Visual Exploration of Omni-Directional Panoramic Scenes

Walter F. Bischof, Nicola C. Anderson, Michael T. Doswell, and Alan Kingstone

University of British Columbia, Vancouver, Canada

Author Note

Walter F. Bischof, Department of Psychology, University of British Columbia, 2136 West Mall,
Vancouver BC, V6T 1Z4, Canada

Abstract

How do we explore the visual environment around us, and how are head and eye movements coordinated during our exploration? To investigate this question, we had observers look at omni-directional panoramic scenes, composed of both landscape and fractal images, using a virtual-reality (VR) viewer while their eye and head movements were tracked. We analyzed the spatial distribution of eye fixations and the distribution of saccade directions; the spatial distribution of head positions and the distribution of head shifts; as well as the relation between eye and head movements. The results show that, for landscape scenes, eye and head behaviour best fit the allocentric frame defined by the scene horizon, especially when head tilt (i.e., head rotation around the view axis) is considered. For fractal scenes, which have an isotropic texture, eye and head movements were executed primarily along the cardinal directions in world coordinates. The results also show that eye and head movements are closely linked in space and time in a complementary way, with stimulus-driven eye movements predominantly leading the head movements. Our study is the first to systematically examine eye and head movements in a panoramic VR environment, and the results demonstrate that a VR environment constitutes a powerful and informative research alternative to traditional methods for investigating looking behaviour.

Keywords: eye movements, head movements, gaze-head relationship, omni-directional panoramic scenes.

Visual Exploration of Omni-Directional Panoramic Scenes

Humans explore the visual environment by moving their eyes several times per second to sample different locations. This is necessitated by the variable-resolution characteristics of the retina and thus the need to project different locations of the environment onto the high-resolution fovea for processing during each fixation. Evidence from many studies on the characteristics of fixations and saccades suggests that these locations are not randomly selected (e.g., Buswell, 1935; Yarbus, 1967). Fixations occur significantly more often near the center of the visual field (central fixation bias, Tatler, 2007), on scene parts that are rated as informative rather than redundant (Mackworth & Morandi, 1967; Underwood & Foulsham, 2006), on objects that are inconsistent with scene expectations (e.g., Loftus & Mackworth, 1978), on items that convey social information, such as faces of people (e.g., Birmingham, Bischof, & Kingstone, 2008, 2009; Yarbus, 1967), and on regions of high contrast in low-level image features (e.g., Parkhurst, Law, & Niebur, 2002).

The non-random nature of selecting fixation positions is also evident from the characteristics of saccades. Several studies have shown systematic tendencies in saccade directions, with saccades executed preferably in the cardinal directions rather than the obliques (Foulsham, Kingstone, & Underwood, 2008; Foulsham & Kingstone, 2010). The bias of saccade directions could be due to oculomotor factors, i.e., the dominance of the muscle or neural apparatus, which preferentially triggers horizontal shifts of the eyes, regardless of the stimulus being viewed. The bias could also be a consequence of the distribution of salient features in pictures and natural scenes. Furthermore, a learned account could predict a horizontal bias based on our experience with pictures and the environment that is acquired over time and initiated top-down, to reflect one's goals and expectations. Finally, the biases in saccade direction could be display-specific, an artefact of laboratory-based eye tracking studies that present scenes on a computer monitor.

To investigate the roots of these saccade biases, Foulsham and colleagues investigated saccadic behavior while viewing images rotated to various extents within frames that remained constant or were also rotated (Foulsham et al., 2008; Underwood, Foulsham & Humphrey, 2009; Foulsham & Kingstone,

2010). For instance, Foulsham et al. presented natural images that were rotated by 0, 45, 90, 135 or 180 degrees. The results showed that saccades in the horizon direction of the images were more frequent and larger in amplitude, independent of the image orientation. For images with indoor scenes (which have an increased presence of vertical edges), the saccades were biased in the cardinal directions (left – right and up – down), again with respect to the image “horizon”. Foulsham and colleagues concluded that these results rule out a simple oculomotor explanation of saccade generation and favored a combination of image-based and frame-based accounts: First, the distribution of image features (edges or salient locations) may guide eye movements in a bottom-up (stimulus-driven) manner towards the image horizon. Second, the gist of the images (Potter, Staub, & O’Connor, 2004; Torralba, 2003) conveyed by the low spatial frequencies of the images may guide the eyes in a top-down manner towards the horizon where the most informative image features are likely to be located. Finally, the increased presence of vertical and horizontal edges in natural images (Coppola, Purves, McCoy, & Purves, 1998) may be the cause for the increased presence of saccades in the cardinal, but not oblique, directions in images.

The frame that the stimuli are presented in also seems to be of critical importance. For example, Foulsham et al. (2008) presented natural scenes within a square window while Foulsham and Kingstone (2010) presented similar natural scenes within a circular window. The two studies yielded very different saccade patterns, which are thus likely due to the difference in frame shape. In a related paper, Foulsham, Teska, and Kingstone (2011) presented natural scenes viewed through gaze-contingent windows of different shapes (square, vertical and horizontal) and found that the saccade distributions were strongly affected by the window shape. Finally, the biases in saccade patterns can also be attributed to the shape of the computer monitor on which the stimuli are presented. These monitors are normally wider than they are high, and pictures are often presented in the landscape orientation filling the screen. It may thus be more efficient to move horizontally than vertically in order to cover the whole area.

To address these possible limitations, we compared the eye movements of people immersed in a virtual environment with 360° panoramic scenes while their head was unrestrained. We asked whether people would distribute their gaze in the same way as past studies, and what role, if any, the head plays in

the visual exploration. This design has several advantages over previous studies. First, following Foulsham and colleagues, we rotated the panoramas around the view axis to study the effect of image rotation on saccade directions (see Figure 1), but our design allowed to study these effects more precisely by eliminating any reference frames defined by the stimulus edge or the monitor frame. Second, the view of the environment was not preselected by the experimenter, but selected by the participant as their head was free to move. This design can thus shed light on both how one explores and inspects the visual environment. Third, the design permitted us to study the relation between, and the coordination of eye and head movements under general, but well controlled conditions. These advantages are discussed in more detail below.

Figure 1: Sketch of the rotated panoramas. Only a central area of the panoramas is shown, but participants could look at the panoramas in all directions.

