likely happened because the audio listener (attached to the AVP’s position) was dominated by the audio source of the object where P21 inserted their head due to the distance attenuation in Unity’s 3D sound settings. In the sound settings, other audio sources attenuate according to a logarithmic distance scale. As a result, sources that are in close proximity, or in this case near the same position, will dominate the audio output. All user responses in the coded group shared similar statements about the ease of perceptibility, except for one, who noted that the stark differences in audio could cause some confusion when mixing:

“I found some inconsistencies with how I wanted the scene to sound in my head when looking at the objects versus how the sound actually rendered out. Moving around an object (walking in a circle around a ball) made me question the placement of other objects in relation to it because the changes in level were semi-stark.” – P27

P27 describes the disorientation they experienced due to conflicting expectations of where they observed their placement versus the sound they heard. The misalignment was exacerbated by walking in a circle, which likely caused the 6DoF audio to pan too sensitively in the user’s headphones, hindering their ability to perceive the audio sources coherently and thus affecting their expected mix. To address this sentiment, it may be worth offering users greater flexibility over the sensitivity of the spatial audio panning.

Visual-Aural Responsiveness. Another subset of participants (N = 7; 4 Non-Expert, 3 Expert) found the alignment of the position of the audio objects’ visual appearance corresponding to changes in the audio directionality, specifically when they would manipulate objects with free-hand gestures, to be both useful and enjoyable. Users remarked the immediate audio and visual feedback in response to their free-hand manipulation as *intuitive* and assisted them in mixing, allowing for easy self-correction of placements:

“I really enjoyed the lack of latency when moving objects around the space. It really felt like the objects responded to my gestures in real-time, and their movements were also accurately reflected by the audio.” – P23

P23’s response explicitly highlights the sense of control in feeling convinced the audio objects reacted visually and aurally to their physical manipulations within a time window, making the physical-virtual interaction believable. Congruent audio-visual feedback has been shown to result in better performance and higher engagement [53, 59]. The blending of senses, also known as “synesthesia”, has been previously linked to excitability and arousal [72], a fact that film, games and VR entertainment creators have been exploiting – by integrating cross-modal stimuli, they evoke a sensation of presence that creates exciting experiences [45, 59], and if spent an extended period of time in the immersive experience, virtual artifacts start feeling like physical [26]. One of our participants expressed a synesthetic experience:

“I quite like how close I can be to the audio objects (it almost feels like there’s a haptic feedback from being close to certain audio objects)” – P16

Another participant echoes this sentiment and also expressed how the immediate visual-aural feedback assisted their mixing process:

“The immediate acoustic feedback made it easy to converge on manipulations that were most effective.”

– P22 (Expert)

Overall, the immediate visual-aural feedback and alignment were complimentary in providing interactions that allow users to feel in control, thereby making their mixing process feel like a *direct* result of their intended placements. The interrelationship of the visual and aural attributes of the interface captured in participant responses highlights the need for proper alignment and immediacy of feedback; otherwise, the mixing experience may be disrupted. Additionally, multimodal synchronization is a key sound design principle necessary to reduce cognitive load when, for example, viewing and listening to cinematic experiences or when developing audio authoring tools [25].

6.4.2 Theme 2: Mapping Sound to AR GUIs. Next, we outline the codes centered on controlling sound attributes in an AR interface. Responses (N = 18; 10 Non-Expert, 8 Expert) in this category generally consisted of the trade-offs of using depth to modify audio attributes (e.g., loudness or reverberation), instead of varying alternative visual properties and physical interactions of the audio objects to have better UI affordances and enable more granular control.

Audio Attribute Level Control. The *Audio Attribute Control* centers on challenges that arose when using the *push-pull* action to vary the loudness parameter coupled with suggestions to modify alternative audio attributes to enable more fine-tuned control over their soundscapes. In this category, users (N = 15; 9 Non-Expert, 6 Expert) noted difficulties in depth perception when pushing objects out to control the loudness parameter of the audio objects:

“Depending on where I was standing in the room, I wished there was some guidance on how far I could push out an object on each XYZ plane before they became inaudible/not part of the mix.” – P27

P27 highlights an issue with recognizing when an audio object becomes “inaudible”. This sentiment is aligned with how multimodal cues are often explored for additional guidance through an interface, and how depth alone may not offer enough control over audio attributes for some users [15, 32]. In this particular case, the user expresses uncertainty when audio objects will become inaudible, and their response suggests the inclusion of additional visual cues, like implementing a visual boundary representing the *inaudible zone* that persists through the duration of the application to identify the distance before an object is silenced quickly.

Participant responses in this category also expressed mostly negative sentiments regarding the “loudness” parameter control implemented in the interface. Users generally felt pushing objects *too far away* made it harder to align their hands with the perceivably smaller objects due to the camera’s perspective view in Unity. P2 describes the difficulty of aligning their hands with the virtual audio object that has been pushed far away from them:

“At one point, I had moved some objects really far away from me, and I was struggling to bring them back. I had to physically walk over to the objects to bring them back.” – P2 (Non-Expert).

Moreover, determining the sensitivity using the forward axis for loudness combined with the challenges in perspective projection presented difficulties in varying the audio attribute. P16 states:

“It would be nice to have a way to tweak the sensitivity concerning distance. It took some trial and error to figure out how far I needed to move the object to get the level of loudness I wanted.” – P16

Other participants echoed this sentiment and typically suggested different features like modifying the sensitivity of the “*push-pull motion*” or alternative methods for level control that do not rely on distance.

Object Appearance. In contrast to the *Audio Attribute Level Control*, the *Audio Object Appearance* categorizes sentiments related to participants' (N = 8; 4 Non-Expert, 4 Expert) preconceived associations with mentally mapping visual elements to audio attributes (e.g., loudness, reverb, or spread). Most of the respondents in this category described preferences of *changing the size of the sphere to increase/decrease loudness* or other audio attributes not implemented in the interface. Additionally, some responses related to the trade-offs of using translucence to represent the loudness:

"I think the size of the audio object should be associated with spread" – P27 (Expert)

Their mixing experience with DAWs may influence P27's sentiment because the spread is typically associated with the size of audio objects in object-based spatial audio metadata, such as Dolby Atmos [22]. In contrast, a response from P22 (Expert) associated audio object size with *draping* and felt translucence could represent *direct-to-reverberant ratio*. The other two experts in this category also stated that size and translucence make sense when representing loudness. While there were mixed sentiments on the preferred visual elements for audio objects, all participants in this category suggested varying visual associations with different audio attributes, indicating that there may be associations from their mixing experience or, in the case of non-experts, there are natural associations with visual properties of the spheres and audio attributes:

"I wasn't sure if I could change the magnitude of the 'orb' or if I need to bring it closer to increase the volume" – P15 (Non-Expert)

P15 expresses confusion regarding the visible size of the "orb" being associated with loudness. The visual parameter, *size* being mistakenly associated with *volume* or loudness was also expressed by P23 and it may be due to the natural associations humans have between audio and visual representations, often referred to as audiovisual perception [24, 29, 46, 53]. This could also contribute to a GUI design strategy focused on how to represent sound visually in the most effective way for the intended task.

