Audio-Augmented Museum Experiences with Gaze Tracking

Jing Yang
Department of Computer Science
ETH Zurich, Switzerland
jing.yang@inf.ethz.ch

Cheuk Yu Chan
Department of Computer Science
ETH Zurich, Switzerland
chanc@ethz.ch

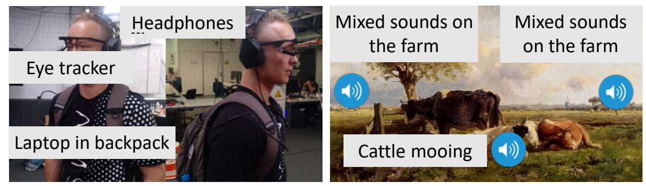

(a) The prototype setup (b) Sounds added to a painting

Figure 1: To experience the gaze tracking-based audio augmentation, a user wears a wearable eye tracker, a pair of unmodified headphones, and a laptop in a backpack, as shown in (a). With this setup, users can hear sounds coming from the painting while viewing it. The sound volume will be adjusted according to the user's gaze. The blue audio icons annotate the virtual sounds that are spatialized with authentic distances and directions relative to the user.

ABSTRACT

In this work, we enrich landscape and genre paintings by spatializing sounds for the drawn objects and scenes, which expands visitors' perception of the paintings and immerses them in the depicted scenarios. Plus, we personalize such spatial audio perception based on visitors' viewing behavior by applying gaze tracking. Through a preliminary user study with 14 participants, we observed that the gaze tracking-based audio augmentation helped people better focus on the areas of interest in the paintings, and enhanced their overall viewing experience.

CCS CONCEPTS

• **Human-centered computing → Mixed / augmented reality**; **Auditory feedback**.

KEYWORDS

Audio augmented reality, Gaze tracking, Museum exhibition, User experience

ACM Reference Format:
Jing Yang and Cheuk Yu Chan. 2019. Audio-Augmented Museum Experiences with Gaze Tracking. In *MUM 2019: 18th International Conference on Mobile and Ubiquitous Multimedia (MUM 2019), November 26–29, 2019, Pisa, Italy.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3365610.3368415

1 INTRODUCTION

Museums and galleries have been motivating researchers to integrate and stimulate multiple human senses, in order to explore novel ways of presenting artworks and to improve visitors' perception and interaction with them [29]. Among these investigations, the auditory sense plays an important role. A typical application is the audio guide that is often used to navigate visitors or give introductions about the exhibits [8, 11]. In recent years, some researchers have enriched these auditory information with spatial sound effects [28] that help to enhance visitors' sense of engagement and spatial awareness.

However, such audio guide only provides "meta data" instead of augmenting the artworks themselves. To this end, some researchers have integrated multiple human senses into the expression of the artwork and thus enhanced visitors' perception. For example, in the Tate Sensorium exhibition [29], paintings were presented together with music, smell, and haptic feelings that were specifically designed to match the artist's intention, which helped visitors perceive the paintings from multiple channels and improved visitors' interaction with the paintings. However, one particular challenge is the lack of personalization that can adapt the enhanced perceptions to the visitor's movements and viewing behaviors.

To facilitate immersive, interactive, and personalized museum experiences, in this work, we enrich landscape and genre paintings by virtually spatializing corresponding sounds for the objects and scenes drawn in the paintings (e.g. cattle mooing, as shown in Figure 1(b)). Plus, we customize visitors' audio perception by tracking their gaze during the viewing process. Specifically, when the visitor focuses on an object or a particular area in a painting, the corresponding sound will be amplified while the other sounds will be attenuated. We implement a prototype using wearable eye tracker for gaze tracking, a pair of normal headphones for audio perception, and a laptop in backpack for gaze calculation, viewing pose estimation, and virtual sound spatialization.

In a user study with 14 participants, people generally reported that the gaze-based audio augmentation helped them better focus on the areas of interest, and the whole pipeline enhanced their experience with the paintings. They could also imagine several other applications with such a gaze-audio augmentation setup, such as effortlessly selecting the sound of interest in an environment with multiple sound sources, seamlessly interacting with smart objects via visual contact, etc.

Our main contributions are:

(1) A user study that explores gaze tracking-based audio augmentation in a museum scenario;

(2) An evaluation of how such gaze-audio augmentation can enhance museum visitors' experience.