To date, the vast majority of eye movement studies have measured eye movements while the observer's head is immobilized, usually by a chinrest, but more extreme measures, such as a bite-bar, were once routine. On the one hand, immobilizing the head not only permits accurate eye movement measurements, but it also enables one to study oculomotor processes in isolation, uninfluenced by head or body movements. On the other hand, this approach means that all the visual information displayed to the observer has been pre-selected by the experimenter, i.e. the observer is a passive recipient of the visual information. To the extent that researchers wish to understand how people explore and inspect the visual information that surrounds them, and not just how people make eye movements to pictures that are presented to them, these limitations are potentially profound. In real life, people turn their heads and bodies to acquire and view new information that is currently unavailable. For example, when we hear a loud noise behind us on the street, we turn our heads to explore and inspect the source of the noise. In other words, the act of selecting a new view by turning one's body and head, and the process of examining the information that is delivered to the eyes by a head turn, are complimentary processes. This was demonstrated, for example, by Foulsham, Walker and Kingstone (2011) who analyzed the visual behavior of participants as they walked across a university campus while wearing a mobile eye tracker. They found that locations around the horizon were selected with head movements, and eye movements were centralized within that head-centered visual field.

When the head is free to move, changes in the line of sight involve simultaneous saccadic eye and head movements. In this situation, the rules defining head-restrained saccadic eye movement have to be altered (Bizzi, Kalil, & Tagliasco, 1971; Freedman, 2008). With attention shifts to sudden, bottom-up cues, for example, the line of sight is adjusted quickly through a saccadic eye movement, followed by the slower head movement combined with eye movement to keep the relevant stimulus projected on the fovea. With top-down control of the line of sight, the head tends to move first to select a newly centered visual field, followed by the eyes (Doshi & Trivedi, 2012). Taken together, adjustments of gaze, i.e. the line of sight, are achieved through a combination of eye and head movements, both with different functions in visual exploration and with different temporal characteristics. As Solman, Foulsham, and

Kingstone (2017) showed using gaze-contingent displays, eye movements are more likely to move in the direction of visible information while head movements showed a complementary pattern, moving primarily in a direction to reveal new information. Moreover, different cognitive effects of head and eye movements were found by Solman and Kingstone (2014). They discovered greater use of memory in head-contingent search compared to gaze-contingent search, which they attribute to the difference in energetic cost of turning the head to orient attention (e.g., having worked hard to turn one's head to see an item it is prudent to commit it to memory).

In the present study participants were immersed in a virtual environment containing panoramic scenes with landscape scenes and fractal scenes. Most of the landscape scenes have a fairly well-defined horizon, which separates the top half (sky) from the bottom half (ground; see Figure 2 top row). To investigate the effect of image content on visual exploration (e.g., the effect of image horizon), we also used fractal panoramas (Figure 2, bottom row), which contain similar spatial frequency variations as landscape scenes but they lack any clear semantically meaningful items including horizons. In addition, we rotated the panoramas around the view axis.

Figure 2: Samples of the panoramic scenes used in the present experiment. Landscape scenes are shown in the top row, fractal scenes in the bottom row.

There is substantial confusion in the eye-movement field regarding the definition of fixations and saccades, in particular when eye movements are combined with free head movements (see e.g., Hessels, Niehorster, Nyström, Andersson & Hooze, 2018). In our study, participants were seated inside a virtual sphere and looked at the panoramic scenes from approximately the center of the sphere. In this context, we now define terms related to gaze, head, and eyes.

A gaze hit point, what we call 'gaze point' for short, is defined as the intersection of the gaze vector with the virtual sphere on which the panoramas are projected. It is defined in world coordinates, with longitudes in the range $[-180, 180]$ degrees (from west to east) and latitudes in the range $[-90, 90]$ degrees (from the south pole to the north pole). Fixations are defined as stable gaze points and are extracted using the IDT algorithm (Salvucci & Goldberg, 2000; Blignaut, 2009; Komogortsev, Gobert, Jayarathna, Koh, & Gowda, 2010) adapted to spherical displays, using a minimum duration of 80 ms and a maximum dispersion threshold of 3° . Saccades are defined as differences between successive fixations (see Appendix 1 for more details).

A head hit point, or 'head point', is defined the intersection of the vector pointing forward from the face with the virtual sphere and is also defined in world coordinates, namely the longitude and latitude of the spherical panorama. Head movements with the VR viewer are normally smooth and do not permit identifying stable points from the raw data. For this reason, we define a head fixation as the average head point during a fixation. Head shifts are defined as the differences between successive head fixations.

Finally, eye hit point, or eye point, is defined as the difference between gaze point and head point. With this definition, the eye point is also expressed in world coordinates, just like gaze and head points.

Methods

Participants

18 participants (ages 18 - 37, $M = 21.3$ years, 12 female) from the University of British Columbia participated in this experiment for course credit. All reported normal or corrected-to-normal vision.

Participants provided informed consent prior to participation, and the study was approved by the ethics board of the University of British Columbia (H10-00527).

Apparatus

The stimuli were presented on an HTC Vive Virtual Reality headset equipped with an SMI eye tracker, controlled by a custom-built desktop computer. The HTC Vive headset has a display resolution of 1080x1200 pixels, a horizontal field of view of approximately 110°, a vertical field of view of approximately 113°, a refresh rate of 90 Hz, and a weight of approx. 635g. The SMI eye tracker recorded eye movements at 250 Hz. To measure head movements, two infrared base stations situated in opposite corners of the room tracked the location of the HTC Vive headset at 60 Hz. Participants were seated in a swivel chair in the center of the room and held a keyboard to provide input throughout the experiment.

Stimuli

The virtual space was created using Unity (Technologies, 2015) and consisted of a sphere around the participant, onto which different omni-directional panoramic scenes were projected, such that the participant appeared to be immersed in the scenes. The scenes consisted of 160 full 360° panoramic images with 80 landscape scenes and 80 computer-generated 3D fractal scenes. The landscape scenes were taken from the SUN360 Panorama Database (Xiao, Hays, Ehinger, Oliva, & Torralba, 2010). The computer-generated fractal images were taken from the Fractal Canyon website (Schizo, 2010). The landscape and fractal images contained similar spatial frequency variations, and all images had a resolution of 4096×2160 pixels. The scenes were presented in eight rotations, rotated counter-clockwise about a horizontal axis (the z-axis of the virtual space) by 0, 45, 90, 135, 180, 225, 270 and 315 degrees (see Figure 1).

