Subtle Gaze Guidance for Immersive Environments

Steve Grogorick
TU Braunschweig
Germany

Michael Stengel
Elmar Eisemann
TU Delft
Netherlands

Marcus Magnor
TU Braunschweig
Germany

ABSTRACT

Immersive displays allow presentation of rich video content over a wide field of view. We present a method to boost visual importance for a selected - possibly invisible - scene part in a cluttered virtual environment. This desirable feature enables to unobtrusively guide the gaze direction of a user to any location within the immersive 360° surrounding. Our method is based on subtle gaze direction which did not include head rotations in previous work. For covering the full 360° environment and wide field of view, we contribute an approach for dynamic stimulus positioning and shape variation based on eccentricity to compensate for visibility differences across the visual field. Our approach is calibrated in a perceptual study for a head-mounted display with binocular eye tracking. An additional study validates the method within an immersive visual search task.

CCS CONCEPTS

• **Computing methodologies → Perception**; **Virtual reality**; *Image processing*;

KEYWORDS

gaze guidance, virtual reality, eye tracking, perceptual study, immersive applications, head-mounted display

ACM Reference format:
Steve Grogorick, Michael Stengel, Elmar Eisemann, and Marcus Magnor. 2017. Subtle Gaze Guidance for Immersive Environments. In *Proceedings of SAP '17, Cottbus, Germany, September 16–17, 2017,* 7 pages.
https://doi.org/10.1145/3119881.3119890

1 INTRODUCTION

The availability of consumer-level wide field-of-view head-mounted displays has triggered the “third wave” in the field of Virtual Reality (VR) and continuously brings more immersive applications to the consumer. Current GPUs and 360° cameras enable richer rendered and captured content than ever experienced before. However, presentation quality of the virtual environment and the resulting perceived impression for the user is still far away from perceiving reality. As a result a user in VR might quickly feel overstrained and miss intended actions and processes within the application, e.g. when directing the attention to a non-intended part of the virtual surrounding.

In this work, we present *immersive subtle gaze guidance*, a real-time image processing technique for wide field-of-view displays which unobtrusively grabs the user's attention and triggers refocusing the view to any selected target location in the 360° environment. Thus, the method can be used to assist the viewer in orientation and navigation in the virtual surrounding without reducing immersion. The method works independent from the presented scene content and is therefore generally applicable. Our method is based on previous work on subtle gaze guidance but includes significant changes in the presented guidance stimulus in order to enable guidance across the full field of view.

In particular, we contribute:

- a model for shape adjustment of the guidance stimulus based on eccentricity
- dynamic stimulus positioning including the user's head rotation and gaze direction
- results from three perceptual studies for calibration and validation of the immersive subtle gaze guidance model.

The paper is structured as follows: First, Section 2 discusses related work on visual gaze guidance. Section 3 gives few details on the gaze direction method we build upon, and then describes our extensions for use in immersive environments. In Section 4, two experiments are conducted to determine perception–driven parameters of the proposed method, in particular eccentricity–based stimulus size. Next, Section 5 presents a user study to validate our findings from the previous experiments in a search task. Section 6 gives a discussion and conclusions on our findings as well as possible directions for future research.

2 RELATED WORK

Under normal circumstances attention is guided by visual features as well as the task of the user. This property is exploited for passive gaze prediction. Strategies for gaze guidance are aiming for steering attention to a specified target location which can differ significantly from the natural fixation location. Therefore, visual gaze guidance requires altering the visible scene content.

A great deal of work exploits gaze direction techniques to guide the users's attention in graphical interfaces [Andrist et al. 2017], training applications [Andrist et al. 2017], driving assistance [Pomarjanschi et al. 2012] and visualization [de Koning and Jarodzka 2017]. Approaches can be classified by the visibility of the used guidance stimulus into *overt* and *subtle* gaze direction methods.

Author's addresses: S. Grogorick and M. Magnor, Computer Graphics Lab, TU Braunschweig; M. Stengel and E. Eisemann, Computer Graphics and Visualization Institute, TU Delft. Steve Grogorick and Marcus Magnor gratefully acknowledge funding by the German Science Foundation (DFG MA2555/15-1 “Immersive Digital Reality” and DFG INST 188/409-1 FUGG “ICG Dome”).

SAP '17, September 16–17, 2017, Cottbus, Germany

2.1 Overt Gaze Direction

Overt Gaze Direction (OGD) makes usage of a global image transformation to trigger a saccade towards the target location. Explored methods make use of blur [Hata et al. 2016; Jarodzka et al. 2013; Lintu and Carbonell 2009], highlighting [Khan et al. 2005], saliency adjustment [Dorr et al. 2008; Sato et al. 2016; Vig et al. 2011], or add a clearly recognizable image overlay such as arrows [Lin et al. 2017], leader lines [Hoffmann et al. 2008], or funnel [Biocca et al. 2007] directing to the target location. Kosara et al. [2002] introduce the semantic depth-of-field for guidance based on the observation that gaze is attracted by high frequencies. Cole et al. [2006] use stylized rendering including desaturation and blur for non-target regions guiding the view in static images. For guidance in immersive 360° videos Lin et al. [2017] present the autopilot technique which rotates the view automatically to the most salient object. Although OGD methods often provide high success rates, they might be not applicable as they disturb the overall viewing experience [Gutwin and Fedak 2004].