MUM 2019, November 26–29, 2019, Pisa, Italy

ACM ISBN 978-1-4503-7624-2/19/11.
https://doi.org/10.1145/3365610.3368415

2 RELATED WORK

The use of visual and auditory senses have been explored in the field of augmented reality (AR) and human-computer interaction (HCI) for a few decades. Researchers have been using the auditory channel, especially with the help of spatialized audio, to redirect a user's attention, to help visually impaired people, and to understand surroundings. Typical applications include navigation [1, 5, 14, 23], notification [3, 12, 13, 24, 27], and audio content creation [21]. By tracking a user's gaze, researchers can understand the user's attention and even intention [2], and then interpret the user's behavior or activate interactions with the environment. Typical application examples include home appliances control [25], clinical practice [4], activity recognition [6, 7], gaze-controlled typing [16], interaction with displays [17], etc.

Researchers have also integrated the auditory sense with the visual sense in several scenarios. For example, some projects help a user understand an urban environment by giving auditory guides to direct the user's attention from current focus [15, 18]; the volume of a character on the screen can be changed based on a user's gaze in a virtual reality scenario [30]; some researchers enhance users' emotional experience with music by providing visual information presentation along with auditory signal [20, 22]; and the combination of these two senses can help to facilitate better human-robot interactions [19, 26].

In museum and art gallery scenarios, the auditory perception has been typically used for navigation and exhibit introductions [8, 11]. Some researchers enriched such auditory information with 3D effects to further improve a user's spatial perception of the environment [10, 28, 31]. The combination of auditory sense and visual sense has also been applied in some museum applications. The LISTEN project [32] personalized a user's audio guide based on their movement in the space and the direction of their gaze. The ARtSENSE system [9] could analyse a visitor's preference profile and then recommend artworks based on the visitor's viewing histories.

In our work, we intend to synthesize and spatialize content-related sounds to enrich the presentation of landscape and genre paintings. We also adjust the visitor's audio perception based on their gaze on the paintings, thus to more actively involve a visitor into the perception and experience with paintings.

3 GAZE TRACKING-BASED AUDIO AUGMENTATION

3.1 Environment Setup and Simulation

In a room of size $6\,m \times 6.5\,m \times 3.4\,m$, we distributed four landscape and genre paintings of which the sizes were around $1\,m^2$. As illustrated in Figure 2, on each painting we used light-yellow tape to frame an object of interest. At the bottom left of the frame and at the corners of each painting we attached ArUco markers. These frames and markers were used for the convenience of the user study and gaze tracking. In the following sections we will give more details.

Figure 2 also shows the virtual sounds attached to the paintings. These audio clips were gathered online from YouTube and Freesound, and spatialized from appropriate locations and distances. We carefully selected these sounds according to the painting contents in order to create authentic soundscape as it might exist in real life.

To model the sound source locations, we made a digital copy of this room in the game engine Unity3D, with the sound sources registered at correct locations in each painting. During the study, participants wore a wearable eye tracker, a pair of unmodified headphones, and a laptop in a backpack, like shown in Figure 1(a). They were asked to view each painting standing at a line that was marked approximately $1.3\,m$ in front of the painting. This distance was decided as a trade-off between smooth gaze tracking and a pleasant painting viewing experience.

3.2 Gaze Tracking and Pose Estimation

We need to estimate users' viewing pose and track their gaze on the paintings for spatializing the virtual soundscape and adjusting the audio perception. To this end, we used a Pupil Labs eye tracker[1]. This eye tracker was equipped with a scene camera that captured the user's field of view (FoV) and one infrared spectrum eye camera that detected the user's pupil. In our setup, the eye camera was mounted for the right eye.

The eye tracker was connected to a laptop running Windows 10 OS for real-time gaze and pose calculation. To calibrate and calculate the viewer's gaze in his/her FoV, we used the plugin hmd-eyes[2] that allows the implementation of Pupil Labs programs in Unity3D. We also used the framework OpenCV for Unity[3] to detect the poses of the ArUco markers that were attached on the paintings. Given the gaze positions and the marker's poses in the user's FoV, we then estimated the user's viewing pose with respect to the painting, and determined whether the user was focusing on the object of interest that was framed in the painting. Accordingly, we could then render and adjust the audio augmentation that corresponded to the user's viewing behavior.

3.3 Spatial Audio Augmentation

Given the user's gaze and pose data in the Unity3D scene, we utilized the Google Resonance Audio SDK to simulate the sound propagation and the spatialized sounds were played to the user via off-the-shelf headphones.