Procedure

At the start of the experiment, participants were familiarized with the VR equipment and instructed to remember the presented scenes for a subsequent memory test. They were handed a keyboard and instructed to use the Enter key for starting a stimulus presentation. The experiment proceeded in two phases. In the first phase, participants viewed 160 scenes with the instruction to remember them for later

testing. In the second phase, participants were presented with 10 old and 10 new, unrotated scenes, and had to indicate whether they had seen the scenes before. The purpose of the memory procedure was merely to ensure that participants engaged in scanning the panoramas, and the delivery of the test was done to ensure that if participants talked to one another outside the study, then future participants would know that the memory instruction was authentic, and not merely a deception. As this small condition is not the focus of our study, and was carried out solely to maintain participant compliance during the initial viewing phase, the data are not discussed.

At the beginning of the experiment and after every 20 trials, the eye tracker was calibrated using a 5-point calibration test, in which the participants had to fixate five locations while their eye movements were recorded. After a successful calibration, the experimental procedure continued, otherwise the calibration was repeated until it was successful.

In the initial viewing phase, each trial began with a fixation cross in the screen center. This allowed the experimenter to monitor the accuracy of the eye tracker. Participants pressed the Enter key to initiate a new trial. Each scene was presented for 10 seconds, long enough to allow slower head movements during the stimulus presentation. There were 160 trials, 80 with fractal scenes and 80 with landscape scenes. Thus, there were 80 unique landscape and 80 unique fractal scenes, and these scenes were shown in one of 8 different orientations, resulting in 10 scenes per orientation and scene type. The scene orientations were balanced across participants to ensure that each scene was presented in every orientation. This phase of the experiment took approximately 20 minutes.

In the Recognition phase, the same procedure was followed, except that, at the end of each trial, participants had to indicate on the keyboard whether they had seen the scene before or whether it was a new scene. Twenty scenes were presented, five 'old' landscape scenes, five 'old' fractal scenes, five 'new' landscape scenes, and five 'new' fractal scenes, each one being presented for 10 seconds. This phase of the experiment took approximately 10 minutes.

Data Analysis

The gaze direction vector was recorded at 250 Hz while head position (i.e., the position of the headset) and head orientation were recorded at a rate of approximately 50 Hz rate, with the exact timing being dependent on the Unity system (Technologies, 2015). Head position and orientation were linearly interpolated to 250 Hz, and a gaze vector was extracted by combining head and eye direction.

The gaze values were transformed into longitude and latitude values in the spherical panorama, with longitudes in the range $[-180, 180]$ degrees (west to east) and latitudes in the range $[-90, 90]$ degrees (south pole to north pole). For display purposes, however, panorama images and data are presented on a flat projection, an equirectangular (or equidistant) projection (Wikipedia, 2019). This projection maps meridians into vertical straight lines of constant spacing, introducing strong distortions near the poles compared to the horizon, hence data in the polar regions have to be interpreted carefully.

Gaze Analysis

The analysis of the gaze data consists of three parts. In the first part, basic fixation data are analyzed. In the second part, the fixation distributions are presented, which differ considerably from previous results obtained with images presented on a screen. In the third part, saccade patterns are analyzed to permit a direct comparison with the results of Foulsham and Kingstone (2010).

Basic Fixation Data

An analysis of the number of fixations showed that participants made on average 3.75 fixations per second for the landscape scenes and 3.23 fixations per second for the fractal scenes, $F(1, 17) = 20.0$, $p < .001$, $\eta^2 = 0.212$. There was no effect of Scene-rotation, $F(7, 17) = 0.48$, $p = 0.851$.

Fixation duration was on average 199.1 ms for the landscape scenes and 234.9 ms for the fractal scenes, $F(1, 17) = 11.1$, $p = .004$, $\eta^2 = 0.159$. There was a weak but significant effect of Scene-rotation, $F(7, 17) =$

2.90, $p = .008$, $\eta^2 = 0.012$. In summary, participants made fewer and longer fixations for the fractal scenes than for the landscape scenes.

Fixation Distribution

The distribution of fixations differed substantially between the landscape and fractal scenes (see Figure 3). The panels on the left show the bivariate fixation distributions for the landscape scenes for the eight scene orientations, those on the right for the fractal scenes. The distribution patterns for the landscape scenes show that participants tended to fixate along the horizon of these images, and a principal component analysis (PCA) shows that the orientations of the fixation distributions (as defined by the angle to the first eigenvector) are all within a range of ± 3.5 degrees of the scene rotation, confirming the close match between the orientation of the fixation distributions and scene rotation. The fixations patterns for the landscape images are closely related to those obtained by Sitzman et al. (2018) who found a similar horizon bias. For the fractal scenes, the fixation patterns are much closer to isotropic and there is no evidence of the fixation patterns being aligned with the virtual horizon of the images.

Figure 3: Distribution of fixations for landscape images (left) and fractal images (right), for all scene rotations in the range 0° - 315° . Darker pixels in these images correspond to higher fixation densities.

As shown in Table 1, the spread of fixations along the scene horizons (longitude) was much larger than the spread in the orthogonal direction (latitude). For fractal scenes, this effect was much reduced.

Table 1

Standard deviations of the fixation distributions, in longitude and latitude and for the two scene types

Scene type	<i>SD</i> Longitude (deg)	<i>SD</i> Latitude (deg)	<i>t</i> (17)	Cohen's <i>d</i>
landscapes	64.82	21.76	11.83**	3.55
fractals	47.10	29.59	10.22**	1.14

Saccade Direction

We now describe the saccade patterns in these scenes. They are closely related to the fixation distributions, but permit a direct comparison with the results of Foulsham and Kingstone (2010). Given the large spread of fixations along the horizon of the landscape images, it is plausible that saccade directions also align with the scene horizons. This is illustrated in Figure 4, which shows polar histograms of saccade directions in world coordinates, for landscape scenes on the left, for fractal scenes on the right, and for all scene rotations. The histograms for the landscape scenes show that most saccades were made along the horizon direction of the panoramas (i.e., left and right for the scene rotation of 0°), thus showing a strong effect of scene rotation. The histograms for the fractal scenes, however, show a very different effect of scene rotation. Saccades were made in all directions, with a slight preference for saccades along the cardinal directions in world coordinates, namely to the left, the right, and up. This dovetails with past work (Foulsham & Kingstone, 2008; Foulsham et al., 2008), and as the fractal images are isotropic on average, with no clear texture orientation along the horizon direction, any response biases likely reflect preferential scanning biases that are structural (i.e., oculomotor) and/or learned.

