

An Evaluation of Radar Metaphors for Providing Directional Stimuli Using Non-Verbal Sound

Brendan Cassidy

Child Computer Interaction Group
University of Central Lancashire
Preston, Lancashire, UK
bcassidy1@uclan.ac.uk

Janet C Read

Child Computer Interaction Group
University of Central Lancashire
Preston, Lancashire, UK
jcread@uclan.ac.uk

I Scott MacKenzie

Department of Computing
York University
Toronto, Canada
mack@yorku.ca

ABSTRACT

We compared four audio-based radar metaphors for providing directional stimuli to users of AR headsets. The metaphors are clock face, compass, white noise, and scale. Each metaphor, or method, signals the movement of a virtual arm in a radar sweep. In a user study, statistically significant differences were observed for accuracy and response time. Beat-based methods (clock face, compass) elicited responses biased to the left of the stimulus location, and non-beat-based methods (white noise, scale) produced responses biased to the right of the stimulus location. The beat methods were more accurate than the non-beat methods. However, the non-beat methods elicited quicker responses. We also discuss how response accuracy varies along the radar sweep between methods. These observations contribute design insights for non-verbal, non-visual directional prompting.

CCS CONCEPTS

• Human Centered Computing → **Auditory Feedback; User Studies; Mixed / Augmented Reality**

KEYWORDS: Visual Impairment; Radar; Augmented Reality; Headset; Accessibility; Directional Prompting; Spatial Audio

ACM Reference format:

Brendan Cassidy, Janet C Read & I Scott Mackenzie. 2019. An Evaluation of Radar Metaphors for Providing Directional Stimuli Using Non-Verbal Sound. In *2019 CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland, UK. ACM, NY, NY, USA. Paper 59, 8 pages. <https://doi.org/10.1145/3290605.3300289>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI 2019, May 4–9, 2019, Glasgow, Scotland, UK.
© 2019 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-5970-2/19/05...\$15.00.
DOI: <https://doi.org/10.1145/3290605.3300289>

1 INTRODUCTION

Augmented Reality (AR) headsets, such as the Microsoft HoloLens, allow placing digital content into the real world. To do this effectively, headsets need to be aware of the environment they are in. For example, a headset needs to know the location of the surface of a table in order to place a virtual object on it. Awareness of position is done through a process commonly known as spatial mapping. With this method a 3D mesh is generated over real-world objects that virtual objects can be occluded by or placed upon. This digital representation of the environment is potentially useful for people who are unable to visually perceive their environment, such as people with a visual impairment. To make this information useful to people with visual impairments, a method is needed to convert the visual representation of the environment into a modality that the user can perceive. For this, sound is a good candidate.




Figure 1: A radar metaphor for directional stimuli

Using sound, one method is to verbalize the user's immediate surroundings based on head position using a ray cast from the headset into the environment. Similar systems using hand-held pointing devices are effective [16, 21]. A drawback is the need to continuously move the pointing/scanning device to search an area, which leads to user fatigue and an increased risk of the user losing awareness of important parts of their environment. An alternative method – successfully evaluated as an assistive technology – is the use of a radar metaphor [8, 15], whereby a virtual radar arm sweeps around a user and activates a stimulus at a point of interest (see Figure 1). We believe audio radar metaphors have potential for AR

headsets. Audio radar sweeps encoded with environmental information extrapolated from spatial mapping data have the potential to provide an “audio snapshot” of visually impaired users’ surroundings. While visual augmentations are of use to visually impaired users, audio augmentations that use the same spatial mapping data could be helpful. This research builds on prior work by investigating alternative methods for delineating the movement of a radar arm when using such a metaphor.

2 RELATED WORK

Advances in AR headsets create opportunities to present the environment to users in new and novel ways. Much research in this area focuses on using headset displays to visualize parts of the environment not normally seen, such as the visualization of sound fields [10] or visualizations for spatially locating sound [11]. Less research exists on presenting local environmental information (such as spatial mapping data) using audio. Early research indicates spatially augmenting sound within an environment for musical appreciation still requires visual accompaniment [13]. The use of radar metaphors, as described in section 1, provide an opportunity to convey spatial information without the need for accompanying visuals. Fujimoto and Turk [8] compared a radar metaphor against other audio and haptic-based directional prompting methods including directional audio, Geiger counter metaphors [9], and haptic pulses [14]. The radar method was most accurate (when used in combination with hard distractor tasks), but significantly more annoying and more difficult to use than the other directional prompting methods. A reason is perhaps that the radar method involved panning the audio from left to right to delineate the passage of the radar arm, which some users found disorientating.