Modularity. The AudioMiXR interface was seen by four experts (N = 8; 4 Non-Expert, 4 Expert) as a component for a potential hybrid system combined with an existing DAW for audio mixing with more spatial awareness. The four experts expressed that designing within the execution environment could fit nicely in-game audio design or XR workflows and foresee the *AudioMiXR* interface integrated into a hybrid system where they have the flexibility to switch between their DAWs on a desktop and into a headset when they require more spatial awareness. For example, P27 states:

"Think this fits nicely within game audio design workflows and specific audio mixing workflows where the output is specifically for AR/XR especially at the ideation phase of the project... DAWs are notoriously cumbersome to use unless you are experienced with hot keys and shortcuts so this was very refreshing to be able to make large changes with little physical exertion." – 27 (Expert)

Respondents thought of *AudioMiXR* as a complement to their existing workflows where they can overcome some of the drawbacks of DAWs but felt they would still want access to their traditional mixing tools. The flexibility of control underscores an optimal medium where all audio attributes may not necessarily need to be mapped to the AR interface and instead can leverage its features when required.

6.4.3 Theme 3: Prototyping with Spatial Placement. Lastly, we present responses related to the potential affordances and drawbacks spatially placing objects has on prototyping mixes. Participants (N = 12; 5 Non-Expert, 7 Expert) expressed suggestions, challenges, and benefits the existing interface has on prototyping their soundscapes.

Vertical Positioning. *Vertical positioning* refers to sentiments centered on audio object placement along the vertical axis affected participants' (N = 7; 5 Non-Expert, 2 Expert) (i.e., placing an object above or below their

line of sight). P5 and P20 both expressed the features was *playing with height* while the rest of the respondents voiced difficulties with placement along the vertical axes:

“Placing object overhead was a bit awkward with the limited field of view. You can also feel the mass of the headset when doing that” – P9 (Non-Expert)

Placement of objects overhead or below the HMD is a unique feature supported by AudioMiXR. Still, the limited FOV forced users to engage in excessive upward head rotation so that the cameras on the headset could capture their hand gestures, thereby disrupting their placements. The difficulties with the vertical placement due to the FOV of the headset may imply that there should be alternatives to object placement. For example, a function that maps subtle movements along the vertical axis may be a potential circumvention to take advantage of the vertical placement unique to an audio mixing interface with 6DoF.

Spatial Exploration. Respondents expressed how the autonomy and *freedom* provided by the AudioMiXR interface encouraged exploration and experimentation of their arrangements while designing their mix (N = 6; 4 Non-Expert, 2 Expert). Responses center on free-hand movement and being able to use *all* of the surrounding space to manipulate objects:

“It allowed me to be creative and place the objects as far or close as I wanted. I felt that it understood what I wanted it to do most of the times.” – P8 (Non-Expert)

P8’s sentiment represents the code group in expressing how the interface feels unrestricted to their movements. P12 describes how the spatial placement enabled them to *remix* audio objects easily and try different combinations of loudness and spatial configurations, thus supporting their creativity. As metaphorized in [65], a mix is a “sonic portrait” and its essence can be conveyed in a myriad of ways. This flexibility may be worth considering for creative workflows within AR Sound Design. Another interesting connection is related to the quantitative results on virtual audio object spread reported in Subsection 6.3, where a significant difference in spread and, thus, spatial exploration between the cinematic and music mix scenes can be observed, which underscores that the amount of spatial exploration is task-dependent.

Shaping of Soundscape. During the experiment, a few participants (N = 5; 3 Non-Expert, 2 Expert) described different strategies of grouping or arranging multiple audio objects, as well as how these objects worked together to create a cohesive soundscape. For instance, P2 would have preferred the ability to “tie objects” to their location instead of the “physical environment” in order for the arrangement of objects in the scene to be relative to the headset which is more aligned with an audio experience with 3DoF. Further, the rest of the participants echoed different types of arrangement functionalities centered on different types of grouping of objects to make the design of their soundscapes more flexible. P27 captures these sentiments well by suggesting moving objects together to simplify their mixing process:

“Grouping objects – link multiple balls together and moving as a unit.” – P27 (Expert)

P27 describes a feature that may embody the concept of a cohesive soundscape better than simply treating each audio object as separate entities at all times. It is possible P27’s experience with DAWs inspired their suggestion as modern DAWs like Logic Pro support users to arrange audio channel groups [3], which is a known technique for professional audio mixers to design complex soundscapes [25, 29]. This would exhibit the Gestalt principles of proximity and common fate [25], where similar sounds – where similarity may be measured in terms of rhythm, timbre, or semantics – would have the tendency to be arranged and moved together. This is something that has been reflected in our quantitative data (Subsection 6.2) as well, where one can clearly see several instances of semantically-related audio objects to be clustered together in space.

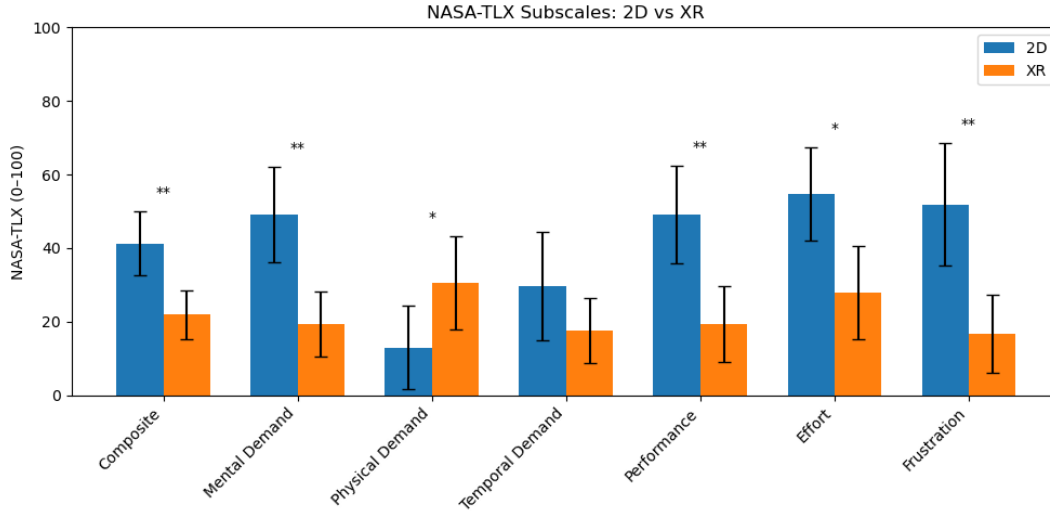


Fig. 16. NASA-TLX workload comparison. XR showed lower cognitive load across key subscales. **Note:** “Composite” refers to the overall distribution across all NASA-TLX categories.