Figure 4: Polar histogram of saccade directions in world coordinates, for the landscape scenes (left), and for the fractal scenes (right), for all scene rotations.

For the statistical analysis of the saccade directions, opposite saccade directions (0° and 180° , 45° and 225° , 90° and 270° , 135° and 315°) and scene rotations (45° and 225° , 90° and 270° , 135° and 315°) were combined to allow a direct comparison with the results of Foulsham et al. (2008) and Foulsham and Kingstone (2010) (see Figure 5). An overall analysis of variance yielded two significant effects, namely the interaction of scene rotation $\times$ saccade direction, $F(12,204) = 64.12, p < .001, \eta^2 = 0.316$, and the interaction of scene type $\times$ scene rotation $\times$ saccade direction, $F(12,204) = 110.89, p < .001, \eta^2 = 0.408$. A more detailed analysis revealed that for landscape scenes (Figure 5 left panel), there was a strong interaction of scene rotation with saccade direction, $F(12,204) = 134.5, p < .001, \eta^2 = 0.851$, with the saccade frequencies along the landscape horizons being approximately twice as high as those in the other directions. For the fractal landscapes (Figure 5 right panel), there was also a strong interaction of scene

rotation with saccade direction, $F(12,204) = 9.43$, $p < .001$, $\eta^2 = 0.174$, with more frequent saccades in the cardinal directions for scene rotations 0° , 90° , 180° , and 270° and more isotropic saccade patterns for the other scene rotations.

Figure 5: Saccade directions organized into four symmetric axes in world coordinates, for landscape scenes (left) and fractal scenes (right). Data show the mean proportions of saccades and standard error bars in each axis across participants.

The present study revealed strong patterns of saccade directions: For all rotations of the landscape scenes, saccades directions were primarily aligned with the scene horizons (Figure 4 left panels and Figure 5 left panel), with 40.5% of the saccade directions within a range of $\pm 22.5^\circ$ of the horizon direction. For fractal scenes, saccades were more frequent in the horizontal and upward direction (in world coordinates) for scene rotations 0° , 90° , 180° , and 270° , an effect that was less pronounced for the other scene rotations. The saccade directions were, however, not aligned with the scene horizon, as was the case for the landscape scenes.

The results obtained for the landscape scenes in the present study are similar to those obtained by Foulsham et al. (2008) who found the saccade directions tended to be aligned with the horizon. For our

fractal scenes, which show biases along cardinal coordinates, the results are convergent with past works (for instance, see Foulsham & Kingstone, 2010, Figures 2 and 3).

Saccade Amplitudes

The fixation distributions for landscape images shown in Figure 3 suggest a) that saccades in the direction of the horizon are more frequent, as discussed in the previous section, and b) that saccade amplitudes may also be larger in that direction. In contrast, the saccade characteristics (direction and amplitude) may be more balanced for fractal landscapes.

An overall analysis of variance of saccade amplitudes yielded two significant main effects, a main effect of scene type, $F(1,17) = 26.85, p < .001, \eta^2 = 0.074$, with saccade amplitudes for landscape panoramas, $M = 13.13$, being larger than for fractal panoramas, $M = 11.54$; and a main effect of saccade direction, $F(3,51) = 57.03, p < .001, \eta^2 = 0.147$. There were also two significant interactions, namely scene rotation $\times$ saccade direction, $F(12,204) = 26.23, p < .001, \eta^2 = 0.093$, and scene type $\times$ scene rotation $\times$ scene direction, $F(12,204) = 25.27, p < .001, \eta^2 = 0.078$. A more detailed analysis revealed that for landscape scenes (Figure 6 left), there was a main effect of saccade direction, $F(3,51) = 64.6, p < .001, \eta^2 = 0.151$, and a significant interaction of scene rotation with saccade direction, $F(12,204) = 37.5, p < .001, \eta^2 = 0.302$, with the saccade amplitudes along the landscape horizons larger than those in the other directions. For the fractal landscapes (Figure 6 right), there was significant main effect of saccade direction, $F(3,51) = 25.38, p < .001, \eta^2 = 0.173$, with higher saccade amplitudes in the horizontal direction (in world coordinates).

Figure 6: Saccade amplitudes along four axes in world coordinates, for landscape scenes (top) and fractal scenes (bottom). Data show the mean amplitudes of saccades and standard error bars in each axis across participants.

For landscape scenes, saccade amplitudes were on average 13.1° and fixation durations were 199.1 ms; for fractal scenes, saccade amplitude were on average 10.8° and fixation durations were 234.9 ms. The inverse relationship between saccade amplitudes and fixation durations has been reported before (e.g., Unema, Pannasch, Joos, & Velichkovsky, 2005; Tatler & Vincent, 2008). It suggests that global scanning - characterized by large amplitude saccades and short fixation durations - was prevalent for the landscape scenes, while local scanning - characterized by smaller amplitude saccades and longer fixation durations - was prevalent for the fractal scenes.

Summary

The results of the gaze analysis show that the fixations and saccades are determined by the content of the panoramas: In landscape panoramas, the fixation distribution and the most frequent saccade direction are aligned with the horizon, even when the panoramas are rotated about the view axis. The alignment is precise, that is, within at most a few degrees of the rotation angle. In fractal panoramas, which lack an

identifiable horizon, the fixation distributions are closer to isotropic, with a tendency of saccade directions along the cardinal directions, and the largest amplitudes to be horizontal, in world coordinates.

Head Analysis

As described earlier, the location and orientation of the VR headset were measured using two infrared base stations situated in opposite corners of the experiment room. From the location and orientation of the headset, one can determine where on the stimulus sphere the head is pointing. This point is referred to as head hit point, or in brief, head point. Head movements are not ballistic in the way that eye movements are, and for this reason, there are, in contrast to gaze, no natural demarcations for head shifts and head fixations. In order to compare head with gaze, we adopted the convention that, for every eye fixation, we define head fixation as the average head point during that fixation. Eye fixations and head fixations for one participant and a landscape stimulus are shown in Figure 7. The red circles show a sequence of eye fixations (numbered from 1 to 45), the black circles show the head fixations during those fixations, and the lines connect fixations with the corresponding head fixations. The lines thus represent differences between gaze and head, and are determined by the eye direction within the headset.

