

Considering Spatial Cognition of Blind Travelers in Utilizing Augmented Reality for Navigation

Paul Ruvolo
Olin College of Engineering
Needham, USA
Paul.Ruvolo@olin.edu

Abstract—Recent work has shown that smartphone-based augmented reality technology (AR) has the potential to be leveraged by people who are blind or visually impaired (BVI) for indoor navigation. The fact that this technology is low-cost, widely available, and portable further amplifies the opportunities for impact. However, when utilizing AR for navigation, there are many possible ways to communicate the spatial information encoded in the AR world to the user, and the choice of how this information is presented to the user may have profound effects on the usability of this information for navigation. In this paper we describe frameworks from the field of spatial cognition, discuss important results in spatial cognition for folks who are BVI, and use these results and frameworks to lay out possible user interface paradigms for AR-based navigation technology for people who are BVI. We also present findings from a route dataset collected from an AR-based navigation application that support the urgency of considering spatial cognition when developing AR technology for people who are BVI.

Index Terms—assistive technology, indoor navigation, augmented reality

I. INTRODUCTION

For people who are blind or visually impaired (BVI), the ability to safely orient oneself in an environment and navigate to a location of interest, a set of skills collectively known as orientation and mobility (O&M), is of vital importance to one's quality of life. For example, individuals with better O&M skills have a higher employment rate [1]–[4]. Technologists have long sought to create devices to allow people who are BVI to augment their own O&M abilities [5]. In particular, smartphones, with their sophisticated sensors and relative ubiquity within the BVI community, have attracted considerable attention due to their great potential to support O&M tasks such as navigation in unfamiliar environments, exploration, and locating objects of interest. For example, sensors such as GPS and Inertial Measurement Units (IMUs) are suitable for low-precision outdoor navigation or indoor navigation over short distances.

In 2017, Apple released ARKit, the first widely available platform for augmented reality (AR) technology. Since the introduction of ARKit (and later Google's ARCore) there has been a surge of interest among assistive technology researchers in leveraging this technology to support people who are BVI. This interest is largely driven by the fact that AR technology enables accurate, low-latency 3D-tracking of the position and orientation of a user's phone. Importantly, these location estimates can be generated in environments where GPS is

unavailable (e.g., indoors) and can be more precise than GPS-based location estimates, which are accurate to about 5m under ideal conditions. Further, smartphone-based AR technology fuses sensory information from the phone's IMU and its camera, providing more accurate motion estimates than either in isolation. While the primary use case of AR technology is to render virtual content visually (e.g., on the phone's camera feed), these same location and orientation estimates also provide a rich and robust sensor for creating assistive technology to empower people who are BVI to more easily perform tasks that require understanding or navigating through spatial environments. For example, a number of research groups have created systems to allow BVI users to navigate using AR [6]–[8]. Although in this paper we are primarily concerned with the use of AR for navigation, AR has also been used as a means to make touchscreen interfaces on appliances accessible to people who are blind [9].

When utilizing AR technology to make information accessible for blind users, there are a number of potential interaction patterns. Hershovitz et al. [10] identified a number of common tasks and provided some UI paradigms for making these accessible to people who are blind. Here, we discuss the specific case of AR-based navigation as it magnifies spatial cognition challenges not fully considered in [10].

When considering the use of AR for navigation, one typically has some sort of digital route or map of the environment that can be parsed by the AR system. Once the AR position tracking of the phone is aligned to the map (a process known as *coordinate system registration*), location-based information (e.g., about a route the user would like to traverse) can be provided. For example, in [6] intermediate route waypoints and the final destination of the route can be mapped into the AR coordinate system. Once mapped, this information can be communicated to the user through non-visual modalities.

II. COORDINATE SYSTEMS FOR REPRESENTING SPATIAL INFORMATION

The spatial cognition literature provides a number of important results and frameworks for assistive technology researchers when considering the coordinate systems in which to represent spatial information to a user. In this section, we will discuss two different distinctions drawn in the literature and how they might manifest themselves in the design of an AR-based navigation app for people who are BVI.