6.5 Study 2: Quantitative Results

We analyzed quantitative differences between the XR interface and the 2D panner using paired-sample t-tests with Holm-corrected p -values to account for multiple comparisons. Measures included perceived workload (NASA-TLX), system usability (SUS), and creativity.

6.5.1 NASA-TLX. Participants reported significantly lower cognitive workload when using the XR interface compared to the 2D panner. Specifically, XR showed significantly lower scores in frustration ($t(11) = -3.73$, $p_{\text{holm}} = 0.0115$, $d = -0.88$), mental demand ($t(11) = -3.19$, $p_{\text{holm}} = 0.0278$, $d = -0.75$), and effort ($t(11) = -2.47$, $p_{\text{holm}} = 0.0576$, $d = -0.58$). The composite workload score also favored XR ($t(11) = -2.95$, $p_{\text{holm}} = 0.0362$, $d = -0.69$). While performance and temporal demand also trended toward improvement, these differences did not reach statistical significance under Holm correction. These results suggest that the XR interface imposes less cognitive load for spatial audio tasks than a 2D panner.

6.5.2 System Usability Scale (SUS). The XR interface was rated as significantly more usable across several SUS dimensions. Participants gave higher ratings for overall usability ($t(11) = 5.31$, $p_{\text{holm}} = 0.0005$, $d = 1.25$), confidence in use ($t(11) = 4.78$, $p_{\text{holm}} = 0.0014$, $d = 1.13$), learnability ($t(11) = 5.48$, $p_{\text{holm}} = 0.0004$, $d = 1.29$), and ease of use ($t(11) = 4.43$, $p_{\text{holm}} = 0.0023$, $d = 1.04$). The XR interface was also seen as significantly less cumbersome ($t(11) = -4.49$, $p_{\text{holm}} = 0.0023$, $d = -1.06$), and more conducive to frequent use ($t(11) = 6.99$, $p_{\text{holm}} < 0.0001$, $d = 1.65$). While some dimensions such as “need for support” and “getting started” did not differ significantly, the large effect sizes observed for many subscales indicate strong usability advantages for XR.

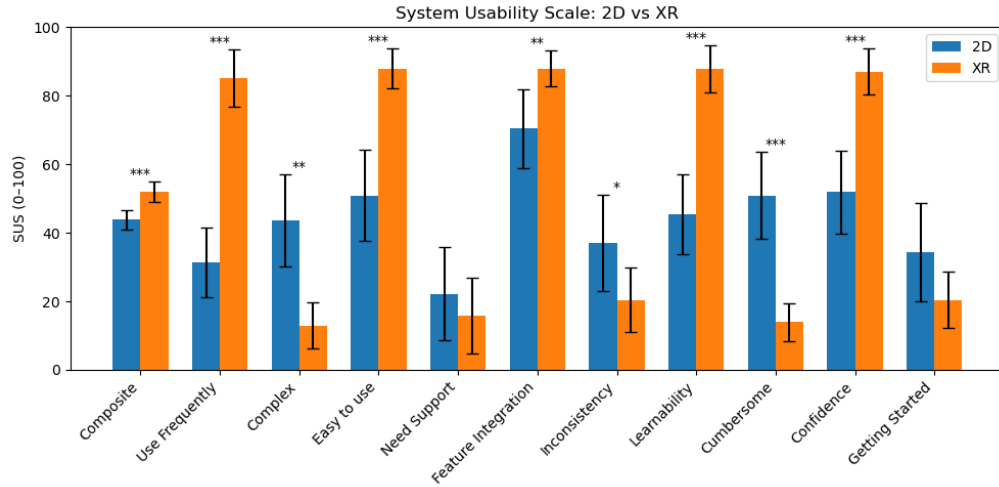


Fig. 17. SUS scale results comparing 2D and XR interfaces (Holm-corrected p -values). **Note:** “Composite” refers to the overall distribution across all SUS categories.

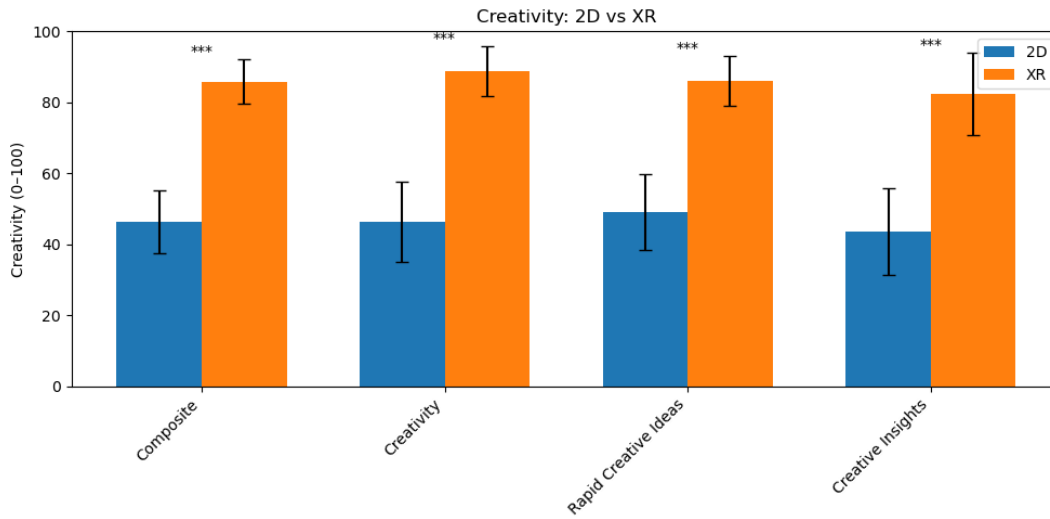


Fig. 18. Creativity scale results. XR outperformed 2D on all creative subdimensions (***) $p < 0.001$. **Note:** “Composite” refers to the overall distribution across all creativity categories.

6.5.3 Creative Output. The XR interface significantly outperformed the 2D panner across all creativity measures. Participants reported higher overall creativity scores with XR ($t(11) = 6.06$, $p_{\text{holm}} < 0.0001$, $d = 1.43$), and reported generating more rapid creative ideas ($t(11) = 5.33$, $p_{\text{holm}} = 0.0001$, $d = 1.26$). They also attributed more creative insight to the XR experience ($t(11) = 6.02$, $p_{\text{holm}} < 0.0001$, $d = 1.42$). The composite creativity score exhibited a large effect size favoring XR ($t(11) = 7.69$, $p_{\text{holm}} < 0.0001$, $d = 1.81$). These results suggest that

the immersive and visual affordances of XR interfaces may substantially enhance creative engagement during spatial audio tasks.