2.2 Subtle Gaze Direction

Subtle Gaze Direction (SGD) tries to keep the visibility of the guidance stimulus to a minimum. Such stimuli are designed to be only perceivable in the periphery of the visual field and usually disappear when directly looking at them. Barth et al. [2006] enable gaze guidance for videos by augmenting the video with small bright red dots appearing at the target location exploiting the fact that sudden object onsets in the periphery attract attention. Sheikh et al. [2016] present gaze guidance for 360° videos using filming techniques such as motion, audio cues or actor gestures which turned out to be effective if applicable for the respective video content. Bailey et al. [2009] introduce a more subtle, yet effective gaze guidance strategy. The authors apply temporally varying image space modulations in the luminance channel of a static image to guide a viewer's gaze through the scene without interrupting their visual experience. The principle has been successfully applied to increase search task performance in desktop settings [McNamara et al. 2009] and for projection-based AR [Booth et al. 2013] as well as to direct gaze in narrative art [McNamara et al. 2012]. A closely related approach applies static image space contrast and lightness modulation in CIE L*a*b space to chosen image parts of video sequences [Veas et al. 2011]. Lu et al. [2014] use a visual cue overlaying the image and test attributes such as transparency, size and shape to evaluate scene-dependent parameter sets based on the present amount of clutter. Waldin et al. [2017] present passive SGD for desktop monitors with high refresh rate displays. The method uses dynamic flicker frequencies calibrated across the visual field so that visibility of a stimulus can be controlled without eye tracking just by its size and flicker rate. In contrast to the latter approach, we model stimulus thresholds for wider eccentricities using immersive displays and include guidance to potentially invisible target locations.

2.3 Immersive Environments

Previous approaches have been mostly used to guide users for non-immersive applications on displays providing a rather narrow field of view. Waldin et al. [2017] explore eccentricities up to 30° in their work. To the best of our knowledge we are the first to explore SGD visibility thresholds for a eccentricities up to 45° horizontally and vertically to optimize SGD for Virtual Reality and Augmented Reality applications. Many existing SGD methods measure the current gaze location of the viewer by eye tracking and control location and visibility of the guidance stimulus accordingly. Such gaze-contingent methods turned out to be very successful for subtle gaze guidance [Booth et al. 2013; McNamara et al. 2009, 2012]. Passive methods replace eye tracking by a saliency estimator [Dorr et al. 2008] to approximate the viewer's current gaze location. Removing the eye tracker from the pipeline reduces guidance precision but is useful when guiding many viewers at the same time [Waldin et al. 2017].

Our approach is designed to guide the viewpoint of a single person wearing an immersive AR/VR headset with high precision and as subtle as possible. Recently, more devices including eye tracking have been accessible [Stengel et al. 2015]. Reasoned by the many advantages given by precise, low-latency eye tracking such as gaze-contingent rendering [Weier et al. 2017], we follow this path when steering the stimulus location and visibility. Hence, our immersive gaze direction technique extends previous SGD techniques [Bailey et al. 2009; Booth et al. 2013] in order to allow guidance to target locations at arbitrary eccentricities in the periphery and even to locations outside the visible visual field.

3 METHOD

The original SGD technique has been developed for a desktop environment [Bailey et al. 2009]. We extend this approach by adding stimulus shape variation in order to enable stimuli presentation at greater eccentricity and dynamic stimulus positioning to allow for guidance to locations outside the current field of view, in the following described as *external target locations*.

3.1 Subtle Gaze Direction Stimuli

In the original method, the guidance stimulus is set to a fixed 2D target location on the screen where the user's gaze should be moved to. The stimulus is shown as a circular region of 0.76° visual angle in which the original color of underlying pixels smoothly alternates between 9.5 % black or white with a frequency of 10 Hz (luminance modulation) [Bailey et al. 2009]. As an alternative a warm–cool color modulation computed from the original color mixed with red/blue has been tested but resulted in a weaker guidance performance. Hence, we make use of luminance modulation using the same frequency and modulation intensity for high success in gaze guidance. A radial gaussian fall-off is used to smoothly blend the stimulus into the surrounding image area. To keep stimuli as subtle as possible, the effect was switched off as soon as the user performed a saccade with a maximum of 10° deviation towards the target. A stationary eye tracking system was used to track the users' gaze in real–time.

3.2 Immersive Subtle Gaze Guidance

Two properties being key for immersive applications prohibit direct usage of the original SGD method in VR:

(1) The enlarged field of view. in immersive environments invalidates usage of simple circular stimuli. Caused by the perspective projection circular objects in the far periphery of the users' sight

would result in thin-shaped ellipses as seen from the users’ point of view.

(2) Increased navigational abilities. of users, i.e. turning the head into different directions and greater eye motion to orient in the scene, introduce the possibility for external target locations.