Figure 7: Example of (gaze) fixations and the corresponding head fixations. The red circles show a sequence of fixations (numbered from 1 to 45). The black circles show the head fixations during each fixation, and the lines connect the two.

Distribution of Head Fixations

The distribution of head fixations differed substantially between the landscape and fractal scenes. This is illustrated in Figure 8, which shows the distribution head fixations for all eight scene orientations, for the landscape panoramas in the left panels and the fractal panoramas in the right panels.

Figure 8: Distribution of head fixations for the landscape panoramas (left) and for the fractal panoramas (right), for scene rotations 0° , 45° , ..., 315° .

An informal comparison of Figure 8 with Figures 3 indicates that the range of eye fixations is larger than the range of head fixations. The statistical analysis in Table 2 shows that the standard deviations of the eye fixations are significantly larger than those of the head fixations, for both scene types, and for longitudes and latitudes. The smaller range of head movements suggests that they are used to center the visual field on a particular point in the environment, and eye movements are used to explore the visual field. This results in the range of fixations exceeding the range of head positions. Apart from these quantitative differences, the distributions of eye fixations and head positions are, however, qualitatively very similar.

Table 2

Fixation and head fixation spreads in longitude and latitude for the two scene types

Scene type	Longitude / Latitude	<i>SD</i> Gaze	<i>SD</i> Head	<i>t</i> (17)	Cohen's <i>d</i>
Landscape	Longitude	64.82	54.54	13.93**	0.585
Landscape	Latitude	21.76	19.51	4.47**	0.659
Fractal	Longitude	47.10	38.36	10.22**	0.426
Fractal	Latitude	29.59	24.30	8.38**	0.739

Direction of Head Shifts

In analogy to the gaze analysis, we now describe the head shift patterns, allowing again a direct comparison with results obtained in previous studies. Figure 9 shows polar histograms of head shift directions, for landscape scenes on the left, for fractal scenes on the right, and for all scene rotations. The histograms for the landscape scenes show that most movements were made along the horizon of the images (i.e., left and right for the scene rotation of 0 degrees), thus showing a strong effect of scene rotation. The histograms for the fractal scenes show no systematic effect of scene rotation. Head shifts were made in all directions, with a slight preference for the cardinal directions in world coordinates, i.e. determined by the gravitational vertical (Cristino & Baddeley, 2009). This pattern of results can be attributed to the fact that the fractal images were isotropic on average, with no clear horizon-defined anisotropy. These results for head shifts are similar to those of the saccades. It is interesting to speculate that shared similarities for the fractal images, which are not unique to the head or eyes, points toward scanning patterns that are learned rather than structural (e.g., mechanisms that move the head or eyes).

Figure 9: Direction distribution of head shifts in world coordinates, for the landscape scenes (left), and for the fractal scenes (right), for all scene rotations.

For the statistical analysis of head shifts, head shift in opposite directions (0° and 180° , 45° and 225° , 90° and 270° , 135° and 315°) and scene rotations (45° and 225° , 90° and 270° , 135° and 315°) were combined to allow a direct comparison with results obtained previously (Foulsham et al., 2008; Foulsham & Kingstone, 2010). An overall analysis of variance yielded three significant effects, namely the effect of head shift direction, $F(3,51) = 6.96, p < .001, \eta^2 = 0.077$, the interaction scene rotation $\times$ head shift direction, $F(12,204) = 88.67, p < .001, \eta^2 = 0.284$, and the interaction scene type $\times$ scene rotation $\times$ head shift direction, $F(12,204) = 106.07, p < .001, \eta^2 = 0.288$.

A more detailed analysis revealed that, for landscape scenes (Figure 10 left panel), there was a significant effect of head shift direction, $F(3,51) = 22.01, p < .001, \eta^2 = 0.081$, and a significant

interaction of scene rotation $\times$ head shift direction, $F(12,204) = 153.2$, $p < .001$, $\eta^2 = 0.770$, with the head shifts along the landscape horizons being approximately three times as frequent as those in the other directions. For the fractal landscapes (Figure 10 right panel), there was a marginal effect of head shift direction, $F(3,51) = 2.73$, $p = .054$, $\eta^2 = 0.116$, with head shift directions in the cardinal directions (in world coordinates) somewhat higher than in the other directions.

Figure 10: Head shift directions organized into four symmetric axes in world coordinates, for landscape scenes (left) and fractal scenes (right). Data show the mean proportions of head shifts and standard error bars in each axis across participants.

For landscape scenes, the frequencies of head shifts along the horizons were more frequent than for the other directions, and there was no systematic difference between the other three directions. This is the same as the effect found for the saccade directions (see Figure 5), but even more biased towards movements along the horizon line: 47.9% of the head shifts were within $\pm 22.5^\circ$ of the horizon line, but only 42.1% of the saccades, $t(142) = 4.204$, Cohen's $d = 0.701$. For the fractal scenes, there was a stronger bias towards movements in the cardinal directions (horizontal and vertical) in world coordinates than was the case for saccades.

Amplitude of Head Shifts

Head shifts were defined as the differences between successive head fixations. An overall analysis of head shift amplitudes yielded several significant effects, namely an effect of scene type, $F(1,17) = 34.3$, $p < .001$, $\eta^2 = 0.128$, with $M = 6.92^\circ$ for the landscape scenes and $M = 4.70^\circ$ for the fractal scenes. Furthermore, there was a significant interaction scene rotation $\times$ head shift direction, $F(12,204) = 25.45$, $p < .001$, $\eta^2 = 0.048$, and a significant interaction scene type $\times$ scene rotation $\times$ head shift direction, $F(12,204) = 18.28$, $p < .001$, $\eta^2 = 0.033$. A more detailed analysis of the landscape scenes revealed a weak effect of scene rotation, $F(4,68) = 3.89$, $p = .007$, $\eta^2 = 0.015$; a significant effect of head shift direction, $F(3,51) = 83.04$, $p < .001$, $\eta^2 = 0.082$; and a significant interaction scene rotation $\times$ head shift direction, $F(12,204) = 32.59$, $p < .001$, $\eta^2 = 0.157$. An analysis of the fractal scenes revealed a significant effect of head shift direction, $F(3,51) = 33.11$, $p < .001$, $\eta^2 = 0.084$; and a weak effect of the interaction scene rotation $\times$ head shift direction, $F(12,204) = 1.78$, $p = .054$, $\eta^2 = 0.006$. The effect of scene type and scene rotation on head movement amplitude is illustrated in Figure 11. Taken together, the results show that, for landscape scenes, the amplitude of the head shifts were largest for saccades in the direction of the horizon while for fractal scenes, the amplitudes were largest for head shifts in the horizontal direction (in world coordinates).