It is possible that user experience and acceptance could be improved by examining other methods for delineating the passage of a virtual radar arm – methods that do not involve directional audio. Another radar method is Rumelin et al.’s NaviRadar [15]. This approach used tactile stimuli to delineate the movement of the radar arm, with a haptic pulse identifying the start of the radar sweep and another identifying the location of the directional stimulus. This required users to track the movement of the radar arm by timing the gap between haptic pulses with no intermittent delineation of the passage of the radar arm. The method reported a mean deviation of 37° from a full 360° radar sweep and achieved perceived usability and navigation performance similar to spoken instructions. This demonstrates use of a radar metaphor in providing

directional stimuli. However, when implemented using sound, especially spatial audio, users reported a poor user experience.

Given the paucity of research on how radar sweep metaphors are delineated, reference is made to studies on the estimation of position with number lines. Whilst only loosely related to the radar metaphor, the work suggests that with more delineations, it is easier to estimate position, whilst at the same time pointing to variability in how different users approach estimation of positions [18,19,20]. Given this interest and mixed messages on how the delineation of a radar sweep affects positioning a sound, the present study explores the same problem in the context of hearing, and then positioning, sounds. This could be due to how the radar sweep is delineated and worthy of further investigation.

3 THE RADAR METAPHORS

Four methods for delineating the movement of a radar arm were developed for evaluation (see Figure 2). Two (clock face, compass) were beat-based [5], wherein a beat delineated the movement/position of the radar arm. The other two (scale, white noise) were non-beat-based. The scale method used a constant tone changing in musical pitch to delineate movement of the radar arm. White noise used a constant white noise sound, containing no intermittent delineations.

3.1 Clock Face Method

The clock face method used a 3/4 waltz-style beat resulting in seven beats in each radar sweep. As the radar arm sweeps from left to right, the timing of the beats and the position of the radar arm map directly to the hours on a clock face. see Figure 2a. The first and last beats (9 o’clock and 3 o’clock) are represented with a distinct sound and timbre that is higher in pitch and slightly longer in length. This was done to help distinguish positions of the radar arm (left, forward, right). As such, the middle beat in the sequence (12 o’clock) is represented with the same distinct sound. A clock face was chosen as it is a directional/orientation metaphor that has been successfully used to convey direction with visually impaired people in assistive technologies [4].

3.2 Compass Method

The compass method used a 4/4 beat with eight beats in each radar sweep. See Figure 2b. The beat maps directly to the cardinal points of a compass, another directional metaphor used in assistive technology [14].




Figure 2. Four audio methods for delineating the passage of a radar arm for directional prompting. (a) clock face, (b) compass, (c) scale, and (d) white noise.

The compass method works in the same way as the clock face method with the beat sounds activating as the radar arm passes over them. The key cardinal points of the compass (west, north, east) are represented with the same distinct sounds as with the clock face method (9 o'clock, 12 o'clock, 3 o'clock).

3.3 Scale Method

The scale method was non-beat-based with the same number of delineations and mappings as the compass method. See Figure 2c. Instead of percussive beats delineating the passage of the radar arm, a rising C major scale was used. With this method, a sound is played while the radar arm is moving, with the pitch of the sound rising note-by-note as the radar arm passes over each delineation seen in Figure 2c. There were no other distinct indicators for left/forwards/right as seen with the compass or clock face methods. Instead, the user is made aware of the start and end of the sweep by the beginning and end of the sound itself. A scale method was chosen to examine the differences between beat and note-based delineations.

3.4 White Noise Method

The white noise method contained no delineations and was simply white noise. See Figure 2d. White noise was played for the duration of the radar sweep. This was chosen as a control, to investigate if delineations are at all useful or if no delineations are just as good.