Our method introduces *dynamic stimulus shape adaptation* to counteract (1) and *dynamic stimulus positioning* to induce head rotations towards external targets (2).

Dynamic Stimulus Shape Adaptation. The problem of augmenting a rendered frame with a simple circular shape is that it would only be perceived as such when being viewed in a perpendicular direction, i.e. in the central region in case of HMDs. In outer regions the shape would appear strongly distorted as depicted in Fig. 1a–b. Therefore, the circular shape of the stimulus must be adapted, i.e. elongated towards the periphery, to instead be perceived circular as in Fig. 1c–d.

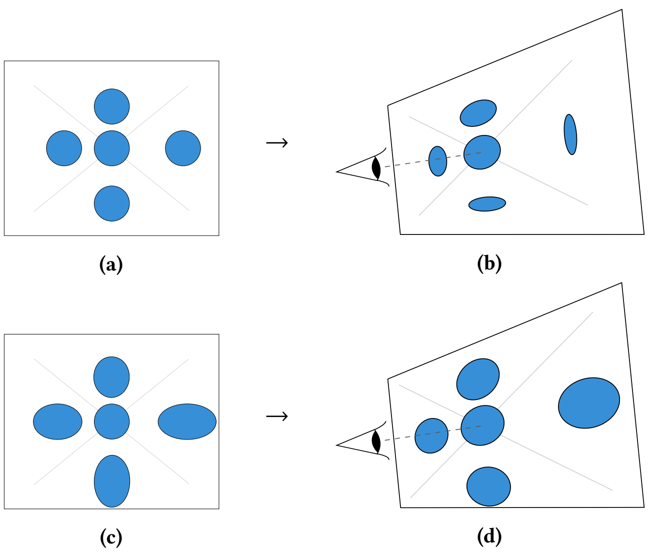

Figure 1: (a) Five equal-sized circles on a plane straight in front of the viewer. (b) With increasing eccentricity the circles appear squeezed when being viewed from the viewer’s perspective. (c–d) The same example with perspectively corrected shapes.

Assuming the stimulus is generated in a fragment shader with the current gaze–to–stimulus angular distance e (aka eccentricity) being available as uniform input, given in radians. First, the angle between eye–to–target direction d_{target} and eye–to–fragment direction $d_{fragment}$ is calculated via dot product. This angle is then scaled by the eccentricity–dependent stimulus size $s(e)$ and smoothed with a sigmoid-like function f to get the fragment’s stimulus intensity

$$i_{stimulus} = f\left(\frac{\arccos(||d_{target}|| \cdot ||d_{fragment}||)}{s(e)}\right) \qquad (1)$$

with f, in our case being the OpenGL function *smoothstep*, that applies Hermite interpolation between 1 and 0. Afterwards, stimulus color and frame color interpolation is applied identical to the original implementation.

Dynamic Stimulus Positioning. For external target locations, e.g. behind the user, it is not sufficient anymore to only steer the user’s eyes. Instead we need to induce a rotation of the head or even the whole body. We, therefore, introduce another modification that dynamically adds motion to the stimulus. More precisely, the stimulus performs a repeated motion towards the screen edge, in which direction the user should turn to reach the target, as depicted in Fig. 2. First, the 3D target location d_{target} is projected on it’s 2D

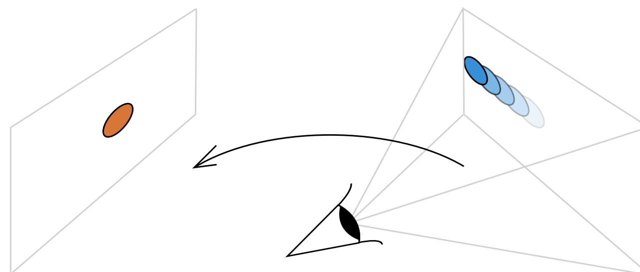

Figure 2: Stimulus (blue) moves out of the user’s field of view to induce a head rotation towards the target (orange).

position p_{target} in the image plane via the current projection matrix, and clamped ($p'_{clamped}$) to the virtual view frustum bounds via division by the maximum of its absolute component values (maxAbsComp). If the target location is behind the virtual camera the value is negated ($p_{clamped}$). We then project $p_{clamped}$ back using the inverse of the projection matrix to get the corresponding 3D direction $d_{clamped}$. Now we select the start direction of the motion d_{start} as the (in camera space) forward direction $d_{forward} = (0, 0, -1)$ rotated ($M_{rotation}$) about an angle α towards $d_{clamped}$. Finally the moving stimulus position $p_{stimulus}$ is a time–driven linear interpolate (lerp) of d_{start} and $d_{clamped}$:

$$p'_{clamped} = p_{target} / \text{maxAbsComp}(p_{target}) \qquad (2)$$

$$p_{clamped} = p'_{clamped} * \text{sign}(d_{forward} \cdot d_{target}) \qquad (3)$$

$$d_{start} = M_{rotation} * d_{forward} \qquad (4)$$

$$p_{stimulus} = \text{lerp}(d_{start}, d_{clamped}, t) \qquad (5)$$

with $t \in [0, 1]$ being the fractional part of the current time in seconds. The rotation axis of the matrix $M_{rotation}$ is derived via cross product of $d_{forward}$ and $d_{clamped}$. A fixed angle of $\alpha = 35°$ turned out to be an appropriate value in all our experiments.

