Shedding Light on Cast Shadows: An Investigation of Perceived Ground Contact in AR and VR

Haley Adams, Jeanine Stefanucci, Sarah Creem-Regehr, Grant Pointon, William Thompson, and Bobby Bodenheimer, *Member, IEEE*

Abstract—Virtual objects in augmented reality (AR) often appear to float atop real world surfaces, which makes it difficult to determine where they are positioned in space. This is problematic as many applications for AR require accurate spatial perception. In the current study, we examine how the way we render cast shadows–which act as an important monocular depth cue for creating a sense of contact between an object and the surface beneath it–impacts spatial perception. Over two experiments, we evaluate people's sense of surface contact given both traditional and non-traditional shadow shading methods in optical see-through augmented reality (OST AR), video see-through augmented reality (VST AR), and virtual reality (VR) head-mounted displays. Our results provide evidence that nontraditional shading techniques for rendering shadows in AR displays may enhance the accuracy of one's perception of surface contact. This finding implies a possible tradeoff between photorealism and accuracy of depth perception, especially in OST AR displays. However, it also supports the use of more stylized graphics like non-traditional cast shadows to improve perception and interaction in AR applications.

Index Terms—Augmented Reality, OST AR, VST AR, VR, Perception, Ground Contact, Shadows, Contrast

◆

1 INTRODUCTION

A common complaint from users of augmented reality (AR) technology is that virtual objects appear to "float" within real world scenes. A detached—or "floaty"—appearance is a strong indicator that the depth cues provided by virtual stimuli are insufficient to accurately locate the stimulus' position within a real world environment. Given that the human visual system integrates information from a variety of cues to interpret depth, the presence of unreliable depth cues can cause unstable depth perception when these cues are not combined in a consistent manner [20, 42]. Drascic and Milgram [17] as well as Adams [1] have pointed to cue conflicts as a potential factor contributing to inaccurate depth perception in augmented reality.

All visual information is produced by structured patterns of light and shadow. As such, one of the most salient visual cues for the layout of objects within a scene is the cast shadow [45]. Cast shadows may be defined as holes in light that occur when an opaque or semi-opaque object blocks the light that falls onto a surface [10]. By providing relative position information between an object and the surface upon which an object's shadow rests, cast shadows provide important depth information in both the real world [45, 68] and in augmented reality [41, 61]. As such, a significant body of work has been conducted on how to best render shadows in graphical displays [43].

According to Gibson's ground theory of space perception, one's perception of space is defined by the layout of surfaces, and the position of an object in space is defined by its relationship to surfaces [24, 25]. Within this framework, cast shadows function as a "visual glue" to attach virtual objects to surfaces [44, 63]. Furthermore, it has been demonstrated that people become more accurate when estimating egocentric distances to objects placed above the ground when they are clearly connected to the ground via shadow [52]. However, it is unclear how to best create this visual glue for augmented reality devices, especially for those devices that rely on additive light displays such as optical see-through devices. Given that this type of AR display cannot remove light—and thereby darken—virtual or real objects, rendering shadows in these devices is a challenge. We are therefore also interested in assessing how non-photorealistic shadows (cf. [8, 65]) affect perceived visual glue and depth perception. This paper describes different options for rendering cast shadows in augmented reality and virtual reality head-mounted displays (HMDs). It evaluates different shading methods in terms of how well people can perceive ground contact when they are used. Our goal is to determine possible methods for creating shadows that aid perception and will generalize across display types when possible, and to understand the current limitations and features specific to each display type. The results of this research should benefit depth perception in immersive displays by allowing for more precise localization of virtual objects through ground contact cues.

In the current work, we evaluate how the manner in which we render shadows affects ground contact perception across two experiments. A priori, we anticipated that rendering shadows that were more consistent with the real world environment–and therefore more perceptually valid–would improve ground contact perception in AR devices. Accordingly, in Experiment 1 (Section 4), we evaluated a variety of shadow shading methods, which included both perceptually motivated methods and a photometrically in-

---

-

Manuscript received November XX, 2020; revised December XX, 2020.

correct shading method. Curiously, our predictions were proven wrong and the photometrically incorrect shading method had the most pronounced effect on people's certainty in estimating ground contact in AR. These findings motivated a second investigation in which we hypothesized that the color contrast between the target object and its shadow was a contributing factor to our results in the first experiment. In Experiment 2 (Section 5) we were able to confirm that object and shadow color contrast had an effect on the likelihood of correct ground contact judgement. However, this effect was different for OST AR, VST AR, and VR devices.

2 RELATED WORK

2.1 Ground Contact Theory

In 1950, Gibson introduced his ground theory of space perception in which he proposed that there is "no such thing as the perception of space without the perception of continuous background surface" [23]. Gibson argued that visual space is not defined by arrays of objects in empty air, but by the layout of surfaces, sets of adjoined surfaces, and entities that are arranged in relation to surfaces [24, 25].

As the field of computer science advanced, Gibson's theories were tested with computer generated graphics, as well. By evaluating people's distance perception, researchers were able to manipulate intermediate platforms and surface discontinuities to measure how they affected one's perceived distance to a target in highly controlled, desktop virtual environments [7, 50]. Through these studies, it became clear that people's perception of distance to objects varied as the optical contact—the location at which the projected image of an object contacts the image of the ground beneath it—between a target and the surface beneath it varied [49, 53, 54]. The same effect of optical contact on distance perception has even recently been demonstrated in immersive augmented reality with the Microsoft HoloLens by Rosales et al. [59]. In their research, they demonstrated that when virtual targets were presented floating above the ground with no cast shadow, people perceived them as on the ground but farther away rather than as floating. As Gibson had formulated, the perception of surface-layout was essential for an individual to determine absolute and relative information for an object's position in space [23, 24].

2.2 Perception of Shadows

Cast shadows provide a strong and salient cue for depth perception by forming a point of contact between an object and adjacent surfaces [29, 44, 63]. This relationship has been demonstrated further by distance perception research in virtual environments in which people's egocentric distance estimates to targets were more accurate when cast shadows were present [52].

Prior research provides evidence that the association between an object and its cast shadow can be surprisingly robust. In addition, the manner in which shadows are shaded has different effects depending on the perceptual task. Although dark shading benefits the perception of shape from shadow [11], even white shadows are beneficial for creating a sense of depth [39]. As demonstrated by Kersten et. al.'s "ball-in-a-box" study, cast shadow shading may be manipulated to unnatural extents and yet still provide a powerful tool to determine spatial location [39]. In this particular study, given a stationary target, lightly shaded (photometrically incorrect) shadows were less effective in producing apparent motion in depth than more traditional, dark shadows. However, in a 3D environment with motion cues, light shadows proved as effective as dark shadows for determining spatial location [39]. Similarly, light shadows have been used to study visual search by various research groups. Visual search investigations have found that light shadows are processed more slowly (a matter of milliseconds) than dark shadows in visual search tasks—providing evidence that a higher-level cognitive process may be required to process shadow shading approaches that do not conform to the darkness constraint of more naturally occurring shadows [18, 32, 56].

And, yet, most of the aforementioned traditional graphics research was conducted via desktop and in completely virtual environments. It is not known how the manner in which shadows are rendered across head-mounted virtual and augmented reality devices affects a viewer's sense of ground contact to improve spatial perception in XR. This may be especially problematic for augmented reality devices, which combine both real and virtual depth cue information. This process often results in conflicting depth information and increased perceptual uncertainty [1, 17, 34].

Accordingly, researchers have begun investigating how graphically provided depth cues must be rendered to enhance spatial perception in augmented reality (AR) devices. While the results of some research evaluating depth cues like shading and texture have been mixed, findings suggest that shadows successfully improve the accuracy of depth perception in both immersive optical see-through (OST) AR [15, 22, 55] and video see-through (VST) AR [6, 61] displays. This literature also suggests that the manner in which shadows are rendered makes a difference for accuracy.

2.3 Shadows in Optical See-through AR

In augmented reality, lighting misalignment—in which the position of real world and virtual lights do not coincide—may adversely affect distance perception [22], unless the misalignment is due to the use of drop shadows [15]. Drop shadows are dark silhouettes ("shadows") that are displayed ("dropped") immediately below an object–regardless of the position of light sources in a scene. In a study of depth perception in immersive OST AR, Diaz et al. [15] found that participants' depth perception was significantly better in a drop shadow condition over a coherent lighting condition. This same work also suggests that the salience of a shadow may affect spatial perception such that more transparent shadows are less effective as depth cues.

It may be possible to leverage the human visual system to create perceptually valid shadows that aid in determining the three-dimensional layout of a scene in these devices. Such an approach holds promise for optical see-through devices, which cannot rely on traditional shading solutions for shadows due to their reliance on additive light for rendering. Thus, Manabe and colleagues [31, 46, 47] have developed a rendering technique that leverages simultaneous contrast to change the visual appearance of two adjacent colors, thereby giving the illusion of a shadow by rendering light near the outer edge of the shadow's shape. In their most recent research, they found that people perceived the illusory cast shadow as a dark color value. This finding provides evidence that simultaneous contrast illusion may be an effective approach for rendering cast shadows in additive light displays [31]. Our first experiment uses a similar technique, which we refer to as the gradient shadow, to render more realistic shadows in optical see-through displays (See Figure 1). However, our technique uses a simple linear falloff to produce a gradient, whereas the algorithms proposed by Manabe, Ikeda, and colleagues [31, 46, 47] produce more complex lighting interactions. For example, Ikeda et al. [31] use a photograph as a texture to estimate the radiance of the surface along with an empirically determined constant to account for viewpoint changes in their falloff algorithm. In contrast, some researchers have also created new display types for optical see-through AR that allow for subtractive rendering, like the recent display prototype designed by Itoh, Kaminokado, and colleagues [33, 36]. However, the design of these displays are currently an active area of research and are not commercially available.

2.4 Shadows in Video See-through AR

In contrast to OST AR devices, video see-through (VST) head-mounted displays (HMDs) are immersive augmented reality devices that may benefit from the ability to use traditional shadow shading techniques since they are able to render dark color values. But conflicts between a shadow's shading method and the real world scene can still adversely affect depth perception as a result of conflicting cue information in these devices [1, 17, 40]. Fortunately, in contrast to optical see-through AR, more research has been conducted in VST AR on cast shadows, especially for mobile AR [4, 27]. This work has mostly focused on approximating real world lighting to obtain more photorealistic shadows. However, there has also been some work evaluating how the presence of shadows, and other monocular cues affect a viewer's depth perception in mobile VST displays, that encourage the pursuit of using monocular cues to improve depth perception in AR [6, 14, 16]. In immersive video see-through displays, Kyto et al. [41] were able to improve depth perception within action space by adding binocular depth cues and relative size depth cues. Meanwhile, Vaziri et al. [64] induced global non-photorealistic effects on the video feed in augmented reality, which interestingly did not induce a significant difference in distance estimation in comparison to their unfiltered video condition.