In order to render a natural gaze-based audio adjustment as well as to compensate accidental gaze movements into and out of an area, we set a $1\,s$ activation/deactivation time and a threshold of 80%, i.e., we would regard that the user was viewing the object of interest if 80% of his/her gaze was located in the frame; and it would take $1\,s$ to gradually amplify the corresponding sound while attenuating the other sounds to pre-defined levels. The sound levels would be kept until the user shifted their visual focus. When the user was viewing the other parts of a painting, all virtual sounds would be played at a balanced level without emphasizing any object or scene.

Note that the implementation described in the whole section was specific to our user study. In different real-world applications, the

[1]https://pupil-labs.com/products/core/
[2]https://github.com/pupil-labs/hmd-eyes
[3]https://assetstore.unity.com/packages/tools/integration/opencv-for-unity-21088

Painting 1: *Harvesting* by Henry H. Parker

Painting 2: *De Melkmeid* by Johannes Vermeer

Painting 3: *Resting Cows* by Anton Mauve

Painting 4: *Homeward Bound* by John Frederick Herring Jr

Figure 2: The four paintings used in our application. Virtual sounds are illustrated with blue icons at corresponding positions. We spatialized the sounds with proper depth (1-200 m). The ArUco markers on the paintings were used for gaze tracking and pose estimation. Inside the light-yellow squares were the designated objects of interest for the user study. On real paintings the frames were marked using tapes. Painting 2 is 0.82 m×0.92 m and the others are 1.4 m×0.85 m. The images are from `http://www.paintinghere.com`.

same operating principle can be applied for other gaze-based audio augmentations.

4 EXPERIMENT & EVALUATION

We conducted a user study to explore whether the gaze tracking-based audio augmentation could provide visitors with intuitive and immersive perception of artworks.

4.1 Experiment procedure

We invited 14 participants (age $\in [19, 29]$, average = 23.5, standard deviation = 2.94, six female) to our experiment that consisted of three parts: gaze calibration, painting viewing, and questionnaire.

Gaze calibration was first implemented to establish a mapping from pupil to gaze coordinates. We used one gaze calibration method provided by Pupil Labs, in which nine points on an ellipse were displayed one after another on a screen. Each point was visible for around three seconds, during which the participant focused on the point until it disappeared and the next one showed up. Participants followed the points by only moving their eyes but keeping the head still. To guarantee that the calibrated gaze tracking could work well when viewing paintings, we used a TV screen of which the size was similar to the painting size, and we asked participants to stand 1.3 m in front of the screen. During gaze calibration, we recorded the pupil detection confidences frame by frame for each point. The confidence refers to a value in the range $[0, 1]$ that indicates the quality of the gaze detection (the higher, the better). Pupil Docs[4] state 0.6 as the threshold for valid gaze detection.

After calibrating the gaze tracking, we asked participants to view all four paintings. To experience the difference, participants viewed each painting twice before moving to the next one, once with the gaze tracking-based audio adjustment as described in Section 3.3 and once without such adjustment, i.e., the sounds would always be played at balanced levels regardless of the user's gaze. The order of these two modes was counterbalanced among participants for each painting. Participants were free to view the paintings in their preferred way with only one restriction — they were asked to view the framed object for a few seconds at least once. This way, the sound levels would be properly adjusted for the participants to experience the change.

After viewing all four paintings, participants answered the following questions on a 5-point Likert scale from "strongly disagree" (1) to "strongly agree" (5): (Q1) My audio perception difference between the two modes of viewing was clear. (Q2) I felt that the gaze-based audio adjustment helped me better focus on the area I was looking at. (Q3) I felt that the gaze-based audio adjustment smoothly matched my visual focus shifts. (Q4) For me, the better painting experience was WITHOUT gaze-based audio adjustment. After finishing the questionnaire, we also talked to the participants for their general feedback and suggestions.

[4]https://docs.pupil-labs.com/

Table 1: The mean, median, and standard deviation values of pupil detection confidence for each gaze calibration point.

