# HiveFive360: Extending the VR Gaze Guidance Technique HiveFive to Highlight Out-Of-FOV Targets


Sophie Kergaßner
Hochschule der Medien
Stuttgart, Germany
Università della Svizzera italiana
Lugano, Switzerland

Nina Doerr
VISUS, University of Stuttgart
Stuttgart, Germany

Markus Wieland
VISUS, University of Stuttgart
Stuttgart, Germany

Martin Fuchs
Institute for Games, Hochschule der Medien
Stuttgart, Germany

Michael Sedlmair
VISUS, University of Stuttgart
Stuttgart, Germany


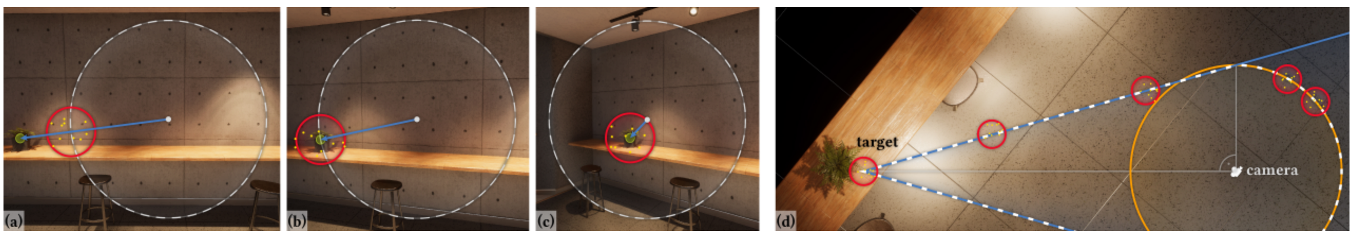


**Figure 1: Overview of our proposed HiveFive360 placement algorithm in action. A swarm of diegetically embedded objects guides the user to the target (plant). (a-c) User's view with schematic representation of the swarm placement (red circles) within the FOV. The dashed circle represents 30° eccentricity. (a) While the target is out-of-FOV, the swarm is positioned at 30° eccentricity in the direction of the target. (b) As the user rotates the head towards the target, the swarm moves toward the object. (c) As soon as the target enters <30° eccentricity, the swarm is placed in front of the target. (d) Top-down parallel rendering with schematic representation of five possible swarm placements on the approximated tangential cone (dashed line), which is constructed from a sphere (orange) and a cone (blue). 3D scene from Unity Technologies[4].**


## ABSTRACT

Modern display technologies, particularly those supporting 360° content, are increasingly used for immersive experiences in a variety of domains. However, information outside of the user's field of view (FOV) may be easily overlooked. To address this, guiding cues can be provided to effectively direct attention. Subtle and diegetic cues are particularly effective in keeping the coherence and immersion of the presented content. HiveFive is one of the few diegetic highlighting techniques. It effectively highlights objects by attracting the user's attention with swarm-like motion. However, HiveFive is restricted to in-FOV target highlighting. This work presents the novel technique HiveFive360, an extension of HiveFive that enables it to guide users to out-of-FOV targets. HiveFive360 is evaluated in a user study against FlyingARrow and Subtle Gaze Direction VR regarding completion time, sense of presence and task load. HiveFive360 was found to effectively guide users in various environments without excessive distraction or task load.



## CCS CONCEPTS

• **Human-centered computing → Empirical studies in HCI**.

## KEYWORDS

Virtual Reality, Visual Gaze Guidance

**ACM Reference Format:**
Sophie Kergaßner, Nina Doerr, Markus Wieland, Martin Fuchs, and Michael Sedlmair. 2024. HiveFive360: Extending the VR Gaze Guidance Technique HiveFive to Highlight Out-Of-FOV Targets. In *Proceedings of Mensch und Computer 2024 (MuC '24), September 01–04, 2024, Karlsruhe, Germany*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3670653.3670662


## 1 INTRODUCTION

With the extension of freely navigatable 360° content, limitations arise from both the devices that display the content and the human visual system. These limitations restrict a user's view to a specific solid angle, referred to as the FOV [21]. Information outside of the user's FOV can easily be overlooked. This poses challenges, especially in the context of 360° movies or virtual reality (VR) games, as information relevant to the story may be missed. In these situations, it is important to provide guidance while ensuring that the



*MuC '24, September 01–04, 2024, Karlsruhe, Germany*



ACM ISBN 979-8-4007-0998-2/24/09
https://doi.org/10.1145/3670653.3670662

highlighting cue is seamlessly integrated into the context of the presented material.

In recent years, numerous highlighting techniques with distinct guidance designs were developed for various application scenarios, for which specific techniques have been found. However, certain use cases remain unresolved. One of them is efficiently guiding users to out-of-FOV targets without using highly overt cues that distract the user from the virtual environment or the story, or too subtle cues that might be overlooked.

HiveFive is a diegetic highlighting technique by Lange et al. [27]. It is a particle system that imitates swarm-like movements, attracting the user's attention and guiding their gaze towards the target object. However, HiveFive can only actively highlight targets that are within the user's FOV. As soon as the target is initialized out-of-FOV, or the user turns away, HiveFive has no ability to gain the user's attention.

In this article, we propose HiveFive360, an extension of HiveFive that enables it to guide users to out-of-FOV targets. We developed and tested three HiveFive360 variants to determine the optimal parameter set. We further evaluated this technique in a user study with FlyingARrow, an arrow based overt highlighting technique [18], and an adaptation of Subtle Gaze Direction (SGD) for immersive environments [13] within three different virtual environments. Participants conducted a search-and-select task, where they searched ten targets in each technique-environment-combination. In summary, our main contributions are:

- We *developed the novel technique HiveFive360,* and optimized it based on the findings of a concept validation study.
- We *evaluated HiveFive360* in a user study and compared it with FlyingARrow [18] and Subtle Gaze Direction VR (SGD-VR) [13].

Our results show that HiveFive360 outperforms SGD-VR in terms of completion time, task load (mental demand and effort), and system usability. However, FlyingARrow guides faster than HiveFive360 and requires less mental demand. Regarding the sense of presence, no differences between the techniques are observed. Among the tested techniques, HiveFive360 blends best into the virtual environments and distracts the least.

## 2 RELATED WORK

HiveFive360 is the first highlighting technique that is capable of highlighting out-of-FOV targets in 360° VR environments with a procedurally generated, diegetic cue. The most relevant related methods that are capable of highlighting out-of-FOV targets are FlyingARrow [18] as an example of an overt technique, and SGD-VR [13] as an example of a subtle technique. Both of which are subject of our study (Section 4). Among the techniques that *either* can not highlight out-of-FOV targets, are not diegetic, or are not generated automatically, we can discern:

*Overt Techniques.* Overt techniques feature highlighting cues that are clearly visible to the user, but are not elements of the environment. Common overt cues include lines and circles [2, 16, 19, 35, 37, 47–49], points and spheres [17, 20, 23, 45, 49], arrows [3, 16, 18, 22, 29, 30, 37, 42, 47], or other structures like bars, maps, or funnels [4, 6, 30, 42, 44]. Although all of the overt techniques mentioned above can highlight out-of-FOV targets, the group of arrow-based methods was particularly interesting for our study, as many of them also continue to highlight the target as it enters the FOV [18, 30, 37, 42, 47]. This characteristic makes arrows a suitable cue for our user study. Additionally, for stereoscopic applications it is beneficial to place the arrow in the correct depth plane [3]. This attribute is exhibited by three of our reviewed arrow designs [18, 42, 47]. Ultimately, we chose FlyingARrow, as it was tested in several studies, proving effective, comfortable and easy guidance [18, 22, 43]. Hu et al. found that adding a trail had a positive effect on completion time [22]. However, participants found their trail design uncomfortable to use and recommended reducing the thickness of the trail and making it less opaque [22]. This was taken into account in our FlyingARrow implementation (Section 4).