2.5 XR HMD Tradeoffs

There are functional tradeoffs between immersive OST and VST augmented reality displays, such as the ability to view the real world unobstructed in optical see-through displays versus the comparatively easy integration of virtual and augmented stimuli in VST [58]. Although direct comparisons of different XR displays can be challenging to understand due to the variability in optics, rendering, and position tracking across experimental studies, these studies are nonetheless worthwhile as they allow us to assess specific cues and perturbations across families of devices. Only by understanding how graphical techniques impact user experience and spatial perception in different devices can researchers begin to create generalizable development guidelines for XR.

However, at present the state of this research is difficult to interpret. For example, in recent distance estimation research that compared people's perception in OST AR and VST AR displays, Medeiros et al. [48] found that OST AR displays resulted in more accurate depth perception over VST AR displays–yet Ballestin et al. found the opposite [3].

Nonetheless, design guidelines have begun to emerge from these studies. Ahn et al. [2] compared three AR devices: OST AR, VST AR, and mobile AR. They found that across devices, people's accuracy and speed was best in a size-matching task when they were presented with a more detailed 3D model (in this case, a 3D scanned object). In contrast, Cidota et al. [12] evaluated how 'diminishing' visual effects—in this case, fade and blur—affected depth perception when reaching to targets in OST AR and VR displays. They found that measured performance was best in VR when visual effects were applied but performance was best in AR when no visual effects were applied.

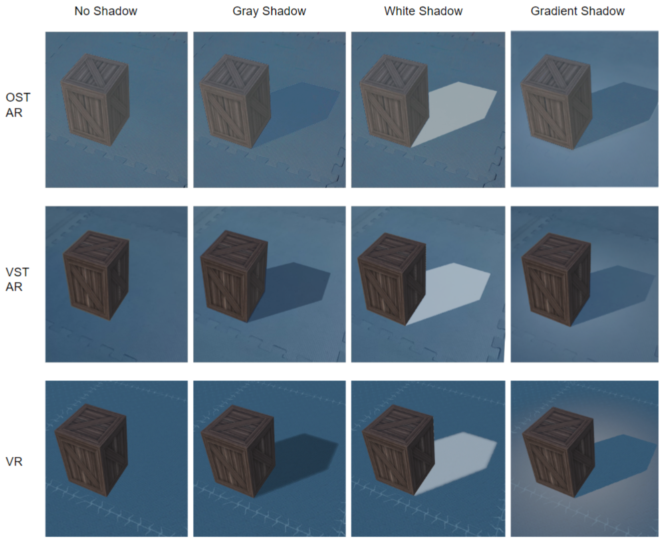

Fig. 1. Close up images of the experimental stimuli used in Experiment 1 are displayed for each XR device. All stimuli are presented on the ground in this image. The same shaders for cast shadows were used across devices for the (1) no shadow, (2) gray shadow, (3) white shadow, and (4) gradient shadow shading conditions. OST AR images were captured with the HoloLens' native mixed reality capture feature, which relies on video input. VST AR images relied on the Zed Mini's video feed. Color correction has been applied to both augmented reality images to better match what participants saw during the experiment.

3 GENERAL METHODS

In order to evaluate how shadow shading methods affect a viewer's certainty in estimating ground contact in head-mounted extended reality displays, we conducted two experiments with three unique devices: an optical see-through augmented reality (OST AR) display, a video see-through augmented reality (VST AR) display, and a virtual reality

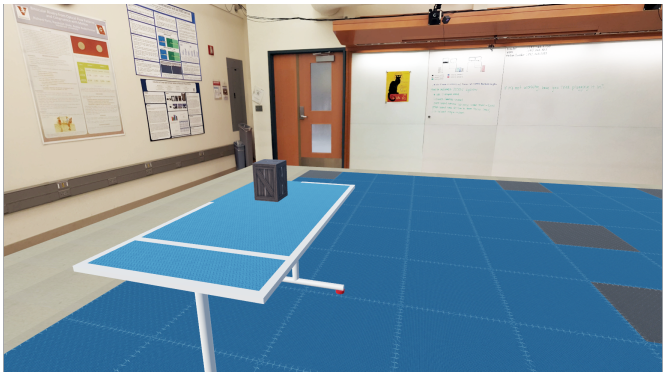

Fig. 2. A screenshot of the virtual environment from the user's perspective is displayed. A target object is placed 1m away on a nearby table. The target is rendered without a cast shadow.

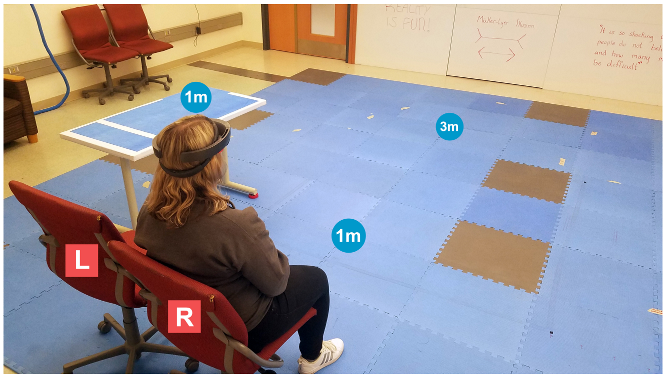

Fig. 3. A participant views experimental stimuli in the Microsoft HoloLens. The image marks the left (L) and right (R) chairs as well as the three distance conditions: 1m table, 1m floor, 3m floor.

(VR) display. The subsequent sections discuss the technical setup of our experiments (Subsection 3.1) as well as the specific solutions used for rendering shadows (Subsection 3.2) and positioning virtual objects based on viewing angle (Subsection 3.3).

3.1 Materials

We employed three immersive head-mounted displays (HMDs) for our investigations. We used the Microsoft HoloLens 1 for our optical see-through display condition, and a wireless HTC Vive Pro was used to render the virtual reality scene. The same Vive Pro was used in conjunction with a Zed Mini stereoscopic camera for the video see-through display condition. Head tracking was used in all conditions to allow natural viewing of experimental stimuli. Applications for each device were developed in Unity version 2017.4.4f1 with the C# programming language.

The Microsoft HoloLens 1 has an approximate per eye resolution of 1268×720 and field of view of $30^\circ \times 17^\circ$. Although the augmented field of view (FOV) of the HoloLens is narrow, outside of this viewing area users' vision is not occluded by the device. This OST display relies on additive light to render images and is therefore unable to render black color values. For our experiments, position tracking was performed using the HoloLens' native inside-out tracking solution.

The virtual reality environment was rendered using a wireless HTC Vive Pro, which has a maximum per eye resolution of 1440×1600 and an approximate field of view of $110^\circ \times 113^\circ$. Position tracking for this condition was performed using the Vive's lighthouse tracking system. In addition, because this condition relied on completely virtual imagery, a virtual model of the real world environment was created. This model included photographed images of the real room along the walls and custom 3D models designed to match the table and foam floor tiles present within the real world environment. An image of the virtual environment can be seen in Figure 2. An image of the real world environment can be seen in Figure 3.

The video see-through device also relied on the HTC Vive Pro for rendering. However, it created an augmented reality environment by combining virtual overlays with real time video footage, which was captured using the Zed Mini stereoscopic camera system. The Zed Mini was affixed to the front of the head-mounted display. The use of the Zed Mini's camera feed constricted the Vive's resolution and field of view to 1280×720 and $90^\circ \times 60^\circ$, respectively. Position tracking was performed using the Zed Mini's inside-out tracking solution, which integrated with the HTC Vive's tracking system.

An HTC Vive tracking puck was used to position virtual target objects within the real world environment for both the VST AR and VR conditions. At the start of an experiment and between each experimental block, the tracking puck was placed near the viewer at predetermined positions in the physical room, and virtual objects were rendered at the puck's position virtually. Similarly, for the OST AR system, this position calibration involved placing a virtual HoloLens spatial anchor at the same predetermined positions.

For the the VST AR and VR conditions, selections were performed using a wireless mouse. With the Microsoft HoloLens, users selected inputs using the HoloLens' clicker in Experiment 1. However, because a different experimental paradigm was used for Experiment 2, inputs were performed with a mouse for all three HMDs. For both experiments, a gaze-directed paradigm was useds to guide selections. A small gaze cursor appeared at the center of the user's vision whenever their forward head orientation pointed towards a user interface element within the HMDs. However, this cursor disappeared when viewing target objects so as not to disrupt their vision when evaluating ground contact.

3.2 Shadows

Shaders to render three distinct hard shadows were programmed using a variant of the HLSL language that is compatible with the Unity game engine. A directional light was positioned so that a target object's shadow would lie behind and to the right of the object. To accomplish this, the orientation of a virtual, directional light was set to 141° along the x axis and -141° along the y axis. A depiction of the three shadow conditions for each device can be seen in Figure 1. It should be noted that the images displayed in the figure do not perfectly match those presented by the immersive HMDs since there are display and capture differences. For example, the HoloLens screen capture relies on a monocular video feed for image capture but the actual user only experiences AR through the stereoscopic optical see-through display.

The dark gray shadow condition represented the most traditional method. It rendered a dark color value within the umbra of the shadow and therefore created a perceptually valid impression of a shadow for most devices.

The white shading condition, which added white light to create a shadow instead of subtracted, represented a photometrically incorrect shadow. Accordingly, it was also the most perceptually incorrect, or non-photorealistic, shading method included in this study.

The gradient shadow condition was designed as another perceptually correct method–especially for OST devices, which are unable to render black. Our gradient shadow method used simultaneous contrast to change the visual appearance of two adjacent colors and give the illusion of a dark shadow by rendering light outside of the shadow's umbra. In our method, the intensity of the light near the edge of the umbra also gradually decreased as the distance from the shadow increases, which created a gradient of light along the ground surface. Within the shadow's umbra nothing was rendered.

3.3 Vertical Displacement

For both studies, in order for participants to judge if targets were in contact with a surface, stimuli had to be presented both on and slightly above the ground for discrimination. Because we evaluated surface contact judgments across multiple distances, we displaced each target vertically based on viewing angle to ensure fair comparisons across distances as much as possible. Participants also viewed stimuli while seated throughout both experiments for consistency. The average eye height of the viewer was calculated by summing the average eye height of a person while seated (i.e., the distance from their bottom to their eyes while seated) from Harrison et al. [28] and the height of the seat of the chairs used in our setup, which resulted in a value of 1.171 m for h_e. In both experiments, some or all stimuli were placed on a table in front of the user (See Figure 3). For these conditions, h_e was adjusted to account for the table by subtracting the table height (0.7461 m) from this value.