A. Egocentric versus Allocentric Coordinate Systems

In the spatial cognition literature, a distinction is drawn between egocentric and allocentric coordinate systems (e.g., see [11] for a detailed discussion). In an *egocentric* coordinate system spatial information is encoded with respect to a coordinate system that places the user's body at its origin and defines coordinate axes based on the orientation of the user's body (e.g., one might speak of objects being "in front of" or "to the left of a user"). In an *allocentric* coordinate system spatial information is encoded relative to objects or features of the environment (e.g., one might say that the bathroom is located "down the main hallway through the third door on the right"). A common experimental finding is that for people who are blind (particularly for those who are congenitally blind), egocentric coordinate systems are more effective [11], [12], although with additional training allocentric coordinate systems can become easier to utilize. Further [11] highlights the importance of training to help the development of the ability to use allocentric coordinate systems.

B. Local Versus Global Spatial Representations

Another important framework from the field of spatial cognition is the that of local and global representations of space. Local spatial representations apply to only a portion of the overall environment (e.g., a particular segment of a navigation route). Multiple local spatial representations might be stitched together to enable navigation over long distances. For example, one commonly used method by which O&M instructors teach navigation of a route to people who are BVI is by using local representations. An O&M instructor might breakdown a long journey into small segments, providing important landmarks and spatial information to help the student easily navigate through each segment of the route and determine when they are ready to switch to the next segment. In this way, many local representations are combined in service of completing a single route. In contrast, a global representation of a space would involve learning a globally coherent map of the environment. In this representation the traveler would have a concept of not only how to navigate along each segment of the route but also how the various segments of the route relate to each other (e.g., when navigating in a multi-story building to a spot on the floor directly above, the traveler would know how the starting and ending location relate to each other in a straight line sense).

According to [11] it has been shown that people who are blind have little trouble with learning routes but that global structure of the environment (e.g., the previously presented example of determining the relative position of a starting and ending location) is harder to ascertain [13]. This difference manifests itself experimentally, for example, in people who are blind having inaccurate sense of global scale or being unable to accurately judge global spatial relationships [13]–[15] (although there are large individual differences, and some people who are blind are able to create very accurate global spatial representations). Due to the fact that local spatial representations are more easily utilized for most people who are BVI, O&M instructors use route learning through local




Fig. 1. A visualization of both the phone and egocentric coordinate systems. The user is depicted in a topdown perspective holding their phone in the left hand and their long cane in the right hand. Differences between the two coordinate systems can be observed in both the position of the origin and the orientation of the coordinate axes.

spatial maps as their primary mode of teaching navigation [5]. As a student progresses to a higher level of expertise in a particular environment, global spatial structure might be introduced into training [5].

III. CHALLENGES AND OPPORTUNITIES IN UTILIZING PHONE COORDINATES

Before discussing how the distinctions of egocentric versus allocentric and local versus global can inform AR-based navigation technology, it helps to discuss one additional component of AR technology: the smartphone. As mentioned earlier, AR systems track 3D motion of the user's phone (more specifically the phone's camera). As a result, when thinking about the design of AR technology, we should consider the coordinate system of the phone in addition to egocentric and allocentric coordinate systems (see Figure 1). For users with typical vision, the distinction between phone and egocentric coordinate systems is not of great importance as these users can easily map between the two systems. These users can perform the mapping using two primary strategies. First, a user with typical vision might hold their phone so that the camera is directly inline with their field of view (making the phone coordinate system and the egocentric coordinate system coincident and interchangeable). Second, a person might use the perceived location of objects in the camera feed and objects in their field of view as a means of computing the offset between phone and egocentric coordinate systems, and thereby allowing information to be mapped between the two.

For people with limited or no functional vision, it is not a straightforward process to reconcile phone and egocentric coordinate systems. Specifically, due to a lack of visual feedback it may be difficult for a person who is blind to either make the phone coincident with their egocentric coordinate system or to ascertain the relationship between these coordinate systems using external landmarks. As a result, it will be more difficult to map information presented in the phone coordinate system (the one easiest to reference in AR applications)

to the egocentric coordinate system. This difficulty presents AR researchers with several choices. First, one might simply ignore the problem and assume that a user with no or limited functional vision can resolve these two coordinate systems *well enough* to utilize AR-based guidance. Second, one might try to convert the phone coordinate system to the egocentric coordinate through some algorithm (see Section V). Third, it may be assumed that these two coordinate systems cannot be reconciled to a high degree of accuracy, but additional information (e.g., references to fixed landmarks or the users prior experience with navigating an environment) might make up for this difficulty.