*Diegetic Techniques.* Diegetic cues are elements or objects that are perceivable for all entities within an environment [39]. The diegetic VR guidance approach by Pausch et al. [34] includes in-scene characters pointing or moving towards the target object. Compared against non-diegetic and half-diegetic cues, participants preferred the diegetic characters in terms of sense of co-presence, interaction quality and usefulness [10]. Likewise, Speicher et al. [45] showed that their technique "Person to Follow" performed second best in terms of accuracy and user experience. Nielsen et al. [33] designed a glowing firefly that guides the user through the scene by following a predetermined path. Similarly, Cao et al. invented a crow that flies through the scene [8]. Unlike the aforementioned approaches, their method automatically spawned the diegetic cue at the right time and position. However, as with all of the techniques mentioned, their technique still required significant work by a skilled 3D artist to seamlessly integrate the cue into the scene.

Lange et al. developed HiveFive [27], a technique that serves as the basis for HiveFive360. HiveFive is based on Reynolds' flocking algorithm [38], which procedurally generates swarm movements. Artists only need to define position and general swarm parameters such as size, speed, extent, and color. However, the swarm's position in space is fixed. It can only highlight in-FOV targets.

*Subtle Techniques.* Subtle Techniques are not consciously perceivable by the user. A well-known technique is SGD [1]. SGD features a flickering stimulus, that is only presented in peripheral vision. It was adapted for various use-cases, such as guidance without eye-tracking [46], and for immersive environments like domes and VR applications [13].

A second group of subtle guiding techniques utilizes stereo rivalry. These techniques involve a dichoptic presentation of two different images to both eyes. Deadeye presents the target object for one eye only, which makes it the first pop-out cue that does not modify the visible characteristics of the target object [26]. It solely highlights in-FOV and was adapted for VR applications successfully [25]. Stereo Inverse Brightness Modulation (SIBM) modulates the brightness for the target's image regions differently per eye [14]. The cue was extended for out-of-FOV guidance as well [15] and was found to effectively guide in static and dynamic environments [14, 15].

Both SGD and SIBM highlight out-of-FOV targets and are therefore suitable for our study. We chose SGD-VR over SIBM, as it is a well-established cue [1, 11, 12, 31, 32, 46]. Additionally, Grogorick et al. have already designed a functional VR adaptation [13].

## 3 DEVELOPMENT OF HIVEFIVE360

The initial swarm design of Lange et al. [27] stays at a fixed position. Thus, it does not have the ability to guide users to out-of-FOV targets. We expand HiveFive by positioning it dynamically, enabling it to highlight out-of-FOV targets. The swarm is spawned within the user's FOV and gradually moves to the out-of-FOV target. In our adaptation, we left the original swarm, e.g. its visual appearance, as unaltered as possible to minimize confounding factors. Thus, we only modified the swarm's placement in 3D space and its behavior to fit our out-of-FOV guidance requirements.

### 3.1 Swarm's Placement in 3D Space

We position the swarm in 3D space so that it is always inside the FOV of the user, with situation-dependent direction (0–360°), eccentricity (0–30°), and distance (Figure 1).

*Direction and Eccentricity.* Based on Schor [41], we positioned the swarm at min. 30° eccentricity (Figure 1a) to trigger a head rotation towards the out-of-FOV target. As soon as the target is in the FOV of the user (<30° eccentricity), the swarm maintains its placement in the FOV in the direction of the target (Figure 1b/c). The swarm entering <30° eccentricity, thus, changing its behavior, indicates the presence of an interesting object in the area.

*Distance.* Simply placing the swarm at a fixed distance (on a spherical surface, Figure 2a) can cause confusion due to binocular vision [3]. Therefore, this approach is not sufficient. To address this, the swarm must leave the spherical surface and fly to the correct depth plane. We defined that it should maintain a minimum distance of 1m constantly, but move towards the target position (i.e., the correct depth plane) at the earliest opportunity. For this, we added an approximated tangential cone to our placement algorithm (Figure 1d/2b). First, the algorithm calculates the direction and eccentricity as described above (Figure 1a–c). Then, it places the swarm on either the cone or the sphere, depending on which position is closer to the target (Figure 1d). The surface resulting from the approximation of the tangential cone is not C1 continuous. This is negligible, as it is not noticeable due to the dynamic expansion and movement of the swarm.

### 3.2 Particle Behavior

In the initial implementation of Lange et al. [27], the swarm particles aim for the user-set center point of the swarm. In our adaptation, the center point is dynamically positioned based on the user's head orientation (Figure 1). However, when translating the center point, the particles move too slowly towards their new target position and are unresponsive to the user's actions in the meantime. We resolve this by accelerating swarm particles that are too far away from the center position. Once the particle returns within the swarm's boundaries, it resumes its default speed. To prevent particles from abruptly decelerating, we add an interpolation between the increased and default speed. The rate of deceleration can be manipulated by setting the size of a buffer zone in Unity, in which the speed is linearly interpolated from high to low speed.

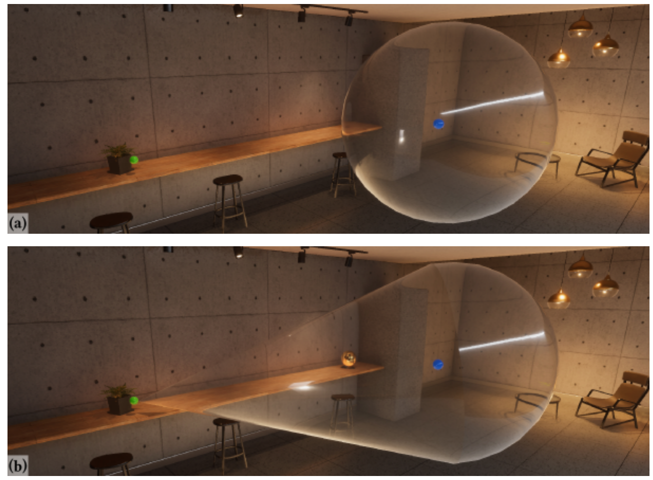

**Figure 2: The glass surface represents all possible swarm positions for the given camera location (blue sphere) and target (green sphere). Two placement options are considered: (a) Placement on a sphere, where the swarm maintains a constant distance of 1m to the camera; and (b) Placement on sphere+cone, where the swarm can leave the spherical surface and fly towards the target, which is at an approx. distance of 3.5m. 3D scene from Unity Technologies[4].**

### 3.3 Swarm Behavior Variants

We design three different swarm behavior variants which we evaluate in terms of subjective intuitivity in a concept validation study (Section 3.4). All three variants inherit the aforementioned changes (Section 3.1/3.2).

*Basic*: No additional changes were made to this swarm variant.