Figure 11: Amplitudes of head shifts along four axes in world coordinates, for landscape scenes (left) and fractal scenes (right). Data show the mean amplitudes of saccades and standard error bars in each axis across participants.

Head Rotation around the View Axis

We observed that some participants tilted their head, (i.e., rotated their head around the view axis to some extent), consistent with the attempt to see the scenes closer to their canonical (upright) orientation. For example, if a scene is rotated by 45° then a head rotation of 45° would result in upright orientation of the scene within the visual field. If this observation is correct then the head rotations should occur in the direction of the scene rotations. This adjustment should happen gradually during the stimulus presentation.

Figure 12 shows individual curves of head rotations around the view axis as a function of head fixation number and relative to the head rotation for the first fixation. The effect is quite strong for the landscape scenes and virtually non-existent for the fractal scenes. For scene orientations 45° and 90° , the head was on average rotated up to about 10 degrees counter-clockwise, and for scene orientations 270° and 315° , the head was rotated on average up to about 7 degrees clockwise. These findings converge with the notion that when head rotation occurs, it is performed to bring the scenes closer to the upright orientation.

Figure 12: Head rotation around the view axis as a function of fixation number. The graphs show individual head rotation curves that are normalized relative to the head rotation for the first fixation. The responses for landscape scenes are shown in the top half and for fractal scenes in the bottom half.

To examine the rate of head rotation, the first 10 head fixations of the rotation data were analyzed using a mixed-effects linear regression analysis, which yielded regression slopes significantly different from zero and a magnitude exceeding $0.4^\circ / \text{fixation}$ for the landscape scenes and for scene rotations 45° , 90° , 135° , 225° , 270° and 315° . For the other scene rotations of the landscape scenes, and for the fractal scenes, the slopes were either non-significant or smaller than $0.4^\circ / \text{fixation}$.

Summary

Head fixations and movements are determined by the content of the panorama, in a manner that is remarkably similar to what was observed for the gaze data. In landscape panoramas, head fixations, direction, and movement amplitude show a pronounced tendency to be in alignment with the horizon, even when the panoramas are rotated along the view axis. This was complemented by the tendency of observers to rotate their heads to bring the landscape horizons somewhat closer to the canonical upright orientation. This pattern of fixations, movements, and head rotations was absent for the fractal images, which lack a well-defined horizon line. Instead, the direction of head movements tended to be in the cardinal directions, and the amplitude of the head movements tended to be in the horizontal direction, in world coordinates.

Gaze and Head

The previous sections analyzed head fixation patterns and eye fixation patterns separately, and as noted, they are remarkably similar. In this section, the focus is on the *relation* between head and gaze. One example of the relation between head and gaze is shown in Figure 3. Several observations are noteworthy: The head positions change only slowly; the head positions vary in a smaller range than the gaze positions (see Table 2); the head and gaze fixations coincide only in rare cases; and finally, the head fixations seem to lag behind the gaze fixations. These observations are explored more formally in this section.

Gaze – Head: Lag Analysis

If gaze and head positions coincide rarely, the question arises whether gaze direction is ahead of head direction or whether it is lagging behind. To answer this question, gaze and head fixations were compared using a minimum distance analysis: Given a (gaze) fixation g_i and a set of head fixations h_j before and after the time point i , the h_j^{min} with minimum distance was determined. If the h_j^{min} occurs before g_i then it is inferred that head is leading gaze, otherwise it was concluded that head is lagging behind gaze. The

results of this minimum distance analysis for landscape images are shown in Figure 13, indicating that most minimum distances occurred at a positive lag, i.e., the gaze positions were most often leading the head positions. Further details of the minimum distance analysis are presented in Appendix 2.

Figure 13: Histogram of gaze-head lags for landscape scenes (top panel) and fractal scenes (bottom panel). If gaze is closer to a later head position, i.e., has a positive lag, then gaze is leading. If gaze is closer to an earlier head position, i.e., has a negative lag, then gaze is trailing. The results show that gaze is leading head in most cases

Previous work on the coordination between eye and head movements suggest that the coordination between head and eye consists of synchronous movement onset with the eyes terminating well in advance of the slower head movements owing to longer contraction times for the neck muscles and the greater inertial forces acting on the head compared to the eye (Bizzi et al., 1971; Freedman, 2008; Gilchrist, Brown, Findlay, & Clarke, 1998). With attention shifts to sudden, bottom-up cues, for example, the line of sight is adjusted quickly through a saccadic eye movement, followed by the slower head movement combined with eye movement to keep the relevant stimulus projected on the fovea. With top-down control of the line of sight, the head tends to move first to select a newly centered visual field, followed by the eyes (Doshi & Trivedi, 2012). Taken together, adjustments of gaze, that is, the line of sight, are achieved through a combination of eye and head movements, both with different functions and with different temporal characteristics. The present work shows gaze to lead lag. This suggests that free visual exploration in panoramic scenes is largely a stimulus-driven bottom-up process, rather than a top-down controlled process.

Gaze – Head: Distance Analysis

Head movements are used to reposition the visual field in different directions of the environment and eye movements serve to explore the visual field defined by the head position. It is thus interesting to explore the range of fixations around the center of the visual field when the head is allowed to move freely. The upper panel of Figure 14 shows a histogram of gaze-head distances averaged over both scene type, all eight scene rotations, and all participants. To be clear, the histogram is based on the distances illustrated in Figure 7, that is, the distances between head positions and gaze positions that correspond in time. The results show that gaze and head positions coincide only rarely, and the most frequent head-gaze distances are found to be in the range of about 10 – 20 degrees, *Mode* = 15, *Md* = 17.3. The comparison between head and gaze positions at the same time point may, however, be somewhat misleading: As shown in the gaze-head lag analysis, head tends to lag behind gaze by a few fixations (Figure 13). It may thus be more appropriate to analyze the minimum distances between gaze and head positions, as explained in the

Section on gaze-head analysis and in Appendix 2. The lower panel of Figure 14 shows the distribution of minimum distances between head and gaze positions, with the most frequent minimum distances in the range 2 – 5 degrees, $Mode = 3$, $Md = 9.5$, which is substantially lower than the distances in the upper panel. Together, these results indicate that a fairly local range of locations close to the center of the head-defined visual field is explored. When points at larger distances are being explored then the head is moved.

