Research Article | Vol. 31, No. 4 / 13 Feb 2023 / Optics Express 6827
Optics EXPRESS

Using augmented reality to cue obstacles for people with low vision

DYLAN R. FOX,[1,*] AHMAD AHMADZADA,[1] CLARA T. FRIEDMAN,[1] SHIRI AZENKOT,[2] MARLENA A. CHU,[1] ROBERTO MANDUCHI,[3] AND EMILY A. COOPER[1]

[1] *University of California, Berkeley, Berkeley, CA, USA*
[2] *Cornell Tech, Cornell University, New York City, NY, USA*
[3] *University of California, Santa Cruz, Santa Cruz, CA, USA*
[*] *Dylan.r.fox@berkeley.edu*

Abstract: Detecting and avoiding obstacles while navigating can pose a challenge for people with low vision, but augmented reality (AR) has the potential to assist by enhancing obstacle visibility. Perceptual and user experience research is needed to understand how to craft effective AR visuals for this purpose. We developed a prototype AR application capable of displaying multiple kinds of visual cues for obstacles on an optical see-through head-mounted display. We assessed the usability of these cues via a study in which participants with low vision navigated an obstacle course. The results suggest that 3D world-locked AR cues were superior to directional heads-up cues for most participants during this activity.

1. Introduction

Over one billion people around the world live with conditions that reduce the clarity of their vision [1]. While many of these people can be helped with simple refractive optical correction (e.g., glasses or contact lenses to correct near- and far-sightedness), many others live with low vision. Low vision comprises a heterogeneous set of visual impairments that cannot be corrected by refractive optics and that interfere with daily activities. Common conditions that cause low vision include macular degeneration, glaucoma, diabetic retinopathy, and retinitis pigmentosa.

For individuals with low vision, emerging augmented reality (AR) systems offer a promising approach to new forms of digital vision correction. These systems can overlay imagery from a display system onto the user's natural vision, and can therefore provide forms of visual enhancement that are not possible with optics alone. For example, AR systems have the potential to boost visibility by highlighting the edges of objects that are hard to see [2], to enhance key information like signs that are difficult to read [3], or to provide customized visual cues that aid in finding things in cluttered environments [4].

One of the notable challenges associated with low vision is safe navigation and obstacle avoidance [5–11]. Modern AR headsets have the capability to gather useful information for helping the user to locate obstacles, for example, via the visual and inertial sensors used to create 3D maps of the surrounding environment [12,13]. As such, several recent studies have proposed to use AR systems to assist people with low vision as they navigate [14–17]. But how obstacles and other spatial information can be most effectively communicated in AR to users with low vision is still an open question. Importantly, AR obstacle cues must be salient and visible to the user, but must not be distracting or interfere with their larger environmental awareness.

One source of underutilized knowledge to help address these challenges can be found in cue design for video games. Game designers have grappled for decades with the task of visually directing the player's attention to particular objects in the game world, and have used a variety of means to do so. One recent framework, for example, classifies visual cues by their purpose ("discover," "look," or "go") and their markedness ("subtle," "emphasized," "integrated," or

#479258
Journal © 2023
https://doi.org/10.1364/OE.479258
Received 31 Oct 2022; revised 13 Jan 2023; accepted 16 Jan 2023; published 9 Feb 2023

Corrected 3 June 2024

Optics EXPRESS

"overlaid"), and proposes several types of AR cues utilizing this framework for navigation guidance [18]. But this work still leaves many design iterations of cues that have not been tested, let alone with low vision users who face unique challenges in viewing and interpreting visual cues.

To address this gap, we compared the effectiveness of two different types of AR visual cues for supporting navigation with low vision: *world-locked* cues that "stick" to real-world obstacles to enhance their visibility in situ, and *heads-up* cues that point in the direction of obstacles outside the user's field of view. By studying how each type of cue varies in effectiveness and user experience, we hope to gain a better understanding of what aspects of AR visual cuing are most valuable and salient for different kinds of obstacles and situations. As AR systems become increasingly powerful and affordable, we hope this work will contribute to their use as a widespread assistive tool for low vision, for example, in the form of *smartglasses* that combine optical refractive correction and digital vision enhancement.

2. Related work

2.1. Low vision adversely affects mobility

Different eye-related conditions cause diverse forms of low vision with varying degrees of severity. Figure 1 shows simulations of the impact of three common causes of low vision. Despite different causes, people with low vision experience many similar challenges. Low vision, for example, can make it difficult for people to read, use computers and mobile devices, and recognize other people [19,20]. Importantly, low vision adversely affects people's ability to walk safely and independently, which is critical for maintaining a healthy and active life. Typically, reduced peripheral vision is the strongest predictor of mobility challenges, since obstacles and elevation changes often lie in our periphery [7,8]. However, reduced contrast sensitivity also makes it difficult to detect obstacles that do not "stand out" [6]. Failure to detect and avoid obstacles can be dangerous, leading to falls and injuries. As a result, people with low vision report being afraid of traveling on their own, restricting their mobility [21,22].

Fig. 1. Illustration of the effects of eye-related conditions on vision. From left to right: macular degeneration results in a loss of central vision, glaucoma results in a loss of peripheral vision and can affect central vision too, and diabetic retinopathy results in a loss or distortion of vision across the visual field. These are meant only to be illustrative examples and do not necessarily reflect the visual perception of people with these conditions.

2.2. AR holds potential for vision enhancement

Assistive technologies that aim to support daily activities like navigation can come in many different varieties, including vision enhancement systems and sensory substitution devices that convey information non-visually (e.g., using sound or haptics) [23,24]. Here, we focus on reviewing wearable vision enhancement platforms.

Early research on wearable electronic vision enhancement focused on designing video see-through AR systems with functionality such as digital image magnification and contrast boosting [25]. Such systems have a form factor similar to virtual reality (VR) devices, but rather than show a virtual world to the user they present a real-time video feed of the physical world in front of the user that has been enhanced in some way. While these systems are promising and can increase functional vision [26], they also supplant the user's natural vision. As a result, they can reduce peripheral vision and occlude the user's eyes during social interactions. Therefore, a range of research has investigated the potential for vision enhancement protocols that run on optical see-through AR systems in which digital imagery is superimposed directly onto the user's natural vision.

A varied set of AR-based vision enhancements have been proposed for optical see-through systems. Peli and colleagues, for example, developed a method that aims to boost visual contrast by superimposing lines over the edges in the visual environment or showing minified edges of a full scene overlaid in the user's limited field. [2,27,28]. The creators of this approach found that users were able to exploit the added visual field information during simple visual search tasks [27,29]. Other groups have investigated AR techniques that overlay distance-dependent brightness and color information over the user's natural view [15,17,30,31]. One such system, used in [15], was associated with improved object recognition, but some users found the system distracting or felt that it made their vision worse.

Both edge-based and distance-based AR enhancements show promising results in lab studies, but each approach has drawbacks when considering potential for real life use. Distance-based enhancements emphasize distance information at the expense of other detail, while edge enhancement techniques may enhance all visual details irrespective of their importance for the task at hand. Importantly, assistive devices tend to be abandoned if they do not adequately consider the needs and preferences of users in the context of real use [32]. Ongoing research in the domain of human-computer interaction is focusing on comparing and contrasting different approaches to AR vision enhancement for specific tasks [4,14,33]. Researchers are also beginning to experiment with minimally-intrusive interfaces which can reduce distraction and cognitive load [34]. As such, a key question becomes how to effectively use AR to cue and enhance information in a way that supports specific tasks, like mobility and obstacle avoidance.

2.3. Visual cue design can be informed by game design

Video game designers have often contended with the challenge of graphically cuing players to important information that is not obvious or salient in the gaming environment, without interfering with the experience of the player. This work focuses on users with typical vision, but can nonetheless provide a useful framework for considering different visual cuing techniques [18,35]. For example, in the game Doom, objects of interest are directly highlighted with a high-contrast coloration to make them "pop out" [36]. On the other hand, games like Food Fight VR [37] and Owlchemy Lab's "Vacation Simulator" [38] have adopted heads up-style indicators to point the user to content that is not visible because it is currently offscreen.

Researchers have begun to leverage cues like these from video games, cinema, and other 2D applications for 3D AR applications, exploring ways of cuing real-world objects using the kinds of visual indicators often found in these media. For example, visualizations that add a halo [39] or a wedge [40] indicator to objects of interest were originally proposed for supporting off-screen targeting on mobile phones (for example, showing off-screen points of interest on maps), then adapted for AR (and VR) technologies [41]. Shapes like arrows and 3D wedges embedded inside a 3D virtual environment have been found to be effective compared to 2D maps and radars, but cue effectiveness varies based on number of targets, environmental occlusion, and relative distance [42,43].

Optics EXPRESS

The visual language of these kinds of cues in AR is not yet fully developed [18], but they nonetheless provide a starting point for considering AR visual cues for users with low vision. Based on this cuing literature, we focus in this report on two visual cue techniques: world-locked cues that are embedded within the environment and directly highlight obstacles in 3D to improve their visibility, and 2D heads-up cues that hover in the user's vision and provide an always-available directional cue towards potential obstacles. Each of these techniques has strengths and drawbacks based on their unique attributes. Our study will examine the usefulness of these two techniques for supporting navigation for users with low vision.