7 DISCUSSION

In this section, we form design lessons derived by combining themes based on participant responses, analyzing our quantitative data, and sharing observations during the experiment while supporting these design lessons with existing literature in audio augmentations in AR, and authoring in XR. To complement these insights, we introduce a final subsection that presents findings from Study 2, a follow-up comparison study evaluating the 6DoF AR interface against a 2D panner baseline. This quantitative comparison strengthens our claims by highlighting where AR offers concrete advantages in usability, cognitive load, and creative potential. Afterward, we describe the types of mixing applications that may benefit most from a 6DoF AR interface, informed by both qualitative and quantitative findings.

7.1 Design Lesson – Proprioception for AR Sound Design

Most users shared positive sentiments related to their experience when “*freely walking around through their soundscapes*”, “*hearing spatial audio from all directions change*” according to their movements, “*immediate audio feedback*” in response to their actions and changes in the audio object positions, and “*being able to explore and be creative*” which may inform how sound design principles should be translated into an AR setup. Theme #1 highlights how users’ bodies influencing the mix made them feel actively engaged in the directed tasks where they noticed their movements and actions directly affected the soundscapes they designed. Their fascination with their body’s influence on the soundscape was primarily attributed to immediate UI response to participants’ actions, position, and orientation, which are inherent to the 6DoF interactions in the interface and can facilitate sound design principles like sound attribute perception, layering and spatial awareness [25, 29]. For instance, the dynamic visual-aural feedback in response to participants’ movements seemed to allow for feelings of self-agency and control of the UI throughout the mixing tasks, as expressed in P23’s statement: “*It really felt like the objects responded to my gestures in real-time, and their movements were also accurately reflected by the audio*” in regards the immediate feedback in response to their physical actions. Similar reactions align well with existing literature on perceiving ILDs and assist with audio localization, which can support the spatial audio mixing process [11, 25, 28]. A 6DoF interface may be able to leverage proprioception in assisting the user in understanding the spatial relationships between their body, actions, audio objects, and the physical space around them.

Moreover, designing dynamic sound design interfaces in AR requires a use case that effectively leverages the space of the virtual environment. In Figure 15, we demonstrate the spatial distribution of all audio objects across all participants in the music mix and cinematic mix tasks – participants generally used significantly more of the physical space to place audio objects for the cinematic blend. Additionally, we have provided a visualization of heat maps (Figures 12 and 13) demonstrating where participants typically placed audio objects in the cinematic scene and uncovered that their placements aligned with real-world expectations of where participants could imagine the audio source (e.g., bird) would be concerning their physical location. The quantitative data demonstrates a possible benefit of using an AR design interface for mixing tasks that have familiar contexts in the real world, suggesting that this may be a reasonable focus for AR sound design UIs with 6DoF interactions. Further, our quantitative findings align well with some of our expert responses where they suggest using the AudioMiXR interface for soundscapes intended for *experimenting with object positions in VR experiences* and *video game design workflows*, which both use cases that involve a ‘perceivably’ physical virtual environment where 3D position and orientation of audio objects would be necessary.

7.2 Design Lesson: Balancing Audio-Visual Modalities in AR GUI

A number of feature requests emerged from our user study that reflect established psychophysical principles related to spatial cues, crossmodal congruences, and audiovisual localization – all of which play a significant role in how participants want to experience sound design in AR. By addressing these features, AudioMiXR can better align with human perception, ultimately enhancing presence, immersion.

Several participants expressed alternatives for parameter control that align with familiar associations between visual effects and audio. Participants desired to control volume by “*making objects larger or smaller*” or the “*level of translucence*” of the audio objects, which are common associations present in traditional DAWs or music production tools [29]. Further, associations of visual and audio attributes are commonly researched in audiovisual perception [53]. They may be coupled with AR design frameworks to design AR sound design interfaces sensitive to humans’ preconceived associations with audio attributes and their visual representations.

Participants consistently expressed the desire to have independent control of an audio object’s loudness, as opposed to just moving it closer or farther away, as reflected by the following responses of P24: “*I’d like to have control over object size and overall object volume at different distances [...] I’d like to specify volume across different distances more precisely*” or P27 “*Expanding and reducing objects to represent volume change rather than just using distance to control levels*”, reflecting the strong cross-modal tendency to associate larger visuals with louder sound and viceversa [25, 29]. Users also noted the challenge of distance-loudness calibration, (P24 “*For example close distances always made the object much louder than I wanted*” or P16 “*It would be nice to have a way to tweak the sensitivity with respect to distance. It took some trial and error to figure out how far I actually need to move the object to get the level of loudness I want*”) e.g., moving an object only a short distance sometimes resulted in an unexpectedly large change in loudness. These observations highlight how individual expectations of sound intensity and distance often vary due to distance compression biases [25, 26], where the weight of the HMD itself can be a contributing factor to this bias [75]. Although the mutisensory integration of both visual and auditory stimuli has been proven to gauge better distance localization due to gain re-balancing between senses [53, 58], it has been also observed that in scenarios of level/size discrepancy, the audio component takes precedence over visuals [53], up to the extent that spatial audio, in particular, can overpower visual perception [18, 34]. Therefore, when dealing with spatial audio one has to be careful to not overdo with audio. Several participants noted the benefit of having “visual boundaries” (i.e. weighting on the visual component) in order to have some guidance on level control. (P27 “*Knowing the boundaries of when an object became visible/invisible aka audible/inaudible. Depending on where I was standing in the room, I was wishing there were some guidance on how far I could push out an object on each XYZ plane before they became inaudible/not part of the mix*”). Previous research has shown that visual boundaries can provide better sensitivity in spatial localization [27]. Additionally, our participants re-iteratively requested a quick way to enable or silence specific virtual audio objects so that they can isolate sounds or compare before-and-after placement. This can be linked to selective attention mechanisms [25]. In particular, the “cocktail party effect” describes this selective attention phenomenon, where listeners can concentrate on one channel in a noisy environment and dynamically switch attention as needed [34, 65]. In a mixing context, muting or soloing audio objects leverages *figure-ground organization* – also rooted in Gestalt principles – allowing foreground sounds to be distinguished from background layers. This is where the quantitative results of spread difference between scenes start making sense: in the cinematic the spatial spread of audio objects is significantly larger than in the music mixing scene, and this can be attributed to the fact that participants had a bias towards considering cinematic audio objects as more ambient-related and tried to magnify the spread each point-source audio object by extending the sonic perimeter of their mix. In contrast, in the music mixing scene, most of the audio objects were considered as “salient” and were required to be loud enough in the mix to contribute its timbral and rhythmic component, which lead to narrower spread, mostly at the hand’s reach.

Beyond loudness, participants requested control over other audio attributes, such as reverberation levels or pitch. Even though loudness is considered a primary cue for distance estimation [26], studies report that reverberation and timbral qualities (e.g., the frequency spectrum) are some other cues that greatly influence distance perception. Participants expected to manipulate these cues in tandem with loudness, similarly, modifying the physical placement of virtual markers for a coherent and holistic mixing experience. There is a lot of previous work relating size of visual markers to pitch [24, 25, 34], vertical positioning to pitch [9, 24, 53], even marker color to pitch [24, 59], some of which converged into heuristics adopted in practice by audio mixing practitioners [20, 29].