4 PERCEPTUAL STUDY

A first study was conducted to estimate the stimulus size distribution $s(e)$ for different users and across the available range of eccentricities in a VR headset. To account for static as well as dynamic scenarios, the study was split into two corresponding parts. For the first part a fully static scene was used, with all objects staying at a fixed position within the users visual field. During the second part, again a mostly static scene was used, but this time participants had to move and turn around in order to complete the given task.

4.1 Static Scenario

The first and static scenario calibrates gaze guidance for such virtual environments that do not require large head movements. 12 participants, recruited at a local university, volunteered to take part in this experiment – 3 female and 9 male, aged 26 to 48. Familiarity with VR was reported very low to not existing for all participants. None of them reported any experience with eye tracking.

A possible example application might be a virtual museum. While virtually staying in front of a painting or sculpture and listening to a description of the exhibit, gaze guidance might be used to assist the viewer by guiding his gaze unobtrusively towards the currently described parts. Therefore, subtle gaze guidance allows to maintain the immersion inside the virtual world and contrasts strongly to highlighting techniques which, for example, add a solid 3D marker to the target region or manipulate the user's orientation [Lin et al. 2017].

Experiment Setup. We used the *SMI Eye Tracking HMD based on the HTC Vive* to present the virtual environments to participants. The contents were shown using a custom built in–house OpenGL renderer, utilizing the *OpenVR SDK*[1] to provide basic HMD functionalities and the *iViewHMD-HTC SDK*[2] for 250 Hz binocular eye tracking. The *iViewHMD-HTC SDK* provides a 5–point calibration procedure (simultaneously for both eyes), which we have integrated and used to calibrate the eye tracking system at the beginning of every experiment. We performed the experiments on a desktop computer with an Intel(R) Core(TM) i7-6700 CPU and an NVIDIA GeForce GTX 970 GPU.

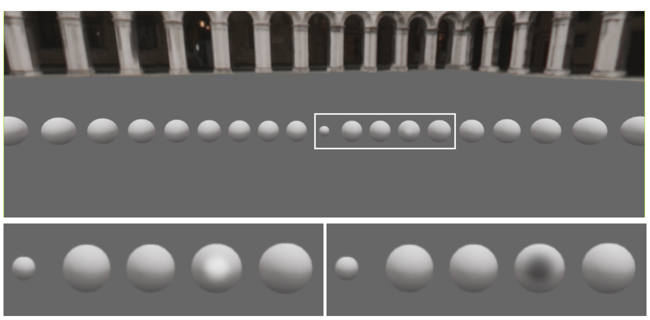

Figure 3: (Top) Virtual environment set-up for static scenario as seen by participants. (Bottom) Insets of exemplary stimulus on the 3rd right sphere, shown for white (left) and black (right) interpolation (effect exaggerated for depiction).

The virtual environment, as depicted in Fig. 3a, contained 19 equal sized floating spheres at fixed positions relative to the participant's head position and orientation. The small sphere in the middle was placed virtually 1.5 m in front of the user's head, i.e. in front of a virtual camera without eye offset. Further spheres were placed every 5° visual angle (until 45°) horizontally in both directions starting from the small sphere, covering 90° horizontal visual field of view. The position of all spheres was kept constant relative to the participant's head during the whole experiment.

Procedure. Participants were seated on a swivel chair, wearing the eye-tracking VR headset. First, the task was explained to the participants: Their goal is to determine the required stimulus size at each sphere's position for which the stimulus becomes just noticeable when looking straight at the small sphere in the center. The stimulus size was initialized with zero at all spheres' positions. The participants have been asked to, for one sphere at a time, slowly increase the stimulus size until the stimulus becomes perceivable. Using a common keyboard, the stimulus diameter could be increased and decreased via *up* and *down arrow keys*, with a step size of 0.02° visual angle. When satisfied with the stimulus size for a certain sphere, they could switch to the next one. Switching between spheres was done via *left* and *right arrow keys*. The stimulus size of a sphere was automatically saved when switching to a different one and restored upon return. Finally, they were reminded to look straight at the small sphere for the whole time. They were allowed to look around and relax their eyes directly after switching to a new sphere, i.e. before again starting to increase stimulus size. The experiment was stopped if they found a threshold for every visible sphere's position. Due to inter–individual differences, especially in interpupillary distance, some participants were not able to see the outermost spheres through the HMD's lenses. We, therefore, excluded the two outermost spheres (45°) from the dataset.

4.2 Dynamic Scenario

The second, dynamic scenario was conducted to cover applications that do require a significant amount of dynamic user movement within a mostly static virtual environment. 17 participants took part in this experiment – 5 female and 12 male, aged 22 to 48 – of which 3 were also part of the first experiment. Familiarity with VR was reported very low to not existing for all participants, again. Also, none of them reported any experience with eye tracking.