3. Prototype design

3.1. Platform

We developed a prototype AR application to support navigating around obstacles using the HoloLens 2 headset and the Unity engine [44]. This prototype was designed to provide visual cues to the user about the location of obstacles in their environment. The HoloLens 2 is an optical see-through device: light reflected from real-world objects is passed directly to the user's eyes and augmented via digital imagery directed to the user's view via waveguides.

3.2. Visual cues

We developed two different types of visual cues to assist users, designed to be used alone or combined. Before finalizing the cue design, we conducted a formative study with a group of four participants with low vision to obtain qualitative feedback on an initial prototype. We used the feedback to improve the visual cues (and to refine the testing paradigm for the main study). Details about the formative study are provided in the Supplement 1.

The AR obstacle cues followed two design principles: world-locked and heads-up. Both cue types were designed to favor the closest obstacles, as these were considered most important as collision hazards for the purposes of this study. The world-locked cues provided a high-contrast 3D box highlighting each obstacle (Fig. 2(A),(B)). This cue type was designed to support obstacle detection when contrast sensitivity is impaired. World-locked cues appeared when the user was within 5 meters of an obstacle, and multiple world-locked cues could appear simultaneously. The heads-up cues, on the other hand, constituted a set of indicator bars directing the users towards the location of an obstacle (Fig. 2(C)). These cues were head-tracking, meaning they were designed to always appear in the same position relative to the user's head. They always appeared in the center of the user's vision and only pointed towards obstacles that were out of view. If the user looked directly at an obstacle, the heads-up cue for that obstacle would disappear. The heads-up cues were designed to support obstacle detection when limitations of peripheral vision made obstacles undetectable.

The heads-up cues underwent a redesign between the formative study and main study. In the formative study, they took the form of arrows that could rotate freely and point towards any number of obstacles at once (one arrow per obstacle). However, participants found it hard to determine which direction the arrows were pointing, and most had difficulty interpreting more than one arrow simultaneously. For the main study, the heads-up cues were adjusted to be rectangular bars with inset arrows that could appear in the up, right, down, or left positions. These discrete positions were aimed at making it easier for participants to identify cue direction. Two bars could appear simultaneously, for example, to indicate that an obstacle was up and to the right. They were also adjusted to point only towards the nearest obstacle in front of the user, at a maximum distance of 2.5 meters. The bar width grew slightly in inverse proportion to the distance between the headset and the obstacle, cluing the user to the distance to the obstacle.

Optics EXPRESS

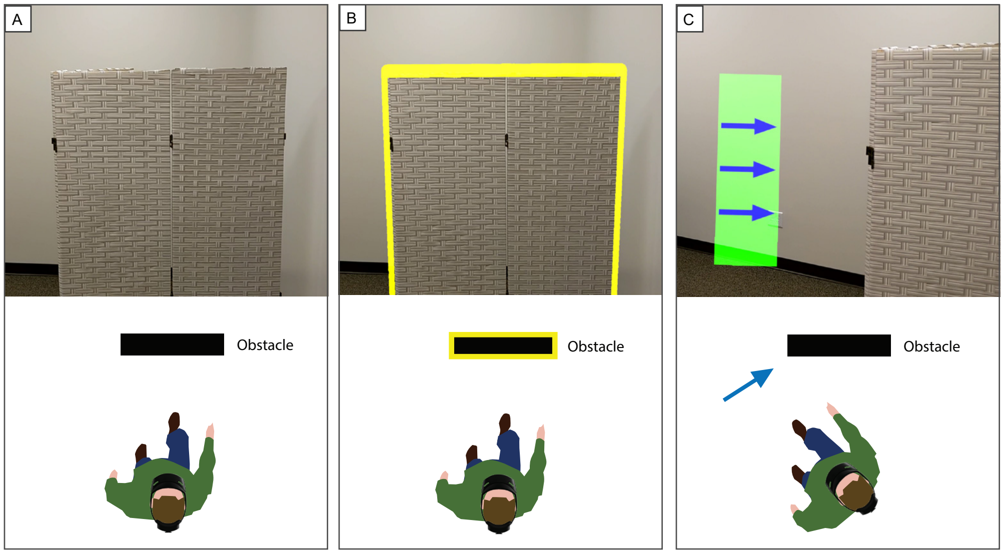

Fig. 2. Visual cue designs for supporting obstacle detection with AR. A) For demonstration purposes, a room divider is used as a wide obstacle. B) The world-locked cues rendered high-contrast 3D yellow boxes around obstacles. C) The heads-up cues were bars that could appear in the up, down, left, and/or right directions, and pointed to one out-of-view obstacle at a time. These images were obtained from the Hololens 2 using the Mixed Reality Capture functionality. Note that in the raw captured output, the headset filled in the faces of the world-locked cue bounding box with black. However, the faces of the box were not rendered for the user, so these were removed in post-processing for visualization.

3.3. Implementation

Our prototype is designed to test the user's experience with the visual cues, and is not a fully implemented obstacle detection and cuing application. That is, we assume that the obstacle locations are known. We had hoped to use the HoloLens 2's built-in spatial mapping capabilities to automatically identify obstacles on the fly, but the current system and implementation were not fast enough to reliably provide cues to users in real time. Thus, in the current implementation, a set of objects were created in Unity that corresponded to the shapes/layouts of physical obstacles in the environment. These objects were loaded by scanning a QR code placed at a known location in the environment, so that the system had ground truth information about the location and shape of each physical obstacle.

To create the world-locked cues, the prototype simply visualized these Unity objects, which are stored as rectangular prisms with yellow edges and transparent faces. Due to the HoloLens 2's limited field of view, the world-locked cues only appeared over the physical obstacles when they were within the display region (43° horizontally by 29° vertically in central vision).

To create the heads-up cues, the Unity objects were hidden with a camera mask, and directional cues were shown instead. Because these cues were only supposed to render when obstacles were outside of central vision, calculating whether to show each directional cue at any point in time required establishing whether an obstacle was somewhere in front of the participant, but not directly in their central vision (we assume that objects behind the observer at not potential obstacles). First, distances were calculated to the center of each obstacle from the camera (participant's head). Second, the horizontal and vertical angles to the edges were calculated based on the camera's right and up axes. The application then determined whether obstacles were “in front of” the participant using a 75° cone projected from the camera. This cone was projected on a horizontal axis in front of the user, so that if they pointed their head downward, the back of the cone would not extend behind their feet.

Once it was determined which obstacle (if any) was closest, within 2.5 m, and in front of the participant, that obstacle was declared the target. To determine which (if any) of the four directional cues should be rendered, the minimum and maximum angles to the obstacle were compared against a threshold angle relative to the center of the user’s head orientation. The threshold determined whether the obstacle was far enough from central vision to trigger a cue. The thresholds for rendering the cues in the down, left and right directions were 9°. The threshold for the up direction was 6°. The lower threshold in the up direction was consistent with the reduced size of the upper visual field in human vision [45]. If the threshold was met or exceeded, the arrow(s) were rendered, and their size multiplied based on the participant’s proximity to the obstacle, up to a maximum of four times the original width of 3.4° for the horizontal and 2.3° degrees for the vertical heads-up cue. Thus, if the user looked directly at an obstacle (a relative angle of 0°), the cue would not appear; if they turned their head so the obstacle was further than 9° to the left or right, the cue would appear pointing to it; and if they continued to turn until the obstacle was more than 75° to the left or right (and thus no longer “in front of” the user), the cue would disappear. By default, the heads-up cues were placed at the edges of the HoloLens 2 screen, but they could be adjusted to the left or right to accommodate participants with vision in only one eye or one side of the visual field.

While the application was in use, data including the position, head rotation, and system status were logged to a csv file every 0.02 seconds via Unity’s Fixed Update system. However, due to a logging error the user position and rotation were not updated at every log entry. This information was updated on average every 0.047 seconds. To the best of our knowledge, this error did not affect the accuracy of results or impact the participants’ experience.

4. Study participants

We recruited 20 adults with low vision to participate in a study aimed at assessing and comparing the efficacy of the AR visual cues and getting feedback on their experiences. Participants were recruited who had moderate visual impairment or worse in both eyes (<20/70 Snellan acuity) and/or a visual field restriction (20 deg or less in diameter), but who retained some functional vision. Participants were also required to be able to repeatedly walk a low-intensity obstacle course (excluding, for example, wheelchair and walker users). Participants were recruited primarily from the Meredith Morgan Eye Center at University of California, Berkeley, though patients from other clinics were also accepted if they learned about the study through word of mouth.