*Agitated Bees*: The swarm gets increasingly agitated as the user fails to look towards the target (i.e., the target is at >30° eccentricity). The particle's speed and rotation speed are slowly increased to max. ×1.5 default speed. We hypothesize that increased particle movement and a change in behavior attracts attention. Furthermore, an increasingly hectic swarm may become uncomfortable, potentially making it too hectic to ignore and prompting the user to take some form of action.

*Periodic Movement*: According to Pratt et al. [36], sudden movement attracts attention. In this variant, the swarm periodically leaves the user's FOV. It remains within the FOV for ~4.5s, before rapidly flying towards the target in a direct path and returning after ~1.5s. This abrupt and unpredictable change in movement should capture attention [36], prompting the user to track the swarm. Additionally, we hypothesize that the larger movements of the swarm leaving the FOV attracts additional attention [46].

### 3.4 Concept Validation Study

We conducted an interview-based concept validation study to evaluate our aforementioned three variants for the adjustment of HiveFive360. The study included 5 participants (3 male, 2 female) with an average age of 24.8 years. All participants had normal or corrected-to-normal vision. Two participants were visualization experts and three participants had VR experience. The concept validation study was conducted in the same study setup as the main study, as described in Section 4. Participants examined all three variants in random order in the DesertTemple environment. They completed the same search-and-select task as in the main study (Section 4). However, they were informed that they could take as much time as they needed to examine the presented variant thoroughly. After each presented variant, the participants were interviewed.

Similar to the results of Lange et al. [27], participants liked the swarm design and its bright color. In general, participants preferred the swarm to remain in their FOV, thus, opting for the *basic* and *agitated* swarm (ID2–5). Four participants explicitly stated that they did not observe differences between the *basic* and *agitated* swarm (ID1, ID3–5). Two participants noted positively that the *agitated* swarm had a high reactivity, which is due to it's higher particle speed (ID2, ID3). Four participants disliked the *periodic* swarm due to its rapid and hectic movements. In addition, they disliked that it left the FOV completely and had to be searched for (ID2–5).

Based on these results, we made the following changes to the swarm for our user study: We decided for the *basic* swarm variant, as participants did not perceive the accelerations of the *agitated* variant. In addition, we enhanced the swarm's reactivity by increasing the particle's default speed by ~42%, and increasing the speed of the particles that have a greater distance to the center point by ~30%

## 4 METHODS

Our goal is to examine whether HiveFive360 is a suitable highlighting technique for out-of-FOV target highlighting, assessing its efficiency, ability to create a sense of presence, and induced task load. To do this, we compare it to two other highlighting techniques in a search-and-select task within three virtual environments.

*Independent Variables.* The study involved two independent variables, each with three levels:

- *highlighting technique* (HiveFive360, FlyingARrow and SGD-VR) and
- *virtual environment* (StreetCorner, DesertTemple and IndoorRoom).

As described in Section 2, we chose FlyingARrow as overt highlighting technique. It features an arrow that repeatedly spawns inside the user's FOV and flies linearly towards the target. For the arrow implementation, we integrated the provided implementation of the GitHub repository[1] by Gruenefeld et al. [18]. As suggested by Hu et al. [22], we also added a faint trail to the arrow to make its trajectory more visible (Figure 3b).

In addition, we used SGD-VR as subtle highlighting technique (see Section 2). Our implementation follows the design and implementation of Grogorick et al. [13] as they already adjusted the original SGD of Bailey et al. [1] for immersive environments (Figure 3c). For this, they adapted the stimulus by moving it inside the user's FOV from the center towards the target direction (max. 35° eccentricity). In our implementation, the stimulus remains in a fixed

[1]https://github.com/UweGruenefeld/OutOfView

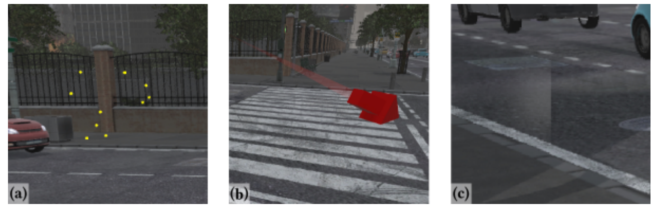


**Figure 3: Screenshots of the three highlighting techniques used in the user study (a) HiveFive360 (b) FlyingARrow + Trail (c) SGD-VR. SGD-VR is only displayed on the left half of the image for better visibility. 3D scene from ReversedInt[2].**

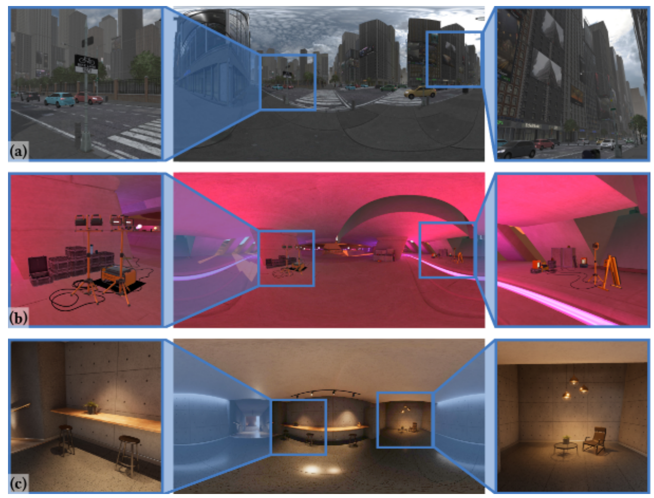


**Figure 4: Equirectangular images of the three virtual environments that were used in the user study. Screenshots of the in-game view are added left and right of the 360° images. (a) StreetCorner (b) DesertTemple (c) IndoorRoom. 3D scenes from ReversedInt[2] and Unity Technologies[3,4].**

position once the target is within 30° eccentricity. The stimulus is deactivated when the gaze of the user is within 10° of the stimulus' surrounding.

We selected three different virtual environments (VEs) for our user study that contained different amounts of visual clutter and consisted of indoor and outdoor scenarios. They did not contain any dynamic distractors or motion [9]. We chose *StreetCorner*[2] as a cluttered outdoor scene, *DesertTemple*[3] as a semi-cluttered and partially covered outdoor scene and *IndoorRoom*[4] as a simple indoor scene (Figure 4).

*Task.* Participants had to find ten consecutively highlighted targets per stimulus. They started by turning towards a fixation cross. As soon as the cross disappeared, the target was highlighted with

[2]*DayScene* from https://assetstore.unity.com/packages/3d/environments/urban/newgen-urban-229501
[3]*The Courtyard* from https://assetstore.unity.com/packages/essentials/tutorial-projects/the-courtyard-49377
[4]Room from Unity's *3D Sample Scene (HDRP)*.

the respective highlighting technique (HT). After successfully selecting the target with a controller, the fixation cross reappeared and the process was repeated. When all ten targets were found, participants answered the four questionnaires.

*Dependent Variables.* We measured seven dependent variables:

- *completion time* (i.e. the time between onset of the HT and selection of the correct target),
- *sense of presence*, measured with the Igroup Presence Questionnaire (IPQ)[5],
- *mental demand*, measured with the NASA Task Load Index (NASA TLX)[6],
- *effort*, measured with the NASA TLX,
- *system usability*, measured with the System Usability Score (SUS) [7], and
- two *additional subjective metrics (ASM)*, both measured on a 5-point scale inspired by Lange et al. [27].