In addition to the challenges of utilizing phone coordinates, there are also opportunities. For instance, the fact that a phone can be easily manipulated to “query” the local environment provides a powerful interaction paradigm for AR systems. For example, in the same way that a user might sweep their cane to understand the tactile character of the surfaces and objects around them, they might also sweep their phone to explore the virtual objects in their environment. For BVI users sweeping their phone in this way is potentially more efficient for querying spatial information than moving their body.

IV. MAPPING SPATIAL COGNITION CONCEPTS TO AR

Given the two frameworks of local versus global representations and egocentric versus allocentric coordinate systems, an important follow-on question for developers of AR-based navigation technology is: what would it look like to utilize these concepts in the design of an AR-based navigation system? In this section we provide some methods for realizing these concepts in AR and some of the challenges for users as well as for the underlying AR technology itself. For now we will not consider the phone coordinate system explicitly, but we will discuss how to incorporate it in Section V.

A. Communicating Spatial Information in an Egocentric Coordinate System

There are a number of methods for communicating spatial information in an egocentric, non-visual manner [16]–[18]. While a full discussion of this topic is outside the scope of this paper, if we consider the three most common non-visual modalities—speech, sounds, and haptic—we can broadly outline some possibilities. For instance, when using speech one might use directions such as “turn left” or “walk straight for 10 meters,” Alternatively, clock directions such as “turn towards your 2 o’clock” might be used (a technique common in O&M training [5]). When considering non-speech sounds, one might use beeps or tones for particular egocentric spatial relationships. For example, one might play a tone to indicate that the user is on the correct path, modulate the pitch of a sound to indicate the relative height of an object, or play different sounds to indicate whether something is to the left or to the right. Further, if the user has stereo headphones, 3D-spatial audio might be used to make sound appear as if it is emanating from a position of interest. Haptic feedback might be utilized in many of the same ways as non-speech sound

with the proviso that for most smartphones there is much less flexibility in the presentation of haptic feedback than non-speech sound (e.g., the only qualities of haptic feedback that might be controlled are frequency and intensity).

As discussed earlier, egocentric representations are typically the easiest for people who are BVI to utilize, and they are also the easiest to generate using AR technology. If we think of the phone coordinates as being coincident with the egocentric coordinate system, then obtaining the egocentric coordinates of mapped positions in the environment is trivial as this functionality is at the core of any commercial AR library. Additionally, expressing information about the local characteristics of an environment is similarly straightforward (e.g., providing notifications when a user is near a point of interest or providing them with the information about their next waypoint along a navigation route).

B. Communicating Spatial Information in an Allocentric Coordinate System

Allocentric coordinate systems are more difficult to represent with AR systems. The principal challenge is the need to have some set of external landmarks to reference when describing spatial information. For libraries like Apple’s ARKit, the information that is available most easily to the application programmer is the 3D path of travel of the phone and the 3D position of horizontal and vertical planar surfaces, which can be identified provided they have suitable visual texture and images have been captured of the surfaces from a range of positions. One choice then is to reference allocentric coordinates to these planar surfaces, however, they may be difficult to detect (e.g., if they lack visual texture, such as in the case of solid-colored walls) and they may not be semantically meaningful to the traveler (e.g., it is crucial to know whether a planar surface represents a door versus a wall when providing spatial information in relationship to this surface). A second potential choice would be to use landmarks that are either detected automatically (e.g., using object recognition) or that have been entered manually by someone mapping a route. This sort of guidance would be useful in the event that the traveler had a sense of where these landmarks were in advance or if they could ascertain whether or not they arrived at them (e.g., using their tactile or auditory sense). Examples of this sort of guidance might include:

- Telling the user to walk through an open area until they reach the far wall.
- Telling the user that there is a water fountain in the next alcove on the left.

C. Local Spatial Representations

When utilizing local representations, the system would provide the user with spatial information about their immediate surroundings or a proximal subgoal when navigating a route. Examples of this sort of guidance might include indicating the direction a user should turn once they reach a route waypoint (e.g., in order to face the next waypoint) or affordances of the local environment that are nearby (e.g., bathrooms or door

ways). Given that the AR system has mapped the environment the user is navigating through, spatial information such as this would be readily attainable for presentation to the user.