The demographics and levels of vision for the study participants are summarized in Table 1. The average age was 46 years (range of 20-87 years). Six participants identified as female and 14 identified as male. Levels of vision were obtained from each participant’s most recent eye exam. For recording visual acuity, we converted to units of logMAR (logarithm of the minimum angle of resolution in units of arcminutes), in which 0 logMAR corresponds to normal visual acuity (resolving an angle of 1 arcminute) and higher numbers correspond to worse acuity. Contrast detection thresholds are indicated as the minimal amount of contrast required to make a target visible, in units of Weber contrast (the luminance of a letter target divided by the luminance of the background, multiplied by 100). Weber contrast values can be range from normal (less than or equal to 2%) to profound loss (greater than or equal to 20%). Peripheral visual field limitations are summarized based on clinical notes, and white cane usage was self-reported by the participants. Informed consent was obtained from all subjects, and they were compensated for their time.

Optics EXPRESS

Table 1. Participant information.[a]

ID	Age (yrs)	Etiology	Acuity (better eye), logMAR	Contrast Threshold	Peripheral Visual Field Limitations	White Cane Used
1	44	Homonymous hemianopsia	0.01 (R)	1.6%	Left side	No
2	29	Optic nerve hypoplasia	2.00 (L)	5%	CFC	Yes
3	34	Achromatopsia	0.90 (R)	4%	No	No
4	61	Retinitis pigmentosa	0.10 (L)	2%	CFC + peripheral islands	Yes
5	35	Oculocutaneous albinism	0.60 (R)	2%	No	No
6	62	Retinitis pigmentosa	CF	-	Peripheral island	Yes
7	32	Retinopathy of prematurity, glaucoma	1.60 (R)	40%	CFC	Yes
8	27	Retinitis pigmentosa	1.60 (R)	63%	CFC + peripheral islands	Yes
9	28	Optic nerve atrophy	1.38 (L)	5%	CFC	No
10	20	Retinitis pigmentosa	CF	-	CFC	Yes
11	87	Age-related macular degeneration	0.90 (L)	12%	No	No
12	59	Leber's hereditary optic neuropathy	1.38 (L/R)	16%	No	No
13	34	Leber's hereditary optic neuropathy	1.30 (R)	-	No	Yes
14	32	Congenital, cause unknown	1.40 (L)	6.3%	Restricted superior field	No
15	66	Stargardt disease	1.08 (L)	19.1%	No	No
16	40	Gyrate atrophy	0.70 (L)	5.8%	CFC	No
17	27	Retinal dystrophy	1.80 (L)	63%	CFC	Yes
18	55	Retinal dystrophy	0.54 (R)	-	CFC	Yes
19	73	Choroidal hemorrhage, glaucoma	0.30 (L)	-	CFC	No
20	83	Myopic degeneration	0.30 (L)	10%	CFC	No

[a] Acuity is indicated as the log of the minimum angle of resolution (logMAR) in the better seeing eye. Contrast threshold is indicated in units of Weber contrast (the luminance of a letter target divided by the luminance of the background, multiplied by 100) and reflects both eyes. A lower acuity value and contrast threshold indicate better vision. Peripheral visual field limitations are summarized based on clinical notes. M/F: Male/Female; CF: counting fingers; CFC: Concentric Field Constriction; L/R: Left/right eye. - indicates that data were unavailable.

5. Study procedure

The study was designed to obtain both qualitative and quantitative assessments of mobility and obstacle avoidance with the system. To this end, we designed an indoor obstacle course and asked participants to navigate the course. Participants provided feedback on their experience in the form of Likert ratings and we quantified mobility performance with a range of measures. Mobility performance with low vision is often gauged with summary measures such as the overall walking speed and counting of "errors" [5,6,46]. However, these conventional measures provide limited information about how people navigate a space, and the precise definition of metrics such as errors can be challenging [47]. As such, we included these measures but also collected continuous 3D position tracking and head rotation data (see below).

Optics EXPRESS

5.1. Obstacle course

A 15-meter-long obstacle course was constructed down the length of a 1.8-meter-wide hallway. This course consisted of three kinds of real, physical obstacles that were augmented with visual cues (see Fig. 3): low (a 2" x 2" dark gray bar of foam spanning the hallway that participants stepped over), high (a ~2" wide white streamer hung horizontally at eye level that participants ducked under) and wide (a 2 m tall x 1.5 m wide folding room divider that participants walked around). These obstacles were chosen based on the formative study and aimed to simulate real-world hazards such as cracks, branches, and door frames that could prove challenging to navigate for individuals with low vision, while minimizing the likelihood of any physical harm or tripping hazard during the experiment. The placement of the high obstacles was adjusted for each participant to match their specific eye level. (Individual eye level was also entered into the AR application to ensure proper alignment of the cues).

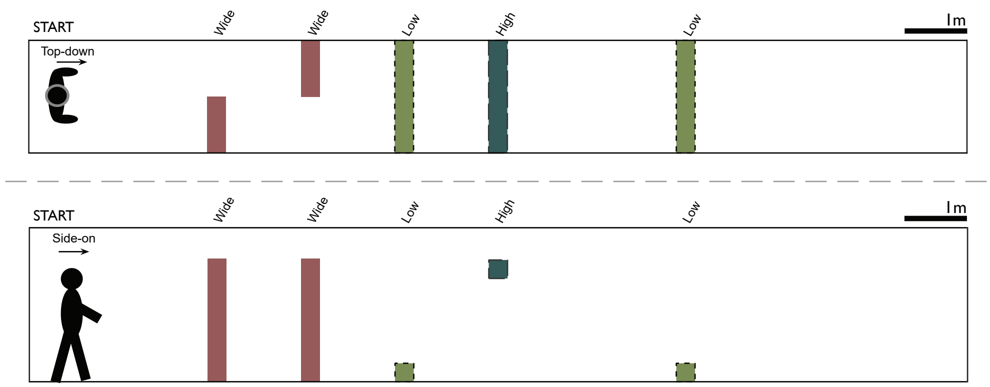

Fig. 3. The 15-meter-long obstacle course was constructed in a hallway with varied illumination and three types of obstacles: wide, high, and low. Pictured here is a sample layout top-down (top) and side-on (bottom). Horizontal dimensions of the obstacles are not drawn to scale.

Obstacles were placed at intervals of 1.5, 3, or 4.5 meters in one of 8 possible layouts. All layouts featured the two wide obstacles first in order to block the participant's view of the rest of the hallway, followed by one or two each of the low and high obstacles in the rest of the hallway. Each layout was used no more than once per participant. The hallway was lit by rectangular fluorescent lights at even intervals. One of the lights at approximately 5 m was off, such that the course was darker towards the beginning and brighter towards the end. Light levels varied from a low of 21 lux to a high of 170 lux.

5.2. Experiment flow

Before asking participants to navigate the obstacle course, we measured their baseline preferred walking speed. This is standard in orientation and mobility experiments, and is used to normalize subsequent walking speed measures as the percentage of preferred walking speed (PPWS). Preferred walking speed was measured walking down a hallway with no obstacles in it, and we took the average of six trials up and down the hallway. White cane users were encouraged to use their cane.

Next, participants were equipped with the HoloLens 2 and the HoloLens' eye calibration procedure was run. Calibration failed for 6 of the 20 main study participants, typically those with central vision loss. However, eye tracking was not used in this experiment so these failures did not affect the experimental procedure. The participant was allowed to adjust the display brightness to a desirable level. They could also shift the heads-up rectangle left or right if needed, condensing the four cues into one side of the HoloLens 2 display.

Participants were then given a full description of each kind of visual cue and encouraged to practice with each cue until they felt comfortable, with a minimum of two walks up and down an obstacle course per cue type. The obstacle course layout used for each participant's practice was not used in subsequent trials. Once the participant reported feeling comfortable with each cue, they were asked to do one final practice run on the same layout without wearing the headset in order to avoid putting the control condition at a disadvantage, and then we commenced the experiment.

The experiment consisted of a set of six trials in which participants walked up and down the obstacle course. Between trials, the course was changed to one of the predetermined layouts, with the layout ordering randomized. On each trial, the participant was timed as they walked forward through the course, stopped, then timed again as they walked back to the start. Walking time was measured via stopwatch for the control condition and via HoloLens 2 logging for the experimental conditions (by subtracting the logged time between the end and beginning of the trial). During each trial, mobility errors were categorized and counted by an experimenter. Errors were defined and recorded as done previously in Leat and Lovie-Kitchin [47]. Classes of errors were: hesitations (defined as when the participant stalls before an obstacle for reasons other than mobility restrictions), object contacts (any contact with obstacles, excluding the use of white canes to detect obstacles), stumbles, high stepping (anticipating a step that is not present), corrections, experimenter interventions (for example, calling the participant's name when in danger of impact), and behavior modifications (such as trailing a hand along a wall). The experimenter who was recording the errors was not blinded to the condition, because it was necessary for them to set the condition on the device. In practice, no experimenter interventions were required. At the end of each trial, the participant was asked a series of questions (with a 7-point Likert scale) about their experience, then given a chance to comment freely on that condition.

The first and last of the six trials were control trials (wearing no headset). The aim of these trials was to assess participants' baseline performance and experience without the headset (first trial), and to measure any possible learning effects (last trial). In between these control trials, we ran four experimental condition trials in a randomized order. These conditions were: no cues (wearing the headset but with no cues displayed), world-locked cues, heads-up cues, and combined cues (world-locked and heads-up together). After all trials were completed, participants were asked to select their favorite cue condition from world-locked, heads-up, or combined.