*Study Design.* We used a 3x3 within-subjects design, evaluating three HTs (HiveFive360, FlyingARrow and SGD-VR) across three VEs (StreetCorner, DesertTemple, and IndoorRoom). Each participant was exposed to each combination of VE and HT once, resulting in 9 conditions per participant. The order of conditions was randomized across subjects to avoid learning effects. The sequence of conditions was randomized by requiring that each HT and VE appeared exactly once in each block of three stimuli. In addition, we ensured that no two VEs or HTs were presented consecutively.

*Hypotheses.* We investigated the following six hypotheses:

- $H_1$: When HiveFive360 is used as a highlighting technique, the completion time is faster than when Subtle Gaze Direction VR is used.
- $H_2$: When HiveFive360 is used as a highlighting technique, the sense of presence (IPQ) is higher than when FlyingARrow is used.
- $H_3$: When HiveFive360 is used as a highlighting technique, the sense of presence (IPQ) is the highest.
- $H_4$: When HiveFive360 is used in an indoor scene (IndoorRoom), the sense of presence (IPQ) is lower than when it is used in an outdoor scene (StreetCorner/DesertTemple).
- $H_5$: When HiveFive360 is used as a highlighting technique, the mental demand (NASA TLX) is the lowest.
- $H_6$: When HiveFive360 is used as a highlighting technique, the effort (NASA TLX) is the lowest.

The project was preregistered at Open Science Framework [24].

*Apparatus and Materials.* The User Study was conducted using the HTC Vive Pro Eye and Unity version 2020.3.21. Eye tracking (used for SGD-VR) was implemented using Tobii XR SDK version 3.0.1.179. The utilized computer features an Intel Core i7-9700 running at 3.60GHz, paired with an NVIDIA GeForce RTX 3070 and 64GB RAM. The utilized questionnaires are listed in Section 4. They were integrated into a single LimeSurvey form[7]. Participants completed the survey on a desktop computer situated beside the study area.

*Procedure.* The procedure was consistent across all participants. Upon arrival, the participants received a verbal introduction. Following this, they read and signed a data privacy and consent form before completing a demographic questionnaire. Participants' VR headsets were adjusted and the instructor explained how to calibrate the eye tracker. The participants received verbal instructions explaining the training task. During the training, they learned how to select nine consecutively color-highlighted targets (cubes) using the VR controller. After successfully completing the training task, the actual study task was explained verbally. The participants saw the initial stimulus. Upon successful detection of all ten targets, participants removed the VR headset and completed the four questionnaires. Once completed, the researcher proceeded with the next stimulus. After testing all nine stimuli and collecting the corresponding questionnaires, participants received a compensation of 12,00€. The duration of participation ranged from 50–70 minutes per participant.

*Participants.* The study included 27 participants (19 male, 8 female) with an average age of 24.9 years. All participants reported normal or corrected-to-normal vision. Six participants reported no prior use of a VR device (rating 1/5), while 13 participants reported only a few previous uses without familiarity (rating 2/5). Six participants self-reported as experienced users (rating 3/5), and two of the participants claimed to be experts (rating 5/5).

## 5 RESULTS

We treated trials that required assistance as outliers and removed them from our evaluation. As assistance, we defined a waiting period of approximately 60 to 180s. By this time, we assumed that most participants clearly expressed that they do not understand the task and we offered subtle help. Furthermore, any attempts that exceeded a completion time (CT) of 60s were eliminated from the evaluation as we do not suspect correct understanding of the task or technique. In total, 17 out of 2430 trials were removed. Of these, 2 trials included HiveFive360 and 15 included SGD-VR as technique. All metrics of the questionnaires were complete and hence, there were no outliers or missing data.

We tested the data for normality using the Shapiro-Wilk test. Normality was not present in most data sets (Table 1). Since ANOVAs are resistant to Type-1 errors, we did not correct to a normal distribution [5]. We tested with Mauchly's sphericity test and in case of a lacking sphericity, applied Greenhouse-Geisser correction. For post-hoc analysis, we used the Games-Howell test. Additionally, data sets that lacked normal distribution were tested with the non-parametric Friedman test (Table 1).

### 5.1 Completion Time

SGD-VR had the highest completion time, followed by HiveFive360 and FlyingARrow (Figure 5a). The rmANOVA and the Friedman test assessing the impact of HT on CT yielded $p < 0.0001$ (Table 1a), indicating significant differences between the techniques. The post-hoc analysis revealed significant differences between all three techniques, with *FlyingARrow < HiveFive360 < SGD-VR.*

[5] http://www.igroup.org/pq/ipq/index.php
[6] https://humansystems.arc.nasa.gov/groups/tlx/
[7] www.limesurvey.org

## 5.2 Sense of Presence

HiveFive360 had the highest overall IPQ presence score, followed by SGD-VR and FlyingARrow. The rmANOVA revealed $p = 0.106$ (Table 1b), indicating no significant difference in IPQ scores between the highlighting techniques. Therefore, no post-hoc tests were performed. HiveFive360 does not outperform any other technique in regard to sense of presence. These results align with the intuitive understanding of the visual presentation in Figure 5b, where the mean values show a high degree of similarity.

Regarding the difference between VE types, both types revealed similar means and standard deviations (Figure 5c). The rmANOVA yielded $p = 0.840$ (Table 1c), indicating no significant differences between the techniques. HiveFive360 induced the same sense of presence for both indoor and outdoor environments.

## 5.3 Task Load

HiveFive360 and FlyingARrow presented a difference in mean by 0.65 for mental demand, while SGD-VR revealed the highest rating (Figure 5d). Both rmANOVA and Friedman test indicated significant differences (Table 1d). A post-hoc tests revealed differences between all techniques (*FlyingARrow* < *HiveFive360* < *SGD-VR*). Thus, the mental demand is the lowest for FlyingARrow.

Regarding effort, the difference in mean was 0.19 between HiveFive360 and FlyingARrow, SGD-VR again revealed the highest rating (Figure 5e). The rmANOVA and the Friedman test (Table 1e), and a consecutive post-hoc test revealed that SGD-VR was rated higher than HiveFive360 and FlyingARrow, with *HiveFive360, FlyingARrow* < *SGD-VR*. The effort is equally as low for both HiveFive360 and FlyingARrow.

## 5.4 System Usability

Based on Sauro and Lewis [40], 68 is an average SUS rating. Both HiveFive360 and FlyingARrow revealed above-average SUS ratings with means of 76.57 and 76.36, respectively (Figure 5f). Both had better usability than 77% of the products in the database of Sauro and Lewis [40] (grade: B [28]). The rmANOVA and the Friedman test (Table 1f), and post-hoc test indicated that they both outperformed SGD-VR, with *HiveFive, FlyingARrow* > *SGD-VR*. SGD-VR showed below-average performance with a mean of 45.99 (percentile rank: 8% [40], grade: F [28]).

## 5.5 Subjective Metrics

To evaluate the highlighting technique's ability to blend into the VE and not distract the user from the VE, we asked two 5-point scale questions inspired by Lange et al. [27]:

*ASM-1*: "How well did the highlighting technique blend into the virtual environment?" (1: It blended very badly into the environment. 5: It blended very well into the environment.)

*ASM-2*: "How much did the highlighting technique distract you from the virtual environment?" (1: It did not distract me at all. 5: It was very distracting.)

