

Fast, Accurate, But Sometimes Too-Compelling Support: The Impact of Imperfectly Automated Cues in an Augmented-Reality Head-Mounted Display on Visual Search Performance

Amelia C. Warden , Christopher D. Wickens , Daniel Rehberg , Francisco R. Ortega ,
and Benjamin A. Clegg 

Abstract—While the visual search for targets in a complex scene might benefit from using augmented-reality (AR) head-mounted display (HMD) technologies by helping to efficiently direct human attention, imperfectly reliable automation support could manifest in occasional errors. The current study examined the effectiveness of different HMD cues that might support visual search performance and their respective consequences following automation errors. A total of 56 participants searched a three-dimensional environment containing 48 objects in a room, in order to locate a target object that was viewed prior to each trial. They searched either unaided or assisted by one of the three HMD types of cues: an arrow pointing to the target, a plan-view minimap highlighting the target, and a constantly visible icon depicting the appearance of the target object. The cue was incorrect in 17% of the trials for one group of participants and 100% correct for the second group. Through both analysis and modeling of both search speed and accuracy, the results indicated that the arrow and minimap cues depicting location information were more effective than the icon cue depicting visual appearance, both overall, and when the cue was correct. However, there was a tradeoff on the infrequent occasions when the cue erred. The most effective AR-based cue led to a greater automation bias in which the cue was more often blindly followed without careful examination of the raw images. The results speak to the benefits of AR and the need to examine potential costs when AR-conveyed information may be incorrect because of imperfectly reliable systems.

Index Terms—Augmented reality (AR), head-mounted display (HMD), imperfect automation, visual attention, visual search.

I. INTRODUCTION

VISUAL search is ubiquitous in the workplace and everyday life [1], [2]. Examples include a driver or pilot searching their forward field of view (FOV) for potential collision hazards [3], the industrial worker or quality control inspector searching a product for flaws [4], [5], the radiologist searching medical imagery for a tumor or fracture [6], [7], the transportation security administration (TSA) inspector searching the X-ray image for a weapon [8], [9], [10], [11], [12], the computer user searching a cluttered screen for a specific file [13], or the warfighter searching a forward scene for a potentially hostile target [14], [15], [16] [17], [18]. When unaided, visual search can be time-consuming and error prone, with a common issue being missing true targets when they are present [19]. This high human performance cost results because complex visual search is often *serial*, effortful, and inefficient, involving the inspection of each item in the search field in turn until one is encountered that is suspected to be the target [20], [21].

Challenges in both speed and accuracy mean there is great potential value in the human performance of a search *cue* that can direct the user's attention to an object [2], [10], [18]. An object might be cued based on an automation or artificial intelligence (AI) inference that the particular object is likely to be a target. The current research addresses the intersection of human performance and computer technology in cued search in two respects: First, we examine the technology of augmented-reality (AR) cueing, and second, we consider the imperfections of automation (e.g., machine vision) that may render the AI-based cueing imperfectly reliable.

The specific contribution of this article is that we demonstrate the differential effectiveness of three different types of head-mounted display (HMD) cueing aids in a comparison that has not been previously reported. HMDs integrate small displays or projection technology into devices, such as eyeglasses, visors, or devices, that are mounted onto a helmet. One fundamental advantage of an HMD, specifically in the case where it displays

Manuscript received 17 January 2023; revised 21 June 2023; accepted 17 July 2023. Date of publication 20 September 2023; date of current version 14 December 2023. This work was supported by the Office of Naval Research under Grant N00014-21-1-2949 and Grant N00014-21-1-2580. This article was recommended by Associate Editor Y. Liu. (Corresponding author: Amelia C. Warden.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Institutional Review Board (IRB) at Colorado State University under Application No. 1818, and performed in line with the the ethical regulations (45 CFR 46) of the United States Department of Health and Human Services.

Amelia C. Warden and Christopher D. Wickens are with the Department of Psychology, Colorado State University, Fort Collins, CO 80523 USA (e-mail: acwarden@colostate.edu; chris.wickens@colostate.edu).

Benjamin A. Clegg is with the Department of Psychology at Montana State University, Bozeman, MT 59717 USA (e-mail: benjamin.clegg@montana.edu).

Daniel Rehberg and Francisco R. Ortega are with the Department of Computer Science, Colorado State University, Fort Collins, CO 80523 USA (e-mail: dan.rehberg@colostate.edu; f.ortega@colostate.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/THMS.2023.3302152>.

Digital Object Identifier 10.1109/THMS.2023.3302152

AR information (e.g., target cueing), is that it reduces the scanning between the display and the real world viewed beyond the display. This allows the system to visually designate a suspected target while the searcher looks directly at the search field in the real world beyond the display.

We provide an underlying psychological model to account for the differences in the effectiveness of the three cueing techniques. In particular, augmented-reality head-mounted display (AR-HMD) cueing based on imperfect automation has rarely been examined. In our examination, we provide a unique comparison revealing that the greater effectiveness of a more realistic, world-centered cue when it is correct may be offset by a greater cost on the infrequent occasions when it is wrong.

II. BACKGROUND

A. Perfect Cueing

In theory, an unaided search through a series of objects to find a target that is similar in appearance to the nontarget objects will take, on average, a search time (ST) of

$$ST = a + \frac{nT}{2} \quad (1)$$

where a is a constant associated with making the response when a target is located but, importantly, it can also involve a time-consuming target confirmation process; n is the number of nontarget objects in the search field; and T is the time to inspect each object and determine that it is *not* a target. The division by 2 represents the fact that, on average, the target will be found after half of the objects have been inspected. This is the *serial self-terminating search* model [20], which has been confirmed in multiple experiments over the past 50 years (see [2] for review). There are multiple exceptions and qualifications to this pure serial search model [21], [22] but, in many applied instances, it provides a reasonable approximation to predicting the ST [23].

Thus, a perfect cue may be able to drastically shorten this time from $(a + nT/2)$ to $(a + T)$, as the single cued object is inspected to confirm that it is the target. A study in our lab confirmed a reduction in ST for 3-D objects located around a room from 10 to 5 seconds (s), associated with one type of perfect cueing [24]. Here, we can calculate that 4.6 s represented the constant a in (1), with $T = 0.29$ s. In hostile, or dynamic environments, saving a few seconds can be of vital importance. Consider, for example, cueing the location of a potential collision hazard while driving fast or flying [25]. Such is an example of one of the many aspects of Wolfe's [21], [22] guided search model, which can shorten ST, where the cue guides attention to a single candidate for a target.

Naturally, the benefits of the cue in reducing ST will also be based on what we label the *cue effectiveness value* (CEV), which accounts for the time required to see and interpret the cue and orient attention from the cue to the cued target [2], [26], [27], [28], [29], an issue of direct interest in the current study. The CEV can be predicted and modeled by considering the human information processing operations involved in the search. One important influence on these operations is the extent to which

the cue is considered *exogenous* or *endogenous* [2], [28]. An endogenous cue presents information about the cued location symbolically but does not appear at the location of the cue itself. For example, an icon cue conveys the appearance of a target but does not directly mark the location where attention should be shifted. Conversely, an exogenous cue orients the user to the location where attention should be shifted, such as a flashing highlight (presented on an HMD) near an object in the world that automatically orients attention toward the location of a target in the search scene.

Basic visual attention research reveals that reliable exogenous (egocentric) cueing is more rapidly processed (greater CEV, smaller CEV penalty) than endogenous cueing. However, while the endogenous–exogenous distinction is widely used in basic research on visual attention, exogenous cues are only effective if the target is in the forward eye field and runs the risk of capturing attention when people are not currently searching for that specific target. In the current research, we have employed more application-relevant terms of *world-referenced* (egocentric) versus *display-referenced* (exocentric) cueing because this maps directly onto two philosophies of HMD cueing, the focus of our investigation.

In HMD design, world-referenced cueing capitalizes on the properties of AR to position and reposition a cue on an HMD such that it always overlays or points directly toward the target, irrespective of how the head is oriented. That is, the cue is presented on the display in world-referenced coordinates that are rapidly updated on the display as the head rotates laterally or vertically. This is one of the cue types to be evaluated in the current study in which an AR *arrow* cue is programmed to always point toward the location of the target.