Using the average eye height of a viewer, denoted as h_e, and the distance to a given target, d_t, we were able to solve for a series of three triangles from which we could extract the degree of vertical displacement, d_v, for target objects placed above the ground. Eye height was calculated by adding the average sitting eye height [28] to the height of the chair used in our study, which resulted in $h_e = 1.171$. In both experiments, when target objects were placed on a nearby table, instead of on the ground, we subtracted our table height from this value to obtain h_e. The trigonometric formulas used for this calculation are shown in the equations below:

$$\sigma = \tan^{-1}\left(\frac{d_t}{h_e}\right) + \omega \tag{1}$$

$$d_v = \left(\frac{\tan(\sigma)\, h_e - d_t}{\tan(\sigma)}\right) \tag{2}$$

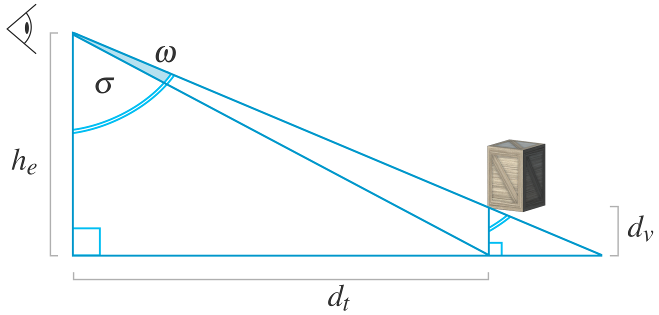

Fig. 4. Visual depiction of trigonometric solution for vertical displacement

For equations 1 and 2, ω represents the degree to which viewing angle was modified and σ represents the updated viewing angle to the vertically displaced target object. Figure 4 shows each variable in context for clarity. In Experiment 1, a viewing angle of $0.3°$ (ω) was selected since it was the height at which people could discern that an object was off the ground more than half the time during preliminary testing. Screenshots of experimental stimuli in the above ground condition can be seen in Figure 5. In Experiment 2, multiple vertical displacements were used.

4 EXPERIMENT 1

Based on our previous discussion of the importance of cast shadows as a cue for surface contact to inform depth perception, we designed our first study to test several hypotheses. First, since prior research has shown that cast shadows are an effective cue for establishing ground contact, we anticipated that the presence of a shadow would significantly affect ground contact perception (**H1**). In addition, given the unique display properties of our two augmented reality devices, we anticipated that the ability of a shadow to create a sense of ground contact would vary depending on the shading technique employed to render it in AR (**H2**). Finally, we anticipated that perceptually valid shading methods would be more beneficial in discerning ground contact over a photometrically incorrect shading method since this method would better match the real world cues for depth given by shadows (**H3**). Accordingly, a priori, we did not anticipate significant differences in ground contact perception for our virtual reality HMD condition as the completely virtual scenes generated by this device benefited from rich and consistent depth cues. Furthermore, we evaluated stimuli in our first experiment across multiple distances to ensure that any effects we found were due to our stimuli and not due to the use of a specific viewing angle, since depth cues can vary in effectiveness depending on the distance of information from the viewer [2, 13].

4.1 Participants

Thirty six individuals (27M, 9F) aged 18–32 from Vanderbilt University were recruited to participate. All participants had normal or corrected to normal vision, and each person was offered donuts for 40 minutes of their time. Our experimental methods were approved by the local institutional review board, and written consent was obtained from all subjects prior to participation.

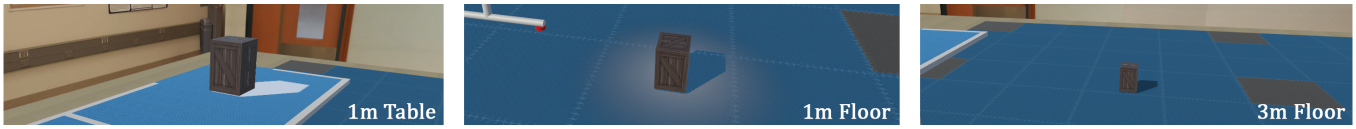

Fig. 5. Target objects in the above ground condition for Experiment 1 are positioned at three distances from the user's perspective. The white shadow, gradient shadow, and gray shadow conditions are displayed from left to right.

4.2 Design

Our protocol was modeled after the one used by Madison et al. [44] in which participants were asked to rate their confidence in perceived contact between a target object and the surface beneath it when virtual stimuli were presented either above or in contact with a surface. For both this prior study and our current research, participants rated their certainty in perceived ground contact for each stimulus using a 5-point response format with verbal anchors where values mapped to: *(1) definitely touching, (2) maybe touching, (3) unsure, (4) maybe above,* and *(5) definitely above*. The input prompt used in the current work is shown in Figure 6.

We used a mixed factorial design for this experiment. Specifically, a 3 (display type) x 4 (shadow type) x 3 (distance) design, with head-mounted display type (OST AR, VST AR, or VR) as a between-subjects variable and shadow type (none, gray, white, or gradient) and distance as within-subjects variables. Section 3.2 discusses the shadow types in further detail. Targets were placed at distances of 1m away from the viewer on a table, 1m away from the viewer on the ground, and 3m away from the viewer on the ground.

To mitigate viewing order effects, our experiment was blocked with respect to distance condition, and the order in which each distance condition was presented was counterbalanced across subjects. Within each block the order of presented stimuli was randomized without repeating within a series of 16 trials. This coincided with an experiment interruption in which participants were prompted to indicate if they required a break.

Within each display condition, participants were exposed to four shadow sub-conditions across three distances. Because we evaluated four shadow shading conditions–that were presented on or above the ground across three distances–there were 24 unique stimuli in total. Additionally, each unique stimulus was viewed 10 times, making the experiment consist of 240 trials total. We used a repeated measures design because it is an effective method for reducing the effect of variance between participants by permitting an individual to act as their own control.

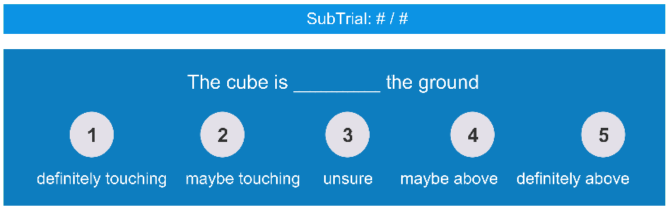
SubTrial: # / #

The cube is ________ the ground

Fig. 6. Input prompt for measuring one's certainty in estimating ground contact

4.3 Procedure

Before beginning the experiment, the participant recorded their basic demographic information and gave written consent. Then, the participant was introduced to one of the three immersive head-mounted displays, and they were instructed on how to wear and interact with the system. During this tutorial, they were also shown how to use a gaze directed interface.

Next, the researcher explained the experimental task to be performed and guided the participant to the experimental setup, which can be seen in Figure 3. In this room, stimuli were presented either on the floor, which was covered in blue foam squares, or on a nearby table. A disposable cloth of a similar blue color was draped along the top of the table with white tape wrapped around its edges. The tape was included to improve the table's salience for the two AR systems, which relied on inside-out tracking. Throughout the experiment participants sat in one of two adjacent chairs in the room. Both the approximate locations and chairs used in the experiment are marked in Figure 3 for clarity.

Virtual target objects were placed in front and slightly to the left of the participant. To view stimuli that were placed on the floor, a participant sat in the right chair, and to view stimuli that were positioned on the table, a participant sat in the left chair. The left chair was included to view targets atop the table since the table was physically offset towards the left of the room. Before each experimental block, a participant was guided to the appropriate chair for viewing. The researcher also gave participants a short break while the system was calibrated for the next distance condition. Calibration was performed to ensure that target objects appeared at the correct position in space. For the OST AR system, this process involved positioning a HoloLens' world anchor at a predetermined position in the room. For the VR and VST AR systems, calibration entailed placing a Vive Tracking puck.

After calibration was complete, the participant was given their head-mounted display. Upon donning the display, the participant saw a single prompt, which asked if they were ready to begin the next portion of the study. The next block of the experiment began once the user selected the 'ready' button below this prompt. Each participant was asked to respond to experimental trials as quickly and as comfortably as possible. The participants viewed one stimulus at a time. Once they determined their answer, they clicked once to reveal an input prompt with the 5 potential confidence in ground contact ratings. This prompt appeared immediately in front of the user to prevent strenuous head movement, and the target object was removed from sight. The next trial began once the user selected a value between 1 to 5 on the input prompt. Clicking anywhere else allowed the user to toggle between viewing the current target and

TABLE 1
Results of Friedman analysis between the no shadow and shadow conditions for each device

	On Ground				Above Ground			
	Avg Rating				Avg Rating			
	No Shadow	Shadow	$\chi^2(2)$	Sig	No Shadow	Shadow	$\chi^2(2)$	Sig
OST	2.5	1.6	4.455	$= 0.035^*$	2.7	3.1	1.600	$= 0.206$
VST	2.2	1.4	8.333	$= 0.004^*$	2.4	3.4	8.333	$= 0.004^*$
VR	2.1	1.5	8.333	$= 0.004^*$	2.8	3.7	12.000	$= 0.001^*$

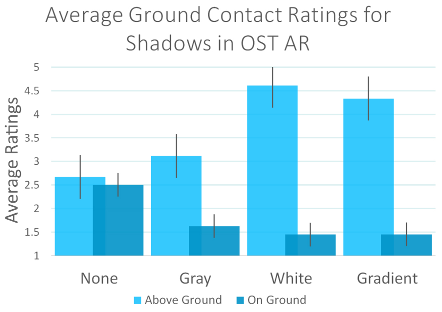

Fig. 7. Average confidence in ground contact rating with 95% CI of ground contact with shadows in OST AR

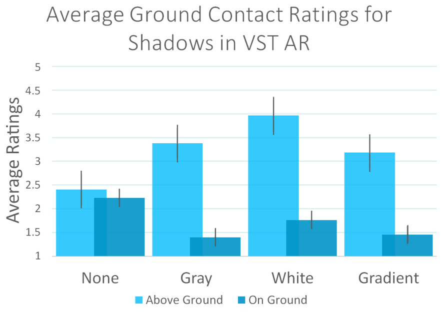

Fig. 8. Average confidence in ground contact ratings with 95% CI of ground contact with shadows in VST AR

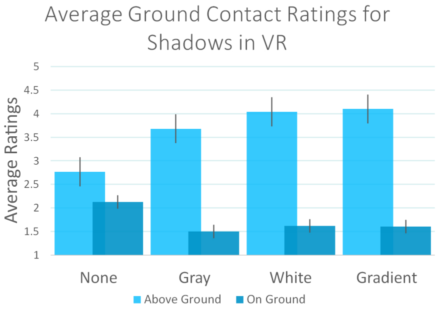

Fig. 9. Average confidence in ground contact ratings with 95% CI of ground contact with shadows in VR

its corresponding input prompt. After every 16 trials, the participant was asked if they needed a break via a virtual prompt within the simulation. After 80 trials, the system was re-calibrated for the next distance condition and the next experimental block began.