The rmANOVA and the Friedman test (Table 1g/h), and post-hoc tests revealed that HiveFive360 outperformed both FlyingARrow and SGD-VR in both ASM questions (Figure 6), indicating that HiveFive360 blended the best into the VEs and distracted the least from the VEs. These results align with the findings of Lange et al. [27].

# 6 DISCUSSION

*Predictable trajectories improve completion time.* Participants reported that FlyingARrow was the fastest due to its immediately predictable trajectory (ID5, ID9, ID12). The participants were able to select the target before the stimulus arrived there. In contrast, HiveFive360's trajectory was less predictable (ID12). ID9 appreciated FlyingARrow's distinct 3D shape and directional clarity, in contrast to HiveFive360's simpler "moving sphere" display that relies solely on changes in position for directional guidance. Compared to SGD-VR, HiveFive360 guided participants significantly faster. Based on these findings, we accept $H_1$.

*In-situ designs guide slower for distant targets.* Participants reported that they were sometimes slowed down by the speed of HiveFive360 or FlyingARrow, having to wait for the HT to reach the target (ID12). This was particularly influential for targets located further away (>6m). Similar effects were reported by the participants of Hu et al. when comparing FlyingARrow to two in-view techniques [22]. Additional data analysis revealed significant effects of the target distance on CT for HiveFive360 and FlyingARrow. The CT for HiveFive360 increased by ~17% for distant targets (>6m) and by ~20% for FlyingARrow, while it decreased by ~9% for SGD-VR.

*Participants remembered target locations and understood the functioning of the techniques.* ID12 mentioned that it was helpful to know the possible target positions beforehand in order to better guess the final object (ID12). When splitting the data between VEs and trials per VE (3 trials per VE), analysis revealed significant differences between first and second trials in the StreetCorner, and all trials in the DesertTemple environment. However, the results can also be explained by the assumption that participants had a better understanding of the overall task in the second trial. Another assumed learning effect might result from understanding the functioning the technique, especially regarding SGD-VR (ID5, ID7, ID9, ID11). We found significant differences between all trials, for all HTs.

*The IPQ yields no significant differences.* Given that overt and subtle cues resulted in the same IPQ scores, no conclusive assertion can be made about HiveFive360's ability to create a sense of presence. We found no significant differences between the techniques, thus, we cannot accept $H_2$ and $H_3$. Likewise, we found no differences between the uses of HiveFive360 in different VE types, so we cannot accept $H_4$ either. In contrast to our findings, Lange et al. [27] measured that HiveFive induced the greatest sense of presence, compared to SGD and Arrow. However, during our study, participants repeatedly expressed the feeling that they were filling in the same values for the IPQ over and over again. Additionally, some participants expressed frustration with the length of the IPQ. As a result, they rushed through the IPQ for the later stimuli and probably did not consider their answers properly. It is unclear whether the IPQ was the correct tool for measuring sense of presence in this particular study design.

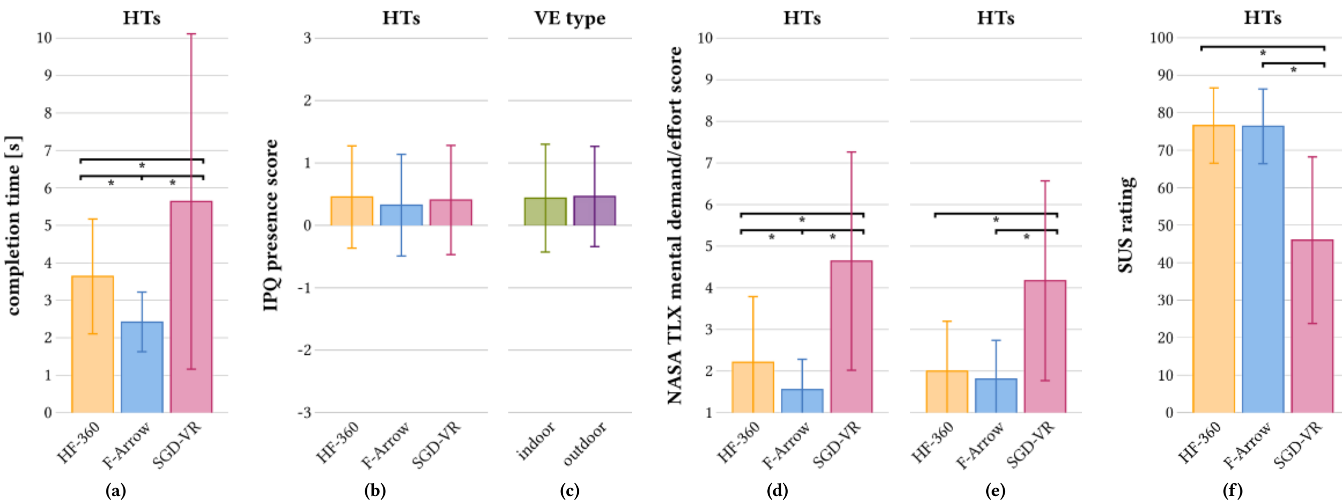


**Figure 5: Statistical evaluation of the comparison between HiveFive360 (ours), FlyingARrow [18], and SGD-VR [13]. Mean and standard deviation are visualized for (a) completion time, (b)–(c) IPQ presence score, (d) NASA TLX mental demand, (e) NASA TLX effort, and (f) System Usability Score (SUS). Significant differences ($p < 0.05$) are denoted by $*$.**

**Table 1: Results of the rmANOVAs and Friedman tests assessing the effect of HTs on (a) completion time (CT) (b) IPQ score (d) NASA TLX mental demand (MD) (e) NASA TLX effort (E) (f) SUS (g) ASM-1 (h) ASM-2, as well as the effect of (c) VE with HiveFive360 as HT on IPQ score.**

| | | HT on CT | HT on IPQ score | VE on IPQ score | HT on MD | HT on E | HT on SUS | HT on ASM-1 | HT on ASM-2 |
|---|---|---|---|---|---|---|---|---|---|
| | Normality | ✗ | ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ |
| rm-ANOVA | DF | 1.1, 28.6 | 2, 52 | 1, 26 | 1.5, 39 | 1.34, 34.84 | 1.32, 34.32 | 1.42, 36.92 | 2, 52 |
| | F | 100.27 | 2.34 | 0.03 | 42.04 | 47.51 | 50.00 | 9.97 | 9.10 |
| | $p$ | <0.0001 | 0.106 | 0.864 | <0.0001 | <0.0001 | <0.0001 | $2{\cdot}10^{-4}$ | $4{\cdot}10^{-4}$ |
| | $\eta^2$ | 0.66 | 0.01 | 0.0001 | 0.43 | 0.39 | 0.51 | 0.21 | 0.16 |
| | $\varepsilon$ | 0.55 | 0.94 | 1.0 | 0.75 | 0.67 | 0.66 | 0.71 | 0.89 |
| Fried-man | DF | 2 | | | 1.93, 50.07 | 1.93, 50.07 | 1.93, 50.07 | 1.93, 50.07 | 1.93, 50.07 |
| | $\chi^2$/F | $\chi^2 = 54$ | | | F = 72.95 | F = 49.27 | F = 36.06 | F = 9.83 | F = 8.47 |
| | $p$ | <0.0001 | | | <0.0001 | <0.0001 | <0.0001 | $3{\cdot}10^{-4}$ | $8{\cdot}10^{-4}$ |
| | | (a) | (b) | (c) | (d) | (e) | (f) | (g) | (h) |