Figure 14: Upper panel: Histogram of gaze-head distances. Lower panel: Histogram of minimum gaze-head distances. The distances in both panels have been averaged across scene types, scene rotations, and participants.

Gaze – Head: Direction Analysis

In a natural environment, visual selection is accomplished by a system of effectors, moving the head and body within space and the eyes within the visual field. Solman et al. (2017) used gaze-contingent displays to analyze whether the principles of selection for these different effectors are the same. They concluded that their results are consistent with a system in which the head favors exploration of unknown regions, while the eye exploits what can be seen with finer-scale saccades. In the present study, we were able to analyze the relationship between gaze and head movement in a setting where first, the visible part of the environment is unconstrained, and second, participants were able to move their head freely.

Polar histograms of the differences between head direction and saccade direction are shown in Figure 15. Recall that head fixations are defined as the average head position during a gaze fixation. For all head shifts $h(t)$, we compare the direction of $h(t)$ with the direction of the gaze saccade 0, 2, 4, or 6 fixations later. The results show that head direction and saccade direction are closely aligned when they occur at the same time (panel a), but are increasingly divergent as time progresses (panel c-d). This suggests that gaze and head have different characteristics, reflecting different roles in visual exploration, a result that confirms similar findings of Fang et al. (2015) and Solman et al. (2017).

Figure 15: Histograms of the directional differences between head direction and saccade direction, in bins of 10 degrees: (a) difference between head direction and gaze direction when they occur at the same time, (b) difference between head direction at time t and gaze direction at time $t+2$, i.e. two fixations later, (c)

difference between head direction at time t and gaze direction at time $t+4$, (d) difference between head direction at time t and gaze direction at time $t+6$.

Summary

The previous sections had analyzed head fixation patterns and eye fixation patterns separately, and as noted, they largely mirror one another. In this section, the *relation* between head and gaze data was examined, and the results suggest a complementary relationship between head and eye. A temporal analysis of gaze and head based on minimum gaze-head distances suggests that, in the visual exploration of panoramic scenes, gaze is almost always leading head. An analysis of the minimum distances between gaze and head suggests that gaze is used to explore a fairly local range within the head-defined visual field. Finally, a comparison of head shifts and saccades indicates that the two are fairly independent, suggesting that gaze and head have different roles in visual exploration (e.g., the head holds steady a particular region of visual space for the eyes to explore).

General Discussion

We investigated visual exploration of omni-directional panoramic scenes in a VR environment where the observers were allowed to freely move their head to inspect different areas of the panorama. We begin by discussing the different approaches to studying eye movement behavior and their advantages and disadvantages, followed by a characterization of head and gaze behavior in omni-directional scenes, as well as its implications for mechanisms of visual exploration.

Approaches to Studying Visual Exploration

In the past, most studies on eye movement behavior have used static images or videos (e.g., Foulsham, Cheng, Tracy, Henrich, & Kingstone, 2010; Foulsham, Walker, & Kingstone, 2011) presented on a screen. A majority of these studies measured eye movements while the observer's head was immobilized, usually by a chinrest. Even in studies that allowed free head movements, the head movements were of

minor importance because the visual field covered most, if not the complete presented stimuli. This study design has major implications for the study of eye movements. First, the visual information presented was pre-selected by the experimenter and was usually available in its entirety, hence it did not require head movements to select new and unexplored areas of the visual environment. Second, as Foulsham, Walker, & Kingstone (2011) have pointed out, gaze behavior without head movements may not reflect the dynamics of gaze selection in the real world, which involves head movements rather than large scanning eye movements (i.e., in real life people tend to look with their head and eyes, not with their eyes alone). Third, the scanning behavior while passively watching static or dynamic scenes shows very different characteristics compared to the scanning behavior while navigating through an environment. In particular, many studies have shown a central tendency of fixations within the visual field (e.g., Tatler, Baddeley, & Gilchrist, 2005; Tatler, 2007; Foulsham & Underwood, 2008). Indeed, a model that simply takes into account where the eyes are normally positioned in the head can outperform other models of looking behaviour (Foulsham & Kingstone, 2017).

In an alternative approach, more recent studies have measured looking behavior as participants navigate through a complex natural environment (e.g. walking through a university campus in Foulsham, Walker, & Kingstone, 2011; see also Foulsham, Chapman, Nasiopoulos, & Kingstone, 2014, t'Hart et al., 2009; t'Hart & Einhäuser, 2012). In this case, looking behavior does not reflect a simple visual exploration of the environment. Observers are faced with multiple tasks that need to be solved at the same time, namely detecting and coding of objects nearby and far away, path planning using far objects, and obstacle avoidance of objects and people nearby. Eye movements were thus more centralized than for scenes presented on a monitor and were often centered on the heading point during walking (Foulsham, Walker, & Kingstone, 2011; Foulsham & Kingstone, 2017). For these reasons, it can be challenging to make inferences about the characteristics of visual exploration from real-world studies using mobile eye trackers. Furthermore, due to physical and other constraints, the visual environment cannot be manipulated easily to test theories of visual exploration.

The present paper focused on visual exploration of omni-directional panoramic scenes in a VR environment, where head movements were freely allowed to select areas of interest in the panoramas. By tracking the head and eyes in VR, we were able to analyze simultaneous head and eye movements and fixations, and study their coordination. Furthermore, the present design also allowed us to control the visual environment for experimental purposes, i.e. to manipulate the content of the entire (360°) visual environment (landscapes and fractals) and to rotate these scenes around the view axis to determine their effects on viewing behavior.