Here, a possible application might be visual inspection training for aircrafts. Gaze guidance could be activated after some time of unsuccessful search for defects. This way, unexperienced users would be supported while retaining a high degree of immersion.

To get a high degree of user movement, we have designed an endless search task that serves as 'background task' during the second experiment. This way, by keeping participants searching (rotating their heads and eyes), we are able to estimate the required stimulus size during such user movements.

Experiment Setup. We used the same hardware and software setup as for the static scenario, described in Sec. 4.1.

The virtual environment was generated as shown in Fig. 4a–b. It was populated by randomly selecting and duplicating four template objects (see Fig. 4c pyramid to sphere) 200 times. Then an additional sphere was added only once and all objects are shuffled, including object position, orientation and color. Fig. 4b shows a view from above that reveals an exemplary arrangement of all objects in a 270° portion of a spherical segment centered at the user. All objects had a virtual size of $20 \times 20 \times 20$ cm and a distance to the observer

[1] https://github.com/ValveSoftware/openvr
[2] https://web.archive.org/web/20161216214330/http://update.smivision.com/iViewHMDHTCscripts/ReleaseNotes.html

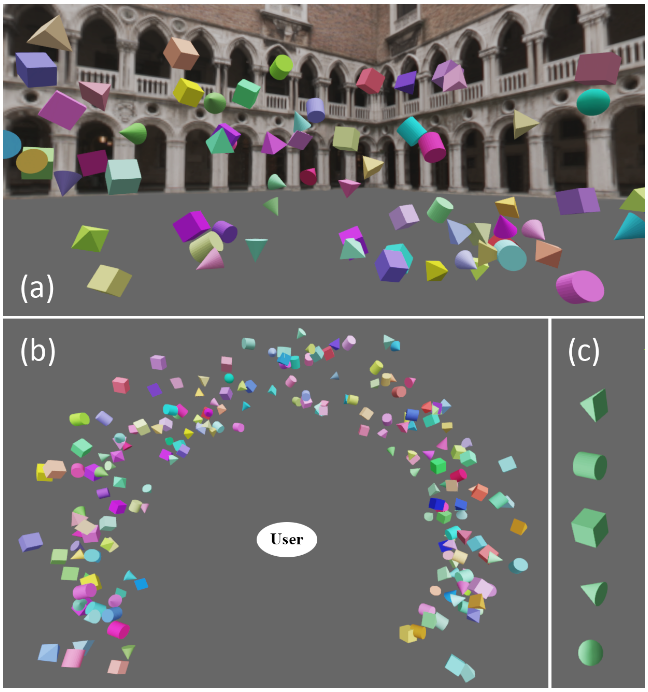

Figure 4: (a) Virtual environment for dynamic scenario from participant's view. (b) The view from above reveals the arrangement around the user. (c) Sphere and further template objects for duplication: cone, cube, cylinder and pyramid.

of 3–4 m. A 90° segment behind the participants was left empty so they do not need to fully turn around during the study.

Procedure. Participants were again seated on a swivel chair, wearing the eye tracking HMD and the 5 object templates were shown as in Fig. 4c. This time their task was twofold: First, they should search for a sphere (shown at the bottom of Fig. 4c) which is hidden somewhere in between multiple duplicates of the other four objects. Second, they were asked to press the *space bar* on the keyboard as soon as they perceive the stimulus either in the left or right part of their peripheral vision. The stimulus was randomly placed left or right of the participants' current tracked gaze direction, starting with an offset of 5°. Its size slowly increased automatically until the *space bar* was pressed. Then the current size was saved together with the current offset and was reset to zero to continue. The same offset was used five times and was then increased by 5°, up to a maximum offset of 40°. The experiment supervisor was able to see the participants' gaze on the desktop screen at real–time. He re–started shuffling all objects' position, orientation and color when a participant found the sphere. This guarantees, that participants were constantly searching, looking around and thus maintaining a high degree of head and eye movement during the experiment. The experiment was stopped when the fifth stimulus for the largest offset was found.

4.3 Results

Static Scenario. Fig. 5a (orange) shows the accumulated measurements of all participants in groups for the tested eccentricities (left and right stimuli are combined). Confirming our expectations based on previous findings on visual acuity [Strasburger et al. 2011] the collected data points reveal a direct linear correlation between stimulus size and eccentricity. Hence, we fit a linear function $s_{static}(e)$ into the data points (mean square error = 0.5965) resulting in the following equation for the eccentricity–dependent stimulus size

$$s_{static}(e) = 0.0098e + 0.1009°. \tag{6}$$

Dynamic Scenario. The measurements of the dynamic scenario are again shown accumulated over all participants and grouped for the distinct tested eccentricities, as shown in Fig. 5a (blue). Similar to Sec. 4.3 the results closely resemble a linear function. We again fit a linear function $s_{dynamic}(e)$ into these collected data points (mean square error = 0,2818) resulting in

$$s_{dynamic}(e) = 0.0076e + 0.289°. \tag{7}$$

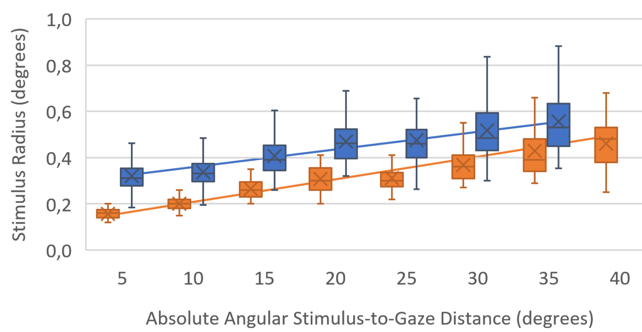

(a)