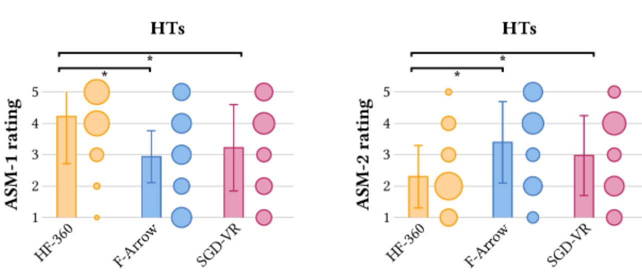


**Figure 6: Mean and standard deviation for our two additional subjective questions (ASM), split by highlighting technique (HT). The bubble chart shows the relative number of responses per rating, coded by area. Significant differences ($p < 0.05$) are denoted by $*$.**

*HiveFive360 has a low task load.* The highest mental demand and effort ratings for SGD-VR are consistent with the high potential for its misinterpretation. The cue was hard to see and hard to understand (ID5, ID7, ID9, ID10, ID11, ID19), resulting in high mental demand and effort ratings. In contrast, mental demand was perceived the lowest for FlyingARrow, likely because of its unambiguous, familiar design. More interpretation was required for HiveFive360. It was clearly visible, however, participants had to understand that it was meant to guide them. Based on these findings, we cannot accept $H_5$. As for effort, both HiveFive360 and FlyingARrow were not challenging. However, HiveFive360 did not outperform FlyingARrow, thus, we cannot accept $H_6$. For SGD-VR, participants had to work harder to complete the task by concentrating on seeing and understanding the cue's guidance.

Compared to the original findings of Lange et al. [27], the mean for mental demand increased only slightly, while the mean for effort remained unchanged. This is an interesting observation, as the additional difficulty introduced by the techniques movement did not result in an increased task load. It is possible that the additional movement may have enhanced the technique's visibility, thereby reducing its perceived task load.

Additional analysis revealed the lowest temporal demand for HiveFive360, indicating that the task was perceived calmest with this technique. However, comments about the pace of HiveFive360 were ambiguous. ID11 found the swarm to be the most relaxed technique, while ID16 stated that the constant movement of the particles and the swarm was "too much movement, too restless".

*Participants understood how to use HiveFive360.* FlyingARrow serves as a distinctive cue with high usability ratings. HiveFive360 was rated as usable as FlyingARrow, despite its higher completion times and mental demand, indicating an easy to understand design.

*HiveFive360 integrates seamlessly into environments.* This is potentially due to its real-world resemblance. Unlike the FlyingARrow, a swarm of insects is plausible in reality, which increases its ability to blend in. However, its responsiveness to user actions was perceived as artificial by one participant (ID9), reducing its environmental integration. The ratings of SGD-VR might be due to visibility issues and its ambiguous guidance (ID5, ID7, ID9–12, ID17, ID19), which shifts focus from the environment to deciphering the technique.

## 7 LIMITATIONS AND FUTURE WORK

*Study Design and Setup.* We assume that the length and repetitions of our study procedure might have lead to some rushing behaviour for the answers of our questionnaires. Additionally, we noted some learning effects in terms of the overall task and the target distribution which might have lead to shorter CTs for the later stimuli. The usage of IPQ as a means to measure the sense of presence may not have been sufficient, given that the participants just spent a few minutes in the VEs. The original website of the IPQ[6] does not provide information on the minimum duration required in the VE for the IPQ to be a suitable questionnaire.

In addition, we noted that participants entangled in the cable, occasionally causing them to turn less efficiently, which might have resulted in a longer CTs. We suppose that they became equally entangled for all stimuli, thus making this effect negligible. Furthermore, the handling of the VR controller could have led to longer CTs. Other studies have addressed this issue by allowing participants to stop the CT measurement by pressing a button and then selecting the target with a selection tool [27], or selecting the correct target by gaze [22].

*HiveFive360's Trajectory.* Regarding HiveFive360's trajectory, collider volumes around objects can be added to avoid collisions. This might improve the diegesis of the swarm by making it behave more realistically. Furthermore, the swarm consistently chose the shortest route, occasionally resulting in guidance via the floor or ceiling (ID12, ID17). This caused participants to search for the target on the floor or ceiling, and might have increased CT due to unnatural and inefficient movements. To address this issue, a more sophisticated pathfinding algorithm could be implemented that prioritizes horizontal guidance.

*Swarm's Visual Appearance.* The swarm's visual appearance impacted both target selection and scene design. The swarm exhibits a vivid yellow color that was not visible against bright textured backgrounds (e.g. marble floor) or yellow objects (e.g. the lounge chair in the IndoorRoom scene). To address this issue, we recolored the affected objects in the scenes. Alternatively, in future applications, the swarm could be visually adapted to the respective scene, which we have avoided for reason of comparability. Additionally, the tiny particles of the swarm were not visible for distant targets due to their small size, so distant objects were excluded as targets. In future applications, the particles could be scaled according to the distance between the swarm and the user to ensure swarm visibility. To increase realism, the swarm's appearance could be additionally tailored to each scene. For instance, in the IndoorRoom scene the swarm could comprise flies or dust particles.

## 8 CONCLUSION

HiveFive360 extends HiveFive [27] to highlight out-of-FOV targets by dynamically placing a swarm in the VE. HiveFive360 is developed in three different variants, which were tested in a qualitative concept validation study. The results indicated that participants preferred the swarm to be constantly visible inside their FOV and quickly react to their actions. The resulting version of HiveFive360 was then tested in a user study against two other highlighting techniques: FlyingARrow [18] and SGD-VR [13].

The results of the user study demonstrate that HiveFive360 guides participants to out-of-FOV targets faster than the other subtle technique, SGD-VR. However, guiding with the overt cue FlyingARrow results in faster completion times than both HiveFive360 and SGD-VR. Furthermore, HiveFive360 outperforms SGD-VR in the NASA TLX dimensions mental demand and effort. In comparison to FlyingARrow, HiveFive360 demonstrates worse performance in mental demand and similar results in effort, though. Regarding the sense of presence, no differences between the techniques are observed.

Further evaluations show that HiveFive360 blends best into the virtual environments and distracts the least, surpassing both SGD-VR and FlyingARrow. HiveFive360 achieves a SUS score of 77 points, denoting good usability. Qualitative feedback indicates that participants appreciate the pace of HiveFive360's guidance but are dissatisfied with its trajectory, which has a tendency to deviate from horizontal guidance. Additionally, they favored FlyingARrow's ability to guide them directly and quickly to the target, while strongly disliking SGD-VR's overall visual and behavioral design.

Overall, HiveFive360 shows promise for efficiently guiding users in various environments without excessive distraction or task load.