5.3. Self-reported experience

Participants were asked to rate their agreement with a series of statements on a scale from 1-7, with 1 being "Strongly Disagree," 4 being "Neither Agree Nor Disagree," and 7 being "Strongly Agree." For the control conditions, these statements were:

1. I felt confident in my ability to navigate this hallway.
2. It was easy to figure out where each obstacle was.
3. It was easy to tell how big each obstacle was.
4. I felt aware of my surroundings while navigating the hallway.

For the experimental conditions, participants were asked to compare their experience to the first control condition. Thus, the statements were:

1. I felt more confident in my ability to navigate this hallway compared to when I was not wearing the HoloLens.

2. It was easier to figure out where each obstacle was compared to when I was not wearing the HoloLens.
3. It was easier to tell how big each obstacle was compared to when I was not wearing the HoloLens.
4. I felt more aware of my surroundings while navigating the hallway compared to when I was not wearing the HoloLens.

For the experimental conditions, the 1-7 scale was shifted to a -3 to 3 scale in data processing to represent the comparative nature of the results, with 0 representing no difference from the control, negative numbers indicating worse than the control, and positive numbers indicating better than the control. For simplicity we refer to each of these ratings as Confidence, Obstacle Location, Obstacle Size, and Awareness.

5.4. Exploratory mobility measures

In addition to PPWS and errors, additional mobility measurements were calculated using the motion tracking data generated by the HoloLens 2 when it was worn. These measurements were used for exploratory analyses of mobility differences across conditions. We used position tracking to measure how far each participant actually traveled, with higher numbers indicating less direct paths through the hallway. We also measured the participant's average speed, calculated by dividing the actual total distance traveled by the time to complete the trial. This differs from the conventional PPWS measure in that it takes total distance traveled into account. To better understand differences in walking speed, we also quantified what percentage of the time participants were slowed (< 0.3 m/s) and stopped (< 0.1 m/s). Using the head rotation tracking, we measured how much each participant rotated their head on average in degrees per second, calculated as total degrees of pitch and yaw divided by their time to complete the course (e.g. turning head right 90° and then back to center would be considered 180° of yaw). Participants 7 and 11 were excluded from these analyses because the HoloLens 2 failed to track their motion for the entire course. All mobility measures, including PPWS and errors, were averaged between the forward and backward portion for each trial.

5.5. Statistical analyses

All statistical analyses were conducted using the R software package. We used Wilcoxon Signed-Rank tests to examine differences in the Likert ratings between the two control trials (first and last), and to compare the Likert ratings associated with each experimental condition against a score of 0. Differences between the Likert ratings associated with different experimental conditions were evaluated with omnibus Friedman tests (to allow comparison of more than two conditions) and paired follow up comparisons also used Wilcoxon Signed-Rank tests, with Bonferroni-corrected p-value thresholds. For the Wilcoxon tests, we also report the test statistic, V, and effect size, r [48].

We used one-way repeated measures ANOVAs to examine differences between the mobility measures associated with each condition. Prior to conducting each ANOVA, participants who had outlier measurements were identified using the IQR method and removed from analysis. This resulted in the removal of three or fewer participants from all analyses, except for the analysis of stop times for which five outliers were removed.

Mauchly's test for Sphericity was applied and in cases when the null hypothesis of sphericity was rejected, we report p-values corrected using the Greenhouse-Geisser method as p_{GG}. For the ANOVAs, we report effect sizes via generalized eta squared, denoted as η^2. Paired follow up comparisons used t-tests with Bonferroni-corrected p-value thresholds. We also report effect sizes in terms of Cohen's d. A p-threshold of 0.05 was used for all tests (with correction for multiple comparisons applied where indicated).

Optics EXPRESS

6. Results

6.1. Subjective reports

First, we consider the overall navigation task difficulty without any AR visual cues. Participants indicated a broad range of difficulties with the task during the first control trial. Responses to the four prompts ranged from 1 to 7, with medians ranging from 5 to 6 (Fig. 4). When participants repeated the control trial at the end of the session, the responses were overall similar, with medians also ranging from 5 to 6. There was no significant change in the responses between the first and last control trials (Confidence: $V = 3$, $p = 0.066$, $r = 0.41$; Obstacle Location: $V = 19$, $p = 0.112$, $r = 0.35$; Obstacle Size: $V = 31$, $p = 0.892$, $r = 0.03$; Awareness: $V = 16.5$, $p = 0.143$, $r = 0.33$).

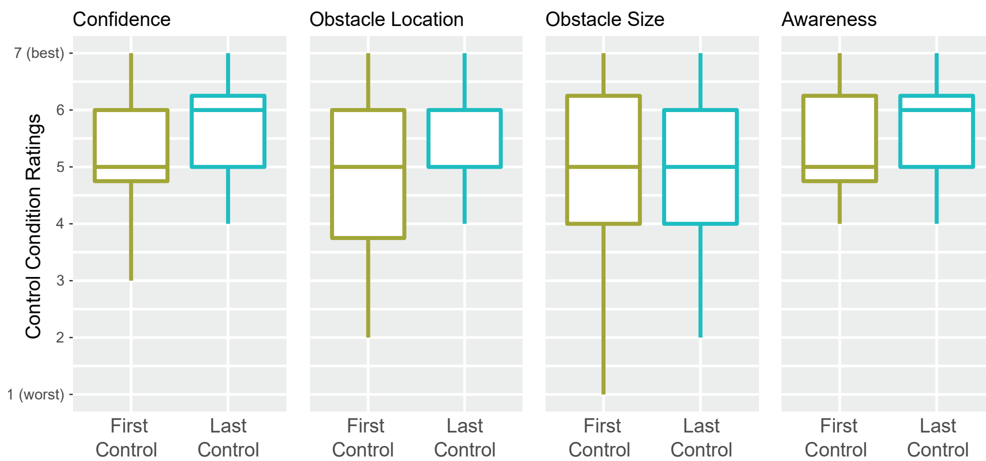

Fig. 4. Control trial Likert ratings for individual questions on a scale from 1 (worst) to 7 (best). Each panel shows the median, interquartile range, and non-outlier upper and lower ranges with a box-and-whisker plot for the first and last control trials. From left to right, panels show results for prompts about confidence, obstacle location, obstacle size, and awareness.

Figure 5 shows the Likert responses comparing each condition to the first control trial, including the no cues condition in which participants wore the headset but saw no cues. Recall that these responses are scaled from -3 to 3, with zero indicating that the condition did not differ from the control. Participants generally rated the no cues condition as worse than the control condition. The responses for this condition were significantly less than zero for all prompts (Table 2), suggesting that wearing the HoloLens 2 hardware alone had a detrimental effect on people's experience navigating. The reasons for these lower ratings may be varied. For example, they may reflect the physical discomfort associated with wearing the device, or possibly the light attenuation associated with the HoloLens 2's visor.

Despite this challenge, participants rated some experimental cue conditions as better than the control (Fig. 5, Table 2). Responses for the world-locked condition were significantly higher than zero for every question (medians ranging from 1 to 2.5). However, responses for the heads-up condition did not differ significantly from zero (medians ranging from 0 to 1). Lastly, responses for the combined cues condition were significantly higher than zero for every question except awareness (medians ranging from 0 to 2).

Next, we asked whether there were significant differences between the ratings associated with the three experimental conditions (Table 3). For all prompts, there was a significant effect of