Screen-referenced exocentric or display-referenced cues are presented at the same fixed screen coordinates independent of head rotation. They can nevertheless designate in some less direct fashion, the object that automation infers to be the target, either by specifying its location or its visual appearance. In our research, a location cue highlights the inferred target, as depicted on a *minimap*, positioned in screen-referenced coordinates on the HMD [24], [30]. The visual appearance cue, which we call the *icon cue*, presents a picture (icon) of the inferred target, continuously visually available, again presented at fixed screen-referenced coordinates. Both of these cue types impose CEV penalties compared with the AR world-referenced arrow cue, but the penalties are of different varieties. For the minimap cue, the penalty is one of the spatial frame-of-reference transformations (FORTs) necessary to determine where the cued target viewed on the minimap (typically a top-down planar display) is located on the forward ego-referenced 3-D view of the search field [31], [32]. For the icon cue, the penalty is inherent in the increase in N because of many objects in the search field that still need to be sequentially examined to determine which one matches the icon. However, some shortening of ST is still preserved here to the extent that salient features of the target icon (e.g., a green color or large size) can discriminate it from many, although not all, of the nontarget objects in the search field. That is, some salient property of the target on any particular trial may help guide the search through a reduced number of items [21], [22].

In summary, comparing the three cue types, their ST, and benefit over the control condition (1) is predicted to be

$$\text{AR Arrow} : (a + T). \quad (2)$$

Here, there is only a single item to be searched for and a single confirmation. Because the cue unambiguously points to the inferred target, CEV is maximum (i.e., having no penalty)

$$\text{Minimap} : (a + T) + (\text{FORT CEV penalty}). \quad (3)$$

This is identical to the predicted ST for the arrow cue, except there is now a penalty to the CEV related to the FORT of mental rotation from the top-down plan-view minimap to the forward 3-D perspective view of the search field

$$\text{Icon} : a' + \frac{((n - x)T)}{2}. \quad (4)$$

The icon cue differs from the arrow cue in two ways. First, as noted above, relative to the control condition, ST may be reduced if certain salient features of some targets reduce the search field from n by a quantity x , as predicted by the guided search model [21]. Second, as noted in the context of the uncued search (1), in all of the equations, the constant a actually contains two subcomponents. One is the time to actually execute the target identification response, which is identical for all four conditions. The other nontrivial component is the time required to confirm that an initially identified candidate is actually the target. This will involve, as participants are requested to do, a close comparison between the candidate target seen in the real world through the HMD and the representation of the target, either visually on the HMD (the icon) or in visual memory (the minimap or arrow). This confirmation time is likely to differ between the arrow and minimap on the one hand (a) and the icon (a') on the other hand. This difference results because, for the two location cues, the comparison must be made with a degraded remembered image due to the passage of time of the true target; whereas for the icon, the image remains directly visible and so invites a careful comparison of the image features with the suspected target in the real world. This may take more time, even if it also improves the accuracy. While the current experiment does not contain the necessary controls to obtain precise estimates of all these terms, it is possible to gauge rough estimates.

Importantly, in an experiment, examining these three types of cues [24], results showed cueing benefits of 5, 4, and 1 s for the arrow, minimap, and icon cues, respectively, indicating that the CEV of location cues (i.e., arrow and minimap) was far more effective than that of visual appearance (i.e., icon). However, the icon cue also provides a benefit to search accuracy not conferred by the other types of cues: that is, the icon cue allows the searcher the capability to confirm with certainty the identity of the target, rather than relying upon the potentially degraded memory of what the target looked like. Thus, Warden et al. [24] found that the accuracy of CEV was just as great (statistically) for the icon cue as for the arrow or minimap cue.

B. Imperfect Cueing and Automation

In [24], the cues were always correct. But in the operational world, a cue must be guided by an automation inference as to what should be the target. Such inference might often be guided by computer vision using machine learning that examines each object and, such as the human, compares the product of this examination with a template of what a true target looks like. As with so many other functions of automation, this inference may, on occasion, be incorrect [10], [11], [33], [34], [35], [36] due to various factors, such as lighting conditions, pattern complexity in the search scene, and the reliability of the automation machine-vision system. Such imperfectly reliable automation also has profound implications for the performance of the human-automation team (HAT), depending upon the degree of dependence or reliance of the human on the automation's decision/diagnosis of what is deemed important to be attended. Indeed, decades of research on human trust in imperfect automation have examined these factors [34], [35], [37], [38]. A smaller number of studies have examined it specifically in the context of visual search and target cueing (e.g., [8], [10], [11], [15], [30], and [36]). However, only one study [15] appears to have done so in the specific context of imperfect AR-HMD target cueing (although see also [30] for a related study).

Three general findings have emerged from the above research on human use of imperfect automation. First, there is a general tendency for humans to depend upon automation recommendation, whether this recommendation is of a course of action, a location to attend to, or a diagnostic state, without adequately examining the raw data upon which that recommendation is based [39]. This dependency is known as the *automation bias* [40], [41] and has the consequence that, on instances when automation fails, which occur with a frequency of $(1.0 - \text{reliability})$, humans are also likely to err in their automation-based judgment. In visual search tasks, such automation failures may be either failing to find a true target (automation miss) or falsely classifying a nontarget as a target (automation false alarm). Somewhat ironically, these failures become more likely, the more reliable automation is [42] (but see [10]), as humans may sometimes express what is called a "perfect automation schema," which is associated with high expectations of automation performance and less forgiving attitudes of automation failures than human failures [43], [44].

Second, as the reliability of automation does decrease, users are somewhat sensitive to this decrease (although more so to automation false alarms than misses [45]); their trust in automation is reduced and their dependence on automation is correspondingly reduced. Consequently, the extent of the automation bias is also reduced. Unlike the previous HMD cueing type comparison [24], the present study included *imperfect* automation cueing at an error rate of 17%, a value that may indeed be typical of many automation systems in use [46] but a level of imperfect reliability that is nevertheless helpful in aiding the HAT performance [47].

Third, there is some evidence that the effect of cue imperfections may interact with the type of cue employed, specifically that the more realistic "compelling" AR arrow cue, while providing greater assistance when it is correct, may lead to greater

human errors caused by the automation bias on the infrequent occasions when the cue is wrong. Some basis for this prediction is provided by the research examining imperfect AR target cueing [15]. In this study, the authors employed an imperfect (85% reliable) cue of a potential hostile enemy target to soldier participants either with a head-down hand-held display to signal possible target location or with an AR-HMD egocentric arrow cue. They observed that the latter was more effective when the cue was correct, but more problematic when the cue was wrong providing a false alarm (i.e., cueing a nontarget). It was as if the more immersive *world-referenced* cue amplified an automation bias.

This tendency aligns with the idea that the more immersive cue produces more *attentional tunneling* [14], [48] at the expense of examining the raw data in the natural world beyond the HMD. In support of this causal assumption, Yeh et al. [14], [15] also observed that the AR-HMD cueing led to a decrease in the detection of other noncued but high-priority threats in the scene. In the current study, we are predicting that all three automation cues will create some degree of automation bias, but that this will be amplified with the arrow cue, under the assumption that the more compelling immersive and easy-to-use display (mitigating FORTs) will amplify dependence on that cue.

This prediction is also consistent with a meta-analysis of the aircraft head-up display (HUD) [49], which revealed that the HUD provided benefits when HUD information was conformal (i.e., a one-to-one mapping between a spatial overlay of a real or virtual object and the far domain) with the world beyond—a concept that is in many ways analogous to AR displays. But that conformal HUD provided a unique cost to detecting events in the world beyond that were not depicted on the HUD; a form of automation failure. In a military study, Wickens et al. [50] found that the more immersive and realistic 3-D display, as with the AR-HMD used by Warden et al. [24], improved performance on tasks depending on the information in the current FOV; but Wickens et al. [14], [50] found that this immersive 3-D display created a sort of attentional tunneling, which degraded detection of events outside of that forward view. Mifsud et al. [30] found that errors in a directional cue presented within an AR-HMD were frequently undetected by the human participants. Again, this is as if this human performance bias was amplified by AR (see also [51]). Neither of these studies nor the prior HUD meta-analysis directly replicate the manipulations or displays used in the current study, but they do point toward certain predicted findings regarding the potential costs of high display realism [51] served by imperfect information.