	P1	P2	P3	P4	P5	P6	P7	P8	P9
mean	0.93	0.86	0.9	0.94	0.77	0.59	0.73	0.87	0.86
median	1	0.97	0.99	1	0.91	0.68	0.85	1	0.98
STD	0.16	0.22	0.2	0.17	0.31	0.34	0.31	0.25	0.23

4.2 Gaze Calibration Results

A good gaze tracking performance was the foundation of the whole working pipeline. Table 1 summarizes the mean, median, and standard deviation values of pupil detection confidence for each calibration point over all participants. Our calibration result was satisfying considering the threshold of 0.6 for valid gaze detection. The results of points 5-7 were obviously worse than the others. This was probably because these three points were towards the left side of the screen while our eye camera was mounted for the right eye. When looking at these three points, the participant's right eye moved towards left, which made it difficult to properly capture the pupil. This problem could influence participants' experience but the impact should not be very serious. One reason was that when viewing the paintings, users would also move their heads when viewing some objects at far angles, so it was less possible that the right pupil would remain difficult to be detected.

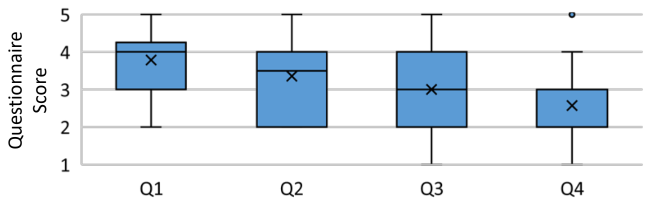

Figure 3: Results of the questionnaire. For Q4, the median value was 2, overlapping the lower limit of the box.

4.3 Experience with the Gaze Tracking-based Audio Augmentation

Figure 3 shows participants' questionnaire answers. As expected, participants could easily distinguish two painting viewing modes (Q1), which indicates that participants could reasonably compare the experience difference between the unchanging audio augmentation and the gaze-based audio adjustment. Regarding Q2, participants leaned towards the opinion that they could better focus on the object of interest when the object sound was amplified over the other sounds. However, some participants reported that they were distracted by such a volume change because the change was a bit too strong so it did not feel very natural to them.

While some participants experienced a smooth audio adjustment that closely matched their visual focus shift, other participants felt the opposite (Q3). One reason could be that the eye tracking system sometimes lost its accuracy due to issues of, for example, viewing angles. Another reason could be traced back to the 1 s activation time as mentioned in Section 3.3. For a few participants this was a bit too long, e.g. *"I had to stare at the object until the sound was adjusted, instead of happening immediately"*. These subjective differences influenced individual's experience.

Finally, regarding Q4, participants tended to agree that gaze-based audio adjustment was better than the unchanging audio augmentation. Some participants stated that the gaze-based audio adjustment smoothly matched their viewing behavior, and *"it felt that the system knew what I wanted to see"*. However, several participants preferred the unchanging audio perception due to various reasons about their subjective feelings. For example, the activation time was a bit long; the sound was amplified too much; and the tracking did not work perfectly, etc. Despite these less pleasant experiences, they all liked the idea to augment a painting with such spatialized content-related sounds.

4.4 Discussion

The user study results show a positive tendency, but are not highly satisfying. We argue that the reason does not lie in the idea of gaze-based audio augmentation, but is more related to technical issues and individual subjective differences. Technically, the wearable eye tracking method implemented in our study was not sufficiently stable. To improve gaze tracking, we can add an eye camera for the left eye. Alternatively, one could also try outside-in eye tracking, with cameras mounted on the paintings. Subjective differences had big influence on people's experiences. This also reflects the general difficulty when stimulating human senses in museum and art galleries to enhance the presentation of artworks. A larger-scale user study can be conducted to collect more opinions from visitors. Since our user study was performed on a relatively young population, it would be reasonable to include people at different ages in future studies. Furthermore, designers can make such a system more flexible for the users to personalize their perception.

Participants also gave us suggestions to improve and extend our system. For example, in addition to adjusting sounds based on gaze, it could also be the other way round — sound change could guide visual attention. This could make use of the spatial feature of spatial sounds.

This gaze-audio augmentation pipeline inspired our participants to imagine several applications outside museum and gallery scenarios. One participant imagined an application of online dictionary — one can hear the pronunciation and explanation of a word when looking at it. Another interesting application scenario is storefront. One can receive information about a product by simply looking at it. Furthermore, participants could imagine selecting the sound of interest in an environment with multiple sound sources, and seamlessly interacting with smart objects via visual contacts and the auditory channel.