Another recurring feature request was replicating and grouping virtual audio objects. One participant described how they wanted to duplicate a single virtual audio object to occupy various positions, increasing its spread. Others recommended grouping virtual audio objects and moving them as a unit or “layer”, distilling core design principles approaches to layering and soundscape arrangement [25]. Grouping resonates again with Gestalt principles, where similar audio sources – especially, those sharing timbral or rhythmic traits, are naturally clustered together, evidenced by the similarity and common-fate principles [25, 65]. It is of extreme interest how our participants positioned virtual audio objects in the space, forming semantic patterns between pairs or groups of sounds, as it can be observed in the heat map analysis (Subsection 6.2). This is especially evident in the cinematic scene, where the task was to create a compelling sonic environment which would resemble a real scene. By supporting object duplication, grouping, and spatial layering, AudioMiXR can help organizing granular virtual audio objects into a coherent and sensible whole more effectively.

7.3 Quantitative Comparison: Usability, Workload, and Creativity

To validate and expand on the design lessons from our exploratory study, we conducted a follow-up comparison study (Study 2) contrasting the 6DoF AudioMiXR interface with a traditional 2D panner UI. This allowed us to directly assess whether the qualitative affordances participants described in Study 1, such as embodied interaction, spatial awareness, and dynamic audio feedback translated into measurable advantages across usability, workload, and creative output.

Usability (SUS). Participants rated the AudioMiXR interface as significantly more usable, with improvements across all SUS subscales (Holm-corrected, $p < 0.05$). They found the AR interface to be less cumbersome and reported greater confidence in completing the mixing tasks. This aligns with earlier feedback on proprioception-driven interaction and perceived control, reinforcing the idea that embodied interfaces can enhance usability by reducing cognitive translation between intent and action. These results suggest that even for non-expert users, a well-designed 6DoF interface can exceed the usability of a familiar 2D control scheme.

Workload (NASA-TLX). We observed significantly lower mental demand, frustration, and perceived performance difficulty when participants used AudioMiXR compared to the 2D panner (Holm-corrected $p < 0.05$). These results suggest that the spatial interaction model of AudioMiXR helped reduce the subjective cognitive and emotional load associated with the task, likely by externalizing spatial reasoning and offering more intuitive control. Although other dimensions like effort, temporal demand, and physical demand trended favorably toward the AR interface, they did not reach statistical significance after correction.

Creative Output. We observed a statistically significant improvement across all subscales of our creativity metrics for the AR condition (Holm-corrected $p < 0.0001$), including originality, elaboration, and expressiveness. These findings support our claim that spatial freedom and bodily agency facilitates divergent thinking and experimental design behavior. In particular, users produced more varied spatial arrangements, including novel object groupings and creative layering approaches, echoing qualitative feedback from Study 1 on the importance of gestural expressiveness and environmental awareness.

Taken together, these results substantiate core design lessons identified in our initial study: that AR interfaces leveraging full-body interaction and real-world spatiality can meaningfully enhance the sound design experience—not only in terms of subjective preference or novelty, but also in concrete usability, workload, and creativity metrics. Study 2 thus provides empirical grounding to our earlier design lessons and helps delineate contexts where 6DoF AR tools may be most advantageous, such as spatially rich or exploratory mixing tasks.

7.4 Future Work and Limitations

AudioMiXR can serve a wide range of sound design tasks by virtue of its object-based workflow 6DoF AR interface. While our current implementation focuses on AR, the system design is equally adaptable to VR and we discuss this next, along with an outline of key avenues for expanding AudioMiXR’s functionality and addressing its current limitations.

7.4.1 Extension to VR. Beyond AR, extending AudioMiXR into VR offers immersive sound design in fully virtual environments. Deploying the same 6DoF interaction model in VR would let users build and experience soundscapes in virtual worlds – useful for VR cinematic productions, interactive media, and game audio production. However, VR introduces unique challenges, such as cybersickness and occlusion of the real environment, calling for careful design of user feedback and locomotion loops. See Figure 19 where we test AudioMiXR within AVP’s *Environments* application.



Fig. 19. Audio mixing in virtual environments. This example illustrates a demo of AudioMiXR running on top of AVP’s native *Environments* application, displaying the virtual Yosemite environment.



Fig. 20. Synchronous multi-user collaborative mixing.

7.4.2 Multi-user Collaboration. Although our prototype focuses on single-user, shared XR mixing session is a very attractive direction of research. Enabling synchronous multi-user collaboration in 3D environments, supporting interactions in both remote and co-located spaces, would allow multiple users wearing headsets to co-edit an audio scene in real time. In co-located setups, where users would share the same space, upon AudioMiXR initialization, users would be able to freely navigate in the *XR origin* position within the space and reset the scene

to return to the origin. The reset would transform virtual audio objects present in the scene relative to their reset position. Resetting the scene would require verbal communication to ensure all users agree and maintain spatial awareness, as changes would affect everyone sharing the workspace (Figure 20). Conversely, in remote collaborations across different physical locations, resetting the workspace would adjust all scene assets, including peer users, relative to the new position. To increase co-presence and spatial awareness in remote AR mixing sessions, visual cues would be used. Each user would be represented by a full-body avatar within the shared workspace, facilitating embodiment and identification. During user-object interactions, such as selection and manipulation, the system would display the user's hands and highlight the manipulated object. Simultaneous selection and manipulation of the same 3D audio object by multiple users would be prevented to maintain interaction integrity.

7.5 Applications

Our experiments relied on synthesized spatial audio based on generic HRTFs. We invite future studies to couple AudioMiXR with personalized HRTFs, which can enhance even further the localization accuracy of the virtual audio objects. Furthermore, AudioMiXR's design is not restricted to binaural rendering; it could be easily extended to loudspeaker stereophony, allowing users to design mixes that target specific speaker arrays – opening up research venues centered on immersion and loudspeaker setups.

7.5.1 Immersive Music Production. We foresee AudioMiXR as a tool for musicians and producers to be used jointly with DAWs by serving as an 6DoF panner for accurate positioning of spatial audio. In live performances (Figure 23) – musicians or DJs could experiment with spatial placement of audio elements while on stage, adjusting soundscapes in real time to match the dynamics of the crowd.




Fig. 21. Snapshots of AudioMiXR users in different environments—AudioMiXR allows for audio mixing in any physical environment.

7.5.2 Cinematic and post-production sound. This speculative use case envisions AudioMiXR as a research probe for professional post-production. Re-recording teams could position sound effects and dialogue within a cinema, theatre, or on-set mock-up, then *walk the room* to verify spatial cues from multiple audience perspectives, reducing reliance on heuristic “sweet-spot” listening. Situated authoring of this kind would allow producers to refine object placement in situ, maximizing narrative and emotional impact (Figure 22). This probe invites controlled



Fig. 22. Application scenario illustrating how AudioMiXR could be applied for cinema-style mixing context. The user positions virtual audio objects throughout the auditorium to explore how spatial cues could be balanced for audiences at different seats.