Optics EXPRESS

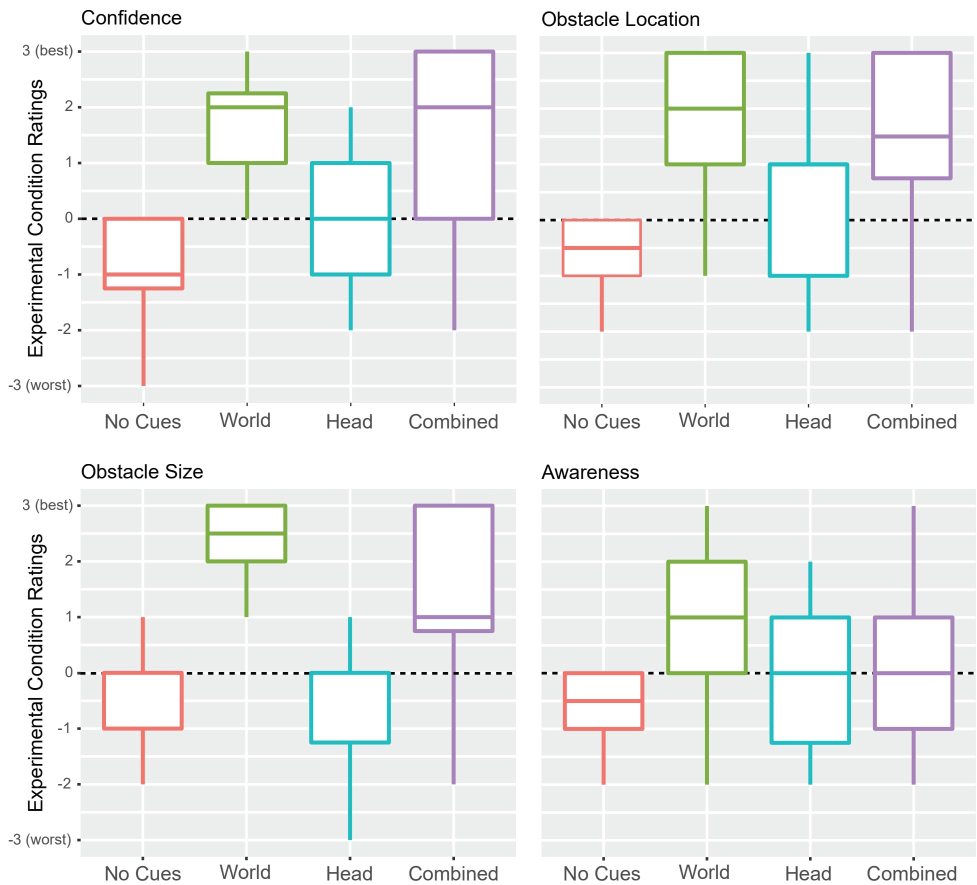

Fig. 5. Likert Ratings for individual prompts, on a scale of -3 to 3, with 0 representing the participant's experience without the headset (first control trial). Each panel shows the median, interquartile range, and non-outlier upper and lower ranges with a box-and-whisker plot for the no cues, world-locked, heads-up and combined cue trials.

cue condition, suggesting that the different cues resulted in consistently different ratings. For the confidence, obstacle location and obstacle size prompts, the pattern of results was largely the same: we found that the world-locked and combined conditions were rated significantly higher than the heads-up condition in all of these categories. The pattern of results for the awareness prompt was different. While the world-locked condition was still rated significantly higher than the heads-up condition, the combined condition was not significantly different from heads-up. We hypothesize that the added heads-up cues in the combined condition detracted from the participant's awareness of their surroundings, which pulled down people's ratings of this condition to be more in-line with their rating for the heads-up cues alone. Consistent with these patterns, when we asked participants at the end of the experiment which cue condition they preferred, the majority of participants (60%) preferred the world-locked cues. By comparison, 25% preferred combined, and 15% preferred heads-up.

6.2. Basic mobility results

Overall, mobility errors were rare and averaged less than one per participant per trial. An ANOVA, including the average errors in the control conditions and all experimental conditions, did not

Optics EXPRESS

Table 2. Comparing Likert ratings for each experimental condition against 0 (first control condition).[a]

Likert Question	Test	Median Rating	V	p value	r
Confidence	No cues	-1.0	0.0	0.002*	0.49
	World-locked	2.0	183.0	<<0.001*	0.57
	Heads-up	0.0	66.5	0.723	0.06
	Combined	2.0	133.0	0.007*	0.43
Obstacle Location	No cues	-0.5	0.0	0.004*	0.46
	World-locked	2.0	164.0	<<0.001*	0.54
	Heads-up	1.0	101.0	0.234	0.19
	Combined	1.5	163.0	0.006*	0.44
Obstacle Size	No cues	0.0	4.0	0.015*	0.39
	World-locked	2.5	184.0	<<0.001*	0.58
	Heads-up	0.0	9.5	0.069	0.29
	Combined	1.0	139.5	0.003*	0.48
Awareness	No cues	-0.5	0.0	0.004*	0.45
	World-locked	1.0	104.0	0.012*	0.40
	Heads-up	0.0	46.0	0.696	0.06
	Combined	0.0	94.5	0.703	0.06

[a]Significance was assessed using single sample Wilcoxon Signed-Rank Tests (using the V test statistic) with p value threshold of 0.05 (no correction for multiple comparisons). Significant p values have an asterisk.

Table 3. Comparing Likert ratings between the three experimental conditions.[a]

Likert Question	Test	Test Statistic (χ^2 /V)	p value	r
Confidence	Omnibus	31.8	<<0.001*	
	World-locked vs heads-up	136.0	<<0.001*	0.56
	World-locked vs combined	55.5	0.564	0.20
	Heads-up vs combined	14.5	0.003*	0.44
Obstacle Location	Omnibus	26.8	<<0.001*	
	World-locked vs heads-up	93.5	0.001*	0.41
	World-locked vs combined	46.5	0.564	0.09
	Heads-up vs combined	3.0	0.004*	0.45
Obstacle Size	Omnibus	37.0	<<0.001*	
	World-locked vs heads-up	188.5	<<0.001*	0.6
	World-locked vs combined	42.0	0.02	0.36
	Heads-up vs combined	3.0	<<0.001*	0.56
Awareness	Omnibus	14.5	0.0023*	
	World-locked vs heads-up	112.0	0.005*	0.47
	World-locked vs combined	74.5	0.166	0.32
	Heads-up vs combined	16.5	0.132	0.25

[a]Omnibus comparisons were conducted using Friedman tests and pairwise follow ups used Wilcoxon Signed-Rank tests. The test statistic in each row is either χ^2 (Friedman/Omnibus) or V (Wilcox). Follow up tests were Bonferroni corrected with a p value threshold of $p < 0.017$. Significant p values have an asterisk.

indicate any significant differences between conditions for number of errors ($F(4, 56) = 2.49$, $\eta^2 = 0.04$, $p_{GG} = 0.072$). Anecdotally, we observed that some behaviors categorized as errors based on the approach used in prior work [47] resulted from intentional strategies adopted by the participants to safely navigate, such as putting their hands in front of their face to detect low-hanging obstacles.

We also compared the percentage of preferred walking speed (PPWS) across conditions, a measure of how quickly the participant moves through the obstacle course compared to their default walking speed. In the control trials (first and last trials), participants tended to walk slower than their preferred walking speed (average PPWS was 61.4% for the first trial and 67.0% for the last trials). This slow down likely reflects the added difficulty associated with navigating around obstacles. We averaged these together to obtain a single control measure that incorporates the expected increase in walking speed over time during the experiment (Fig. 6, far left, red). Descriptively, all other conditions were associated with slower walking speeds compared to this control (average PPWS no cues = 61.2%, world-locked = 60.3%, heads-up = 51.7%, combined = 55.6%). An ANOVA indicated a significant difference between conditions ($F(4, 64) = 12.27$, $\eta^2 = 0.12$, $p_{GG} << 0.001$). Follow up tests indicated that all conditions except for the world-locked were associated with a significantly slower PPWS than the control (Table 4). The heads-up condition was also significantly slower than all other conditions except the combined condition. Thus, the results suggest that the world-locked cues were both rated more highly and were less disruptive to walking speed during navigation, whereas the heads-up cues were not rated well and were disruptive. The combined cues tended to fall in between. Next, we used the Hololens 2 tracking data to explore how the cues affected people's movement, with a focus on understanding the negative effects of the heads-up cues on navigation.

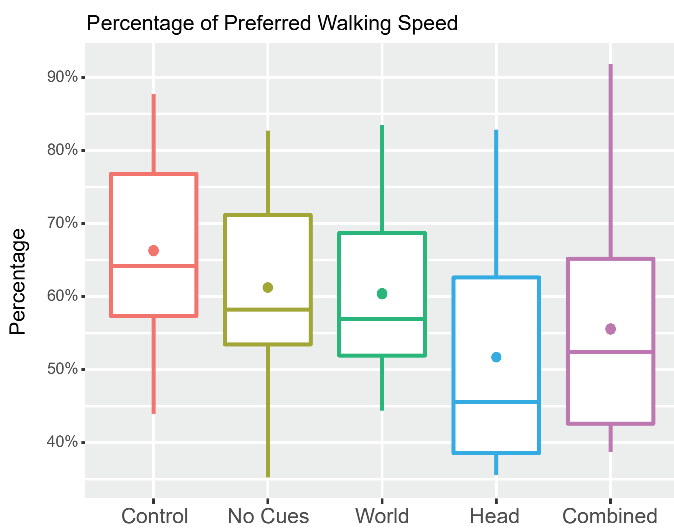

Fig. 6. Percentage of preferred walking speed (PPWS) by condition. Box and whisker plots represent the median, interquartile range, and non-outlier upper and lower ranges. Dots represent means.

6.3. Motion tracking

The HoloLens 2's spatial localization abilities allowed us greater insight into how the participants' behavior was affected by each kind of cue. We conducted a range of exploratory analyses on these data, and here we highlight the results of the analyses that were most revealing. Note that

Table 4. Pairwise comparisons between conditions for PPWS.[a]

Test	t statistic	p value	Cohen's d
Control vs no cues	3.4	0.004*	0.38
Control vs world-locked	2.3	0.035	0.47
Control vs heads-up	5.4	<<0.001*	0.98
Control vs combined	4.2	0.001*	0.73
No cues vs world-locked	0.4	0.72	0.07
No cues vs heads-up	4.5	<<0.001*	0.65
No cues vs combined	2.6	0.021	0.40
World-locked vs heads-up	3.4	0.003*	0.62
World-locked vs combined	2.0	0.059	0.35
Heads-up vs combined	-3.2	0.006	0.24

[a]Each row indicates the results of a paired t-test, with degrees of freedom = 16. The Bonferroni-corrected p-value threshold was set at 0.005. Comparisons associated with significant p values have asterisks.

as the participants did not wear the headset in the control condition, this data is available only for the experimental conditions. Figure 7 shows example walking trajectories for one participant when using the world-locked cues. Pink circles represent the measured location at regular time intervals (every ~500 ms), such that closely spaced circles indicate slower speeds and loosely spaced circles indicate faster speeds. In the top-down view, the participant can be seen weaving around the two wide obstacles and slowing down in between them (A). In the side view, the participant can be seen ducking below the high obstacle (B) and backing up to step over the second low obstacle (C).