4 METHOD

A user study was conducted to compare the four radar metaphors just described.

4.1 Participants

Thirty-six sighted adult participants (18-56 yr) were recruited (19 male, 17 female). Participants were recruited from the local university campus. None of the participants were involved in the design of the research. Participants were asked about their experience playing a musical instrument, but this was not a factor in participation. Fifteen people were experienced in playing a musical instrument; 21 people reported no musical instrument experience.

4.2 Apparatus

Participants completed the study on a 17.5 inch Intel i7 laptop with GTX 1080 graphics card, 32gb RAM, and Windows 10. They each wore identical low-cost headphones to receive sound stimuli (impedance 32 ohms, sensitivity 104 dB, frequency response 20 Hz – 20 kHz). Study data were captured by custom software and saved to a .csv file on the laptop. Preference data were captured via a paper-based debrief questionnaire at the end of the study.

The experiment software was written in C# using the Unity 3D game engine. This software comprises a setup screen for the experimenter to use, a briefing screen for the

each of the four radar methods, a practice screen where participants can familiarize themselves with each method, and the testing screens. In the testing screens, participants are first presented with a blank screen on which they click to trigger a demo run of the radar sweep sound. With this, they hear one of the four radar sounds (depending on the current method) as well as the sound of a location stimulus in a random location along the sweep. The duration of each sweep was 3 seconds for all methods. For each trial the location stimulus was placed at a random location along the radar arc between -90° (beginning of the radar arc) and 90° (end of the radar arc). Response accuracy was calculated as the difference in angle in degrees between the response location and the location stimulus. If the response location was provided before the stimulus location, this would result in a negative response accuracy (e.g., -8°), if the response location was after the stimulus location on the radar arc, this would result in a positive response accuracy (e.g., 10°).

Participants see a copy of the arc (Figure 3) and then estimate where they thought the location stimulus occurred along the radar sweep. Their estimate is provided through a drag and drop interface. Figure 3 illustrates the response screen of the study software.

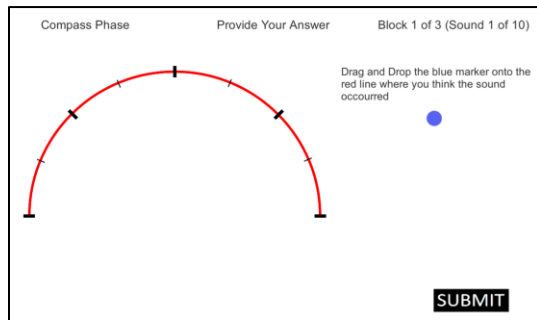


Figure 3. Study software example answer screen

4.3 Procedure

Before the study began, the necessary consents were obtained, and participants were asked if they had any experience playing a musical instrument(s). This was to examine if musical training improved performance with the different radar methods. According to Zhao et al. [22], people with musical experience have a better understanding of metrical structures. Each participant completed stimulus location tasks for all four of the radar delineation methods sequentially. To counter learning effects, order of presenting radar methods was counterbalanced using a 4×4 Latin square.

There was a separate briefing for each radar method before each set of location tasks began. Each briefing

session was in three parts and on three separate screens. The first screen presented a visualization of the radar arm moving while the radar delineation sounds are playing. This allowed the user to gain an understanding of how the sound represents the movement of the radar arm. Participants could listen to the radar method as many times as they liked, up to a maximum of two minutes; none required the full two minutes before moving to part two of the briefing. Figure 4 provides an example of how the radar arm is visualized during the briefing.




Figure 4. Part 2 of a radar method briefing

In part two of the briefing, an additional sound is introduced in the radar sweep to provide the location stimulus. Users were again permitted to play the radar arm as required (up to two minutes) to listen to the location stimulus embedded in the radar sweep. For all radar methods, the location stimulus is a wooden “knock” sound with its location represented as a blue dot on the radar arc (see Figure 4). The sound was similar to a knock on a wooden tabletop. This was chosen for its brevity (110 milliseconds) and unique timbre as this is important in making sounds more distinguishable in computer systems [12]. As it is difficult to distinguish any sound in white noise, during the white noise radar method, the white noise paused when the location stimulus was activated and resumed immediately after the sound ended. This was seamless and helped prevent the location stimulus from being drowned out by the large frequency range of the white noise. The final part of the briefing was a single example location stimulus task. During testing, the radar arc and arm are not visible. The user clicks a button to start the task, playing the radar and embedded location stimulus once. The radar arc is then made visible and the user places a marker on the radar arc where they thought the location stimulus was located. Figure 3 shows an example of the screen participants use to provide the response location.