C. Cue Location and Search Strategy

In addition to the effectiveness of cue type, a second issue examined in the current study is the effect of cue location within the visual field. Of course, the fundamental premise of using an HMD for attentional cueing is to reduce the information access effort required to move attention between the far domain and the displayed cue location [52]. While a head-down display, such as a tablet, was not employed for comparison in the current study (see for example [14]), our study did include a manipulation of cue location on the HMD. We hypothesized a benefit for

placing the screen-referenced cue directly in the center of the HMD FOV compared with positioning it at the bottom because of the reduced information access effort at the former location, as observed in our perfect cueing study [24]. However, we also recognize that any such advantage could be partially or fully offset by two factors.

- 1) The increased clutter that results when a screen-referenced cue (icon or minimap) is directly superimposed on the foveal view of the outside world where all search objects can be seen when the direction of gaze is straight ahead. This offset is expected to grow as the visual complexity of the cue increases and be larger with the more visually complex minimap, a finding observed in the perfect cueing study [24].
- 2) The particular *search strategy* employed: A strategy that proceeds from the top down will likely produce a relatively greater clutter-induced compromise of target detection in the lower search field when the cue is presented downward. Performance using a strategy that proceeds from the center outward on the other hand will be little affected by vertical cue location. We have a particular interest in the basic search strategy employed in this paradigm in the control condition when there is no cue present.

Thus, the current experiment examines the search for 3-D target objects in a room, similar to that employed by Warden et al. [24], where participants were either unaided or assisted by an HMD cue of either of the three types described above. For half the participants, the cueing was imperfect, where in 17% of the trials, an object was cued that was not the target that they had studied prior to the start of the search. Based upon our review of the literature above, we hypothesized the following.

- 1) H_1 : Any cueing would provide a CEV benefit to ST as measured by response time (RT) relative to the control condition, given the long history of human factors research on cueing effectiveness.
- 2) H_2 : This RT benefit would be maximum for the AR arrow cue [24], reduced for the minimap cue because of its spatial transformation requirement that reduces CEV, and least for the icon cue, which relies upon visual appearance rather than location information and forces some degree of time-consuming serial search, as described above. This is based upon both the prior findings of [24] and the information processing analysis of each different cue type described above in (2)–(4).
- 3) H_3 : Cueing would provide a CEV benefit for accuracy relative to the control condition with no cue, based on the history of cueing research (see [24]).
- 4) H_4 : The profile of accuracy benefits across cue types would differ from that shown by RT because of the improved accuracy CEV for the icon cue, which allows memory-free confirmation of the cued target. This hypothesis is based on the known frailties of human working memory when a current image must be compared against a remembered image, which is the case with the arrow and minimap cues.
- 5) H_5 : Imperfect cueing automation would impair the overall search performance because of the contribution of high

error human rate in those 17% of the trials when the cue was wrong [15]. This is based on the large database of research on the effects of imperfect automation on HAT team performance.

- 6) H_6 : This impairment would be proportional to the benefit offered when the cue was correct. This is based on the corresponding finding of the other AR-HMD cueing study available in the literature [15], [30].
- 7) H_7 : Positioning a cue in the center of the HMD FOV rather than downward will generally improve cueing effectiveness because of reduced information access effort, particularly for the visually simple arrow cue whose clutter is low. Ample experimental evidence, as reviewed in [52], predicts some costs in performance when two sources of information that need to be integrated are displayed further apart.

III. EXPERIMENT

A. Method

Participants: A total of 57 participants in an introductory psychology course at Colorado State University received course credit in exchange for completing the experiment. The experiment was approved by the Institutional Review Board (IRB). Before starting the experiment, participants gave their informed consent and completed an online version of the Ishihara color blindness test to rule out any unknown red–green deficiencies. Previous work indicates that online color blindness tests are valid measures for screening for general color deficiencies [53], [54]. All participants had normal or corrected-to-normal vision, and all passed the Ishihara electronic color blindness test.

Apparatus and Stimuli: Participants completed the experiment using a HoloLens 2, an AR-HMD. The HoloLens 2 is an AR optical see-through headset that overlays virtual content onto the real world. The FOV of the device is 43° in the lateral direction and 29° in the vertical direction. The cueing aids used during the visual search task were created in Unity, version 2019.3.1f1. We systematically constructed 48 real-world objects using multicolored megablocks. Each object consisted of five blocks pointing in the $\pm x$, $\pm y$, and $+z$ planes and varying in color and shape. In addition, we created eight objects that were not in the search scene but were similar to eight objects in the search scene (foils). These foil objects differed from their real-world counterpart in two dimensions (e.g., color and shape). These foil objects were used for the imperfect reliable condition only.

Task: Participants completed a 180° visual search task using an AR-HMD. Participants were seated in a stationary chair in front of the search field (see Fig. 1) within which they were to search for a designated target on each trial. In all conditions, the HMD displayed an image of the target object to be searched at the beginning of each trial. Participants had a total of 5 s to study the image until the search began. Once the trial began, they scanned the search field for the designated object. Once they felt that they had located the target, they fixated directly on the object. A gaze-sensitive algorithm then rendered a box around the target. To make their selection, they pressed the “A” button




Fig. 1. Image of the search scene depicted here from a top-down perspective (top) and a forward-facing perspective (bottom). Objects were placed at three heights (shelf, table, and floor) within a 180° search field.



Fig. 2. Example images of each cue with the hit boxes surrounding the object. Hit boxes denote when the participant’s gaze was on that specific object.

on an X-Box controller. If the item they selected was the correct item (e.g., the target image at the beginning of the trial matched the object selected in the real world), the response was recorded correctly. Otherwise, the response was recorded as incorrect. For the imperfectly reliable condition, participants could select the practice block object if they believed the target object at the beginning was not present. This response was recorded as correct.

Participants completed four cueing conditions during the experiment. As shown in Fig. 2, participants were given three types of cueing aids: an arrow cue, a minimap cue, and an icon cue. In addition, they completed a no-cue (control) condition. The arrow cue was a green 2-D AR world-referenced arrow that pointed in the direction of the target as the participant turned their head. The arrow cue turned into a green circle once the participants’ gaze was on the designated target (see left panel of Fig. 2). The minimap cue used screen-referenced coordinates that continuously displayed the position of all objects in the

search scene from a top-down bird's eye perspective. The target object's location was indicated with a yellow box around the object, as can be seen in the center panel of Fig. 2. The icon cue continuously displayed the visual appearance of the target object in screen-referenced coordinates on the HMD, which can be seen in the black box to the lower left of the right panel in Fig. 2. Note that the general similarity of the cued target icon with various search objects (three of which are seen on the ledge above) requires careful comparison to judge which search object matches the icon.

Cues were positioned in either the center or 12.6° downward from the center of the AR-HMD display: this display location represents the cue-location variable in the experiment. As shown in Fig. 1, the objects in the real world were distributed across the 180° search field and located at three different height locations (floor, table, and shelf) that were separated by 28 inches the vertical direction. All objects were deemed to be equally visible based on pilot testing. A total of 48 objects were uniformly distributed across two halves of the room. The target object could appear at any location within this 3-D array.

Design: The experiment was a 2 (reliability) \times 3 (cue) \times 2 (cue location) \times 3 (object vertical location) mixed-subject design, with the reliability as the between-subject variable and cue type, cue location, and object location as the within-subject variables. A fourth no-cue condition was included, but, by definition, this did not include cue location as a variable.