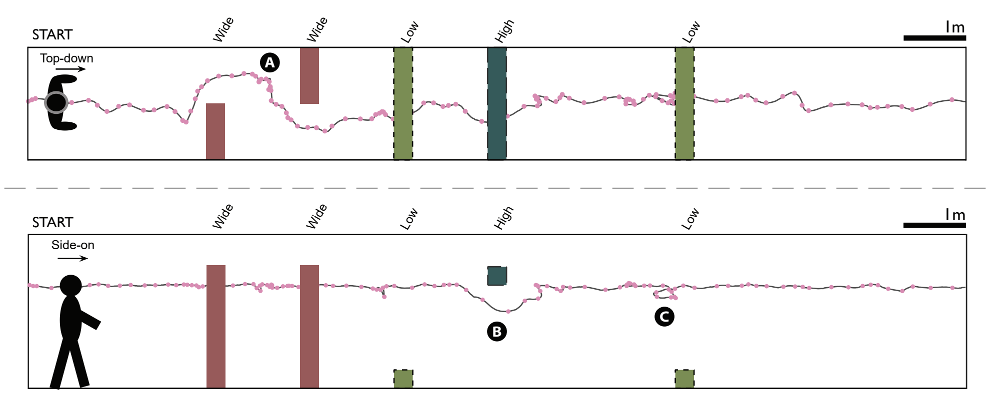

Fig. 7. Top-down (above) and side-on (below) views of participant 10's location during their forward world-locked cue trial. Colored blocks represent obstacles, and pink dots are placed along the path at time intervals of about 500 ms. The width of the blocks is not representative of the actual obstacle width. Clusters of dots indicated where the participant stopped or moved backward after colliding with an obstacle.

The PPWS measure indicated that participants walked more slowly when using the heads-up cues. However, because this measure is averaged over the entire course, it is possible that people actually walked the same speed but just took a more indirect path. So one question we sought to answer was how direct a path each participant took through the obstacle course, by examining the total distance traveled. An ANOVA indicated a significant difference in distance traveled per condition ($F(3, 45) = 4.06$, $\eta^2 = 0.04$, $p = 0.012$); however, all mean distances were between 15.6 and 15.8 meters (no cues = 15.6 m, world-locked = 15.6 m, heads-up = 15.8 m, combined = 15.8 m). None of the follow up tests indicated a significant difference between any two conditions (Table 5). Preliminarily, this suggests that the participants took a fairly direct path to the goal regardless of condition, and thus the slower PPWS in the heads-up condition is due primarily to walking more slowly and not due to zig-zagging across the course.

Table 5. Pairwise comparisons between conditions for total distance traveled.[a]

Test	t statistic	p value	Cohen's d
No cues vs world-locked	-0.4	0.671	0.08
No cues vs heads-up	-2.9	0.011	0.42
No cues vs combined	-2.5	0.026	0.42
World-locked vs heads-up	-2.0	0.066	0.39
World-locked vs combined	-2.4	0.032	0.39
Heads-up vs combined	0.0	0.988	0.00

[a]Each row indicates the results of a paired t-test, with degrees of freedom = 15. The Bonferroni-corrected p-value threshold was set at 0.008.

Additional insight can be gained by closely examining how the speeds of participants changed over the course of their trials. While the change in percentage of time spent stopped between conditions was not significant (Fig. 8 left; $F(3, 36) = 0.364$, $\eta^2 = 0.02$, $p = 0.779$), the change in percentage of time spent slowed (i.e., moving slower than 0.3 m/s) was significantly different between conditions (Fig. 8 right; $F(3, 42) = 5.08$, $\eta^2 = 0.15$, $p = 0.004$). Though none of the follow up tests indicated a significant pairwise difference (Table 6), qualitatively participants seemed to spend more time moving slowly while using the heads-up cues compared to the other conditions. We hypothesized that people may have tended to slow down when a heads-up cue appeared and take some time to look around more for the obstacle being indicated.

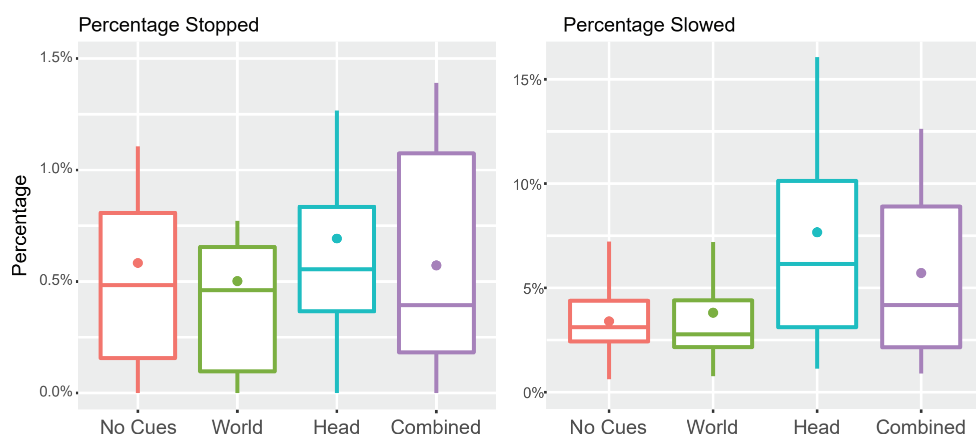

Fig. 8. Percentage of time spent stopped (< 0.1 m/s) or slowed (<0.3 m/s) by condition. Each panel shows the median, interquartile range, and non-outlier upper and lower ranges with a box-and-whisker plot. Dots represent means. Note that the y-axis ranges in the two plots differ.

Lastly, we asked whether the head rotation data contained any indication that people were spending more time looking around in the heads-up condition. We conducted ANOVAs to

Table 6. Pairwise comparisons between conditions for percentage of time slowed.[a]

Test	t statistic	p value	Cohen's d
No cues vs world-locked	-0.5	0.656	0.16
No cues vs heads-up	-2.7	0.017	0.93
No cues vs combined	-2.3	0.035	0.73
World-locked vs heads-up	-2.5	0.028	0.78
World-locked vs combined	-1.9	0.078	0.52
Heads-up vs combined	1.7	0.111	0.37

[a] Each row indicates the results of a paired t-test, with degrees of freedom = 14. The Bonferroni-corrected p-value threshold was set at 0.008. Comparisons associated with significant p values have asterisks.

compare the average speed of head rotation in the vertical (pitch) and horizontal (yaw) directions (Fig. 9 left and right, respectively). There was no significant effect of condition associated with pitch ($F(3, 45) = 0.715$, $\eta2 = 0.02$, $p = 0.548$), nor with yaw ($F(3, 45) = 2.26$, $\eta2 = 0.03$, $p = 0.095$). Preliminarily, these data do not support the notion that participants were necessarily looking around more when using the heads-up cues, and suggest that other behavioral modifications may the causing them to take more time during the task.

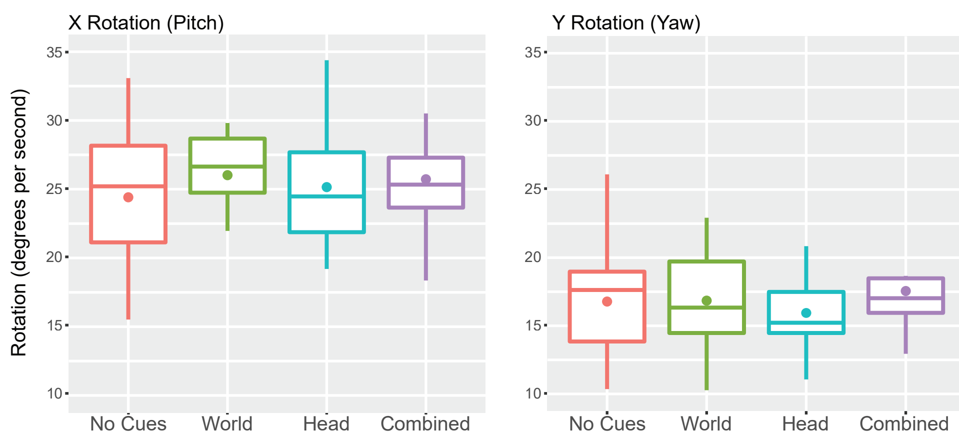

Fig. 9. Rotation in degrees per second for pitch (left) and yaw (right). Each panel shows the median, interquartile range, and non-outlier upper and lower ranges with a box-and-whisker plot. Dots represent means.