During the study, each radar method was evaluated with 30 trials, split into three consecutive blocks of 10 trials. At the end of each block the user was shown their mean

accuracy, so they could see if they were improving. At the end of the third block, the participant would move to the briefing stage of the next radar method. Participants were given unlimited time to respond. Response time was measured from when the radar arc was visible (after the radar sweep sound had played) to when the user clicked the submit button. Participants were directed to proceed at a pace they felt comfortable with. When the trials were complete, participants were asked to complete a short questionnaire on user preference. The questionnaire used a 7-point Likert scale and asked participants to rate the difficulty of each radar method from very easy to very hard. Participants then ranked the radar methods in order of preference.

4.4 Design

The study used a $4 \times 3 \times 10$ within-subjects design with the following independent variables and levels:

Radar method (*clock, compass, white noise, scale*)

Block (*1, 2, 3*)

Trial (*1, 2, 3, 4, 5, 6, 7, 8, 9, 10*)

As well, group was a between-subjects factor for counterbalancing the radar method condition. There were four groups with nine participants in each group. The dependent variables were response accuracy ($^{\circ}$), absolute response accuracy ($^{\circ}$), and response time (s).

5 RESULTS AND DISCUSSION

5.1 Response Accuracy

The effect of radar method on response accuracy was statistically significant ($F_{3,96} = 5.23$, $p < .005$). The extra delineation in the compass method appeared to help reduce overall inaccuracy compared to the clock face method. The effect of block on response accuracy was not statistically significant ($F_{2,64} = 0.26$, ns). This was likely because participants were given a good opportunity to familiarize themselves with the radar methods in the briefings. The Radar Method $\times$ Block interaction effect on response accuracy was not statistically significant ($F_{6,192} = 1.74$, $p > .05$). Mean response accuracy values indicated that on average people provided responses earlier in the radar sweep for beat-based methods (clock: -1.80° , compass: -0.58°) This is in contrast to the non-beat-based methods, where participants were observed to respond later in the radar sweep than the beat-based methods (noise: 2.07° scale: 2.12°) See Fig 5.




Figure 5. Response accuracy by radar method. Error bars show ± 1 SE.

There is a possibility that these results associate with work from line estimation studies where it has been previously noted that individuals determine a general area quite quickly and then position within a small region around that choice, and, in this case, decide whether to swing slightly left or slightly right [1, 20]. Speaking with our participants, it appears that in the beat-based methods there was some post-proportioning which encouraged the left-leaning (negative) response accuracy in the beat-based methods as the individuals used the two adjacent beats for reference points. Positions overall were selected earlier than the location stimulus (-0.89°) for participants with musical experience compared to later (1.42°) for participants with no musical experience.

5.2 Absolute Response Accuracy

To get a clearer understanding of response accuracy, absolute response accuracy rates were also analyzed. As response accuracy can be both positive and negative (depending on if the user responded before or after the location stimulus), responses on either side of the location stimulus balance out over time. By examining absolute response accuracy (where positive and negative values are treated the same) we see how many degrees inaccurate (left or right) users were between radar methods. The effect of radar method on absolute response accuracy was statistically significant ($F_{3,96} = 5.49$, $p < .005$). Figure 6 illustrates absolute response accuracy.

The compass method had the best mean absolute response accuracy of 9.89° . The scale method had the second-best absolute response accuracy at 10.19° . It is worth noting that both these methods contained the maximum number of eight delineations.




Figure 6. Absolute response accuracy by radar method. Error bars show ± 1 SE.