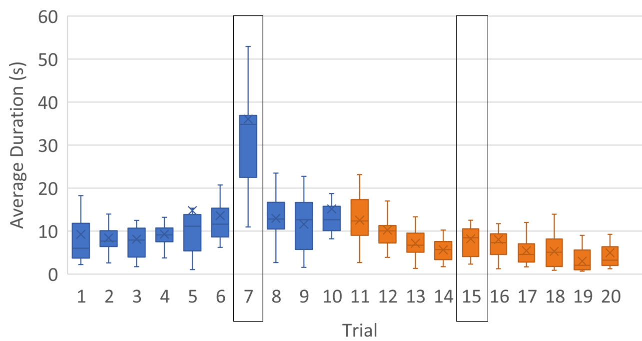

(b)

Figure 5: (a) Distribution of measured stimulus size threshold for static (orange) and dynamic (blue) scenarios, overlaid with corresponding linear fits $s_{static}(e)$ and $s_{dynamic}(e)$. (b) Trial duration distribution of visual search task without (blue) and with (orange) active gaze guidance. Black frames mark challenging cases. Both figures show accumulated data from all participants, as box–and–whisker plots including minimum, maximum, median (–) and mean (×).

5 USER STUDY

The final study evaluates the effectiveness of our method within a visual search task. For this experiment the gaze guidance stimulus size could no longer be manipulated by the participants. Instead, to validate our previous findings, it was configured based on the results of the previous experiment in Sec. 4.3, given in Eq. 7. We used the same virtual environment as for the dynamic scenario (Sec. 4.2). Again, the hardware and software setup was identical to previous studies (Sec. 4.1). Our participants in this third experiment were the same 17 people as in the second (dynamic) experiment.

5.1 Procedure

Participants were again seated on a swivel chair, wearing the eye tracking HMD and the 5 object templates were shown. The participants' task was to find the sphere and hit the *space bar* as soon as they have found it. For this experiment pressing the *space bar* directly started shuffling all objects' position, orientation and color. Overall the participants had to find 20 spheres. The ordering of the 20 trials as well as the state (position, orientation, color) of the objects within each trial, were identical for all participants. During the first 10 trials gaze guidance was completely deactivated and no stimulus was shown at any time. During the last 10 trials gaze guidance was activated and always targeted at the sphere, i.e. the stimulus was superimposed at the sphere's center. Participants were informed at the start of the first trial with activated gaze guidance. Both sequences contained a challenging trial as shown in Fig. 6. For each frame of each of the 20 trials the time since experiment start, gaze direction and gaze-to-target angular distance was saved.

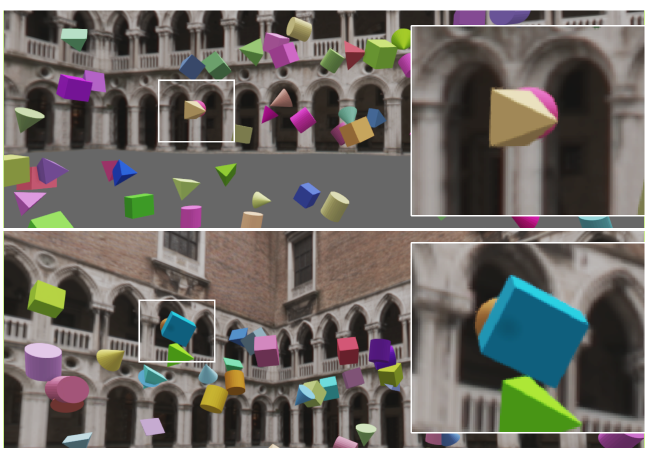

Figure 6: Challenging trials of the visual search task due to occlusion. (top) Pink sphere behind yellow pyramid, without gaze guidance and (bottom) yellow sphere behind blue cube, with gaze guidance.

5.2 Results

The collected data is then analyzed in terms of how long participants needed to find the individual spheres during the visual search task, for which the resulting distribution is shown in Fig. 5b. For the general case, only a slight improvement in detection speed can be identified. Nevertheless, the data reveals great benefit when guiding users towards "hidden" objects such as in Fig. 6. For these challenging trials (7 vs. 15) we can see significant improvement in Fig. 5b (black frames). Without gaze guidance, a multiple of the average time was spent to find the target. For trial 7 the average duration was 36,07 s which is more than two standard deviations (σ = 10, 16 s) above the mean duration of all trials (10,48 s). Trial 15 took only 8,20 s in average, which is even slightly below the overall mean.