4.4 Results

Our study used ordinal subjective assessments to investigate differences in perceived ground contact across shadow shading conditions. Participants gave ratings of their confidence in perceived surface contact using a 5-point response format like the one used in Madison et al. [44]. Due to the use of an ordinal response format, the resulting data were not normally distributed and nonparametric statistical analyses were employed for correct interpretation. First, we employed Friedman tests to determine if there were differences in confidence ratings between experimental conditions. For each participant we evaluated the average confidence rating for perceived contact across 10 repeated trials. We then used Wilcoxon signed-rank tests in post-hoc analyses with Bonferroni correction to understand specific effects of shadow and distance conditions. Bonferroni correction is recommended for evaluations with multiple comparisons to compensate for an increased chance of Type I error. The significance level used for each analysis is included within the statistical reporting for each experimental finding.

Although we conducted our experiments with three XR displays, we did not make direct, statistical comparisons of people's perception between these displays due to the high variance in ergonomic, optical, and graphic properties of each device. Our hypotheses and analyses instead focused on how shadow shading affects perception in each display.

4.4.1 *Does having a shadow matter?*

Based on an abundance of prior research in psychology and computer graphics, which demonstrates that shadows provide an effective cue for establishing ground contact [22, 44, 52, 61], we anticipated that the presence of a shadow would result in a significant difference in people's confidence in perceived contact between the no shadow and the other shadow conditions. Furthermore, we predicted that this trend would hold in all three XR display types. To evaluate our hypothesis, for each display we ran a Friedman test on the average confidence of contact ratings between the no shadow condition and all other shadow conditions when collapsed. Table 1 summarizes the average confidence ratings across participants as well as the results of this analysis.

For the OST AR display, when target objects were placed on a surface, Friedman tests showed a significant difference in people's confidence of contact ratings between the shadow conditions when collapsed together and the no shadow condition ($\chi^2(2) = 4.45$, $p = 0.035$). However, when objects were placed above the ground, there was no significant difference in confidence ratings between the no shadow and shadow conditions ($\chi^2(2) = 1.60, p = 0.206$). This indicates that people's confidence in perceiving when an object was placed above the ground was higher when a cast shadow was present—but only when the target object was truly placed on the ground.

For both the VST AR and VR display conditions, people were more confident in assessing surface contact in the presence of a shadow, regardless of whether the target object was placed on the ground or above it. For the VST AR device, people's confidence differed between the no shadow and shadow conditions for targets with the same degree of significance for both on and above ground targets ($\chi^2(2) = 8.33, p = 0.004$). There was a difference in people's

TABLE 2
Results of Friedman analysis between the shadow conditions for each device

	On Ground					**Above Ground**				
	Avg Rating					Avg Rating				
	Gray	White	Gradient	$\chi^2(2)$	Sig	Gray	White	Gradient	$\chi^2(2)$	Sig
OST	1.6	1.4	1.5	1.317	$= 0.518$	3.1	4.6	3.8	18.500	$< 0.001^*$
VST	1.4	1.8	1.5	7.787	$= 0.020^*$	3.4	4.0	3.4	12.667	$= 0.002^*$
VR	1.5	1.6	1.6	1.227	$= 0.541$	3.7	4.0	3.9	4.667	$= 0.097$

confidence of contact given the presence of a shadow for the VR display, as well, for both the on ground target objects ($\chi^2(2) = 8.33, p = 0.004$) and the above ground target objects ($\chi^2(2) = 12.00, p = 0.001$).

As expected, we found significant differences in confidence of contact between the no shadow and shadow conditions for all devices–a result which confirms our hypothesis (**H1**). However, it was curious that confidence ratings for virtual objects in OST AR were significant only when placed on the ground. To better understand this finding, we conducted a post-hoc analysis on the average confidence ratings for above ground objects for each shadow condition in the OST AR device. Post -hoc, Wilcoxon signed rank tests with Bonferroni corrected significance ($p < 0.0167$) indicated that people's ratings for the no shadow condition were significantly different from the white ($Z = -3.1, p = 0.002$) and gradient ($Z = -2.8, p = 0.005$) shadow conditions, but not the gray shadow condition. These findings were reasonable, given that dark color values are known to appear more transparent in additive light displays like those used in OST AR. If a shadow is too transparent for a viewer to discern when an object is placed above the ground, then it may adversely affect the viewer's confidence in their ability to determine surface contact such that their confidence ratings more closely resemble the ratings expressed in the no shadow condition.

4.4.2 *Which type of shadow affects perception of ground contact?*

Next, we evaluated differences in confidence of contact for on and above ground targets for the different shadow shading methods via Friedman tests for each device. We then ran post-hoc Wilcoxon signed rank tests with Bonferroni corrected significance ($p < 0.0167$). In the OST AR device, Friedman tests across the three rendered shadow conditions revealed no significant differences in people's confidence ratings between the methods when an object was placed on the ground. However, the same analysis did find a difference when objects were placed above the ground ($\chi^2(2) = 18.500$, $p < 0.001$). Post-hoc tests revealed that confidence of contact ratings for all three shadow conditions were significantly different from each other. Namely, ratings for the white shadow ratings differed from the gray shadow ($Z = -3.1, p = 0.002$) and the gradient shadow ($Z = -2.8, p = 0.005$); the ratings for the gradient shadow differed from the gray shadow ($Z = -2.9, p = 0.004$). For each shadow condition the average confidence of contact ratings were: 3.1, 4.6, and 3.8 for the gray, white, and gradient shadows, respectively. Figure 7 further illustrates the differences in these ratings, where the white shadow shading condition is given the highest confidence rating and the gray shadow shading condition is given the lowest confidence rating. For OST AR, brighter–and therefore more salient–shadows appear to greatly influence people's confidence of surface contact when objects are placed above the ground.

For the VST AR device, we found a significant difference in confidence of contact ratings between shadow conditions for both on ground ($\chi^2(2) = 7.787, p = 0.020$) and above ground ($\chi^2(2) = 12.667, p = 0.002$) target objects. The average ratings for the no shadow conditions and the three, shaded shadow conditions are visualized in Figure 8. Post-hoc analyses revealed that when target objects were on the ground, the white shadow was significantly different from the gray shadow condition ($Z = -2.432, p = 0.015$). Furthermore, for above ground objects, the white shadow was significantly different than both the gray ($Z = -2.472, p = 0.013$) and the gradient shadow condition ($Z = -2.590, p = 0.010$). Interestingly, for both on and above ground targets, the white shading method for shadows generally resulted in higher average confidence of contact ratings for target objects when compared to the other shading methods (See Table 2).

In VR there was no significant difference in confidence of contact between the three shadow shading methods (See Figure9 and Table 2). Although confidence was unaffected by shadow shading method in VR, confidence ratings in AR displays proved quite sensitive, which confirmed our second hypothesis (**H2**) that shadow shading method would influence people's confidence in surface contact perception in augmented reality displays. However, we were unable to confirm (**H3**), which predicted that people would make more confident surface contact ratings given perceptually valid cast shadows over photometrically incorrect shadows in AR. This outcome becomes particularly curious when we examine the photometrically incorrect (the white shadow) shading condition's performance between the OST AR and VST AR devices. In the OST AR display, people's confidence was significantly higher when target objects were placed above the ground. Their confidence matched the ground truth of the target's position in space. However, in VST AR, people's confidence ratings were higher when presented with the white shadow, regardless of whether the object was placed on the ground or above it, which may be an undesirable outcome for establishing ground contact.

4.4.3 *Does distance to the object affect the perception of ground contact with shadows?*

To better understand how our results may be influenced by viewing conditions within personal and actions spaces,

TABLE 3
Results of Friedman analysis between the distance conditions for each device–Friedman's Q

	On Ground					Above Ground				
	Avg Rating					Avg Rating				
	1m Table	1m Floor	3m Floor	$\chi^2(2)$	Sig	1m Table	1m Floor	3m Floor	$\chi^2(2)$	Sig
OST	1.5	1.4	1.7	3.268	= 0.195	3.4	3.7	4.5	16.468	< 0.001*
VST	1.4	1.8	1.4	2.227	= 0.328	3.0	3.3	4.4	12.667	= 0.002*
VR	1.7	1.7	1.7	0.304	= 0.859	2.5	4.2	4.1	18.766	< 0.001*

for each device we evaluated shadow conditions across viewing distances from the observer, since depth cues can vary in effectiveness across distances [13]. We first conducted Friedman tests across the 1m table, 1m floor, and 3m floor conditions to determine if there were differences in confidence of surface contact ratings across the viewing distances. Table 3 shows the average ratings and significance values across conditions for this analysis. We then ran post-hoc Wilcoxon signed rank tests with Bonferroni corrected significance ($p < 0.0167$) to compare confidence ratings between specific distance conditions.

For the OST AR display, we found no significant difference in confidence of contact ratings across viewing distance conditions when the object was in contact with the ground. However, there was a significant effect of distance when objects were placed above the ground ($\chi^2(2) = 16.468$, $p < 0.001$). Post-hoc tests showed that people's confidence in contact ratings for the 3m floor distance were significantly higher than the ratings for the 1m floor ($Z = -2.8, p = 0.006$) and 1m table conditions ($Z = -3.1, p = 0.002$).

We also found no difference in confidence of contact ratings across viewing distances for the VST AR when objects were placed on the ground. However, there was a significant difference between viewing distance conditions for above ground objects ($\chi^2(2) = 12.667, p = 0.002$). In a similar pattern to that found in the OST AR condition, we found that people were more confident that target objects were placed above the ground for the 3m floor condition than either the 1m floor condition ($Z = -2.7, p = 0.008$) or the 1m table condition ($Z = -3.1, p = 0.002$).