The scale method had a better absolute response accuracy than the clock face method (10.99°) despite not having a distinct central delineation in the middle of the radar sweep. The noise method had the worst absolute response accuracy with a mean of 12.97° .

These results suggest that audio delineations during a radar sweep help improve accuracy when identifying location stimulus. The effect of block on absolute response accuracy was not statistically significant ($F_{2,64} = 1.06$, $p > .05$). Results for absolute response accuracy by radar method and block are illustrated in Table 1. The Radar Method $\times$ Block interaction effect on absolute response accuracy was also not statistically significant ($F_{6,192} = 2.21$, $p < .05$).

Table 1. Absolute response accuracy by radar method and block

Radar method	Block		
	1	2	3
Clock	11.45	10.31	11.21
Compass	9.90	10.13	9.65
Scale	10.91	10.45	9.19
White noise	13.55	12.26	13.09

Overall participants with musical experience had a 16% better mean absolute response accuracy (9.91°) compared with participants with no musical experience (11.8°). This confirms the idea that musicians are better than non-musicians at processing metrical structures [22].

5.3 Response Time

The effect of radar method on response time was statistically significant ($F_{3,96} = 62.1$, $p < .005$). The non-beat-based methods elicited the fastest response times (noise:

4.48 s, scale: 4.8 s) compared to beat-based methods (clock: 4.93 s, compass: 5.03 s) Figure 7 illustrates.




Figure 7: Response time by radar method. Error bars show ± 1 SE.

A reason for the 4.8% increase in response time between compass and scale (which both have eight delineations) could be reflective of different strategies identifying the position of the location stimulus. This idea is in line with facilitators observations of participants counting the beat-based sounds (clock, compass) in their heads or on fingers before providing responses. This was in contrast to the non-beat-based methods (scale, white noise) where participants responded based on their feel for where the radar arm was. This explanation would account for both the faster response time of non-beat-based methods and the better accuracy of beat-based methods.

5.4 Response Accuracy by Sweep Location

An analysis of absolute response accuracy by location stimulus indicated mean absolute response accuracy was worse for all radar methods in the second half of the radar sweep (clock: 11.71° , compass: 10.47° , scale: 10.57° , noise: 13.18°) when compared to the first half (clock: 10.24° , compass: 9.34° , scale: 9.81° , noise: 12.76°). The analysis was expanded to examine performance in eight equal sections along the radar sweep corresponding to the delineations of the compass and scale method. Accuracy increased towards the beginning and end of the radar sweep for all methods. The beginning and end of the radar sweep are unambiguous reference points for estimating the position of the radar arm. If a participant misjudges the speed of the radar arm, the margin for error will be less at the beginning as less time has passed. Similarly, if only a small amount of time passes from when the location stimulus is activated to the end of the radar sweep, participants can readjust any inaccuracies they initially had in identifying location stimulus to nearer the end of the sweep. Figure 8 illustrates

mean absolute response accuracy in degrees by location stimulus. The figure is separated into beat-based (top) and non-beat-based methods (bottom).




Figure 8. Absolute response accuracy by location stimulus

Note that the non-beat-based methods (white noise, scale) reveal a decrease in performance during the middle of the radar sweep as this is farthest from the beginning/end reference points. Beat-based methods have the best accuracy at the beginning and end of the sweep but show a small decrease in response accuracy at the center. This is likely because of the distinctive central delineation in beat-based methods (clock: 12 o'clock, compass: north). These findings are supported in the line estimation research which found that locations close to delineation points are less likely to be misplaced (i.e., show better accuracy) [17, 20].

5.5 Subjective Results

Finally, an analysis of user preference data was conducted. 35 of 36 participants completed the debrief questionnaire. The clock face method was most preferred (Table 2). This is despite it being less accurate than the compass method. One explanation for this, which was supported by observations and comments from participants during the study, was that, with the increased number of delineations in the compass method, users felt there was “too much going on” when listening to the radar sweep. Participants found the delineations useful (as is reflected in the accuracy data) but there was a greater cognitive load with more beat-based delineations. This could also explain the longer response

time of the compass method compared to the clock face method (5.03 vs. 4.93 s).