Fig. 23. Application scenario illustrating how AudioMiXR could be applied during a live DJ performance. AudioMiXR being used during a DJ set, illustrating how multiple spatial audio objects can be visually placed and manipulated above the turntables and out into the audience. By leveraging 6DoF interaction, DJs can “grab”, rearrange, and adjust each virtual audio object in real time.

comparison studies with expert mixers (e.g., evaluating localization accuracy, mix completion time), and creative satisfaction against conventional workflows and thereby advances the broader investigation of embodied spatial-audio tools for cinematic sound.

7.5.3 Speculative live-performance scenario. Figure 23 sketches a potential use case in which a DJ employs AudioMiXR’s 6DoF controls to place and sculpt virtual audio objects above the turntables and out into the audience, enabling expressive, real-time spatial mixing. We present this vignette as a research probe rather than as a capability of the current prototype: it surfaces open design questions about embodied control in high-tempo, performance contexts. Positioning AudioMiXR in this way connects our work to broader research venues on immersive authoring tools, creative-workflow support, and collaborative performance environments, and it outlines a concrete agenda for future investigations into spatial-audio interaction.

8 CONCLUSION

In this work, we presented AudioMiXR, to the best of the authors’ knowledge, the first 6DoF AR interface for sound design, enabling direct free-hand manipulation of spatial audio. Therefore, no guidelines for sound design in XR currently exist. We developed two design lessons: (1) Proprioception for AR Sound Design and (2) Balancing Audio-Visual Modalities in AR GUI. The first design lesson underscores how utilizing the body when mixing audio supports perceptibility and self-localization in relation to both the physical environment and the manipulated audio objects. The second design lesson highlights the importance of carefully balancing audio and visual modalities to foster a cohesive mixing experience and provide more intuitive representations of sound.

To substantiate these lessons, we conducted a follow-up comparison study that quantitatively evaluated AudioMiXR against a 2D panner interface. This study confirmed that our AR interface significantly improved usability (SUS), reduced frustration and mental workload (NASA-TLX), and enhanced creative output across multiple dimensions. These results empirically reinforce the qualitative insights from our first study, demonstrating that 6DoF interaction not only feels more engaging and intuitive but also offers measurable benefits in terms of user experience and creative efficacy.

Overall, our findings suggest that embracing 6DoF interaction in XR audio mixing can significantly improve usability, reduce cognitive burden, and expand creative possibilities with spatial audio across domains such as music production, cinematic soundscapes, and interactive media.

Acknowledgments

We would like to thank the staff at Dolby Laboratories for their continuous support and guidance throughout this research.

- [27] Francesca C. Fortenbaugh, Shradha Sanghvi, Michael A. Silver, and Lynn C. Robertson. 2012. Exploring the edges of visual space: The influence of visual boundaries on peripheral localization. *Journal of Vision* 12, 2 (02 2012), 19–19. <https://doi.org/10.1167/12.2.19> arXiv:[https://arvojournals.org/arvo/content_public/journal/jov/932802/jov-12-2-19.pdf](https://arxiv.org/abs/https://arvojournals.org/arvo/content_public/journal/jov/932802/jov-12-2-19.pdf)
- [28] Tom Francart and Jan Wouters. 2007. Perception of across-frequency interaural level differences. *The Journal of the Acoustical Society of America* 122, 5 (2007), 2826–2831.
- [29] David Gibson. 1997. *The art of mixing: a visual guide to recording, engineering, and production*. MixBooks.
- [30] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field methods* 18, 1 (2006), 59–82.
- [31] SG Hart. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Human mental workload/Elsevier* (1988).
- [32] Amal Hashky, Benjamin Rheault, Ahmed Rageeb Ahsan, Brett Benda, Tyler Audino, Samuel Lonneman, and Eric D Ragan. 2024. Multi-Modal User Modeling for Task Guidance: A Dataset for Real-Time Assistance with Stress and Interruption Dynamics. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 544–550.
- [33] Marek Hatala and Ron Wakkary. 2005. Ontology-Based User Modeling in an Augmented Audio Reality System for Museums. *User Modeling and User-Adapted Interaction* 15, 3 (01 Aug 2005), 339–380. <https://doi.org/10.1007/s11257-005-2304-5>
- [34] James Hereford and William Winn. 1994. Non-Speech Sound in Human-Computer Interaction: A Review and Design Guidelines. *Journal of Educational Computing Research* 11, 3 (1994), 211–233. <https://doi.org/10.2190/MKD9-W05T-YJ9Y-81NM> arXiv:<https://doi.org/10.2190/MKD9-W05T-YJ9Y-81NM>
- [35] Jane Hoffswell, Wilmot Li, and Zhicheng Liu. 2020. Techniques for Flexible Responsive Visualization Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376777>
- [36] Zainab Iftikhar, Yumeng Ma, and Jeff Huang. 2023. “Together but not together”: Evaluating Typing Indicators for Interaction-Rich Communication. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [37] Image-Line Software. 2025. FL Studio [Computer software]. <https://www.image-line.com> Version 20.8.
- [38] Markus Jelonek. 2023. VRtoER: When Virtual Reality leads to Accidents: A Community on Reddit as Lens to Insights about VR Safety. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). Association for Computing Machinery, New York, NY, USA, Article 315, 6 pages. <https://doi.org/10.1145/3544549.3585783>
- [39] Steven Jiang, Lawrence Lim, and Misha Sra. 2023. Spatializing Music in Virtual Reality (SUI '23). Association for Computing Machinery, New York, NY, USA, Article 42, 3 pages. <https://doi.org/10.1145/3607822.3618011>
- [40] Steven Jiang, Lawrence Lim, and Misha Sra. 2023. Spatializing Music in Virtual Reality. In *Proceedings of the 2023 ACM Symposium on Spatial User Interaction*. 1–3.
- [41] Mohamed Kari, Reinhard Schütte, and Raj Sodhi. 2023. Scene Responsiveness for Visuotactile Illusions in Mixed Reality. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–15.
- [42] Angelika C Kern and Wolfgang Ellermeier. 2020. Audio in VR: Effects of a soundscape and movement-triggered step sounds on presence. *Front. Robot. AI* 7 (Feb. 2020), 20.
- [43] You-Jin Kim, Myungin Lee, Marko Peljhan, JoAnn Kuchera-Morin, and Tobias Höllerer. 2024. Spatial Orchestra: Locomotion Music Instruments through Spatial Exploration. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–5.
- [44] Tobias Langlotz, Holger Regenbrecht, Stefanie Zollmann, and Dieter Schmalstieg. 2013. Audio stickies: visually-guided spatial audio annotations on a mobile augmented reality platform. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration* (Adelaide, Australia) (OzCHI '13). Association for Computing Machinery, New York, NY, USA, 545–554. <https://doi.org/10.1145/2541016.2541022>
- [45] Matthew Lombard and Theresa Ditton. 1997. At the Heart of It All: The Concept of Presence. *Journal of Computer-Mediated Communication* 3, 2 (09 1997), JCMC321. <https://doi.org/10.1111/j.1083-6101.1997.tb00072.x>
- [46] Sandra Malpica, Ana Serrano, Marcos Allue, Manuel G Bedia, and Belén Masia. 2020. Crossmodal perception in virtual reality. *Multi-media Tools and Applications* 79 (2020), 3311–3331.
- [47] Mark Marrington et al. 2017. Composing with the digital audio workstation. *The singer-songwriter handbook* (2017), 77–89.
- [48] Charles Patrick Martin, Zeruo Liu, Yichen Wang, Wennan He, and Henry Gardner. 2020. Sonic sculpture: Activating engagement with head-mounted augmented reality. *arXiv preprint arXiv:2012.02311* (2020).
- [49] Justin Mathew, Stéphane Huot, and Brian F. G. Katz. 2017. Survey and Implications for the Design of New 3D Audio Authoring Tools. *Journal on Multimodal User Interfaces* 11, 3 (April 2017), 277–287. <https://doi.org/10.1007/s12193-017-0245-z>
- [50] Davide A. Mauro, Rufael Mekuria, and Michele Sanna. 2013. Binaural spatialization for 3D immersive audio communication in a virtual world. In *Proceedings of the 8th Audio Mostly Conference* (Piteå, Sweden) (AM '13). Association for Computing Machinery, New York, NY, USA, Article 8, 8 pages. <https://doi.org/10.1145/2544114.2544115>
- [51] David McGookin, Yolanda Vazquez-Alvarez, Stephen Brewster, and Joanna Bergstrom-Lehtovirta. 2012. Shaking the dead: multimodal location based experiences for un-stewarded archaeological sites. In *Proceedings of the 7th Nordic Conference on Human-Computer*