Participants were randomly assigned to one of the two reliability groups. One group of participants ($N = 28$) completed the 100% reliability condition for which the image of the target object always matched the object in the real world where the cueing aid was located. To encourage participants not to rely totally on the cue, the instructions indicated that *“The cue is always correct. However, before selecting the object that is identified by the cue, you should try to assure that the selected object is, in fact, the same as the one that you saw at the beginning of the trial. If you do not see the object in the search field, you can select the blue practice object block to indicate ‘object not present.’”*

The second group of participants ($N = 28$) completed the 83% reliability condition for which the image of the target object occasionally did not match the real-world object that the cueing aid designated (i.e., the cueing aid located the incorrect object 17% of the time). On the incorrect trials, the cue instead designated the foil; that is, the object that was matched to look similar, but not identical, to the object shown at the beginning of the search. Such similarity captured the kind of error that a machine-learning system might make. Participants in this condition were explicitly instructed that *“The cueing aids may NOT BE PERFECTLY RELIABLE, meaning you may not fully rely on them. If you do not see the object in the search field, you can select the blue practice object block to indicate ‘object not present’.”* For the arrow and minimap cues, the spatial properties of the cue designated the object (foil) that looked like the originally viewed target. For the icon cue, the icon representation of the target object remained on the HMD, but, as with the two spatial cueing conditions, that object was not present in the search field, being replaced instead by a similarly looking object.

Participants completed 12 trials for each cue condition in a blocked, counterbalanced order; six trials when the cue was located in the center of the display and six trials when the cue was located 12.6° downward from the center. These two cueing locations were blocked, and the cue conditions were counterbalanced. Participants completed 12 trials for the no-cue (control) condition. The entire experiment consisted of 48 trials. Four different sets of trials were randomly generated such that each cue type was presented at one of the three heights on the left and right sides of the room. For example, on the left side of the room, three trials of the 2-D arrow in the center location of the display would cue the participant to three objects located at the shelf, table, and floor heights, each at one of the three regions in the lateral direction of the left side of the room (see Fig. 1). Of the eight incorrect target images (two for each cue condition), each was paired with a real-world object that was similar but not identical. These foils were presented at the three vertical levels on the two sides of the room based on the locations within the four different sets of trials described above.

Procedure: Participants completed the institutional review board (IRB) consent form, and then completed seven practice trials to familiarize themselves with each cue condition. They completed two trials for each of the three cueing aids (arrow, minimap, and icon): one with the cue in the center of the display and one with the cue 12.6° downward from the center. They completed only one practice trial for the no-cue condition. They searched for the same practice block consisting of only blue megablocks for all practice trials. Next, they completed the test trials where they searched for 48 different target objects for each cueing condition. All participants in both groups were instructed to carry out the task as rapidly and accurately as possible. Participants did not receive feedback. After completing the visual search task, they took a Qualtrics survey to assess which cue they thought was the hardest, easiest, most helpful, and least helpful. We also collected demographic information (e.g., age, gender, AR/VR experience, etc.).

IV. RESULTS

Outlier criteria were based on whether the accuracy was at or below chance performance for all cueing conditions. Of the 57 participants, one participant's data were removed from the analysis. First, we analyzed the data in R using 3 (cue type) \times 2 (reliability level) \times 2 (cue location) mixed ANOVAs to examine the effect of cue type (arrow, minimal, icon, and no cue), cue reliability level (100% and 83%), and location (up versus down) on performance. Cue type was submitted as the within-subject factor and reliability level as the between-subject factor. We log-transformed RT data because the assumptions of normality and homogeneity were violated.

A. Overall Cueing Effectiveness

RT: The effects of cue type and reliability level on RT (measured in s) are shown in Fig. 3. To ease the interpretation, we plotted raw RT rather than the log-transformed RT. The single value of mean RT in the control (no cue) condition is plotted

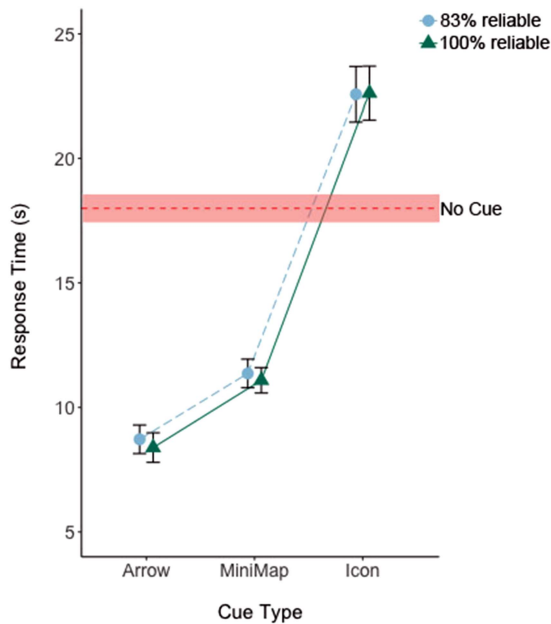


Fig. 3. Mean RT plotted as a function of cue type and reliability. Light blue dashed and green solid lines represent 83% and 100% reliability, respectively. Error bars are the SEM. The horizontal dashed red line represents the mean RT for the no-cue (control) condition. The shaded region indicates the SEM.

as the red dashed horizontal line, and the red shaded region represents 1 standard error of the mean (SEM).

To examine H_1 , the overall benefits of cueing to RT and separate planned comparisons were made between the no-cue and the three cueing conditions. These revealed significant benefits of both the arrow cue ($t(55) = -8.69, p < .001$, and $d = 1.45$) and the minimap cue ($t(55) = -5.66, p < .001$, and $d = 0.87$) with large effect sizes, but participants were actually slower than the uncued condition in finding the target with the icon cue ($t(55) = 3.73, p = .0005$, and $d = 0.44$). Thus, H_1 was supported in two of the three cases.

To examine H_2 , specifically regarding the differences in cue effectiveness for RT, a 3 (cue type) \times 2 (reliability) mixed analysis of variance (ANOVA) was carried out on the log-transformed RT. The assumption of sphericity was violated; therefore, Greenhouse–Geisser corrections for log-transformed RT data are reported here. This ANOVA revealed a significant, large effect of cue type on log-transformed RT, $F(1.52, 82.13) = 105.89, p < .001$, and $\eta_p^2 = 0.66$. As suggested by the similarity of the two curves in Fig. 3, the reliability level did not impact log-transformed RT, $F(1, 54) = 0.05, p = .82$, and $\eta_p^2 < 0.1$. The interaction was also not significant ($p = .42$). Specific contrasts revealed, as shown in Fig. 3, that all three cueing conditions differed from each other, with the arrow supporting the fastest performance, and the icon supporting the slowest performance. As the figure suggests, the minimap was also much more effective than the icon ($t(55) = 11.01, p < .001$, and $d = 1.40$), although still significantly slower than the arrow cue ($t(55) = -6.30, p < .001$, and $d = 0.82$). Thus, these results fully support the ordering of cueing effectiveness on search speed proposed in H_2 .

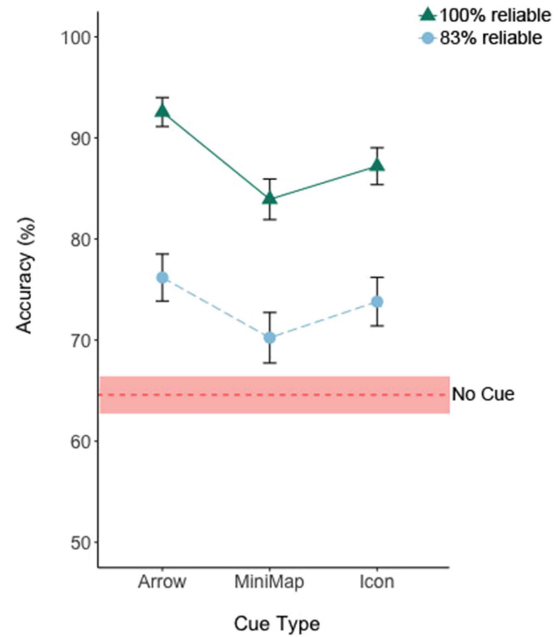


Fig. 4. Mean accuracy plotted as a function of cue type and reliability. Green solid and light blue dashed lines represent 100% and 83% reliability, respectively. Error bars represent the SEM. The horizontal dashed red line represents the mean accuracy for the no-cue (control) condition. The shaded region indicates the SEM.

Accuracy: Fig. 4 presents the accuracy data in an identical format to the RT data of Fig. 3. In contrast to the RT data, these data indicate that all six cueing conditions, whether perfectly reliable or not, improved accuracy over the no-cue (control) condition, supporting H_3 .