D. Global Spatial Representations

As AR systems maintain a globally consistent map of an environment, from a technological point of view it is no more difficult to obtain global spatial information than it is local spatial information. The sorts of global information that might be presented to a user in the context of an AR system might be the straight line distance from the user's current position to the starting or end of a route, nearby branching points or shortcuts and their relationship to the current route, and the relative location of various landmarks in the environment.

V. MAPPING PHONE COORDINATES TO EGOCENTRIC COORDINATES

As shown in Figure 1, the discrepancy between the phone and the egocentric coordinate system presents potential challenges to AR technology creators. In this section we outline some ways to approach these challenges.

A. Allow the User to Do the Mapping

The simplest solution to the challenge of reconciling these coordinate systems is to adopt the perspective that performing this mapping is solely the responsibility of the user. While this is expedient, anecdotally, it is important to note that in our experience developing the Clew app [6] we found that the ability to do this mapping was highly inconsistent across users. In Section VI, we provide some additional data that suggests that there is indeed large variability in users' abilities to perform this mapping effectively.

B. Automatic Registration of Egocentric and Phone Coordinate Systems Using AR

Another approach to the problem is to try to automatically provide feedback in the egocentric coordinate system by computing the transform between the phone coordinate and egocentric coordinate system. Unfortunately, without additional hardware (in addition to the smartphone) there is no straightforward way to do this precisely. There are, however, some ways in which approximate transformations can be computed. Here we outline some general approaches.

1) *Use Movement to Compute the Angle Offset:* If we are content to only calculate the angle offset between the phone and the egocentric coordinate frame and we make the assumption that this offset is constant over time (i.e., rotations in the phone coordinate system are caused by the user rotating their body and not by rotating the phone independently of the body), then we can use periods where the user is walking to try to estimate this offset. If we assume that periods of translation of the phone are caused by the user walking forward (straight ahead in the egocentric frame), we can compute the angle between this translation direction and the axis emanating from the phone coordinate system to determine the offset between the egocentric and phone coordinate systems. When

implementing this approach care should be taken to avoid using periods where the user is not walking forward in order to estimate this offset (e.g., when the user is stepping to the side or moving their phone back and forth while standing still). Another limitation is that this approach could not determine translational offset between the phone and egocentric coordinate systems. However, when navigating along a route, which involves reasoning about landmarks or virtual waypoints that are often meters away, these coordinate systems may be "close enough" in a translational sense even if their alignment in an orientation sense must be precise to be effective.

2) *Detecting Edges When Shore Lining:* Another approach one might take is to use periods of movement when the user is shore lining (meaning following an edge in the environment). If the movement of the phone can be correlated with an edge detected on the ground plane in the camera view, then the relative orientation of the phone and the egocentric coordinate system could be computed.

3) *Using a Sweeping Motion:* A third approach is to avoid using the orientation of the phone directly and instead use a vector emanating from the user's elbow to the phone's location in 3D. If the app instructs the user to scan the phone back and forth by keeping their elbow at their side and rotating their forearm in an arc parallel to the ground plane, then one could infer the location of the elbow from this arc (e.g., by fitting a parametric model to the trajectory of phone positions) and consequently determine the relationship of the egocentric coordinate system to the phone coordinate system.

C. Physical Alignment of Egocentric and Phone Coordinate Systems

A low-tech solution to the issue of coordinate system alignment is to provide the user with instructions for holding their phone that bring the two coordinate systems into closer alignment. For instance, users may be instructed to brace the phone against their bodies so that the phone is approximately at the midpoint of the ribcage and the camera's optical axis points perpendicular to the user's shoulders. While it is not likely that this positioning would be realized exactly, the two coordinate systems would be much more inline than when the user is holding the phone away from their body freely. One downside of this strategy beyond potential ergonomic issues for the user is that the phone in this configuration would be subject to greater vibration, which might degrade the quality of the AR pose tracking.

D. Avoid the Egocentric Frame

Finally, an AR system might eschew the egocentric frame altogether, instead only providing spatial information in terms of external landmarks or features of the scene (allocentric). In such a system, the distinction between phone and egocentric coordinate systems would be irrelevant.