Table 2. Radar method preference

Radar method	1st	2nd	3rd	4th
Clock	14	8	6	7
Compass	10	12	7	6
Scale	7	9	14	5
White noise	4	6	8	17

It is interesting to note that the opposite is true when considering non-beat-based delineations. Participants prefer at least some method of delineation over none. The least preferred method by participants was white noise. This method contained no delineations on the radar arc, so this result is hardly surprising. What is of more interest is that 31% of participants rated it as their first or second most preferred method. This suggests that a significant number of people find it difficult to parse delineated radar sweeps when identifying location stimulus. We speculate that with increased experience, accuracy could improve with non-beat-based methods. Scale was the only method to display a continuous improvement in accuracy as the study progressed, which may be a small indicator. As experience increases there could be more users also showing a preference for non-beat-based approaches, despite lower accuracy. Other factors such as how “agreeable” the radar audio sound is and how often it is activated must also be considered.

6 CONCLUSIONS AND FURTHER WORK

On average, participants were able to perceive location stimulus to within 13° across all methods. With the compass method, this improved to below 10°. Beat-based delineations had better accuracy when identifying location stimulus. Non-beat-based delineations (using constant tones) elicited quicker responses from participants. There was a disparity between user preference and accuracy for beat-based methods, with participants preferring fewer delineations in the radar sweep. If accuracy is flexible, designers of radar metaphors should consider exploring non-beat-based methods of delineation to reduce cognitive load.

There was variation between participants in both performance and preference for each method. This suggests providing options to personalize a radar sweep system to the method users are most comfortable with would be helpful, but (without further work) at the expense of accuracy in some cases. Participants with musical training performed better in all areas than those without. They were quicker and more accurate. This suggests that others

who have enhanced auditory skills may also perform well, and this gives encouragement to the idea that visually-impaired users could find these metaphors useful.

Investigation into the effect of multiple location stimuli on a single radar arm is the next step towards creating a metaphor that can provide a bigger picture of the user's surroundings rather than a single location stimulus. It is possible that a method that works well for a single location stimulus breaks down when it is used with multiple stimuli. Similarly, methods that performed less well with a single stimulus may better support multiple prompts. The stimulus location tasks in this study were all conducted in isolation, further work could also include simultaneous tasks to assess the effects of additional cognitive load on accuracy of sweep methods.

There is also the possibility to investigate how well a user can identify and distinguish different types of location stimuli, for example, using a different timbre for each location stimulus. Integration of effects, such as low-pass filters, could help encode extra information, such as proximity to objects. This has been used before effectively in other location stimuli systems [7]. Adjusting cutoff/resonance of the radar delineation sound could also be useful in conveying information such as proximity to walls. Further work with different participant groups will also help provide insight into the preferences of each group, for example, with children or visually-impaired people. This study was a lab-based study carried out on a laptop. Another logical extension to the work is to implement the appropriate radar methods on an AR headset and evaluate with live spatial mapping data.

ACKNOWLEDGEMENTS

This research was part funded by UK Research Council EPSRC Grant No. EP/L027658/1. 'WE ARE ABLE: Displays and Play'.