The 3 (cue type) \times 2 (reliability) mixed ANOVA conducted on the accuracy data revealed a significant effect of cue type, $F(2, 108) = 4.83, p < .01$, and $\eta_p^2 = 0.08$. In addition, there was a significant and large cost of imperfect reliability level on accuracy, $F(1, 54) = 27.51, p < .001$, and $\eta_p^2 = 0.34$, discussed below. There was no significant interaction between cue type and reliability level, $F(2, 108) = 0.24, p = .78$, and $\eta_p^2 < 0.01$.

The analysis revealed that, as with the RT, the arrow cue was most effective in supporting accuracy. However, the ordering for the icon and minimap cue was reversed from that shown for RT, with the icon cue showing somewhat, although not significantly, greater accuracy than the minimap cue, ($t(55) = 1.46, p = .15$, and $d = 0.20$). Thus, H_4 is supported, replicating the general findings of [24].

B. Automation Bias: Automation Error Trials

The data addressing H_5 and the effects of imperfection were described in the two analyses above. In the analysis of RT, imperfection of cueing had no effect. In the analysis of accuracy, consistent with H_5 , it had a large effect, degrading accuracy across all cue types by approximately 13%. This degradation could result from either or both of two causes: First, following the automation bias, participants blindly followed the cue, whether right or wrong, and hence, when it was wrong, suffered the corresponding increasing error rate; second, participants, knowing




Fig. 5. Performance in the speed–accuracy space. Mean RT (x-axis) by mean accuracy (y-axis) for the arrow (green triangle, solid line), minimap (blue circle, dashed line), and icon (red square, dotted line) cues when the cue was correct (top three points) and incorrect (bottom three points). The open circle represents the no-cue (control) condition. Horizontal and vertical error bars represent the SEM for RT and accuracy, respectively.

that the cue could be wrong on any trial, relied less upon the automation and more on their own imperfect performance level. Because when they were unaided, this performance level at 65% (see horizontal line of Fig. 4) was less than with the cueing aids, such increased self-reliance would also pull the overall average performance downward.

The data to examine the automation bias are presented in the speed–accuracy tradeoff space in Fig. 5. In this, we have plotted the data for all three conditions of cue effectiveness when the cues were accurate and when they were wrong. Our statistical analysis is performed only on the two most different conditions in cueing effectiveness when the cue was correct, specifically the arrow and the icon cue. These two represent very different information processing mechanisms [see (2) and (4)]. Inspection of the figure also reveals that the performance with the two location cues, the arrow and minimap, responds almost identically to the failure.

Each line in Fig. 5 connects the three different cue types: the green triangle (arrow cue) and blue circle (minimap cue) toward the left, reflecting a rapid RT, and the square red (icon cue) farther toward the right, reflecting their longer RT. Each data point represents the speed (RT) and accuracy of *correct cueing* (top) and *wrong cueing* (bottom). The larger standard error bars for the wrong-cueing trials result from the fact that their N was one-sixth the size of the N for correct automation. For comparison purposes, the unfilled circle represents the data from the no-cue (control) condition.

Within this speed–accuracy space, high cue effectiveness is represented in the upper left corner: short RT and high accuracy. The data when the cue is correct show the pattern described above in Figs. 3 and 4. The arrow cue is unambiguously more effective on both speed and accuracy variables. However, when

the cue is wrong (lower data points), the pattern essentially reverses. There is a delayed RT in all cue types, suggesting that the participants are somewhat or sometimes aware of the cueing error [i.e., not always blindly following the cue, but sometimes just more hesitant in responding, which represents an increase in the constant a in (2)–(4)]. There is a large loss in accuracy for all three cue types suggesting an automation bias. Critically, however, this loss is much greater for the arrow cue, where accuracy drops from 89% to 13%, than it is for the icon cue, where accuracy drops from 80% to 43%. That is, when using the icon cue, participants correctly chose the blue practice block object (target not present) in 57% of the trials. The significant difference in accuracy between the two cue types when the cue is wrong showed a large effect size, ($t(27) = 9.50$, $p < .001$, and $d = 2.59$), thereby supporting H_6 . It should be noted here that the positioning of the minimap cue in the speed–accuracy space when the cue is wrong is nearly identical to that of the arrow cue. That is, both location cues showed the same pattern.

Above, we described the two different kinds of human errors that could occur when the automation was wrong: the participant’s decision to use their fallible judgment and the blind following of the automation (the automation bias). The prevalence of the former can be estimated from the error rate in the control condition, which is approximately 30%. Hence, we can assume that, of the 60% error rate in the icon condition, half of these can be attributed to the *automation bias errors*, and for the two location cue conditions, the approximate 92% error rate can be predominately attributed to the *automation bias errors*. The reduced automation bias of the icon condition would appear to reflect the more thorough examination of the target object during the confirmation stage, and hence, the much longer RT delay that was observed as reflected by the constant $a' > a$.

Another way to look at the automation bias is to examine the RTs for correct and incorrect responses on trials when the automation was wrong. The results revealed that RTs for the arrow, minimap, and icon cues were much longer when participants were correct ($M_A = 40.95$ s, $M_{MM} = 38.25$ s, and $M_I = 52.1$ s, respectively), and hence, they “over-ride” the automation, compared to when they were incorrect ($M_A = 11.26$ s, $M_{MM} = 15.03$ s, and $M_I = 37.55$ s, respectively) and did not override. These findings suggest that they spent more time examining the cued object and deciding whether to select the practice block on trials where they realized the automation was wrong. However, this only occurred in 23% of the trials. On the remaining 77% of the trials, RTs were shorter and similar to their values when the automation cues were perfectly reliable, suggesting that the participants blindly followed the automation on the majority of the trials when the cue was incorrect.

C. Location Effects

In addition to our primary interest in cue type and reliability, a secondary interest included the influence of cue location. Our prior research [24] indicated that the performance was somewhat disrupted due to increased visual scanning between the object view and the cue view when the cue was positioned downward. In the current experiment, we were also interested in any possible

interaction between cue location and the location of the target object in the wider visual field. The latter variable could provide us insight into the visual search strategies employed in unaided search and how they might be modified by the different cueing types. To conduct the analyses of both cue and target location effects, we separated the target location by its vertical (shelf, table, and floor; see Fig. 1) and lateral positioning. Because the targets were continuously distributed in the lateral field, we coded lateral location into four categories: left versus right of initial fixation and toward the center or the periphery. In this way, we could discern one of the two different search strategies often employed: left-to-right and center outward. We analyzed the three levels of vertical target positioning, separately from the four levels of lateral positioning, to ensure sufficient data in each. We also analyzed the uncued trials separately from the cued trials, only the latter analyses, including the variable of cue location.

The effects of both cue and target location were, for the most part, muted and nonsignificant. In particular, cue location (center or downward) showed no significant effects nor interactions with either RT or accuracy.

The only prominent effects of vertical target location were for RT. Results showed faster RT for the top ($M_{\text{shelf}} = 13.43$ s) and middle ($M_{\text{table}} = 13.65$ s) locations than the bottom ($M_{\text{floor}} = 15.29$ s), $F(2, 110) = 9.45$, $p < .001$, and $\eta_p^2 = 0.15$, suggesting some combination of a top to bottom, and a middle (table) outward (to shelf or floor) search strategy. This effect was also significant, but smaller in magnitude, and of the same pattern for uncued trials, $F(2, 110) = 3.62$, $p = .03$, and $\eta_p^2 = 0.06$. There was no consistent trend for the lateral position in either cued or uncued trials.

For accuracy (on the cued trials), there was an interaction between the vertical location of targets and cue location, $F(2, 110) = 5.06$, $p = .008$, and $\eta_p^2 = 0.08$. For the floor objects, the greatest accuracy occurred when the cue was presented downward ($t(55) = 2.92$, $p = .005$, and $d = 0.50$), a benefit for downward cueing not conferred in other vertical target locations. The effect of lateral target location (on cued trials only) showed greater accuracy when the target was in the center, $F(1, 55) = 8.42$, $p = .005$, and $\eta_p^2 = 0.13$. There were no consistent effects of lateral target position on the accuracy for uncued trials. The absence of any main effect of cue location fails to confirm H_7 .