5 CONCLUSION & FUTURE WORK

In this work, we implemented a gaze tracking-based audio augmentation pipeline to enhance visitors' perception of paintings in museums and art galleries. A user study indicated that such a gaze-audio augmentation could emphasize the object of interest that was being viewed by the visitor, which seamlessly and intuitively enhanced their focus on a painting.

Although the implementation in this user study was not perfect, it demonstrated the potential of this gaze-audio pipeline in enhancing visitors' museum experiences. Further studies with more participants can be conducted to investigate better paradigms of enriching the presentation of artworks. We also recommend more flexible system design to better personalize visitors' preference and perception.

In addition to museum scenarios, it is also interesting to explore the application of such a gaze-audio pipeline in other everyday applications as discussed in Section 4.4. We believe that the combination of the visual and auditory senses — the most important perceptions for most human beings — can enhance more intuitive human-object interactions, and eventually pushes forward the vision of ubiquitous AR and HCI.

REFERENCES

[1] Robert Albrecht, Riitta Väänänen, and Tapio Lokki. (2016). Guided by Music: Pedestrian and Cyclist Navigation with Route and Beacon Guidance. *Personal and Ubiquitous Computing* 20, 20:1, 20(1), 121–145.

[2] Mihai Bâce, Philippe Schlattner, Vincent Becker, and Gábor Sörös. 2017. Facilitating Object Detection and Recognition through Eye Gaze. In *ACM MobileHCI.*

[3] Amit Barde, Matt Ward, William S. Helton, Mark Billinghurst, and Gun Lee. 2016. Attention Redirection Using Binaurally Spatialised Cues Delivered Over a Bone Conduction Headset. In *Human Factors and Ergonomics Society Annual Meeting*. SAGE Publications, 1534–1538.

[4] K. Bartl-Pokorny, F. Pokorny, S. Bölte, A. Langmann, T. Falck-Ytter, T. Wolin, C. Einspieler, J. Sigafoos, and P. Marschik. 2013. Eye Tracking in Basic Research and Clinical Practice. *Klinische Neurophysiologie* 44, 3 (2013), 193–198.

[5] Simon Blessenohl, Cecily Morrison, Antonio Criminisi, and Jamie Shotton. 2015. Improving Indoor mobility of The Visually Impaired with Depth-based Spatial Sound. In *IEEE ICCV.*

[6] Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Tröster. 2008. Robust Recognition of Reading Activity in Transit using Wearable Electrooculography. In *International Conference on Pervasive Computing*. Springer, 19–37.

[7] Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Tröster. 2009. Eye Movement Analysis for Activity Recognition. In *ACM UbiComp*.

[8] Areti Damala, Pierre Cubaud, Anne Bationo, Pascal Houlier, and Isabelle Marchal. 2008. Bridging the Gap between the Digital and the Physical: Design and Evaluation of a Mobile Augmented Reality Guide for the Museum Visit. In *ACM DIMEA*.

[9] Areti Damala, Tobias Schuchert, Isabel Rodriguez, Jorge Moragues, Kiel Gilleade, and Nenad Stojanovic. (2013). Exploring the Affective Museum Visiting Experience: Adaptive Augmented Reality (A2R) and Cultural Heritage. *International Journal of Heritage in the Digital Era* 2, 1, 2(1), 117–142.

[10] Marcia de Borba Campos, Jaime Sánchez, Anderson C. Martins, Régis Schneider Santana, and Matías Espinoza. 2014. Mobile Navigation through A Science Museum for Users Who Are Blind. In *UAHCI*. 717–728.

[11] Sandra Gebbensleben, Jana Dittmann, and Claus Vielhauer. 2006. Multimodal Audio Guide for Museums and Exhibitions. In *Multimedia on Mobile Devices II*, Vol. 6074. International Society for Optics and Photonics.

[12] Florian Heller and Jan Borchers. 2014. AudioTorch: Using A Smartphone as Directional Microphone in Virtual Audio Spaces. In *ACM MobileHCI*.

[13] Florian Heller, Jayan Jevanesan, Pascal Dietrich, and Jan Borchers. 2016. Where Are We?: Evaluating The Current Rendering Fidelity of Mobile Audio Augmented Reality Systems. In *ACM MobileHCI*.

[14] Florian Heller and Johannes Schöning. 2018. NavigaTone: Seamlessly Embedding Navigation Cues in Mobile Music Listening. In *ACM CHI*.