Fig. 7, exemplarily, shows recorded gaze-to-target angular distance of a single user over time (all 20 trials). As can be seen, the angular distance starts roughly between 100° and 200° at the beginning of each trial, and then decreases towards zero (target found). Besides the overall slightly reduced search times, trials with guidance follow a mostly monotonic drop once the direction towards the target is found, i.e. the stimulus successfully influenced the viewers attention. In contrast, trials without guidance show some local minima, corresponding to overlook and miss a target at a certain distance. This can be seen especially for the first, not gaze guided, challenging case (left framed trial in Fig. 7).

6 DISCUSSION AND CONCLUSION

The results of the final user study validate applicability of subtle gaze guidance in immersive environments, but possible gain in user performance strongly depends on the present scenario (static vs. dynamic) as well as the given task (general vs. hidden object search). Already the results of the perceptual study show that parameters need to be carefully adapted to scene complexity. Also motion within the scene will most probably influence guiding performance, as movements – especially motion onsets – are known to strongly affect attention.

The slight improvements for "general search" compared to the significant performance increase for "hidden object search" reveals a strong dependency on the selected task. One reason for the slight improvement for moderately visible search targets might be limited complexity of the provided test scenario, so that the target object could be spotted more efficiently than expected. Although the chosen environment involves of a certain amount of complexity, there are certainly more complex and (in terms of attention) more challenging scenarios.

In conclusion, we have presented an extended version of the subtle gaze direction method to open up this technique to a wide field of view and full 360° surroundings, within immersive virtual environments. First, we proposed a model for dynamic shape adjustments of the gaze guidance stimulus, to compensate for perceived distortions due to perspective projection of far peripheral image regions. Second, we introduced dynamic stimulus positioning to also guide the user's head orientation towards external target locations. We, afterwards, conducted and evaluated two perceptual studies to derive appropriate stimulus size parameters for static as well as more dynamic scenarios. Finally, a user study was conducted to validate our method in a search task context. The final results show that gaze guidance in immersive environments can be achieved with our presented method.

In future work we would like to investigate scenes with a higher degree of dynamic movements, to see to which extent a flickering

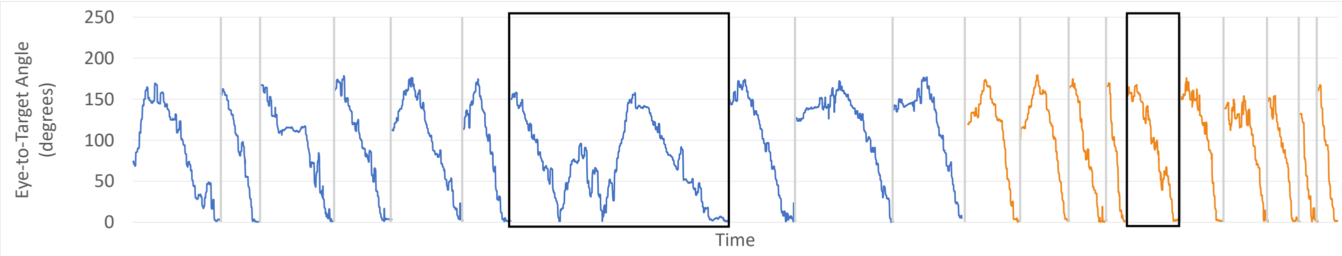

Figure 7: Exemplary data of a single participant showing angular distance from gaze to target over time. 19 thin peaks mark onset of a new object for visual search without (blue) and with (orange) active gaze guidance.

stimulus can bear up against moving objects in the close surrounding. Besides this, the method's performance should be evaluated in scenarios that include objects which by themselves already strongly attract viewers' attention. Also further applications await their chance to take advantage of this new opportunity, e.g. to replace visual distractors by a subtle stimulus during reorientation phases in the context of redirected walking. Additionally, applicability to other virtual environment systems, such as dome or cave environments need to be evaluated.

ACKNOWLEDGMENTS

Steve Grogorick and Marcus Magnor gratefully acknowledge funding by the German Science Foundation (DFG MA2555/15-1 "Immersive Digital Reality" and DFG INST 188/409-1 FUGG "ICG Dome").

REFERENCES

Sean Andrist, Michael Gleicher, and Bilge Mutlu. 2017. Looking Coordinated: Bidirectional Gaze Mechanisms for Collaborative Interaction with Virtual Characters. *Computer Human Interaction (CHI)* (2017).

Reynold Bailey, Ann McNamara, Nisha Sudarsanam, and Cindy Grimm. 2009. Subtle gaze direction. *ACM Transactions on Graphics (TOG)* 28, 4 (2009), 100.

Erhardt Barth, Michael Dorr, Martin Böhme, Karl Gegenfurtner, and Thomas Martinetz. 2006. Guiding the mind's eye: improving communication and vision by external control of the scanpath. In *Electronic Imaging 2006*. International Society for Optics and Photonics, 60570D–60570D.