Similarly, in the VR display there was no effect of viewing distance for the on ground objects; however, there was a significant difference between viewing distance conditions when objects were placed above the ground ($\chi^2(2) = 18.766, p < 0.001$). In VR, people's confidence of surface contact was significantly higher for the 3m floor and 1m floor conditions when compared to the 1m table condition, with the same degree of significance for both comparisons ($Z = -3.1, p = 0.002$). The average confidence ratings for each viewing distance condition are displayed in Table 3. The average ratings for above ground objects were higher in the 1m floor and 3m floor conditions with scores of 4.2 and 4.1, respectively. In contrast, the average rating for the 1m table distance condition was only 2.5.

For all devices, we found significant differences in confidence of surface contact ratings as a function of viewing distance for target objects positioned above the ground, but not for those positioned on the ground. In addition, confidence ratings for above ground objects were significantly higher for the 3m floor condition than either one or both of the 1m floor and 1m table conditions for all devices. A possible explanation for this finding is that the 3m floor condition permitted the viewer to see underneath the target object, thereby making it easier for the viewer to determine when objects were positioned above the ground since their view of the cast shadow underneath the object was more clear. However, people's ratings for the 3m floor and 1m floor conditions did not significantly differ in the VR condition, so this finding may be unique to AR devices.

4.5 Discussion

Overall, the results of our comparisons between the no shadow condition and the shadow conditions when collapsed re-affirmed that shadows provide a powerful cue for ground contact (**H1**). However, they caution against the use of dark color values for shadows in OST AR devices as they may be less effective than other, more visible, techniques for establishing ground contact. We did not find a significant effect of shadow condition for the OST AR display when objects were placed above the ground. However, post-hoc analyses revealed that the gray shadow performed similarly to the no shadow condition, which likely influenced the outcome of our Friedman test across all shadow conditions.

A priori, we also anticipated that shadow shading techniques would vary in effectiveness for both AR devices (**H2**). We confirmed this hypothesis, but we were surprised to find that the photometrically incorrect shading method–the white shadow–generally resulted in higher confidence of surface contact when objects were placed above the ground for both AR devices. This outcome was particularly unexpected since we had predicted that more perceptually valid shadow methods, like the gray and gradient shadows, would be more beneficial for establishing ground contact. As a result, we could not confirm our third hypothesis (**H3**). This counterintuitive result encouraged further evaluation of the effect of high contrast, and therefore more visible, object and shadow shading conditions for our second study.

It should be noted that we saw different effects of shading conditions between the OST AR and VST AR devices. This can most likely be attributed to differences in the displays as both AR devices combine real and virtual images in very different ways. Whereas people's confidence of surface contact in the OST AR device proved highly sensitive to all shading methods when objects were placed above the ground, people's confidence of contact in the VST AR condition was sensitive to a single shading condition–the white shadow–regardless of whether an object was placed on or above the ground.

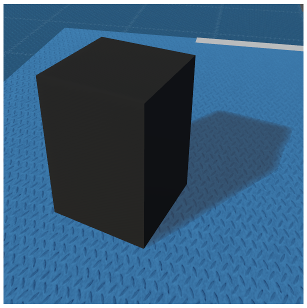

Fig. 10. Dark object with dark shadow condition (DODS)

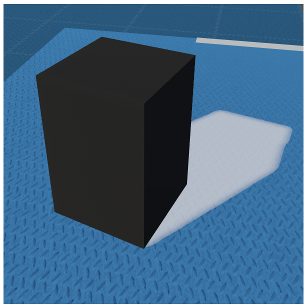

Fig. 11. Dark object with light shadow (DOLS)

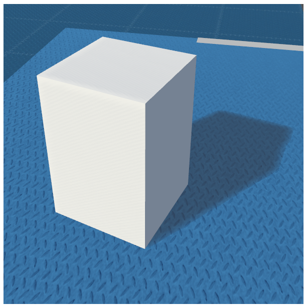

Fig. 12. Light object with dark shadow (LODS)

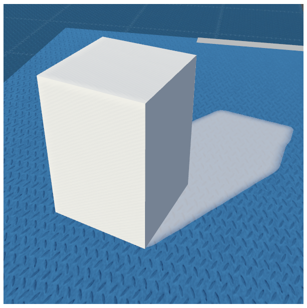

Fig. 13. Light object with light shadow (LOLS)

In addition, for the OST AR device we found that the white shadow resulted in the highest confidence and the gray shadow resulted in the lowest confidence. The gradient shadow's confidence ratings fell between the two. One interpretation of these results is that more salient shading techniques are more effective for determining ground contact in OST AR displays. For example, the white shadow may have been more prominent against the darkly textured target object used in our first evaluation. This provided additional inspiration to evaluate the relationship between object and shadow shading in a second experiment.

5 EXPERIMENT 2

The results of our first experiment indicated that shadow shading plays an important role in determining ground contact in augmented reality. Interestingly, we found that people performed better with white shadows—or photometrically incorrect shadows—than with the other shading methods when discerning ground contact for both OST AR and VST AR devices, especially when objects were placed above the ground. Given these results, we suspected that the high contrast of the white shadow's color value against the darkly textured target object used in Experiment 1 allowed participants to more confidently rate ground contact. In the interest of better understanding how object shading may have influenced our initial findings, we evaluated different object and shadow shading conditions in Experiment 2. Specifically, we evaluated four unique combinations of object and shadow shading conditions to parse out what aspects of perceived surface contact were affected by shadow shading method alone versus the contrast of an object and its shadow shading method. Objects were then presented at multiple heights to permit greater sensitivity in measuring perception of ground contact.

Based on the results of our first experiment, we developed three hypotheses for this experiment. First, we anticipated that high color contrast between objects and their shadows would improve the likelihood of correct assessment of ground contact in augmented reality displays (**H1**). We also predicted that white shadow shading methods would improve participants' ability to discern surface contact in augmented reality (**H2**). However, given the results of our first study and our prior discussion, we did not anticipate any significant effects of shadow shading method on task performance for the virtual reality head-mounted display (**H3**).

5.1 Participants

Six individuals in total (3M, 3F) aged 20–45 from Vanderbilt University volunteered to participate in the second experiment, which used a psychophysical paradigm. Psychophysics is a class of psychological methods that quantitatively measures perceptual responses to changes in physical stimuli [21].

These methods rely on a small number of participants to make a large number of simple, behavioral responses that reveal underlying perceptual processes. This family of methods has proven highly replicable since they employ judgments or adjustments with low individual variance [19, 60]. Although psychophysical paradigms typically rely on a smaller number of participants, six participants were required to counterbalance the presentation order of the three immersive displays. All participants had normal or corrected to normal vision. Our experimental methods were approved by the local institutional review board, and written consent was obtained from all subjects prior to participation.

5.2 Design

For Experiment 2, we evaluated the relationship between object and shadow shading using the same OST AR, VST AR, and VR devices that were used in Experiment 1. We also used the same testing environment with a medium blue background so that we could draw comparisons between the two studies. In this second experiment, all objects were positioned 1m away from the viewer on the same table that was used in Experiment 1.

However, in this experiment we employed a psychophysical approach to evaluate how light and dark color values affect the relationship between a target object and its shadow when determining ground contact. By using this approach we were able to restructure our experiment as a within-subjects evaluation for further experimental control, and we were able to efficiently evaluate ground contact perception for target objects at multiple heights for a more sensitive measure of performance. Specifically, we used a two-alternative forced choice (2AFC) design with method of constant stimuli—a classic approach [26]. We used a within-subjects 2 x 2 x 6 factorial design in which two levels of shading for a target object and its shadow were presented, and targets were presented at 6 different vertical displacements.

The shading levels contained light and dark color values such that two high contrast and two low contrast conditions were created. In this context, we refer to contrast as a difference in color—rather than luminance—since we are unable to directly compare luminance values using traditional methods over the three unique display types. Grayscale color values were used to inform shaders for both the target object and its shadow. The low contrast conditions were: [light object x light shadow] and [dark object x dark shadow]. The high contrast conditions were: [light object x dark shadow] and [dark object x light shadow]. High contrast conditions had a difference of 200 RGB color values and low contrast conditions had a difference of 30 RGB color values. Specifically, we use grayscale RGB color values of 250 and 20 to inform the target object shader and we use grayscale RGB color values of 220 and 50 for the shadow shader in this experiment. The same shadow shaders that were developed for Experiment 1 were also used for Experiment 2 (See Subsection 3.2). This design resulted in 4 unique combinations of experimental stimuli.

For each stimulus pair in our temporal 2AFC protocol, the participant was asked "Which stimulus is closer to the ground?" Each stimulus was presented for 600 msec. In between each stimulus pair, there was an interval of 800 msec in which a random pattern was presented to avoid visual aftereffects. Participants responded using the left mouse button to indicate the first stimulus and the right mouse button to indicate the second stimulus. One stimulus was presented on a surface and the other was presented at some vertical displacement above the surface.

Objects were presented at one of six heights between 0 and 3mm at regular intervals of 0.03 degree changes in viewing angle (See Section 3.3). This change in visual angle corresponded to approximately a 0.06mm change in height per step. Each height comparison was presented 20 times each, except for when both the first and second object heights were both presented on the ground with 0 vertical displacement. In this case, the height conditions were only presented 10 times each.

The experiment was blocked by device and counterbalanced across subjects. Each device block consisted of 440 trials, resulting in each participant completing 1320 trials throughout the experiment. The experiment took approximately two hours to complete, with participants taking around 30 minutes to complete the experimental task in each device. Within each experimental block, stimulus pairs were presented pseudo-randomly such that there were no repetitions of any unique stimulus combination before all other unique stimuli were presented once. The experiment was self-paced. Both the user's response and their response time were recorded for each trial. The next trial began 1000 ms after the participant responded to the previous trial–unless the experiment was paused. In total, across all subjects we collected 7,920 datapoints.

5.3 Procedure

This experiment was conducted after the Covid-19 outbreak. Therefore, special considerations (e.g., social distancing) had to be implemented to protect both the experimenter and the participants. First, the researcher explained the experimental protocol to the participant and gave them an informed consent form. Throughout the experiment, participants were informed that they could take a break at any time, and the programs used to run the study asked participants after every 44 trials if the they required a break from the experimental task.

Participants were asked to calibrate the equipment themselves–instead of the experimenter–to avoid the spread of germs through shared head-mounted displays. Otherwise, the calibration phase for the current experiment was the same as in Experiment 1. After completing the trials in one head-mounted display, the participant filled out a short post survey before continuing onto the next head-mounted display used for the experiment and thus the next block of trials. At the end of the experiment, participants filled out a short final survey.