VI. RESULTS FROM AN AR NAVIGATION APP

In [6] we presented the design and evaluation of our smartphone app, Clew, for AR-based navigation for people who are

TABLE I

THE EFFECT OF THE ESTIMATED MEDIAN ANGLE OFFSET BETWEEN THE EGOCENTRIC AND PHONE COORDINATE SYSTEM AND THE SUCCESS PROBABILITY. SUCCESS IS BASED ON WHETHER THE USER RATED THE ROUTE NAVIGATION AS THUMBS UP (SUCCESSFUL) OR THUMBS DOWN (UNSUCCESSFUL).

Median Offset	N Routes	Success Probability	95% CI
< 20°	2,548	0.843	[0.828, 0.857]
20° to 26°	2,547	0.787	[0.771, 0.803]
26° to 36°	2,547	0.731	[0.713, 0.748]
36° to 62°	2,547	0.633	[0.614, 0.652]
> 62°	2,547	0.543	[0.523, 0.562]

BVI (for a full discussion of the app, refer to the original paper). The app lays down a trail of virtual breadcrumbs (3D positions by the AR library), which can then either be used to immediately generate automatic, non-visual guidance back to the starting location or saved so that the user may receive automated guidance along that route in the future.

In [6] we presented an analysis of the effect of various features of the route (e.g., route length, number of motion tracking errors) on whether or not the app was able to give satisfactory guidance to the user. The definition of satisfactory here is based on whether a user precesses the thumbs up or thumbs down button following the completion of the navigation guidance. Importantly, this data may be biased with respect to the overall user base of that app as users can opt out of rating their routes by changing the app settings (we suspect that users who are more proficient and use the app more frequently are more likely to opt out of rating each route they complete). In total our dataset consists of 12,736 routes.

In the original analysis of this data we did not consider the alignment of the phone and the egocentric coordinate system. Here, we estimate the angle offset between the phone and egocentric coordinate system using the method described in Section V-B1 and evaluate its effect on the probability of the user indicating that the app's guidance was satisfactory at the completion of a route. Importantly, Clew does not attempt to reconcile the egocentric and phone coordinate systems, instead relying on the user to be able to successfully interpret location guidance in the phone coordinate system in terms of how they should walk along the route.

To estimate the angle offset between the two coordinate systems, we filtered logged route data to focus on segments where the user appeared to be walking (as defined by the phone moving along a straight path with a speed of at least 0.3 m/s). For any segment that matched this criterion we computed the offset between the phone's translation direction (assumed to be aligned with straight ahead of the user) and the phone's orientation. We computed the median value across the entirety of the route to summarize the degree to which the phone and the egocentric coordinate systems were misaligned.

We separated the routes based on which quintile their median angle offset value fell into (e.g., 0-20%, 20-40%) and computed the success probability for each quintile. The results are given in Table I (note that since the median angle

offset values are separated into quintiles, each row of the table corresponds to an equivalent number of total routes). As predicted, the probability of the user rating the guidance as successful falls off dramatically as the angle offset increases. The success probability falls from a high of 84.3% for the quintile with the smallest angle offset to 54.3% for the quintile with the largest angle offset. While the correlation of success rate and angle offset, of course, does not imply causation, this finding indicates that more work should be done to understand this effect.

VII. FUTURE WORK

While we have provided some initial work in drawing lessons from the spatial cognition literature, more work needs to be done. First, the difficulties that people who are BVI face in utilizing certain forms of spatial information (e.g., global or allocentric) highlights opportunities for the development of assistive technology the helps to *develop* such skills (rather than to replace traditional O&M skills with technology). Second, more work should be done to allow automatic registration of egocentric and phone coordinate systems. Especially with the emergence of LIDAR technology for depth sensing on the latest smartphones, there are many opportunities to improve upon the initial methods presented here. Finally, spatial cognition researchers should directly assess the abilities of people who are BVI to map between phone and egocentric coordinate systems as the quantification of the distribution of this capability is vital for guiding future research in the field of AR-based assistive navigation technology.