## REFERENCES

[1] Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle Gaze Direction. *ACM Trans. Graphics (TOG)* 28, 4 (2009), 1–14. https://doi.org/10.1145/1559755.1559757

[2] Patrick Baudisch and Ruth Rosenholtz. 2003. Halo: A Technique for Visualizing off-Screen Objects. In *Proc. Conf. Human Factors in Computing Systems (CHI)*. ACM, 481–488. https://doi.org/10.1145/642611.642695

[3] Nicola Binetti, Luyan Wu, Shiping Chen, Ernst Kruijff, Simon Julier, and Duncan P. Brumby. 2021. Using visual and auditory cues to locate out-of-view objects in head-mounted augmented reality. *Displays* 69 (2021), 1–9. https://doi.org/10.1016/j.displa.2021.102032

[4] Frank Biocca, Arthur Tang, Charles Owen, and Fan Xiao. 2006. Attention Funnel: Omnidirectional 3D Cursor for Mobile Augmented Reality Platforms. In *Proc. Conf. Human Factors in Computing Systems (CHI)*. ACM, 1115–1122. https://doi.org/10.1145/1124772.1124939

[5] María José Blanca, Rafael Alarcón Postigo, Jaume Arnau Gras, Roser Bono Cabré, and Rebecca Bendayan. 2017. Non-normal data: Is ANOVA still a valid option? *Psicothema* 29, 4 (2017), 552–557. https://doi.org/10.7334/psicothema2016.383

[6] Felix Bork, Christian Schnelzer, Ulrich Eck, and Nassir Navab. 2018. Towards Efficient Visual Guidance in Limited Field-of-View Head-Mounted Displays. *IEEE Trans. Visualization and Computer Graphics (TVCG)* 24, 11 (2018), 2983–2992. https://doi.org/10.1109/TVCG.2018.2868584

[7] John Brooke. 1995. SUS: A quick and dirty usability scale. *Usability Eval. Ind.* 189 (1995), 1–6. https://doi.org/10.1201/9781498710411-35

[8] Chong Cao, Zhaowei Shi, and Miao Yu. 2020. Automatic Generation of Diegetic Guidance in Cinematic Virtual Reality. In *IEEE Int. Symp. Mixed and Augmented Reality (ISMAR)*. 600–607. https://doi.org/10.1109/ISMAR50242.2020.00087

[9] Nina Doerr, Katrin Angerbauer, Melissa Reinelt, and Michael Sedlmair. 2023. Bees, Birds and Butterflies: Investigating the Influence of Distractors on Visual Attention Guidance Techniques. In *Ext. Abstr. Conf. Human Factors in Computing Systems (CHI EA)*. ACM, 1–7. https://doi.org/10.1145/3544549.3585816

[10] Matt Gottsacker, Nahal Norouzi, Kangsoo Kim, Gerd Bruder, and Greg Welch. 2021. Diegetic Representations for Seamless Cross-Reality Interruptions. In *Proc. IEEE Int. Symp. Mixed and Augmented Reality (ISMAR)*. 310–319. https://doi.org/10.1109/ISMAR52148.2021.00047

[11] Steve Grogorick, Georgia Albuquerque, and Marcus Magnor. 2018. Comparing Unobtrusive Gaze Guiding Stimuli in Head-Mounted Displays. In *Proc. IEEE Int. Conf. Image Processing (ICIP)*. 2805–2809. https://doi.org/10.1109/ICIP.2018.8451784

[12] Steve Grogorick, Georgia Albuquerque, Jan-Philipp Tauscher, and Marcus Magnor. 2018. Comparison of Unobtrusive Visual Guidance Methods in an Immersive Dome Environment. *ACM Trans. Applied Perception (TAP)* 15, 4 (2018), 1–11. https://doi.org/10.1145/3238303

[13] Steve Grogorick, Michael Stengel, Elmar Eisemann, and Marcus Magnor. 2017. Subtle Gaze Guidance for Immersive Environments. In *Proc. ACM Symp. on Applied Perception (SAP)*. 1–7. https://doi.org/10.1145/3119881.3119890

[14] Steve Grogorick, Jan-Philipp Tauscher, Georgia Albuquerque, Marc Kassubeck, and Marcus Magnor. 2019. Towards VR Attention Guidance: Environment-Dependent Perceptual Threshold for Stereo Inverse Brightness Modulation. In *Proc. ACM Symp. on Applied Perception (SAP)*. 1–5. https://doi.org/10.1145/3343036.3343137

[15] Steve Grogorick, Jan-Philipp Tauscher, Nikkel Heesen, Susana Castillo, and Marcus Magnor. 2020. Stereo Inverse Brightness Modulation for Guidance in Dynamic Panorama Videos in Virtual Reality. *Computer Graphics Forum* 39, 6 (2020), 542–553. https://doi.org/10.1111/cgf.14091

[16] Uwe Gruenefeld, Abdallah El Ali, Wilko Heuten, and Susanne Boll. 2017. Visualizing Out-of-View Objects in Head-Mounted Augmented Reality. In *Proc. Int. Conf. Human-Computer Interaction with Mobile Devices and Services (MobileHCI)*. ACM, 1–7. https://doi.org/10.1145/3098279.3122124

[17] Uwe Gruenefeld, Dag Ennenga, Abdallah El Ali, Wilko Heuten, and Susanne Boll. 2017. EyeSee360: Designing a Visualization Technique for out-of-View Objects in Head-Mounted Augmented Reality. In *Proc. Symp. Spatial User Interaction (SUI)*. ACM, 109–118. https://doi.org/10.1145/3131277.3132175

[18] Uwe Gruenefeld, Daniel Lange, Lasse Hammer, Susanne Boll, and Wilko Heuten. 2018. FlyingARrow: Pointing Towards Out-of-View Objects on Augmented Reality Devices. In *Proc. ACM Int. Symp. Pervasive Displays (PerDis)*. 1–6. https://doi.org/10.1145/3205873.3205881

[19] Sean Gustafson, Patrick Baudisch, Carl Gutwin, and Pourang Irani. 2008. Wedge: Clutter-Free Visualization of Off-Screen Locations. In *Proc. Conf. Human Factors in Computing Systems (CHI)*. ACM, 787–796. https://doi.org/10.1145/1357054.1357179

[20] Sean Gustafson and Pourang Irani. 2007. Comparing Visualizations for Tracking Off-Screen Moving Targets. In *Ext. Abstr. Human Factors in Computing Systems (CHI EA)*. ACM, 2399–2404. https://doi.org/10.1145/1240866.1241014

[21] Carl Gutwin, Andy Cockburn, and Ashley Coveney. 2017. Peripheral Popout: The Influence of Visual Angle and Stimulus Intensity on Popout Effects. In *Proc. Conf. Human Factors in Computing Systems (CHI)*. ACM, 208–219. https://doi.org/10.1145/3025453.3025984

[22] Sathaporn "Hubert" Hu, Joseph Malloch, and Derek Reilly. 2021. A Comparative Evaluation of Techniques for Locating Out-of-View Targets in Virtual Reality. In *Proc. Graphics Interface (GI)*. Canadian Information Processing Society, 202–212. https://doi.org/10.20380/gi2021.32

[23] Hyungeun Jo, Sungjae Hwang, Hyunwoo Park, and Jung-hee Ryu. 2011. Mobile Augmented Reality: Aroundplot: Focus+Context Interface for off-Screen Objects in 3D Environments. *Computers & Graphics* 35, 4 (2011), 841–853. https://doi.org/10.1016/j.cag.2011.04.005

[24] Sophie Kergaßner, Nina Doerr, Markus Wieland, and Michael Sedlmair. 2023. Highlighting Out-Of-View Targets in XR Applications. www.osf.io/6dtqr

[25] Andrey Krekhov, Sebastian Cmentowski, Andre Waschk, and Jens Krüger. 2020. Deadeye Visualization Revisited: Investigation of Preattentiveness and Applicability in Virtual Environments. *IEEE Trans. Visualization and Computer Graphics (TVCG)* 26, 1 (2020), 547–557. https://doi.org/10.1109/tvcg.2019.2934370