5.4 Results

We analyzed our data using a binomial mixed model to understand the influence of shading condition on participants' judgments in each display. Generalized linear mixed models (GLMMs) are a form of generalized regression that is appropriate for repeated-measures designs and for non-normally distributed outcomes.

Because participants in our study were asked to make binary decisions about which object was positioned closer to the ground (first or second object) in a 2AFC task, our GLMMs were fitted by specifying binomial errors and a logit link function using the glmer function in the lme4 library [5] in R [30]. To be more specific, we modeled binary outcomes (correct or incorrect) for our input variables (predictors), which included object shading (2 levels: dark, light) shadow shading (2 levels: dark, light), and vertical displacement (5 levels). Although method of constant stimuli dictated that both targets be positioned at the same location for some comparisons, we did not evaluate this condition in our model since there was no incorrect answer for this condition. As such, we evaluated five instead of six levels for vertical displacement in our model. For the GLMM analysis, positive effects indicate an increased likelihood of correct ground contact judgement in our 2AFC task.

For brevity, we will refer to our object shading and shadow shading combinations with the following acronyms throughout the results discussion: 1) dark object with dark shadow (DODS), 2) dark object with light shadow (DOLS), 3) light object with dark shadow (LODS), and 4) light object with light shadow (LOLS).

For all three devices, we found a main effect of shading condition such that participants viewing the dark object with light shadow (DOLS) condition had an increased likelihood of correctly judging a target's ground contact in comparison to the other shading methods. In Figure 14 participants' forced choice response data has been plotted with a psychometric function for each display condition for clarity.

In the video see-through augmented reality device, the DOLS condition resulted in an increased likelihood of correct response in comparison against the DODS ($\beta = 1.397, SE = 0.149, p < 0.0001$), the LODS ($\beta = 1.41, SE = 0.149, p < 0.0001$), and the LOLS ($\beta = 1.25, SE =$

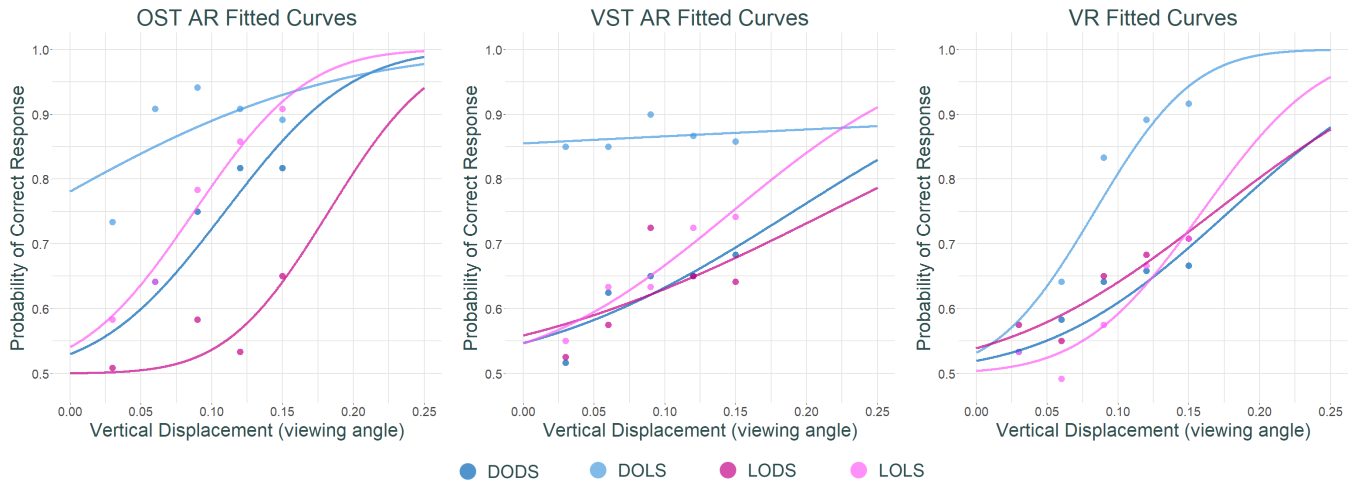

Fig. 14. Response data from the 2AFC task in Experiment 2 has been fitted with a cumulative normal curve using a generalized linear model. The figure shows the psychometric curves for response data for the OST AR display (left), the VST AR display (middle), and the VR display (right) conditions when targets were placed above the ground.

$0.150, p < 0.0001$) shading conditions. For the virtual reality device, the DOLS condition performed similarly against the other conditions with an increased likelihood of correct response in comparison to the DODS ($\beta = 0.807, SE = 0.13, p < 0.0001$), the LODS ($\beta = 0.658, SE = 0.132, p < 0.0001$), and the LOLS ($\beta = 0.830, SE = 0.131, p < 0.0001$) conditions.

From the results of our first experiment, we expected more nuanced relationships between object and shadow shading methods to be revealed for the OST AR device, and this expectation was met. Each shading condition performed significantly different from each other such that the order of highest to lowest likelihood of correct response was: 1) dark object with light shadow (DOLS), 2) light object with light shadow (LOLS), 3) dark object with dark shadow (DODS), and 4) light object with dark shadow (LODS).

In a similar fashion to the results found in the other displays, the highest likelihood of correct response was found in the DOLS shading condition. Statistically speaking, people's responses given the dark object with light shadow condition (DOLS) were more accurate than when they were presented with any of the other shading conditions: LOLS ($\beta = 0.894, SE = 0.1613, p < 0.0001$), DODS ($\beta = 1.215, SE = 0.158, p < 0.0001$), LODS ($\beta = 1.919, SE = 0.155, p < 0.0001$). People also performed better in the light object with light shadow condition (LOLS) than both of the dark shadow conditions: DODS ($\beta = 0.321, SE = 0.136, p = 0.0181$) and LODS ($\beta = 1.0253, SE = 0.132, p < 0.0001$). And they performed better in the dark object with dark shadow condition than the LODS ($\beta = 0.704, SE = 0.127, p < 0.0001$) condition.

5.5 Discussion

For all devices, the dark object with light shadow (DOLS) condition resulted in an increased likelihood of correct judgement of distance to the ground. This result echoes the findings of our first experiment in which the white shadow condition resulted in significantly different responses for both augmented reality devices. Prior research in computer graphics has shown us that light, photometrically incorrect cast shadows may be as effective as dark, perceptually correct cast shadows in spatial location tasks [39]. Our current research finds a similar trend in that people often perform better with the white shadow condition than with other perceptually valid shading approaches—especially in our employed additive light display.

However, it is interesting to note that while the DOLS condition resulted in a significant increase in performance in the VR condition for Experiment 2, we did not see this pattern of results in the first experiment. We believe this is due in part to the difference in paradigm. Whereas our first evaluation was useful for uncovering multiple cues that may be interacting with a viewer's ability to perceive ground contact, in general, our second evaluation used a more sensitive paradigm to allow us to isolate to what degree cast shadow shading cues were affecting a viewer's sense of ground contact. Thus, we may have been better able to isolate subtle differences in performance in Experiment 2. In all devices, participants generally expressed that the dark object with light shadow (DOLS) shading condition was the easiest condition to see the shadow, which explains why it performed pointedly well.

The DOLS method was the only one to result in a significant difference in performance from the other conditions in both the video see-through and virtual reality device (DOLS $>$ LOLS , DODS , LODS) . However, the optical see-through device proved highly sensitive to all shading conditions. In order of highest likelihood of correct response to lowest likelihood, the shading conditions performed as follows: DOLS $>$ LOLS $>$ DODS $>$ LODS. It is important to note that the two high contrast conditions (dark object with light shadow and light object with dark shadow) resulted in highly polarizing performance. This is likely due to the HoloLens' reliance on additive light to render objects. In addition, both light shadow conditions resulted in more accurate ground contract perception over the dark shadow shading methods. While participants expressed that the DOLS condition was easy to interpret since the shadow was very visible, participants complained that the opposite was true for the LODS condition. Participants reported that the bright appearance of the object made it difficult to see the dark shadow underneath it.

In summary, we were unable to confirm our first hypothesis, in which we predicted that people would perform better in the high contrast color conditions than the low contrast conditions when determining ground contact in augmented reality (**H1**). However, our results did indicate that a viewer's perception of ground contact is highly sensitive to high color contrast between objects in their shadows in optical see-through AR. We were also only able to partially confirm our second hypothesis in which we predicted that the white shadow condition would improve surface contact perception in augmented reality displays (**H2**). We found this relationship to be true in the OST AR device; however, in VST AR the white shadow only resulted in an increased likelihood of correct response when the target object was dark (DOLS). Finally, unlike in our first study, we found a significant effect of shading condition in the virtual reality condition where the DOLS shading condition performed significantly better than all other conditions, which means that were unable to confirm our third hypothesis (**H3**).

6 GENERAL DISCUSSION

Over two experiments we evaluated how the appearance of a shadow through shading influences one's perception of ground contact in immersive displays. Cast shadows provide an important depth cue as they allow viewers to perceptually attach objects in space to the surfaces and other objects near them. The three head-mounted displays we evaluated—optical see-through augmented reality, video see-through augmented reality, and virtual reality—each had unique rendering properties that could affect how monocular depth cues, such as cast shadows, were utilized by the viewer. Of particular interest was the optical see-through display, which could not rely on traditional shading techniques for shadows due to its reliance on additive light technology to display virtual overlays and its subsequent inability to present dark color values.

To establish a baseline understanding of people's ability to perceive ground contact from shadows in these devices, our first study evaluated people's ability to discriminate the position of both on and above ground objects using a paradigm employed originally by Madison et al. [44]. The results of our first experiment were surprising in that for both augmented reality conditions, the most photometrically incorrect shading method allowed users to better discriminate when an object was placed above the ground, as opposed to either of the two perceptually motivated shading methods. We believed this was due, in part, to the color contrast between the virtual object that was rendered with a dark crate texture, and its shadow, which was white. The use of a darkly shaded test object may have made it easier for viewers to detect the presence of a shadow beneath the object when it was placed above the ground. Another finding was that there was a significant difference in people's confidence of surface contact across shadow conditions for AR but not for VR. We suspected that these results may have been influenced by display differences between VR and AR or they may have been due to the fact that our paradigm only used two object heights for evaluation of perceived ground contact.

Both of these findings motivated the design of our second study. To better understand the emergent behaviors of our first evaluation, we varied both shadow and object color to understand how the differences in these color values influenced people's ability to accurately perceive ground contact. In addition, we elected to use a psychophysical paradigm, which allowed us to evaluate each shading condition at multiple heights for a more sensitive evaluation of people's ability to discern if an object was in contact with the ground.