VIII. CONCLUSION

We have provided frameworks and results from the field of spatial cognition that are relevant for the creation of AR technology to aid people who are BVI with navigation. Additionally, we have provided a largescale quantitative picture of how failing to adequately considering these factors can result in degraded performance of AR technology in real-world conditions. While the field of AR for navigation is in its infancy, with the rapidly growing technological capabilities of mobile phones, and the consideration of the spatial cognition literature, we hope that the positive impact of such technology for people who are BVI will be large.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CISE-2007824 and by The Peabody Foundation.

REFERENCES

- [1] A. Crudden, L. W. McBroom, A. L. Skinner, and J. E. Moore, "Comprehensive examination of barriers to employment among persons who are blind or visually impaired." *Mississippi State: Rehabilitation Research and Training Center on Blindness and Low Vision, University of Mississippi*, 1998.
- [2] A. Crudden and L. W. McBroom, "Barriers to employment: A survey of employed persons who are visually impaired," *Journal of Visual Impairment and Blindness*, vol. 93, pp. 341–350, 1999.

- [3] R. Leonard, T. D'Allura, and A. Horowitz, "Factors associated with employment among persons who have a vision impairment: A follow-up of vocational placement referrals," *Journal of Vocational Rehabilitation*, vol. 12, no. 1, pp. 33–43, 1999.
- [4] B. O'Day, "Employment barriers for people with visual impairments," *Journal of Visual Impairment & Blindness*, vol. 93, no. 10, 1999.
- [5] W. R. Wiener, R. L. Welsh, and B. B. Blasch, *Foundations of orientation and mobility*. American Foundation for the Blind, 2010, vol. 1.
- [6] C. Yoon, R. Louie, J. Ryan, M. Vu, H. Bang, W. Derksen, and P. Ruvo, "Leveraging augmented reality to create apps for people with visual disabilities: A case study in indoor navigation," in *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, 2019, pp. 210–221.
- [7] Y. Liu, N. R. Stiles, and M. Meister, "Augmented reality powers a cognitive assistant for the blind," *ELife*, vol. 7, p. e37841, 2018.
- [8] G. Fusco, S. A. Cheraghi, L. Neat, and J. M. Coughlan, "An indoor navigation app using computer vision and sign recognition," in *International Conference on Computers Helping People with Special Needs*. Springer, 2020, pp. 485–494.
- [9] A. Guo, X. Chen, H. Qi, S. White, S. Ghosh, C. Asakawa, and J. P. Bigham, "Vizlens: A robust and interactive screen reader for interfaces in the real world," in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, 2016, pp. 651–664.
- [10] J. Herskovitz, J. Wu, S. White, A. Pavel, G. Reyes, A. Guo, and J. P. Bigham, "Making mobile augmented reality applications accessible," in *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, 2020, pp. 1–14.
- [11] N. A. Giudice, "Navigating without vision: Principles of blind spatial cognition," in *Handbook of behavioral and cognitive geography*. Edward Elgar Publishing, 2018.
- [12] S. Millar, *Understanding and representing space: Theory and evidence from studies with blind and sighted children*. Clarendon Press/Oxford University Press, 1994.
- [13] S. Ungar, *Cognitive Mapping without Visual Experience*. Psychology Press, 2000.
- [14] J. M. Loomis, R. L. Klatzky, R. G. Golledge, J. G. Cicinelli, J. W. Pellegrino, and P. A. Fry, "Nonvisual navigation by blind and sighted: assessment of path integration ability," *Journal of Experimental Psychology: General*, vol. 122, no. 1, p. 73, 1993.
- [15] M. Gori, G. Cappagli, G. Baud-Bovy, and S. Finocchietti, "Shape perception and navigation in blind adults," *Frontiers in psychology*, vol. 8, p. 10, 2017.
- [16] D. Ahmetovic, F. Avanzini, A. Baratè, C. Bernareggi, G. Galimberti, L. A. Ludovico, S. Mascetti, and G. Presti, "Sonification of pathways for people with visual impairments," in *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, 2018, pp. 379–381.
- [17] D. Dakopoulos and N. G. Bourbakis, "Wearable obstacle avoidance electronic travel aids for blind: a survey," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 40, no. 1, pp. 25–35, 2009.
- [18] G. Presti, D. Ahmetovic, M. Ducci, C. Bernareggi, L. Ludovico, A. Baratè, F. Avanzini, and S. Mascetti, "Watchout: Obstacle sonification for people with visual impairment or blindness," in *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, 2019, pp. 402–413.