[26] Andrey Krekhov and Jens Krüger. 2019. Deadeye: A Novel Preattentive Visualization Technique Based on Dichoptic Presentation. *IEEE Trans. Visualization and Computer Graphics (TVCG)* 25, 1 (2019), 936–945. https://doi.org/10.1109/tvcg.2018.2864498

[27] Daniel Lange, Tim Claudius Stratmann, Uwe Gruenefeld, and Susanne Boll. 2020. HiveFive: Immersion Preserving Attention Guidance in Virtual Reality. In *Proc. Conf. Human Factors in Computing Systems (CHI)*. ACM, 1–13. https://doi.org/10.1145/3313831.3376803

[28] James R. Lewis and Jeff Sauro. 2021. Item Benchmarks for the System Usability Scale. *Journal of User Experience* 13, 3 (2021), 158–167. https://uxpajournal.org/item-benchmarks-system-usability-scale-sus/

[29] Kunpeng Li, Ziyan Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. 2018. Tell Me Where to Look: Guided Attention Inference Network. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*. 9215–9223. https://doi.org/10.1109/cvpr.2018.00960

[30] Daniela Markov-Vetter, Martin Luboschik, ABM Tariqul Islam, Peter Gauger, and Oliver Staadt. 2020. The Effect of Spatial Reference on Visual Attention and Workload during Viewpoint Guidance in Augmented Reality. In *Proc. ACM Symp. Spatial User Interaction (SUI)*. 1–10. https://doi.org/10.1145/3385959.3418449

[31] Ann McNamara, Reynold Bailey, and Cindy Grimm. 2009. Search Task Performance Using Subtle Gaze Direction with the Presence of Distractions. *ACM Trans. Applied Perception (TAP)* 6, 3 (2009), 1–19. https://doi.org/10.1145/1577755.1577760

[32] Ann McNamara, Thomas Booth, Srinivas Sridharan, Stephen Caffey, Cindy Grimm, and Reynold Bailey. 2012. Directing Gaze in Narrative Art. In *Proc. ACM Symp. Applied Perception (SAP)*. 63–70. https://doi.org/10.1145/2338676.2338689

[33] Lasse T. Nielsen, Matias B. Møller, Sune D. Hartmeyer, Troels C. M. Ljung, Niels C. Nilsson, Rolf Nordahl, and Stefania Serafin. 2016. Missing the Point: An Exploration of How to Guide Users' Attention during Cinematic Virtual Reality. In *Proc. ACM Conf. Virtual Reality Software and Technology (VRST)*. 229–232. https://doi.org/10.1145/2993369.2993405

[34] Randy Pausch, Jon Snoddy, Robert Taylor, Scott Watson, and Eric Haseltine. 1996. Disney's Aladdin: First Steps Toward Storytelling in Virtual Reality. In *Proc. Ann. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM, 193–203. https://doi.org/10.1145/237170.237257

[35] Patrick Perea, Denis Morand, and Laurence Nigay. 2017. Halo3D: A Technique for Visualizing off-Screen Points of Interest in Mobile Augmented Reality. In *Proc. Conf. l'Interaction Homme-Machine (IHM)*. ACM, 43–51. https://doi.org/10.1145/3132129.3132144

[36] Jay Pratt, Petre V. Radulescu, Ruo Mu Guo, and Richard A. Abrams. 2010. It's Alive!: Animate Motion Captures Visual Attention. *Psychological Science* 21, 11 (2010), 1724–1730. https://doi.org/10.1177/0956797610387440

[37] Patrick Renner and Thies Pfeiffer. 2017. Attention Guiding Techniques Using Peripheral Vision and Eye Tracking for Feedback in Augmented-Reality-Based Assistance Systems. In *Proc. IEEE Symp. 3D User Interfaces (3DUI)*. 186–194. https://doi.org/10.1109/3dui.2017.7893338

[38] Craig W. Reynolds. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. *ACM Computer Graphics (SIGGRAPH)* 21, 4 (1987), 25–34. https://doi.org/10.1145/37402.37406

[39] Sylvia Rothe, Daniel Buschek, and Heinrich Hußmann. 2019. Guidance in Cinematic Virtual Reality-Taxonomy, Research Status and Challenges. *Multimodal Technologies and Interaction* 3, 1 (2019), 1–23. https://doi.org/10.3390/mti3010019

[40] Jeff Sauro and James R. Lewis. 2016. *Quantifying the User Experience: Practical Statistics for User Research*. Morgan Kaufmann.

[41] Clifton M. Schor. 2011. Neural Control of Eye Movements. In *Adler's Physiology of the eye*. Saunders Elsevier, 220–242.

[42] Bjorn Schwerdtfeger and Gudrun Klinker. 2008. Supporting order picking with Augmented Reality. In *Proc. IEEE/ACM Int. Symp. Mixed and Augmented Reality (ISMAR)*. 91–94. https://doi.org/10.1109/ISMAR.2008.4637331

[43] Arne Seeliger, Gerrit Merz, Christian Holz, and Stefan Feuerriegel. 2021. Exploring the Effect of Visual Cues on Eye Gaze During AR-Guided Picking and Assembly Tasks. In *IEEE Int. Symp. Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 159–164. https://doi.org/10.1109/ismar-adjunct54149.2021.00041

[44] Teresa Siu and Valeria Herskovic. 2013. SidebARs: Improving Awareness of off-Screen Elements in Mobile Augmented Reality. In *Proc. Chilean Conf. Human-Computer Interaction (ChileCHI)*. ACM, 36–41. https://doi.org/10.1145/2535597.2535608

[45] Marco Speicher, Christoph Rosenberg, Donald Degraen, Florian Daiber, and Antonio Krúger. 2019. Exploring Visual Guidance in 360-degree Videos. In *Proc. Int. Conf. Interactive Experiences for TV and Online Video (TVX)*. ACM, 1–12. https://doi.org/10.1145/3317697.3323350

[46] N. Waldin, M. Waldner, and I. Viola. 2017. Flicker Observer Effect: Guiding Attention Through High Frequency Flicker in Images. *Computer Graphics Forum* 36, 2 (2017), 467–476. https://doi.org/10.1111/cgf.13141

[47] Jonathan Wieland, Rudolf C. Hegemann Garcia, Harald Reiterer, and Tiare Feuchtner. 2022. Arrow, Bézier Curve, or Halos? : Comparing 3D Out-of-View Object Visualization Techniques for Handheld Augmented Reality. In *IEEE Int. Symp. Mixed and Augmented Reality (ISMAR)*. 797–806. https://doi.org/10.1109/ismar55827.2022.00098

[48] Jason Woodworth, Andrew Yoshimura, Nicholas Lipari, and Christoph Borst. 2023. Design and Evaluation of Visual Cues for Restoring and Guiding Visual Attention in Eye-Tracked VR. In *IEEE Conf. Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 442–450. https://doi.org/10.1109/VRW58643.2023.00096

[49] Polle T. Zellweger, Jock D. Mackinlay, Lance Good, Mark Stefik, and Patrick Baudisch. 2003. City Lights: Contextual Views in Minimal Space. In *Ext. Abstr. Human Factors in Computing Systems (CHI EA)*. ACM, 838–839. https://doi.org/10.1145/765891.766022