In this second study, we evaluated four unique shading conditions in which a dark target object was rendered with either a dark shadow (DODS) or a light shadow (DOLS), and a light target object was rendered with either a dark shadow (LODS) or a light shadow (LOLS). In all devices, the dark object with light shadow (DOLS) condition resulted in a significantly increased likelihood of correct judgment of ground contact. This condition is most similar to the white shadow condition of the first experiment, which reinforces our findings related to the photometrically incorrect shading condition in Experiment 1.

Our results have interesting implications for the design of both virtual and augmented reality applications because they encourage the use of non-photorealistic shading methods for improving surface contact perception. Specifically, developers may use non-photorealistic shading methods for shadows to more effectively establish ground contact, which is an especially desirable outcome for augmented reality displays given the common complaint of AR technology adopters that virtual objects in real world spaces appear to "float."

It is important to note that the optical see-through condition in our experiment was especially sensitive to the relationship between object and shadow shading. All four shadow shading methods were significantly different from each other. As previously stated, the dark object light shadow (DOLS) shading was significantly better for estimating ground contact compared to all other methods. But following this performance, the next best method for determining ground contact was the light object, light shadow condition (LOLS). The optical see-through device's reliance on additive light is likely a contributing factor given that people performed better in both light shadow conditions than in either of the dark shadow conditions. The results from our first and second experiments caution against the use of dark color values to render shadows in OST AR devices.

These insights are useful for improving spatial perception in AR, but they may be alarming for those seeking to create more photorealistic graphics for applications, given that our results imply a tradeoff between realism and spatial perception enhanced by ground contact. By using a light color value to render shadows, we were able to improve people's accuracy in detecting subtle changes in position at near distances. This improvement in accuracy may be especially useful for applications that require high precision like the spinal surgery application mentioned at the start of our paper [67].

It should be noted that our application is far from the first to manipulate the appearance of shading or shadows; these applications have typically been developed to achieve a stylized effect [37]. However, little work has investigated using stylized graphics as a tool to improve spatial perception [12, 64, 65], even though colored shadows have been used previously to create numerous interesting experiences and interactions for photography and interactive projections [35, 38, 51]. It may be interesting to compare how colorful shadow effects influence surface contact perception between real and virtual or augmented environments.

6.1 Limitations

There are clear limitations to the current study. In both evaluations, we used a very simple, rectilinear shape as our target object. This was a deliberate choice for our initial investigation into the evaluation of shadow shading methods because we did not want complex geometry to introduce additional depth cue information. However, moving forward, it will be important to evaluate different geometric shapes since geometry may influence the strategy a viewer uses to interpret a shadow. Evidence that object geometry interacts with color and luminance to affect depth perception was provided in recent mobile AR research by Do et al., [16], although the exact influence of these depth cues remains unclear. We also use the same background surface (a medium blue color) across all conditions in our study. Given that prior research has shown that background information can influence shadow perception [31, 56], it will be important to consider different backgrounds for testing surface contact perception in future work. It will also be important to consider evaluating the effects of shadows at farther distances than those used in the current study, since shadows can vary in usefulness across distances [13].

In addition, because we use one display to represent each type of immersive reality device (e.g., the Microsoft HoloLens was evaluated to better understand optical see-through displays), we are limited in our ability to generalize to other displays within the same family of devices. For example, we might find different results if we were to evaluate ground contact perception in another OST AR HMD like the Magic Leap One. Even though both devices are optical see-through displays, the manner in which they render images differs since they rely on unique hardware and software infrastructures. For example, in the current study, the Microsoft HoloLens has a substantially higher pixel density than the HTC Vive Pro or the HTC Vive Pro in conjunction with the Zed Mini. Although prior research has not shown the quality of computer graphics to have an impact on depth perception in virtual reality [62]—it is possible that this may not be the case in AR displays. Likewise, the video stream from the Zed mini is offset along the gaze axis from the HoloLens [57], and we did not test whether this affected perception to a significant degree.

The challenge of direct evaluations between displays can even be seen in the variance of results found in distance perception studies in virtual reality [9], which is a comparatively more mature technology. However, we believe that the current work as well as other prior work that evaluates spatial perception over different display types are essential for identifying salient underlying factors such as the effect of color contrast of shadows seen in the current study.

7 CONCLUSIONS

Our results provide evidence that photorealistic shadows are not essential for creating surface contact in either augmented reality or virtual reality. This is perhaps a counter-intuitive finding, since researchers have previously argued that consistent depth cue information—like the depth information provided by a cast shadow—is essential for accurate spatial perception. However, beyond research, we can see applications of this idea already in practice in other forms of media. Photographers use colored light and shadows for a stylish effect. And in the absence of sufficient lighting and depth information, video game developers often rely on colorful visual markers to indicate where characters or items of interest are located in space. Our findings are especially promising for optical see-through and other additive light displays. The Microsoft HoloLens' inability to render black or subtract light makes it nearly impossible to render realistic shadows, which can be defined as the absence of light. For this family of devices, our research presents an opportunity: the opportunity to enhance surface contact perception in XR without computationally expensive, photorealistic shadow rendering.

Although our current evaluation provides strong evidence for the influence of shadow shading method on one's sense of ground contact in immersive displays, future research would benefit from the evaluation of different target object geometries to see if these effects replicate in the presence of more complex stimuli. In addition, the perception of surface contact has been shown to influence depth perception [52]; future work should evaluate the extent to which different shadow shading methods influence absolute measures of depth perception in immersive technologies across a range of distances. It may also be of interest to evaluate how stylized graphics affect the ability of immersive AR to elicit emotional responses, such as conservative action judgments in the face of virtual danger seen in Wu et al. [66], if non-photorealistic graphics are to be used in commercial and entertainment applications.

ACKNOWLEDGMENTS

We thank Sarah Miele for her insightful guidance and for inspiring us with her dauntless perseverance throughout this project. This material is based upon work supported by the Office of Naval Research under grant N00014-18-1-2964.

REFERENCES

[1] H. Adams, "Resolving cue conflicts in augmented reality," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 2020, pp. 547–548.

[2] J.-g. Ahn, E. Ahn, S. Min, H. Choi, H. Kim, and G. J. Kim, "Size perception of augmented objects by different ar displays," in *HCI International 2019 - Posters*, C. Stephanidis, Ed. Cham: Springer International Publishing, 2019, pp. 337–344.

[3] G. Ballestin, F. Solari, and M. Chessa, "Perception and action in peripersonal space: A comparison between video and optical see-through augmented reality devices," in *2018 IEEE International Symposium on Mixed*

and Augmented Reality Adjunct (ISMAR-Adjunct), 2018, pp. 184–189.

[4] J. Barreira, M. Bessa, L. Barbosa, and L. Magalhães, “A context-aware method for authentically simulating outdoors shadows for mobile augmented reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 3, pp. 1223–1231, 2018.

[5] D. Bates, “lme4: Linear mixed-effects models using s4 classes. r package version 0.999375-34,” *http://cran. r-project. org/package= lme4*, 2010.

[6] M. Berning, D. Kleinert, T. Riedel, and M. Beigl, “A study of depth perception in hand-held augmented reality using autostereoscopic displays,” in *2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2014, pp. 93–98.

[7] Z. Bian, M. L. Braunstein, and G. J. Andersen, “The ground dominance effect in the perception of 3-d layout,” *Perception & Psychophysics*, vol. 67, no. 5, pp. 802–815, 2005.

[8] S. Bruckner and E. Gröller, “Enhancing depth-perception with flexible volumetric halos,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1344–1351, Nov 2007.

[9] L. E. Buck, M. K. Young, and B. Bodenheimer, “A comparison of distance estimation in hmd-based virtual environments with different hmd-based conditions,” *ACM Trans. Appl. Percept.*, vol. 15, no. 3, pp. 21:1–21:15, Jul. 2018. [Online]. Available: http://doi.acm.org/10.1145/3196885

[10] R. Casati and P. Cavanagh, *The visual world of shadows*. MIT Press, 2019.

[11] P. Cavanagh and Y. G. Leclerc, “Shape from shadows.” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 15, no. 1, p. 3, 1989.

[12] M. A. Cidota, R. M. S. Clifford, S. G. Lukosch, and M. Billinghurst, “Using visual effects to facilitate depth perception for spatial tasks in virtual and augmented reality,” in *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*, 2016, pp. 172–177.

[13] J. E. Cutting and P. M. Vishton, “Perceiving layout and knowing distance: The integration, relative potency and contextual use of different information about depth,” in *Perception of Space and Motion*, W. Epstein and S. Rogers, Eds. New York: Academic Press, 1995, pp. 69–117.

[14] A. Dey, A. Cunningham, and C. Sandor, “Evaluating depth perception of photorealistic mixed reality visualizations for occluded objects in outdoor environments,” in *2010 IEEE Symposium on 3D User Interfaces (3DUI)*, ser. VRST ’10. New York, NY, USA: Association for Computing Machinery, 2010, p. 211–218. [Online]. Available: https://doi.org/10.1145/1889863.1889911

[15] C. Diaz, M. Walker, D. A. Szafir, and D. Szafir, “Designing for depth perceptions in augmented reality,” in *2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Oct 2017, pp. 111–122.

[16] T. Do, J. J. LaViola Jr., and R. McMahan, “The effects of object shape, fidelity, color, and luminance on depth perception in handheld mobile augmented reality,” in *2020 IEEE International Symposium on Mixed and Augmented Reality(ISMAR)*, 2020.

[17] D. Drascic and P. Milgram, “Perceptual issues in augmented reality,” in *Stereoscopic Displays and Virtual Reality Systems III*, M. T. Bolas, S. S. Fisher, M. T. Bolas, S. S. Fisher, and J. O. Merritt, Eds., vol. 2653, International Society for Optics and Photonics. SPIE, 1996, pp. 123 – 134. [Online]. Available: https://doi.org/10.1117/12.237425

[18] J. H. Elder, S. Trithart, G. Pintilie, and D. MacLean, “Rapid processing of cast and attached shadows,” *Perception*, vol. 33, no. 11, pp. 1319–1338, 2004.

[19] S. Epstein, “The stability of behavior: Ii. implications for psychological research,” *American Psychologist*, vol. 35, no. 9, pp. 790–806, 09 1980.

[20] M. O. Ernst and H. H. Bülthoff, “Merging the senses into a robust percept,” *Trends in cognitive sciences*, vol. 8, no. 4, pp. 162–169, 2004.

[21] B. Farell and D. G. Pelli, “Psychophysical methods, or how to measure a threshold and why,” *Vision research: A practical guide to laboratory methods*, vol. 5, pp. 129–136, 1999.