7. Discussion

7.1. Preference for world-locked over heads-up cues

The clearest result to come from our study is the overall preference for world-locked cues over heads-up ones for cueing obstacles during navigation. This preference was also supported by mobility measures, which suggested that the heads-up cue condition was associated with slower walking.

World-locked cues may have been generally preferred because these cues require less interpretation and processing compared to heads-up cues. World-locked cues in AR effectively behave like real obstacles: as soon as participants observe them, they stay in one real-world position. In comparison, the heads-up cues appear and disappear depending on the participant's orientation relative to the obstacle, and require the participant to process their meaning, turn, and find the obstacle to which they refer. Some participants stated that they had to look back and forth in order to cause the heads-up cues to reappear so they could better interpret them. The necessity to process the heads-up cues' location and shape in order to interpret them also caused difficulty for those with low acuity: some participants mistook the rectangular ceiling lights for heads-up cues. The world-locked cues also give more information about the obstacles because they indicate the size and shape of each obstacle, which is useful for navigating around them.

We initially hypothesized that world-locked cues might be more useful for those with central vision loss and heads-up cues might be more useful for those with peripheral vision loss. However, ultimately only three participants indicated that the heads-up cues were their preferred condition, and these participants did not differ notably from the other participants in terms of their peripheral vision loss. In another exploratory analysis, we examined differences in cue preferences for the participants with relatively high and relatively low acuity (performing a median split). Overall, the pattern of responses across these two groups was highly similar. Indeed, we know that the effects of low vision on mobility performance are multifaceted [6–8], so it is likely that there are similarly complex interactions between the type of low vision someone has and the most effective assistive technology in this domain. With 20 study participants and so many preferring the same cue, the current dataset is not robustly powered for examining potential relationships between individual differences in vision and cue preferences. However, we speculate that the world-locked cues were useful for diverse types of vision loss. For example, these cues may support participants with low contrast sensitivity in central vision by boosting contrast, and may also make the visual scanning strategies of people with limited peripheral vision more effective, enabling them to easily identify the presence of obstacles by quickly sweeping their vision across the environment.

It should be noted that the design of the user study may have artificially limited the performance of the heads-up cues. Because the obstacle course was a straight hallway, this meant that obstacles were less likely to appear far to the side of the participants. In addition, all obstacles were stationary. For a different task or environment, particularly one with incoming obstacles such as pedestrians, the peripheral vision-boosting effects of the heads-up cue could be more helpful. As another example, when navigating less constrained environments like sidewalks, park trails, and large indoor spaces, it may also be valuable to be alerted to potential obstacles off to the left or right for the purposes of path planning. Indeed, navigation represents a range of different situations and subtasks which may be best supported by different styles of AR cues that prioritize information differently.

In theory, the combined cues should offer the best of both worlds: the heads-up cues could direct users to out-of-view obstacles, at which point they could use the world-locked cues to narrow in on them. Several participants did find this combination helpful. However, the low awareness scores for the combined cues condition point to the main flaw: having multiple cues may lead to additional cognitive load and visual clutter for participants. In addition to the processing required by the heads-up cues, participants had to spend additional thought distinguishing the different cue types from one another. The different coloration on the cues (yellow for world-locked, green for heads-up) was not enough to distinguish them for many participants. Several reported simply trying to ignore the heads-up cues and focus only on the world-locked instead.

7.2. Technology limitations

There are many areas of improvement for AR systems in general that would also benefit AR-based visual assistance applications, such as better ergonomics, wider field of view, increased display brightness, and improved environmental scanning speed and accuracy. For example, while we had intended to utilize the HoloLens' 2 scanning ability to identify and cue obstacles in real time this was ultimately not a feasible approach: at a normal walking speed, a user would walk into an obstacle before it was processed by the hardware. For typical AR applications, Microsoft suggests creating a "scanning experience" in which the user helps the application construct a complete scan [12]. However, this technique is not viable for the type of assistive application investigated here. Additionally, while our study focused on static obstacles, many of our participants described moving obstacles such as people in a grocery store as more of a challenge. Whether via improved hardware capabilities or supplementation with additional devices (one could imagine an airport with built-in LIDAR scanning devices feeding information to user headsets in real time), this capability will need to be improved for viable mobility assistance applications of AR. Such improvements would be especially beneficial if they also increase the ability to detect and quickly cue people to small and low-contrast collision hazards. Different environmental conditions may also pose more or less of a challenge to the system for obstacle detection. For example, detection algorithms based on camera images may be more susceptible to performance decreases under low lighting conditions, whereas approaches based on active infrared illumination should be more robust.

Finally, the HoloLens 2 built-in eye calibration software failed to work for 30% of our participants. Accurate and precise eye tracking calibration for users with low vision, particularly central vision loss and/or nystagmus, is certainly a challenge. However, usable eye tracking procedures for users with low vision has a great potential to expand AR-based visual assistive technology. For example, future AR applications may aim to use eye tracking to allow users to select specific regions of the environment for enhancement just by looking at them, but this application would only be useful if eye tracking is relatively accurate.

7.3. Recommendations for AR cue design

Based on the results of the main study and the interviews conducted during the formative study, here we outline several considerations for designing AR cues for users with low vision.

Legibility of directional cues: In order for cues to be helpful for those that need them most, they must be legible even at very low levels of acuity. Our heads-up cues had a rectangular shape that, though effective for utilizing the HoloLens 2's limited field of view, made left/right and up/down difficult to distinguish except by position. At the same time, the smaller chevron shaped arrows used in the formative study were also difficult for participants to interpret. Thus, directional cues that are large and stable, but that also convey global shape information (e.g., curved bars) may be more successful.

Considering behavioral strategies: Our results also highlighted the interactions between AR cue design and people's existing orientation and mobility strategies, such as visual scanning. The design of AR cue should be tested directly with diverse groups of users that have low vision in order to understand how best to support visual tasks in concert with other behavioral strategies people may use. Cues that work well in a video game context, like heads-up cues, may not translate well to applications in visual assistive technology if they are not synergistic with other tools and strategies that the user employs. Cues should be customizable to support the behavioral strategies each user favors and the types of obstacles they need help with, including what objects are classified as obstacles.

Naturalistic design: Where possible, cues that function like real-world objects may enjoy an advantage over ones that are more abstract. Our world-locked cues seemed to be preferred in part due to their simplicity, acting simply to add contrast to existing obstacles instead of creating new visual items for participants to process.

Depth information: Additional information on the distance to obstacles was one of the most requested features by our participants. Though world-locked cues naturally grow larger as one approaches, and the heads-up cues expanded slightly based on the user's distance, this was too subtle to be noticeable. Clear indication of distance to a given obstacle is likely to be a high priority for AR systems.

8. Conclusion

AR vision enhancement systems have great potential to support people with visual impairments that cannot be corrected by conventional glasses. However, research into the technical requirements and visual design of these systems is still in its infancy. As the current work has demonstrated, many factors related to the user, the system design, and the task can influence the effectiveness of these systems. Fruitful avenues for future research on AR support for obstacle avoidance include examinations of visual cues for moving obstacles, comparisons of different environmental scanning technology, assessments of eye tracking and gaze-mediated interactions, explorations of different kinds of environments, and combining visual cues with other forms of feedback such as haptics and sound. As commercially available AR systems become more common, we hope this work can serve as part of the foundation to support a next generation of smartglasses that incorporate digital vision enhancement.

Funding. National Science Foundation (1942693, 2041726); CITRIS and the Banatao Institute at the University of California.

Acknowledgements. The authors would like to thank Martin Banks for lending an additional HoloLens 2 and for providing space for the obstacle course.

Disclosures. The authors disclose no conflicts of interest.

Data Availability. Data underlying the results presented in this paper are available in Ref. [49]

Supplemental document. See Supplement 1 for supporting content.