V. DISCUSSION

In this experiment, we have joined two human-computer interaction (HCI) issues in the study of AR-HMD target cueing: the capability of HMD cueing to exploit the properties of AR to offer world-referenced cues and the effects of imperfect reliability of automation-generated cues.

A. Overall Cueing Effectiveness

In a task requiring the integration of near-domain (displayed cue) information and far-domain (search field) information, we found AR cueing (i.e., the arrow) to provide a pronounced 10-s benefit relative to uncued targets (H_1). We also found a

benefit for cues that were not presented in world-referenced AR coordinates compared with uncued targets. Specifically, the minimap showed a 4-s benefit and the icon cue showed a 15% accuracy benefit relative to no cue, confirming H_2 .

The time benefit of AR arrow cueing over the minimap was small (2 s) but significant. Both techniques provided precise location cueing, thus eliminating the $nT/2$ from predictive (2) and (3). However, the minimap cue imposed some added time cost due to the FORT, when transforming a top-down 2-D map to the forward 3-D view of the search scene [31], [48]. This penalized the CEV by approximately 2 s (3). While no spatial transformation was required for the icon cue, its time penalty relative to the arrow cue was quite large, at approximately 15 s: a 200% increase.

In the context of the serial self-terminating model of visual search, as discussed in Section I [20], assume that, with the icon cue, each object is examined and compared with the icon cue in order to establish that it is not the target. This search strategy continues until a match is found, a comparison process of time T , which we estimate from the current data to be about 0.40 s. In the context of the equations presented in Section I, this estimate assumes that, for the precise location cues, only one comparison is required ($n = 1$), whereas for the icon cue, on average, $48/2 = 24$ comparisons are required. (We note that, in [24], the estimate of T was quite similar—approximately 0.33 s—to the 40 s observed here). The much longer ST of the icon cue could be due to some combination of any of three factors that differ between the current icon cue condition, and its rendering in the previous research and the arrow cue in the current research.

- 1) n : The requirement to search through more objects here (48) than in the earlier research (32), or 8 more objects on average.
- 2) T : A longer, more careful inspection of each object made in the current research because of the explicit request to do so in the instructions (see *Methods* above). This instruction was offered to both groups.
- 3) d' : A longer time is required for the final confirmation of the candidate believed to be the target, compared with both the prior experiment and the arrow condition. In the latter case, we assume that this comparison, based on memory versus perception, was more rapid and, consequently, more error prone when the cue was wrong (see Fig. 5).

According to the guided search model [21], features of an object, such as salience, help guide search via attentional selection. Different types of attention guidance combine to create a priority map for the information selected by attention [21]. In the context here, features (e.g., shape, color, and orientation) of the icon cue and the objects in the search scene play a role in guiding attention in terms of what to inspect. However, the current data do not allow us to determine the precise value of CEV for the icon cue because we have no way to estimate “ x ,” in (4), the number of items that do not require inspection because of the guided search [21].

While this model accounts for the differences in ST between the cueing conditions and supports the advantage of location over identity cueing, the accuracy differences result from a

different mechanism. In particular, the accuracy of the icon was now elevated to be highly similar to that of the minimap and nonsignificantly greater. The reason for this equivalence is assumed to lie in the target confirmation process (a process reflected in the RT function by the intercept a or a'). For the icon, this confirmation is a direct *visual* comparison between the viewed target object and the icon. However, for the minimap, it is a comparison between the visual image of the target, and the *remembered* image of the target, which was viewed at the beginning of the trial. For the minimap, this time delay, averaging approximately 12 s (see Fig. 3), is of sufficient magnitude such that it degrades the quality of the image in spatial working memory [55] and produces some loss to the accuracy of the comparison. This loss for the minimap is greater than it is for the arrow cue simply because the minimap time delay is about 3 s longer than that of the arrow (see Fig. 3). That time difference is great enough to account for the diminished accuracy of the minimap relative to the arrow in the quality of the remembered target image. With a lower quality image, accuracy will be compromised. We note that, in all cases, the accuracy is far below perfect. That is because the megablock targets were intentionally designed to be highly similar to many of the nontargets in their 3-D shape and color combination (see Fig. 2 right panel). This might represent, for example, a vehicle with a certain differentiating feature (e.g., 1 ft longer) compared with many other nontarget vehicles sharing many similar features to the target (e.g., shape, color, number of doors, etc.).

Despite the pronounced advantage of location (arrow and minimap) over appearance (icon) cueing, it is possible to envision circumstances in which the advantage is reduced or even reversed. For example, in some circumstances, initial location uncertainty may be quite small, as in searching a line up of only five ($N = 5$) suspects for a crime to identify the remembered target previously viewed at the crime scene; but there may be subtle differences in facial appearance of the true target and the nontargets, greatly impinging on accuracy when there is no image for direct visual comparison.

B. Imperfection of Cueing

The second intersecting issue was the 17% imperfection of AI, which degraded the overall HAT accuracy and CEV by 15% supporting H_5 . This effect of automation reliability has been well documented in much of the HAI research in general [56] as well as that specific to target search [10], [11], [15], [36]. These findings indicate that imperfection of a cue will almost always pull HAT performance downward relative to perfect cueing, as long as the accuracy (reliability) of the human alone is less than that of the imperfect cue alone, as was the case in the present experiment (see Fig. 4). But at the same time, imperfect reliability in automation will typically aid human performance accuracy above an unaided manual baseline, as is the case here ([11], [36], [44], see Fig. 4). In this context, it is also important to note that, even with the best cue of the perfect cueing condition, the mean accuracy of 92% is still less than if participants in that group only followed the cue, which was correct 100% of the time. Such a finding is not unexpected, given that other studies have observed that the performance of highly reliable

automation typically remains below that of automation itself ([11], [36], see [2] for a summary).

Our more detailed examination of the data indicated that the consequence of cueing imperfection was not always a blind adherence to the automation bias [40], [41]. If this had been the case, then the speed of responding would have been unaffected by whether the cue was or was not correct. But, as shown in Fig. 5, speed was somewhat affected. The response was significantly slower on those occasions when the cue was wrong, even for the most effective cues (i.e., arrow), which might have been predicted to show the greatest automation bias. We suspect that such hesitation may have been due to the increased overall caution with the participant knowing that cueing errors could occur or the increase in confirmation time on those trials where doubt might have occurred in the participant's mind. This would produce an increase in the ST (RT) constant a in (2) and (3), or a' in (4).

Nevertheless, strong evidence was provided for the presence of automation bias in all cueing conditions. It is here that the two human-automation interaction (HAI) issues, that of AR and automation imperfection, intersect. When location cueing is most intuitive (i.e., with the AR arrow) in world-referenced coordinates, it is most likely to be followed without using the raw data of the target object in the search scene, supporting H_6 . Mifsud et al. [30] also found this automation bias with a different form of AR target cueing, and Yeh et al. [14] found it to be the case when comparing AR-HMD cueing with non-AR cueing. As described in Section I, the more realistic the cue and perhaps the more powerful the response it elicits, the more likely it is to be followed when it is wrong: for example, in this case, pointing an arrow at a potential target that the user must then designate for the response. This can be seen as a specific manifestation of the more general "Lumberjack" tradeoff principle of HAI: the more effective automation is when it is correct, the greater the consequences to the HAT's performance when it is wrong [38], [59]. (That is, the higher the tree, the harder it falls). This leads us to consider two aspects of HAI. On the one hand, consider mitigating solutions to overdependence on automation and AI when their information is rendered in AR. On the other hand, consider the frequency and the kind of errors committed by AI when it is based on actual machine vision operating in real time under time-sensitive conditions of potentially hostile behavior detection or dynamic hazard avoidance [46].

Our findings regarding location effects were muted. Differences in the location of the cue, center or downward on the HMD, had little effect on performance. Thus, the information access penalty of a small downward scan to see any of the cues were quite small, such as findings in another context [52]. It may be that the downward position of the cue location (at 12.6°) remained sufficiently within the eye-head field range to mitigate any cost of information access effort that we might, otherwise, see if the cues were out of the FOV of the HMD entirely (e.g., glanceable AR [57]). Correspondingly, there was not the predicted interaction of cue location with target position that might have suggested a clutter effect. This clutter effect refers to the greatest performance decrement in two circumstances: First, when the cue was in the center and the target on the table, and second, when the cue was positioned downward, and the target

was on the floor. This clutter effect was muted possibly because the minimap was a smaller size than the prior study [24], creating less overlay clutter. Alternatively, the performance penalties for increased clutter due to overlaying target and cue were offset by the benefits of reduced information access between the target and cue location.