[22] Y. Gao, E. Peillard, J.-M. Normand, G. Moreau, Y. Liu, and Y. Wang, “Influence of virtual objects’ shadows and lighting coherence on distance perception in optical see-through augmented reality,” *Journal of the Society for Information Display*, 2019.

[23] J. J. Gibson, *The perception of the visual world*. Westport, Conn.: Greenwood Press, 1950.

[24] ——, *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin, 1979.

[25] ——, *The implications of experiments on the perception of space and motion*. US Department of Commerce, National Technical Information Service, 1975.

[26] D. M. Green, J. A. Swets *et al.*, *Signal detection theory and psychophysics*. Wiley New York, 1966, vol. 1.

[27] L. Gruber, J. Ventura, and D. Schmalstieg, “Image-space illumination for augmented reality in dynamic environments,” in *2015 IEEE Virtual Reality (VR)*, 2015, pp. 127–134.

[28] C. R. Harrison and K. M. Robinette, “Caesar: Summary statistics for the adult population (ages 18-65) of the united states of america,” Air Force Research Lab Wright-Patterson AFB OH Human Effectiveness Directorate, Tech. Rep., 2002.

[29] H. H. Hu, A. A. Gooch, W. B. Thompson, B. E. Smits, J. J. Rieser, and P. Shirley, “Visual cues for imminent object contact in realistic virtual environment,” in *IEEE Visualization*, 2000, pp. 179–185.

[30] R. Ihaka and R. Gentleman, “Development core team (2009). r: A language and environment for statistical computing. r foundation for statistical computing, vienna, austria,” *URL http://www. R-project. org.(Accessed 16 October 2009)*, 1996.

[31] S. Ikeda, Y. Kimura, S. Manabe, A. Kimura, and F. Shibata, “Shadow induction on optical see-through head-mounted displays,” *Computers & Graphics*, vol. 91, pp. 141 – 152, 2020.

[32] T. Imura, N. Shirai, M. Tomonaga, M. K. Yamaguchi, and A. Yagi, “Asymmetry in the perception of motion in depth induced by moving cast shadows,” *Journal of*

Vision, vol. 8, no. 13, pp. 10–10, 2008.

[33] Y. Itoh, T. Langlotz, D. Iwai, K. Kiyokawa, and T. Amano, "Light attenuation display: Subtractive see-through near-eye display via spatial color filtering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 5, pp. 1951–1960, 2019.

[34] R. A. Jacobs, "What determines visual cue reliability?" *Trends in Cognitive Sciences*, vol. 6, no. 8, pp. 345 – 350, 2002.

[35] C. Jacquemin, G. Gagneré, and B. Lahoz, "Shedding light on shadow: Real-time interactive artworks based on cast shadows or silhouettes," in *Proceedings of the 19th ACM International Conference on Multimedia*, ser. MM '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 173–182. [Online]. Available: https://doi.org/10.1145/2072298.2072322

[36] T. Kaminokado, Y. Hiroi, and Y. Itoh, "Stainedview: Variable-intensity light-attenuation display with cascaded spatial color filtering for improved color fidelity," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 12, pp. 3576–3586, 2020.

[37] S. Kasahara, S. Higa, and A. Komori, "Fragment shadow: Generating fragmented shadows with multi-projectors geometry and color calibration," in *ACM SIGGRAPH 2019 Studio*, ser. SIGGRAPH '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: https://doi.org/10.1145/3306306.3328003

[38] H. Kato, T. Naemura, and H. Harashima, "Textured shadow," in *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.*, 2003, pp. 352–353.

[39] D. Kersten, P. Mamassian, and D. C. Knill, "Moving cast shadows induce apparent motion in depth," *Perception*, vol. 26, no. 2, pp. 171–192, 1997.

[40] E. Kruijff, J. E. Swan, and S. Feiner, "Perceptual issues in augmented reality revisited," in *2010 IEEE International Symposium on Mixed and Augmented Reality*, 2010, pp. 3–12.

[41] M. Kytö, A. Mäkinen, T. Tossavainen, and P. T. Oittinen, "Stereoscopic depth perception in video see-through augmented reality within action space," *Journal of Electronic Imaging*, vol. 23, no. 1, pp. 1 – 11, 2014. [Online]. Available: https://doi.org/10.1117/1.JEI.23.1.011006

[42] M. S. Landy, L. T. Maloney, E. B. Johnston, and M. Young, "Measurement and modeling of depth cue combination: In defense of weak fusion," *Vision research*, vol. 35, no. 3, pp. 389–412, 1995.

[43] N. Liu and M. Pang, "A survey of shadow rendering algorithms: Projection shadows and shadow volumes," in *2009 Second International Workshop on Computer Science and Engineering*, vol. 1, Oct 2009, pp. 488–492.

[44] C. Madison, W. Thompson, D. Kersten, P. Shirley, and B. Smits, "Use of interreflection and shadow for surface contact," *Perception & Psychophysics*, vol. 63, no. 2, pp. 187–194, 2001.

[45] P. Mamassian, D. C. Knill, and D. Kersten, "The perception of cast shadows," *Trends in cognitive sciences*, vol. 2, no. 8, pp. 288–295, 1998.

[46] S. Manabe, S. Ikeda, A. Kimura, and F. Shibata, "Casting virtual shadows based on brightness induction for optical see-through displays," in *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2018, pp. 627–628.

[47] ——, "Shadow inducers: Inconspicuous highlights for casting virtual shadows on ost-hmds," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, March 2019, pp. 1331–1332.

[48] D. Medeiros, M. Sousa, D. Mendes, A. Raposo, and J. Jorge, "Perceiving depth: optical versus video see-through," in *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*, 2016, pp. 237–240.

[49] J. Meng and H. Sedgwick, "Distance perception across spatial discontinuities," *Perception & Psychophysics*, vol. 64, no. 1, pp. 1–14, 2002.

[50] J. C. Meng and H. Sedgwick, "Distance perception mediated through nested contact relations among surfaces," *Perception & Psychophysics*, vol. 63, no. 1, pp. 1–15, 2001.

[51] Y. Minomo, Y. Kakehi, and M. Iida, "Transforming your shadow into colorful visual media: Multi-projection of complementary colors," in *Proceedings of the 2005 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*, ser. ACE '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 61–68. [Online]. Available: https://doi.org/10.1145/1178477.1178485

[52] R. Ni, M. L. Braunstein, and G. J. Andersen, "Perception of scene layout from optical contact, shadows, and motion," *Perception*, vol. 33, no. 11, pp. 1305–1318, 2004, pMID: 15693673. [Online]. Available: https://doi.org/10.1068/p5288

[53] R. Ni, M. Braunstein, and G. Andersen, "Distance perception from motion parallax and ground contact," *Visual Cognition*, vol. 12, no. 6, pp. 1235–1254, Aug. 2005.

[54] K. Ozkan and M. L. Braunstein, "Background surface and horizon effects in the perception of relative size and distance," *Visual cognition*, vol. 18, no. 2, pp. 229–254, 2010.

[55] J. Ping, B. H. Thomas, J. Baumeister, J. Guo, D. Weng, and Y. Liu, "Effects of shading model and opacity on depth perception in optical see-through augmented reality," *Journal of the Society for Information Display*, 2020.

[56] R. A. Rensink and P. Cavanagh, "The influence of cast shadows on visual search," *Perception*, vol. 33, no. 11, pp. 1339–1358, 2004.

[57] J. P. Rolland, F. A. Biocca, T. Barlow, and A. Kancherla, "Quantification of adaptation to virtual-eye location in see-thru head-mounted displays," in *Proceedings Virtual Reality Annual International Symposium'95*. IEEE, 1995, pp. 56–66.

[58] J. P. Rolland, R. L. Holloway, and H. Fuchs, "Comparison of optical and video see-through, head-mounted displays," in *Telemanipulator and Telepresence Technologies*, H. Das, Ed., vol. 2351, International Society for Optics and Photonics. SPIE, 1995, pp. 293 – 307. [Online]. Available: https://doi.org/10.1117/12.197322

[59] C. S. Rosales, G. Pointon, H. Adams, J. Stefanucci, S. Creem-Regehr, W. B. Thompson, and B. Bodenheimer, "Distance judgments to on- and off-ground objects in augmented reality," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, 2019, pp. 237–243.

[60] P. L. Smith and D. R. Little, "Small is beautiful: In defense of the small-n design," *Psychonomic bulletin & review*, vol. 25, no. 6, pp. 2083–2101, 2018.

[61] N. Sugano, H. Kato, and K. Tachibana, "The effects of shadow representation of virtual objects in augmented reality," in *Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality*, ser. ISMAR '03. USA: IEEE Computer Society, 2003, p. 76.

[62] W. B. Thompson, P. Willemsen, A. A. Gooch, S. H. Creem-Regehr, J. M. Loomis, and A. C. Beall, "Does the quality of the computer graphics matter when judging distances in visually immersive environments," *Presence: Teleoperators and Virtual Environments*, vol. 13, pp. 560–571, 2004.

[63] W. B. Thompson, P. Shirley, B. Smits, D. J. Kersten, and C. Madison, "Visual glue," *University of Utah Technical Report UUCS-98-007*, 1998.

[64] K. Vaziri, P. Liu, S. Aseeri, and V. Interrante, "Impact of visual and experiential realism on distance perception in vr using a custom video see-through system," in *Proceedings of the ACM Symposium on Applied Perception*, ser. SAP '17. New York, NY, USA: Association for Computing Machinery, 2017. [Online]. Available: https://doi.org/10.1145/3119881.3119892

[65] V. Šoltészová, D. Patel, and I. Viola, "Chromatic shadows for improved perception," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering*, ser. NPAR '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 105–116. [Online]. Available: https://doi.org/10.1145/2024676.2024694

[66] H. Wu, H. Adams, G. Pointon, J. Stefanucci, S. Creem-Regehr, and B. Bodenheimer, "Danger from the deep: A gap affordance study in augmented reality," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2019, pp. 1775–1779.

[67] J.-R. Wu, M.-L. Wang, K.-C. Liu, M.-H. Hu, and P.-Y. Lee, "Real-time advanced spinal surgery via visible patient model and augmented reality system," *Computer Methods and Programs in Biomedicine*, vol. 113, no. 3, pp. 869 – 881, 2014.

[68] A. Yonas, L. T. Goldsmith, and J. L. Hallstrom, "Development of sensitivity to information provided by cast shadows in pictures," *Perception*, vol. 7, no. 3, pp. 333–341, 1978.