References

1. World Health Organization, "Blindness and Vision Impairment," https://www.who.int/news-room/fact-sheets/detail/blindness-and-visual-impairment.
2. A. D. Hwang and E. Peli, "An Augmented-Reality Edge Enhancement Application for Google Glass," Optom. Vis. Sci. **91**(8), 1021–1030 (2014).
3. J. Huang, M. Kinateder, M. J. Dunn, W. Jarosz, X.-D. Yang, and E. A. Cooper, "An Augmented Reality Sign-Reading Assistant for Users with Reduced Vision," PLoS One **14**(1), e0210630 (2019).
4. Y. Zhao, S. Szpiro, J. Knighten, and S. Azenkot, "CueSee: Exploring Visual Cues for People with Low Vision to Facilitate a Visual Search Task," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '16* (ACM, 2016).
5. T. Kuyk and J. L. Elliott, "Visual Factors and Mobility in Persons with Age-Related Macular Degeneration," J. Rehabil. Res. Dev. **36**(4), 303–312 (1999).
6. J. A. Marron and I. L. Bailey, "Visual Factors and Orientation-Mobility Performance," Am. J. Optom. Physiol. Opt. **59**(5), 413–426 (1982).
7. A. Black, J. E. Lovie-Kitchin, R. L. Woods, N. Arnold, J. Byrnes, and J. Murrish, "Mobility Performance with Retinitis Pigmentosa," Clin. Exp. Optom. **80**(1), 1–12 (1997).
8. S. E. Hassan, J. E. Lovie-Kitchin, R. L. Woods, and G. P.-Y. Soong, "Orientation and Mobility in Age-Related Macular Degeneration," Vis. Impair. Res. **1**(3), 175–180 (1999).
9. P. Ramulu, "Glaucoma and Disability: Which Tasks are Affected, and at What Stage of Disease?" Curr. Opin. Ophthalmol. **20**(2), 92–98 (2009).
10. D. J. Taylor, A. E. Hobby, A. M. Binns, and D. P. Crabb, "How Does Age-Related Macular Degeneration Affect Real-World Visual Ability and Quality of Life? A Systematic Review," BMJ Open **6**(12), e011504 (2016).
11. G. E. Legge, D. Yu, C. S. Kallie, T. M. Bochsler, and R. Gage, "Visual Accessibility of Ramps and Steps," Journal of Vision **10**(11), 8 (2010).
12. Microsoft, "Spatial mapping," https://learn.microsoft.com/en-us/windows/mixed-reality/design/spatial-mapping.
13. Magic Leap, "Magic Leap 2," https://www.magicleap.com/magic-leap-2.
14. Y. Zhao, E. Kupferstein, B. V. Castro, S. Feiner, and S. Azenkot, "Designing AR Visualizations to Facilitate Stair Navigation for People with Low Vision," in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (ACM, 2019).
15. M. Kinateder, J. Gualtieri, M. J. Dunn, W. Jarosz, X.-D. Yang, and E. A. Cooper, "Using an Augmented Reality Device as a Distance-Based Vision Aid: Promises and Limitations," Optom. Vis. Sci. **95**(9), 727–737 (2018).
16. C. Hu, G. Zhai, and D. Li, "An Augmented-Reality Night Vision Enhancement Application for See-Through Glasses," in *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)* (IEEE, 2015).

17. A. N. Angelopoulos, H. Ameri, D. Mitra, and M. Humayun, "Enhanced Depth Navigation Through Augmented Reality Depth Mapping in Patients with Low Vision," Sci. Rep. **9**(1), 11230 (2019).
18. K. R. Dillman, T. T. H. Mok, A. Tang, L. Oehlberg, and A. Mitchell, "A Visual Interaction Cue Framework from Video Game Environments for Augmented Reality," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (ACM, 2018).
19. V. R. Cimarolli, K. Boerner, M. Brennan-Ing, J. P. Reinhardt, and A. Horowitz, "Challenges Faced by Older Adults with Vision Loss: a Qualitative Study with Implications for Rehabilitation," Clin. Rehabil. **26**(8), 748–757 (2012).
20. S. Szpiro, Y. Hashash, and S. Zhao, "How People with Low Vision Access Computing Devices: Understanding Challenges and Opportunities," in *18th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS'16)* (ACM, 2016).
21. D. L. Rudman and M. Durdle, "Living with Fear: the Lived Experience of Community Mobility Among Older Adults with Low Vision," J. Aging Phys. Act. **17**(1), 106–122 (2009).
22. M. Y. Wang, J. Rousseau, H. Boisjoly, H. Schmaltz, M.-J. Kergoat, S. Moghadaszadeh, F. Djafari, and E. E. Freeman, "Activity Limitation Due to a Fear of Falling in Older Adults with Eye Disease," Invest. Ophthalmol. Vis. Sci. **53**(13), 7967–7972 (2012).
23. Y. Zhao, E. Kupferstein, H. Rojnirun, L. Findlater, and S. Azenkot, "The effectiveness of visual and audio wayfinding guidance on smartglasses for people with low vision," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (ACM, 2020).
24. H.-C. Wang, R. K. Katzschmann, S. Teng, B. Araki, L. Giarre, and D. Rus, "Enabling independent navigation for visually impaired people through a wearable vision-based feedback system," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017).
25. R. W. Massof and D. L. Rickman, "Obstacles Encountered in the Development of the Low Vision Enhancement System," Optom. Vis. Sci. **69**(1), 32–41 (1992).
26. L. E. Culham, A. Chabra, and G. S. Rubin, "Clinical Performance of Electronic, Head-Mounted, Low-Vision Devices," Ophthalmic Physiol. Opt. **24**(4), 281–290 (2004).
27. E. Peli, G. Luo, A. Bowers, and N. Rensing, "Applications of Augmented Vision Head-Mounted Systems in Vision Rehabilitation," J. Soc. Inf. Disp. **15**(12), 1037–1045 (2007).
28. P. Fernando, "Augmented-View for Restricted Visual Field: Multiple Device Implementations," Optom. Vis. Sci. **79**(11), 715–723 (2002).
29. G. Luo and E. Peli, "Use of an Augmented-Vision Device for Visual Search by Patients with Tunnel Vision," Invest. Ophthalmol. Vis. Sci. **47**(9), 4152–4159 (2006).
30. S. L. Hicks, I. Wilson, L. Muhammed, J. Worsfold, S. M. Downes, and C. Kennard, "A Depth-Based Head-Mounted Visual Display to Aid Navigation in Partially Sighted Individuals," PLoS One **8**(7), e67695 (2013).
31. J. J. Rheede, I. R. Van, R. I. Wilson, S. M. Qian, C. Downes, and S. L. Kennard, "Improving Mobility Performance in Low Vision With a Distance-Based Representation of the Visual Scene," Invest. Ophthalmol. Visual Sci. **56**(8), 4802–4809 (2015).
32. B. Phillips and H. Zhao, "Predictors of Assistive Technology Abandonment," Assist. Technol. **5**(1), 36–45 (1993).
33. H. Min Htike, T. H. Margrain, Y.-K. Lai, and P. Eslambolchilar, "Augmented Reality Glasses as an Orientation and Mobility Aid for People with Low Vision: A Feasibility Study of Experiences and Requirements," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (ACM, 2021).
34. A. Seeliger, R. P. Weibel, and S. Feuerriegel, "Context-Adaptive Visual Cues for Safe Navigation in Augmented Reality Using Machine Learning," Int. J. Hum. Comput. Interact., 1–21 (2022).
35. J. Wuertz, S. A. Alharthi, W. A. Hamilton, S. Bateman, C. Gutwin, A. Tang, Z. Toups, and J. Hammer, "A Design Framework for Awareness Cues in Distributed Multiplayer Games," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (ACM, 2018).
36. Bethesda Software, "Doom," Video game (2016).
37. Health Impact Studio, "Visual Indicators," https://foodfight.healthimpact.studio/blog/food-fight-development-blog-indicators/.
38. Owlchemy Labs, "See What We Said? Vacation Simulator Subtitle Update + Language Support!" https://owlchemylabs.com/blog/vacation-simulator-subtitle-update-language-support.
39. P. Baudisch and R. Rosenholtz, "Halo: A Technique for Visualizing Offscreen Location," in *Proceedings of the Conference on Human Factors in Computing Systems CHI'03* (1993), pp. 5–10.
40. S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani, "Clutter-Free Visualization of Off-Screen Locations," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2008), 98, pp. 787–796.
41. U. Gruenefeld, A. E. Ali, S. Boll, and W. Heuten, "Beyond Halo and Wedge: Visualizing Out-Of-View Objects on Head-Mounted Virtual and Augmented Reality Devices," in *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI '18 No. 40* (ACM, 2018), pp. 1–11.
42. T. Schinke, N. Henze, and S. Boll, "Visualization of Off-Screen Objects in Mobile Augmented Reality," in *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services - MobileHCI '10* (ACM, 2010).
43. D. Yu, H.-N. Liang, K. Fan, H. Zhang, C. Fleming, and K. Papangelis, "Design and Evaluation of Visualization Techniques of Off-Screen and Occluded Targets in Virtual Reality Environments," IEEE Trans. Vis. Comput. Graph. **26**(9), 2762–2774 (2020).

44. AROA Unity Prototype, "AROA Unity Prototype," GitHub, 2022 https://github.com/eacooper/AROA-UnityProject-Public.
45. R. H. Spector, "Visual Fields," in *Clinical Methods: The History, Physical, and Laboratory Examinations* (Butterworths, 1990).
46. S. E. Hassan, J. E. Lovie-Kitchin, and R. L. Woods, "Vision and Mobility Performance of Subjects with Age-Related Macular Degeneration," Optom. Vis. Sci. **79**(11), 697–707 (2002).
47. S. J. Leat and J. E. Lovie-Kitchin, "Measuring Mobility Performance: Experience Gained in Designing a Mobility Course," Clin. Exp. Optom. **89**(4), 215–228 (2006).
48. C. O. Fritz, P. E. Morris, and J. J. Richler, "Effect size estimates: current use, calculations, and interpretation," J. Exp. Psychol. Gen. **141**(1), 2–18 (2012).
49. AROA Data and Analysis, "AROA Data and Analysis," GitHub, 2022 https://github.com/eacooper/AROA-UserStudy-Public.