- Interaction: Making Sense Through Design* (Copenhagen, Denmark) (NordiCHI '12). Association for Computing Machinery, New York, NY, USA, 199–208. <https://doi.org/10.1145/2399016.2399048>
- [52] Oussama Metatla, Nuno N Correia, Fiore Martin, Nick Bryan-Kinns, and Tony Stockman. 2016. Tap the ShapeTones: Exploring the effects of crossmodal congruence in an audio-visual interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 1055–1066.
- [53] Oussama Metatla, Nuno N. Correia, Fiore Martin, Nick Bryan-Kinns, and Tony Stockman. 2016. Tap the ShapeTones: Exploring the Effects of Crossmodal Congruence in an Audio-Visual Interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1055–1066. <https://doi.org/10.1145/2858036.2858456>
- [54] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. 1995. Augmented reality: A class of displays on the reality-virtuality continuum. In *Telemanipulator and telepresence technologies*, Vol. 2351. Spie, 282–292.
- [55] George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* 63, 2 (1956), 81.
- [56] Jörg Müller, Matthias Geier, Christina Dicke, and Sascha Spors. 2014. The boomRoom: mid-air direct interaction with virtual sound sources. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 247–256.
- [57] David Murphy and Flaithri Neff. 2010. Spatial Sound for Computer Games and Virtual Reality. (01 2010). <https://doi.org/10.4018/978-1-61692-828-5.ch014>
- [58] Brian Odegaard, David R Wozny, and Ladan Shams. 2015. Biases in visual, auditory, and audiovisual perception of space. *PLoS computational biology* 11, 12 (2015), e1004649.
- [59] Patrick O'Toole, Maurizio Mancini, and Ian Pitt. 2024. Sound and Colour: Evaluating Auditory-Visual Tests in Virtual Reality and Traditional Desktop Settings. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces*. 1–5.
- [60] Alvaro Pastor and Pierre Bourdin-Kreitz. 2024. Comparing episodic memory outcomes from walking augmented reality and stationary virtual reality encoding experiences. *Scientific Reports* 14, 1 (30 Mar 2024), 7580. <https://doi.org/10.1038/s41598-024-57668-w>
- [61] Jayesh Pillai, Priyanshu Thakur, and Tushar Vaid. 2024. Belonging: Exploring Audio-Visual Cues in Real-Time 6DoF VR Narrative. In *Proceedings of the 14th Indian Conference on Human-Computer Interaction*, Anirudha Joshi and Gavin R. Sim (Eds.). Springer Nature Singapore, Singapore, 171–178.
- [62] Roy A. Ruddell, Ekaterina Volkova, and Heinrich H. Bühlhoff. 2011. Walking improves your cognitive map in environments that are large-scale and large in extent. *ACM Trans. Comput.-Hum. Interact.* 18, 2, Article 10 (July 2011), 20 pages. <https://doi.org/10.1145/1970378.1970384>
- [63] Dariusz Rumiński. 2015. An experimental study of spatial sound usefulness in searching and navigating through AR environments. *Virtual Reality* 19, 3 (01 Nov 2015), 223–233. <https://doi.org/10.1007/s10055-015-0274-4>
- [64] Kaili Shan, Tiger Sun, Jarrod Tart, Brandon Woodard, Irene Humer, and Christian Eckhardt. 2024. Work-in-Progress—Virtual Learning Laboratories for High School Chemistry Lab: An Immersive Learning User Study. *Immersive Learning Research-Academic* (2024), 61–71.
- [65] MT Shelvock. 2016. Gestalt theory and mixing audio. *Innovation in Music II* (2016), 1–14.
- [66] J.L. Sinclair. 2020. *Principles of Game Audio and Sound Design: Sound Design and Audio Implementation for Interactive and Immersive Media*. Taylor & Francis. <https://books.google.com/books?id=EDkPEAAQBAJ>
- [67] Mel Slater, Bernhard Spanlang, Maria V. Sanchez-Vives, and Olaf Blanke. 2010. First Person Experience of Body Transfer in Virtual Reality. *PLOS ONE* 5, 5 (05 2010), 1–9. <https://doi.org/10.1371/journal.pone.0010564>
- [68] Jaka Sodnik, Saso Tomazic, Raphael Grasset, Andreas Duenser, and Mark Billinghurst. 2006. Spatial sound localization in an augmented reality environment. In *Proceedings of the 18th Australia Conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments* (Sydney, Australia) (OZCHI '06). Association for Computing Machinery, New York, NY, USA, 111–118. <https://doi.org/10.1145/1228175.1228197>
- [69] SPARTA Team. 2025. SPARTA [Computer software]. <https://www.sparta-project.org> Version 1.0.
- [70] Christoph Stahl. 2007. The roaring navigator: a group guide for the zoo with shared auditory landmark display. In *Proceedings of the 9th International Conference on Human Computer Interaction with Mobile Devices and Services* (Singapore) (MobileHCI '07). Association for Computing Machinery, New York, NY, USA, 383–386. <https://doi.org/10.1145/1377999.1378042>
- [71] Jacob H. Steffen, James E. Gaskin, Thomas O. Meservy, Jeffrey L. Jenkins, and Iopa Wolman. 2019. Framework of Affordances for Virtual Reality and Augmented Reality. *Journal of Management Information Systems* 36, 3 (2019), 683–729. <https://doi.org/10.1080/07421222.2019.1628877> arXiv:<https://doi.org/10.1080/07421222.2019.1628877>
- [72] Devin Blair Terhune, Sarah Tai, Alan Cowey, Tudor Popescu, and Roi Cohen Kadosh. 2011. Enhanced Cortical Excitability in Grapheme-Color Synesthesia and Its Modulation. *Current Biology* 21, 23 (06 Dec 2011), 2006–2009. <https://doi.org/10.1016/j.cub.2011.10.032>
- [73] Wen-Jie Tseng, Petros Dimitrios Kontrazis, Eric Lecolinet, Samuel Huron, and Jan Gugenheimer. 2024. Understanding Interaction and Breakouts of Safety Boundaries in Virtual Reality Through Mixed-Method Studies. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE Computer Society, Los Alamitos, CA, USA, 482–492. <https://doi.org/10.1109/VR58804.2024.00069>