Frank Biocca, Charles Owen, Arthur Tang, and Corey Bohil. 2007. Attention issues in spatial information systems: Directing mobile users' visual attention using augmented reality. *Journal of Management Information Systems* 23, 4 (2007), 163–184.

Thomas Booth, Srinivas Sridharan, Ann McNamara, Cindy Grimm, and Reynold Bailey. 2013. Guiding attention in controlled real-world environments. In *Proceedings of the ACM Symposium on Applied Perception*. ACM, 75–82.

Forrester Cole, Douglas DeCarlo, Adam Finkelstein, Kenrick Kin, R Keith Morley, and Anthony Santella. 2006. Directing Gaze in 3D Models with Stylized Focus. *Rendering Techniques* 2006 (2006), 17th.

Björn B de Koning and Halszka Jarodzka. 2017. Attention guidance strategies for supporting learning from dynamic visualizations. In *Learning from Dynamic Visualization*. Springer, 255–278.

Michael Dorr, Eleonora Vig, Karl R Gegenfurtner, Thomas Martinetz, and Erhardt Barth. 2008. Eye movement modelling and gaze guidance. In *Fourth International Workshop on Human-Computer Conversation*.

Carl Gutwin and Chris Fedak. 2004. A comparison of fisheye lenses for interactive layout tasks. In *Proceedings of Graphics interface 2004*. Canadian Human-Computer Communications Society, 213–220.

Hajime Hata, Hideki Koike, and Yoichi Sato. 2016. Visual Guidance with Unnoticed Blur Effect. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*. ACM, 28–35.

Raphael Hoffmann, Patrick Baudisch, and Daniel S Weld. 2008. Evaluating visual cues for window switching on large screens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 929–938.

Halszka Jarodzka, Tamara van Gog, Michael Dorr, Katharina Scheiter, and Peter Gerjets. 2013. Learning to see: Guiding students' attention via a model's eye movements fosters learning. *Learning and Instruction* 25 (2013), 62–70.

Azam Khan, Justin Matejka, George Fitzmaurice, and Gordon Kurtenbach. 2005. Spotlight: directing users' attention on large displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 791–798.

Robert Kosara, Silvia Miksch, and Helwig Hauser. 2002. Focus + context taken literally. *IEEE Computer Graphics and Applications* 22, 1 (2002), 22–29.

Yen-Chen Lin, Yung-Ju Chang, Hou-Ning Hu, Hsien-Tzu Cheng, Chi-Wen Huang, and Min Sun. 2017. Tell Me Where to Look: Investigating Ways for Assisting Focus in 360° Video. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2535–2545.

Andrei Lintu and Noelle Carbonell. 2009. Gaze Guidance through Peripheral Stimuli. (2009).

Weiquan Lu, Henry Been-Lirn Duh, Steven Feiner, and Qi Zhao. 2014. Attributes of subtle cues for facilitating visual search in augmented reality. *IEEE transactions on visualization and computer graphics* 20, 3 (2014), 404–412.

Ann McNamara, Reynold Bailey, and Cindy Grimm. 2009. Search task performance using subtle gaze direction with the presence of distractions. *ACM Transactions on Applied Perception (TAP)* 6, 3 (2009), 17.

Ann McNamara, Thomas Booth, Srinivas Sridharan, Stephen Caffey, Cindy Grimm, and Reynold Bailey. 2012. Directing gaze in narrative art. In *Proceedings of the ACM Symposium on Applied Perception*. ACM, 63–70.

Laura Pomarjanschi, Michael Dorr, and Erhardt Barth. 2012. Gaze guidance reduces the number of collisions with pedestrians in a driving simulator. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1, 2 (2012), 8.

Yoichi Sato, Yusuke Sugano, Akihiro Sugimoto, Yoshinori Kuno, and Hideki Koike. 2016. Sensing and Controlling Human Gaze in Daily Living Space for Human-Harmonized Information Environments. In *Human-Harmonized Information Technology, Volume 1*. Springer, 199–237.

Alia Sheikh, Andy Brown, Zillah Watson, and Michael Evans. 2016. Directing attention in 360-degree video. (2016).

Michael Stengel, Steve Grogorick, Martin Eisemann, Elmar Eisemann, and Marcus A Magnor. 2015. An affordable solution for binocular eye tracking and calibration in head-mounted displays. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 15–24.

Hans Strasburger, Ingo Rentschler, and Martin Jüttner. 2011. Peripheral vision and pattern recognition: A review. 11, 5 (2011), 13–13.

Eduardo E Veas, Erick Mendez, Steven K Feiner, and Dieter Schmalstieg. 2011. Directing attention and influencing memory with visual saliency modulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1471–1480.

Eleonora Vig, Michael Dorr, and Erhardt Barth. 2011. Learned saliency transformations for gaze guidance. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 78650W–78650W.

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

Martin Weier, Michael Stengel, Thorsten Roth, Piotr Didyk, Elmar Eisemann, Martin Eisemann, Steve Grogorick, André Hinkenjann, E Kruijff, M Magnor, et al. 2017. Perception-driven Accelerated Rendering. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 611–643.

Received May 2017; revised July 2017