REFERENCES

- [1] Hilary Barth, Emily Slusser, Dale Cohen, and Annie Paladino. 2011. "A sense of proportion: Commentary on Opfer, Siegler, and Young." *Developmental Science*: 1205.
- [2] Susanne Brandler and Thomas H. Rammsayer. 2003. "Differences in mental abilities between musicians and non-musicians." *Psychology of music* 31.2: 123-138.
- [3] Stephen A. Brewster, Peter C. Wright, and Alistair DN Edwards. 1993. "An evaluation of earcons for use in auditory human-computer interfaces." In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*, pp. 222-227. ACM.
- [4] Brendan Cassidy, Gilbert Cockton, and Lynne Coventry. 2013. "A haptic ATM interface to assist visually impaired users." In *Proceedings of the 15th international ACM SIGACCESS conference on computers and accessibility*, p. 1. ACM.
- [5] Grosvenor Cooper and Leonard B. Meyer. 1963. *The rhythmic structure of music*. University of Chicago Press
- [6] Saravanan Ergovan Nicole Payne, Jacek Smurzynski, and Marc Fagelson. 2018. "Musical Training Influences Auditory Temporal Processing." *Journal of Hearing Science* 6, no. 3: 37-44.
- [7] Richard Etter and Marcus Specht. 2008. *Melodious Walkabout—Implicit Navigation with Contextualized Personal Audio Contents*. Adjunct Proc. 3rd Int. Conf. Pervasive Computing, Munich, Germany, May 8–13, 2005, pp. 43–49
- [8] Emily Fujimoto and Matthew Turk. 2014. "Non-visual navigation using combined audio music and haptic cues." In *Proceedings of the 16th International Conference on Multimodal Interaction*, pp. 411–418. ACM
- [9] Simon Holland David R. Morse, and Henrik Gedenryd. "AudioGPS: Spatial audio navigation with a minimal attention interface." *Personal and Ubiquitous computing* 6, no. 4: 253-259.
- [10] Atsuto Inoue, Kohei Yatabe, Yasuhiro Oikawa, and Yusuke Ikeda. 2002. "Visualization of 3D sound field using see-through head mounted display." In *ACM SIGGRAPH 2017 Posters*, p. 34. ACM, 2017.
- [11] Dhruv Jain, Leah Findlater, Jamie Gilkeson, Benjamin Holland, Ramani Duraiswami, Dmitry Zotkin, Christian Vogler, and Jon E. Froehlich. 2015. "Head-mounted display visualizations to support sound awareness for the deaf and hard of hearing." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 241-250. ACM
- [12] David K. McGookin and Stephen A. Brewster. 2004. "Understanding concurrent earcons: Applying auditory scene analysis principles to concurrent earcon recognition." *ACM Transactions on Applied Perception (TAP)* 1, no. 2: 130-155.
- [13] Shoki Miyagawa, Yuki Koyama, Jun Kato, Masataka Goto, and Shigeo Morishima. 2018. "Placing Music in Space: A Study on Music Appreciation with Spatial Mapping." In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*, pp. 39-43. ACM
- [14] Martin Pielot, Benjamin Poppinga, Wilko Heuten, and Susanne Boll. 2011. "A tactile compass for eyes-free pedestrian navigation." In *IFIP Conference on Human-Computer Interaction*, pp. 640-656. Springer, Berlin, Heidelberg
- [15] Sonja Rümelin, Enrico Rukzio, and Robert Hardy. 2011. "NaviRadar: a novel tactile information display for pedestrian navigation." In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pp. 293-302. ACM
- [16] Simon Robinson, Parisa Eslambolchilar, and Matt Jones. 2009. "Sweep-Shake: finding digital resources in physical environments." In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*, p. 12. ACM
- [17] Michael Schneider, Angela Heine, Verena Thaler, Joke Torbeyns, Bert De Smedt, Lieven Verschaffel, Arthur M. Jacobs, and Elsbeth Stern. 2008. "A validation of eye movements as a measure of elementary school children's developing number sense." *Cognitive Development* 23, no. 3: 409-422.
- [18] Robert S. Siegler and John E. Opfer. 2003. "The development of numerical estimation: Evidence for multiple representations of numerical quantity." *Psychological science* 14, no. 3: 237-250.
- [19] Robert S. Siegler and Clarissa A. Thompson. 2014. "Numerical landmarks are useful—except when they're not." *Journal of experimental child psychology* 120: 39-58. (2)
- [20] Jessica L. Sullivan, Barbara J. Juhasz, Timothy J. Slattery, and Hilary C. Barth. 2011. "Adults' number-line estimation strategies: Evidence from eye movements." *Psychonomic bulletin & review* 18, no. 3: 557-563.
- [21] John Williamson, Simon Robinson, Craig Stewart, Roderick Murray-Smith, Matt Jones, and Stephen Brewster. 2010. "Social gravity: a virtual elastic tether for casual, privacy-preserving pedestrian rendezvous." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1485-1494. ACM
- [22] Christina Zhao, HT Gloria Lam, Harkirat Sohi, and Patricia K. Kuhl. 2017. "Neural processing of musical meter in musicians and non-musicians." *Neuropsychologia* 106: 289-297.