In this regard, the significant interaction observed between cue and target vertical location suggested that the lower targets were best served by the bottom cue location: close spatial proximity between cue and target fosters best performance. This finding is a prediction of the proximity compatibility principle [52], [58], which states that, when two sources of information require integration, those sources should be placed closer together in space.

VI. LIMITATIONS

One limitation of the current research, in generalizing to cued search in the real world, is that the latter includes many cases in which the objects need not be compared with an image to be detected. For example, consider searching for a weapon in an X-ray image at airport security. The searcher has a well-formed image in the long-term memory of what a “typical weapon” looks like. Yet, the aforementioned example of searching for a particular vehicle that may have a similar shape and color to other vehicles but differs in its length, or instances where specific targets may be less familiar and are not stored in long-term memory, demonstrate cases more typical of those studied here where objects cannot be easily discriminated from foils.

A second limitation is that the arrow, such as any AR cue, could be subject to parallax errors. The magnitude of this effect may be quite dependent on the proximity between the target and adjacent items. If search items are close together, parallax errors could easily result in an arrow cue pointing to an adjacent item rather than the true target. In the current research, we assured sufficient spacing between search items that this did not occur, but it certainly remains a concern in other applications.

A third limitation is that we did not measure how accurate participants thought the system was. In doing so, we would be able to assess how calibrated participants were to the system’s reliability. Perceptions of system reliability influence trust in the system and, consequently, how people perform with the system [60]. Therefore, future work should assess how well participants are calibrated to the system’s reliability.

Finally, we did not use expert participants who engage in visual search tasks, such as those in a military combat role for which this paradigm might be highly relevant. We did not have a sufficient number of expert participants available to gain adequate statistical power. However, we note that, in other AR-HMD studies, our subject matter experts have stated that a major concern of the forward observer is that of confusing a relevant from a nonrelevant object on the ground when searching for the former. Therefore, we believe that the current research addresses a highly relevant applied issue.

VII. CONCLUSION AND FUTURE WORK

Overall, we found that cues depicting location information, the arrow, and minimap cues led to faster and, in the case of the arrow cue, more accurate visual searches compared with the

icon cue, which only depicted information about the appearance of a target. However, when imperfect cues were erroneous, the most effective cues for supporting search performance led to worse outcomes. This is indicative of an automation bias where people blindly follow the cue without sufficient, independent confirmation for themselves that the identified object is the correct target. These findings highlight an important prospective use case for AR support during visual search tasks and also shed light on the costs associated with imperfectly reliable cueing aids. A tendency toward greater automation bias may be a critical drawback in some high-stakes’ contexts. Such findings strongly imply that designers will need to take a system view of error management in supporting visual search rather than just optimizing performance under the assumption of perfectly reliable support.

In the future, several experiments are warranted to explore the interaction between the two HMD automation issues. For example, different levels of automation reliability should be explored to assess if the AR cost for the more realistic world-referenced cue diminishes as cueing reliability also decreases. Also, the current findings, when automation errors are scripted by the experiment, should be extended to circumstances in which actual machine vision determines target cueing in real time. Correspondingly, effort should be taken to systematically manipulate target versus nontarget similarity to establish if advantages to the accuracy of the icon cue begin to emerge, particularly as the size of the search field n is reduced. Finally, more research should be conducted when the screen-referenced cue is further displaced from the search field, thereby increasing the information access effort.

ACKNOWLEDGMENT

Dr. P. Squire was the scientific/technical monitor.

REFERENCES

- [1] G. Dury C, “Inspection,” in *International Encyclopedia of Ergonomics and Human Factors*, vol. 2, W. Karwowski, Ed. Boca Raton, FL, USA: CRC Press, 2006.
- [2] C. D. Wickens, J. S. McCarley, and R. S. Gutzwiller, *Applied Attention Theory*. Boca Raton, FL, USA, CRC Press, 2022.
- [3] J. R. Helleberg and C. D. Wickens, “Effects of data-link modality and display redundancy on pilot performance: An attentional perspective,” *Int. J. Aviation Psychol.*, vol. 13, no. 3, pp. 189–210, 2003.
- [4] K. Gramopadhye A, G. Drury C, X. Jiang, and R. Sreenivasan, “Visual search and visual lobe size: Can training on one affect the other?,” *Int. J. Ind. Ergonom.*, vol. 30, no. 3, pp. 181–195, 2002.
- [5] C. G. Dury, “Visual search in industrial inspection,” *Appl. Ergonom.*, vol. 22, no. 3, pp. 263–276, 1991.
- [6] C.-C. Wu and J. M. Wolfe, “Eye movements in medical image perception: A selective review of past, present and future,” *Vision*, vol. 3, no. 2, 2019, Art. no. 32.
- [7] H. Kundel and C. Nodine, *Studies of Eye Movements and Visual Search in Radiology*, J.W. Senders, D.F. Fisher, R.A. Monty, Eds, 1978.
- [8] D. Wiegmann, J. S. McCarley, A. F. Kramer, and C. D. Wickens, “Age and automation interact to influence performance of a simulated luggage screening task,” *Aviation, Space, Environ. Med.*, vol. 77, no. 8, pp. 825–831, 2006.
- [9] D. Huegeli, S. Merks, and A. Schwaninger, “Benefits of decision support systems in relation to task difficulty in airport security X-ray screening,” *Int. J. Human-Comput. Interaction*, pp. 1–16, 2022.
- [10] J. Goh, D. A. Wiegmann, and P. Madhavan, “Effects of automation failure in a luggage screening task: A comparison between direct and indirect cueing,” in *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, vol. 49, pp. 492–496, 2005.