[15] Yi-Ta Hsieh, Valeria Orso, Salvatore Andolina, Manuela Canaveras, Diogo Cabral, Anna Spagnolli, Luciano Gamberini, and Gamberini Jacucci. 2018. Interweaving Visual and Audio-Haptic Augmented Reality for Urban Exploration. In *ACM DIS*.

[16] Mohamed Khamis, Mariam Hassib, Emanuel von Zezschwitz, Andreas Bulling, and Florian Alt. 2017. GazeTouchPIN: Protecting Sensitive Data on Mobile Devices using Secure Multimodal Authentication. In *ACM ICMI*.

[17] Mohamed Khamis, Axel Hoesl, Alexander Klimczak, Martin Reiss, Florian Alt, and Andreas Bulling. 2017. Eyescout: Active Eye Tracking for Position and Movement Independent Gaze Interaction with Large Public Displays. In *ACM UIST*.

[18] Tiffany CK. Kwok, Peter Kiefer, Victor R. Schinazi, Benjamin Adams, and Martin Raubal. 2019. Gaze-Guided Narratives: Adapting Audio Guide Content to Gaze in Virtual and Real Environments. In *ACM CHI*.

[19] Stéphane Lathuilière, Benoît Massé, Pablo Mesejo, and Radu Horaud. 2019. Neural Network based Reinforcement Learning for Audio-Visual Gaze Control in Human-Robot Interaction. *Pattern Recognition Letters* 118 (2019), 61–71.

[20] Yaxiong Ma, Yixue Hao, Min Chen, Jincai Chen, Ping Lu, and Andrej Košir. 2019. Audio-Visual Emotion Fusion (AVEF): A Deep Efficient Weighted Approach. *Information Fusion* 46 (2019), 184–192.

[21] Jörg Müller, Matthias Geier, Christina Dicke, and Sascha Spors. 2014. The Boom-Room: Mid-air Direct Interaction with Virtual Sound Sources. In *ACM CHI*.

[22] Fada Pan, Li Zhang, Yuhong Ou, and Xinni Zhang. 2019. The Audio-Visual Integration Effect on Music Emotion: Behavioral and Physiological Evidence. *PloS one* 14, 5 (2019).

[23] Spencer Russell, Gershon Dublon, and Joseph A. Paradiso. 2016. HearThere: Networked Sensory ProsThetics through Auditory Augmented Reality. In *ACM AH*.

[24] Eldon Schoop, James Smith, and Bjoern Hartmann. 2018. HindSight: Enhancing Spatial Awareness by Sonifying Detected Objects in Real-Time 360-Degree Video. In *ACM CHI*.

[25] Jeffrey S. Shell, Roel Vertegaal, Aadil Mamuji, Thanh Pham, Changuk Sohn, and Alexander W. Skaburskis. 2003. Eyepliances and Eyereason: Using Attention to Drive Interactions with Ubiquitous Appliances. In *Extended Abstracts of UIST*.

[26] Zhihao Shen, Armagan Elibol, and Nak Young Chong. 2019. Inferring Human Personality Traits in Human-Robot Social Interaction. In *ACM/IEEE HRI*.

[27] Titus JJ. Tang and Wai Ho Li. 2014. An Assistive Eyewear Prototype That Interactively Converts 3D Object Locations into Spatial Audio. In *ACM ISWC*.

[28] Yolanda Vazquez-Alvarez, Matthew P. Aylett, Stephen A. Brewster, Rocio von Jungenfeld, and Antti Virolainen. 2014. Multilevel Auditory Displays for Mobile Eyes-free Location-based Interaction. In *ACM CHI*.

[29] Chi Thanh Vi, Damien Ablart, Elia Gatti, Carlos Velasco, and Marianna Obrist. 2017. Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition. *International Journal of Human-Computer Studies* 108 (2017), 1–14.

[30] Margarita Vinnikov, Robert S. Allison, and Suzette Fernandes. 2017. Gaze-contingent Auditory Displays for Improved Spatial Attention in Virtual Reality. *ACM Transactions on Computer-Human Interaction* 24, 3 (2017), 19.

[31] Ron Wakkary and Marek Hatala. (2007). Situated Play in A Tangible Interface and Adaptive Audio Museum Guide. *Personal and Ubiquitous Computing* 11, 3, 11(3), 171–191.

[32] Andreas Zimmermann and Andreas Lorenz. 2008. LISTEN: a User-Adaptive Audio-Augmented Museum Guide. *User Modeling and User-Adapted Interaction* 18, 5 (2008), 389–416.