- [74] Stefan Weber, David Weibel, and Fred W Mast. 2021. How to get there when you are there already? Defining presence in virtual reality and the importance of perceived realism. *Frontiers in psychology* 12 (2021), 628298.
- [75] Peter Willemsen, Mark B. Colton, Sarah H. Creem-Regehr, and William B. Thompson. 2009. The effects of head-mounted display mechanical properties and field of view on distance judgments in virtual environments. *ACM Trans. Appl. Percept.* 6, 2, Article 8 (March 2009), 14 pages. <https://doi.org/10.1145/1498700.1498702>
- [76] Brandon Woodard, Melvin He, Mose Sakashita, Jing Qian, Zainab Iftikhar, and Joseph J. LaViola Jr. 2025. Cam-2-Cam: Exploring the Design Space of Dual-Camera Interactions for Smartphone-based Augmented Reality. arXiv:2504.20035 [cs.HC] <https://arxiv.org/abs/2504.20035>
- [77] Muzhe Wu, Yi Fei Cheng, and David Lindlbauer. 2024. New Ears: An Exploratory Study of Audio Interaction Techniques for Performing Search in a Virtual Reality Environment. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 386–395. <https://doi.org/10.1109/ISMAR62088.2024.00053>
- [78] Jing Yang, Amit Barde, and Mark Billinghurst. 2022. Audio augmented reality: A systematic review of technologies, applications, and future research directions. *journal of the audio engineering society* 70, 10 (2022), 788–809.

A Survey for Professional Audio Mixers

Questions for professional audio mixers:

- (1) Can you describe your general process when mixing audio tracks with spatial audio objects?
- (2) What tools or software do you typically use for spatial audio localization, mixing, and panning? How do these tools support your workflow?
- (3) How do you approach spatialization of audio objects within a 3D environment? What are your considerations when positioning audio objects spatially?
- (4) Can you discuss any specific techniques or strategies you employ to achieve optimal spatialization and panning effects in your projects?
- (5) In your experience, what are the main challenges or difficulties you encounter when working with spatial audio? How do you typically address these challenges?
- (6) How often do you collaborate with others (e.g., other sound engineers, content creators) during the spatial audio production process? Can you describe how collaboration influences your workflow and decision-making?
- (7) Have you encountered any notable successes or breakthroughs in your approach to spatial audio that you'd like to share? How did these impact your projects or clients?
- (8) Looking ahead, what advancements or improvements do you envision in spatial audio technology or tools that would benefit your work?
- (9) Can you describe a recent project where spatial audio played a critical role? What were the specific challenges and successes in that project?
- (10) How do you evaluate the effectiveness of spatial audio mixes? Are there specific criteria or metrics you use to assess quality?
- (11) What audio attributes do you think can benefit from spatialized interaction? (eg. in a audio panner to manipulate, volume, reverb, equalization etc.)
- (12) Finally, based on your experience, what advice would you give to someone new to spatial audio mixing and localization?

B Demographic Questions

- (1) Age
- (2) Sex (M/F)
- (3) Glasses prescription
- (4) Experience with DAWs (1 = no experience, 2 = beginner, 3 = intermediate, 4 = advanced, 5 = expert)
- (5) Experience with Augmented reality prior to the experiment 1-5 (1 = no experience, 2 = beginner, 3 = intermediate, 4 = advanced, 5 = expert)

C NASA TLX Questions for Both Non-experts and Experts

- (1) **Mental Demand:** Using a scale from 1 to 7, where 1 means very low workload and 7 means very high workload, please rate the mental workload involved in manipulating virtual audio objects in the XR system. How mentally demanding was the task? (1 = very low, 7 = very high)
- (2) **Physical Demand:** How physically demanding was it to interact with virtual audio objects using free-hand gestures? How physically demanding was the task? (1 = very low, 7 = very high)
- (3) **Temporal Demand:** How hurried or rushed was the pace of the task? (1 = very low, 7 = very high)
- (4) **Performance:** How accurately were you able to position and manipulate virtual audio objects using free-hand gestures in the XR environment? How successful were you in accomplishing the task? (1 = perfect, 7 = failure)

- (5) **Effort:** How hard did you have to work to accomplish your level of performance? (1 = very low, 7 = very high)
- (6) **Frustration:** To what extent did you experience frustration or confusion while using the XR system to manipulate audio objects? How insecure, discouraged, irritated, stressed, and annoyed were you? (1 = very low, 7 = very high)
- (7) **Efficiency:** Rate the efficiency of completing tasks (e.g., moving, resizing, rotating) with virtual audio objects in the XR system. (1 = very low, 7 = very high)
- (8) **Overall Satisfaction:** On a scale from 1 to 7, how satisfied are you with the overall usability of the XR system for manipulating audio objects? (1 = very low, 7 = very high)

D Short-answer survey

Questions for non-experts:

- (1) How intuitive was the process of navigating and interacting with virtual audio objects using free-hand manipulation?
- (2) Did you encounter any difficulties or frustrations while attempting to manipulate the virtual audio objects? Please describe.
- (3) Were there any features or functionalities that were unclear or difficult to understand during your interaction with the system?
- (4) How comfortable did you feel using the XR system to manipulate audio objects in your physical environment?
- (5) What aspects of the system did you find most useful or enjoyable? Why?
- (6) Were there any specific improvements or additional features you would suggest to enhance the usability of the system?
- (7) (Experts only) Where can you see this interface being applied to your audio mixing workflow?