- [11] M. M. Boskemper, M. L. Bartlett, and J. S. McCarley, "Measuring the efficiency of automation-aided performance in a simulated baggage screening task," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 64, no. 6, pp. 945–961, 2022.
- [12] L. Swann, V. Popovic, A. Blackler, and H. Thompson, "Airport security screener problem-solving knowledge and implications," *Hum. Factors*, vol. 62, no. 8, pp. 1265–1285, 2020.
- [13] M. D. Fleetwood and M. D. Byrne, "Modeling the visual search of displays: A revised ACT-R model of icon search based on eye-tracking data," *Hum.-Comput. Interaction*, vol. 21, no. 2, pp. 153–197, 2006.
- [14] M. Yeh, J. L. Merlo, C. D. Wickens, and D. L. Brandenburg, "Head up versus head down: The costs of imprecision, unreliability, and visual clutter on cue effectiveness for display signaling," *Hum. Factors*, vol. 45, no. 3, pp. 390–407, 2003.
- [15] M. Yeh and C. D. Wickens, "Display signaling in augmented reality: Effects of cue reliability and image realism on attention allocation and trust calibration," *Hum. Factors*, vol. 43, no. 3, pp. 355–365, 2001.
- [16] M. Yeh, C. D. Wickens, and F. J. Seagull, "Target cuing in visual search: The effects of conformality and display location on the allocation of visual attention," *Hum. Factors*, vol. 41, no. 4, pp. 524–542, 1999.
- [17] S. P. Devlin, N. M. Moacdieh, C. D. Wickens, and S. L. Riggs, "Transitions between low and high levels of mental workload can improve multitasking performance," *IISE Trans. Occup. Ergonom. Hum. Factors*, vol. 8, no. 2, pp. 72–87, 2020.
- [18] M. E. Frame, A. M. Maresca, A. Christensen-Salem, and R. E. Patterson, "Evaluation of simulated recognition aids for human sensemaking in applied surveillance scenarios," *Hum. Factors*, pp. 1–12, 2022.
- [19] J. M. Wolfe, T. S. Horowitz, and N. M. Kenner, "Rare items often missed in visual searches," *Nature*, vol. 435, no. 7041, pp. 439–440, 2005.
- [20] S. Sternberg, "Memory-scanning: Mental processes revealed by reaction-time experiments," *Amer. Scientist*, vol. 57, no. 4, pp. 421–457, 1969.
- [21] J. M. Wolfe, "Guided search 6.0: An updated model of visual search," *Psychon. Bull. Rev.*, vol. 28, no. 4, pp. 1060–1092, 2021.
- [22] J. M. Wolfe and T. S. Horowitz, "Five factors that guide attention in visual search," *Nature Hum. Behav.*, vol. 1, no. 3, 2017, Art. no. 0058.
- [23] M. Yeh and C. D. Wickens, "Attention filtering in the design of electronic map displays: A comparison of color coding, intensity coding, and decluttering techniques," *Hum. Factors*, vol. 43, pp. 543–562, 2001.
- [24] A. C. Warden, C. D. Wickens, D. Mifsud, S. Ourada, B. A. Clegg, and F. R. Ortega, "Visual search in augmented reality: Effect of target cue type and location," in *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, vol. 66, no. 1, pp. 373–377, 2022.
- [25] L. C. Thomas and C. D. Wickens, "Display dimensionality, conflict geometry, and time pressure effects on conflict detection and resolution performance using cockpit displays of traffic information," *Int. J. Aviation Psychol.*, vol. 16, no. 3, pp. 321–342, 2006.
- [26] S. Yantis, "Stimulus-driven attentional capture," *Curr. Directions Psychol. Sci.*, vol. 2, no. 5, pp. 156–161, 1993.
- [27] M. I. Posner, "Orienting of attention," *Quart. J. Exp. Psychol.*, vol. 32, no. 1, pp. 3–25, 1980.
- [28] M. I. Posner, C. R. Snyder, and B. J. Davidson, "Attention and the detection of signals," *J. Exp. Psychol., Gen.*, vol. 109, no. 2, 1980, Art. no. 160.
- [29] M. I. Posner, "Orienting of attention: Then and now," *Quart. J. Exp. Psychol.*, vol. 69, no. 10, pp. 1864–1875, 2016.
- [30] D. Mifsud, C. Wickens, M. Maulbeck, P. Crane, and F. R. Ortega, "The effectiveness of gaze guidance lines in supporting JTAC's attention allocation," in *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, vol. 66, pp. 2198–2201, 2022.
- [31] C. D. Wickens, M. Vincow, and M. Yeh, "Design applications of visual spatial thinking," in *The Cambridge Handbook of Visuospatial Thinking*. Cambridge, U.K.: Cambridge Univ. Press, 2005, pp. 383–425.
- [32] C. D. Wickens, J. W. Keller, and R. L. Small, "Left. No, right! Development of the frame of reference transformation tool (FORT)," *Proc. Hum. Factors Ergonom. Soc. Annu. Meeting*, vol. 54, pp. 1022–1026, 2010.
- [33] B. Schneiderman, *Human-Centered AI*. Oxford, U.K.: Oxford Univ. Press, 2022.
- [34] J. Lee and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, no. 10, pp. 1243–1270, 1992.
- [35] K. A. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," *Hum. Factors*, vol. 57, no. 3, pp. 407–434, 2015.
- [36] M. L. Bartlett and J. S. McCarley, "Benchmarking aided decision making in a signal detection task," *Hum. Factors*, vol. 59, no. 6, pp. 881–900, 2017.
- [37] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 46, no. 1, pp. 50–80, 2004.
- [38] L. Onnasch, C. D. Wickens, H. Li, and D. Manzey, "Human performance consequences of stages and levels of automation: An integrated meta-analysis," *Hum. Factors*, vol. 56, no. 3, pp. 476–488, 2014.
- [39] J. E. Bahner, A.-D. Huper, and D. Manzey, "Misuse of automated decision aids: Complacency, automation bias and the impact of training experience," *Int. J. Human-Comput. Stud.*, vol. 66, no. 9, pp. 688–699, 2008.
- [40] K. L. Mosier, L. J. Skitka, S. Heers, and M. Burdick, "Automation bias: Decision making and performance in high-tech cockpits," in *Decision Making in Aviation*. Evanston, IL, USA: Routledge, 2017, pp. 271–288.
- [41] R. Parasuraman and D. H. Manzey, "Complacency and bias in human use of automation: An attentional integration," *Hum. Factors*, vol. 52, no. 3, pp. 381–410, 2010.
- [42] L. Bainbridge, "Ironies of automation," in *Analysis, Design and Evaluation of Man-Machine Systems*. Amsterdam, The Netherlands: Elsevier, 1983, pp. 129–135.
- [43] S. M. Merritt, J. L. Unnerstall, D. Lee, and K. Huber, "Measuring individual differences in the perfect automation schema," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 57, no. 5, pp. 740–753, 2015.
- [44] T. Rieger and D. Manzey, "Understanding the impact of time pressure and automation support in a visual search task," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, 2022.
- [45] S. R. Dixon, C. D. Wickens, and J. S. McCarley, "On the independence of compliance and reliance: Are automation false alarms worse than misses?," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 49, no. 4, pp. 564–572, 2007.
- [46] J. P. Bliss, "Investigation of alarm-related accidents and incidents in aviation," *Int. J. Aviation Psychol.*, vol. 13, no. 3, pp. 249–268, 2003.
- [47] C. D. Wickens and S. R. Dixon, "The benefits of imperfect diagnostic automation: A synthesis of the literature," *Theor. Issues Ergonom. Sci.*, vol. 8, no. 3, pp. 201–212, 2007.
- [48] C. Wickens, "Attention: Theory, principles, models and applications," *Int. J. Human-Comput. Interaction*, vol. 37, no. 5, pp. 403–417, 2021.
- [49] S. Fadden, C. Wickens, and P. Ververs, "Costs and benefits of head up displays—An attention perspective and a meta-analysis," in *Proc. World Aviation Conf.*, 2000, Art. no. 5542.
- [50] C. D. Wickens, L. C. Thomas, and R. Young, "Frames of reference for the display of battlefield information: Judgment-display dependencies," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 42, no. 4, pp. 660–675, 2000.
- [51] C. D. Wickens and M. Yeh, "Display compellingness: A literature review," *Cherokee CRC, LLC*, Tulsa, OK, USA, 2018.
- [52] A. C. Warden, C. D. Wickens, D. Rehbert, B. Clegg A, and F. Ortega R, "Information access effort: Are head movements 'cheap' or even 'free'?", in *Proc. 66th Annu. Hum. Factors Soc. Annu. Meeting*, 2022, pp. 2203–2207.
- [53] B. Rashid et al., "Comparing the validity of electronic color vision test with conventional Ishihara pseudoisochromatic plate test: Validity of electronic color vision test," *Pakistan Biomed. J.*, vol. 4, no. 2, pp. 156–159, 2021.
- [54] D. van Staden, F. N. Mahomed, S. Govender, L. Lengisi, B. Singh, and O. Aboobaker, "Comparing the validity of an online Ishihara colour vision test to the traditional Ishihara handbook in a South African university population," *Afr. Vis. Eye Health*, vol. 77, no. 1, 2018, Art. no. a370.
- [55] R. H. Logie, "Visuo-spatial working memory," Psychol. Press, 1995.
- [56] C. D. Wickens, W. S. Helton, J. G. Hollands, and S. Banbury, *Engineering Psychology and Human Performance*. Evanston, IL, USA: Routledge, 2021.
- [57] F. Lu, S. Davari, L. Lisle, Y. Li, and D. A. Bowman, "Glanceable AR: Evaluating information access methods for head-worn augmented reality," in *Proc. IEEE Conf. Virtual Reality 3D User Interfaces*, 2020, pp. 930–939.
- [58] C. D. Wickens and C. M. Carswell, "The proximity compatibility principle: Its psychological foundation and relevance to display design," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 37, no. 3, pp. 473–494, 1995.
- [59] A. Sebok and C. D. Wickens, "Implementing Lumberjacks and Black Swans into model-based tools to support human-automation interaction," *Hum. Factors*, vol. 59, pp. 189–202, 2017.
- [60] R. Sargent, B. Walters, and C. Wickens, "Meta-analysis qualifying and quantifying the benefits of automation transparency to enhance models of human performance," in *Proc. Int. Conf. Hum.-Comput. Interaction*, 2023, pp. 243–261.