Characterization of Gaze and Head Behavior in Omni-directional Scenes

The analysis of gaze behavior revealed different fixation patterns for landscape scenes and fractal scenes: For landscape scenes, fixations were concentrated along the scene horizons, a finding that confirmed earlier results of Foulsham, Walker, and Kingstone (2011) and Sitzman et al. (2018) for gaze distributions during walking. For fractal scenes, however, the fixation distributions were close to being isotropic, a result that can be attributed to the statistically non-oriented nature of the fractal scenes used. The analysis of saccades reflected the fixation distributions: For landscape scenes, a majority of saccades were made in the direction of the horizon, with those saccade amplitudes being larger than in the other directions (see also Foulsham, Teska, & Kingstone, 2011). For fractal scenes, the direction distribution of saccade was closer to isotropic, with a tendency of saccades in the cardinal directions in world coordinates, and saccade amplitudes being somewhat larger in the horizontal direction. A comparison of landscape and fractal scenes reveals that saccade amplitudes were larger and fixation durations shorter for the landscape scenes compared to the fractal scenes, suggesting that global scanning was more prevalent for landscape scenes, and local scanning more prevalent for fractal scenes.

The analysis of head movements showed that the variance of the head fixations was smaller than for the (gaze) fixations, suggesting that the eyes were used to extend the range of gaze positions. An analysis of head shifts yielded results that paralleled those of gaze: For landscape scenes, most head shifts were made in the direction of the scene horizons, with head shift amplitudes being larger than in the other

directions. For fractal scenes, head shifts were more likely in the cardinal directions and the largest amplitudes in the horizontal direction (in world coordinates). An analysis of head rotations (around the view axis) showed that the head tended to be rotated in counter-clockwise direction for scene rotations 45° to 135° , and clockwise for scene rotations 225° to 315° , in an attempt to bring the landscape horizons somewhat closer to the canonical upright orientation. For fractal scenes, there was no significant head rotation for any of the scene rotations.

Collectively, these data for both head and eye reveal that both systems are affected by the content of the 360° panorama. For landscape scenes, both effector systems seek out the horizon, even when the images are rotated. And when these scenes are rotated, the head often rotates to return the landscape horizons closer to their canonical upright orientation. For fractals, which have an isotropic texture, the behaviour of the eye and head systems change dramatically but in a similar way to one another. They are both insensitive to image rotation and display a preferential bias to move in the world centered cardinal coordinates, with the greatest amplitude being horizontal.

Finally, the analysis of the relationship between gaze and head movements indicates that while the two movements are similar, their functional relationship appears to be complementary. For landscapes, eye movements tended to lead head movements. This could be due to a bottom-up guided visual exploration (see e.g., Pelz, Hayhoe, & Loeber, 2001; Freedman, 2008; Doshi & Trivedi, 2012). An analysis of the distance between gaze and head positions indicated that gaze was used to explore a fairly local range close to the center of the visual field, consistent with the observations of Foulsham, Walker, and Kingstone (2011) who found that most eye movements occurred, relative to the head, within the center of the visual field relative (see also Foulsham & Kingstone, 2017). Finally, a directional analysis showed that head and saccade directions were closely aligned when they occurred at the same time, but diverge increasingly as time progresses. This suggests that gaze and head have complementary roles in visual exploration, confirming similar findings of Fang et al. (2015) and Solman et al. (2017).

Implications for Visual Exploration

In everyday life, we rely on visual input for a multitude of reasons, such as to detect persons and objects of interest and to interact with them, to plan navigational paths through the environment and execute these plans while avoiding obstacles, to recognize and respond to potentially dangerous event, and to explore and remember the structure of the environment around us. The present study is focused exclusively on the last aspect, namely the visual exploration of the environment unencumbered by other tasks. To this effect, we studied the basic exploration characteristics in a VR environment, which allows one to not only have full control over the environment that observers are exploring, but to also study how their head and eye movement systems are interacting during their exploration. To our knowledge this is the first study to systematically investigate the relationship between eye and head movements while observers view fully immersive panoramic scenes. We have found that when people are free to explore these scenes, they move both their head and their eyes. While this may raise concerns for past work that has focused almost exclusively on head-fixed viewing, we found that head movements closely mirror eye movement behaviour. Yet, the head movement data reveal significant, and new findings. First, the head seems to act in service of the eye. Not only does it follow the eye to keep the eye centered in the head, it rotates in order to bring the landscape images into a more canonical upright position for exploration. Second, dynamic patterns of eye and head movements that differ between landscape and fractal images reveal specific scanning strategies that are exquisitely sensitive to the scene content.

In sum, our results demonstrate that a VR environment constitutes a very useful and informative alternative to traditional eye movement research approaches, to those methods that use static and dynamic scenes presented on a monitor, and to investigations that use mobile eye trackers during navigation through the real world. The present study is only a first step in this direction, and the outlook for future studies with static and dynamic VR environments looks very promising indeed.

Acknowledgements

This work was supported by Discovery Grants from the Natural Sciences and Engineering Research Council of Canada (WFB: RGPIN-2015-05560, AK: Grant 12R80338, NCA: Postdoctoral Fellowship).

Appendix 1: Circular statistics

The gaze tracker values were transformed into longitude and latitude of the spherical panorama, with longitude in the range $[-180, 180]$ degrees and latitude in the range $[-90, 90]$ degrees. To cope with the wrap-around at ± 180 degrees longitude, means and dispersion of longitude values were computed using circular statistics (Batschelet, 1981; Mardia & Jupp, 2000).

Circular mean: Given a set of longitude values $\lambda_1, \lambda_2, \dots, \lambda_n$ in the range $[-180, 180]$ degrees, the mean longitude L is defined as

$$L = \tan^{-1} \left(\frac{\sum_{i=1}^n \sin \lambda_i}{\sum_{i=1}^n \cos \lambda_i} \right)$$

Circular range: Given a set of longitude values $\lambda_1, \lambda_2, \dots, \lambda_n$ in the range $[-180, 180]$ degrees, the range R is defined as follows: Let $\lambda_{(1)}, \lambda_{(2)}, \dots, \lambda_{(n)}$ be the linearly ordered $\lambda_1, \lambda_2, \dots, \lambda_n$. Define $T_i = \lambda_{(i+1)} - \lambda_{(i)}$ for $i = 1, \dots, n-1$ and $T_n = 360 - \lambda_{(n)} + \lambda_{(1)}$. Then the circular range ω is defined as $\omega = 360 - \max(T_1, T_2, \dots, T_n)$.

Saccades are defined as differences between successive fixations, more precisely in terms of great circles on the spherical panorama: Given are two fixations $f_1 = (\lambda_1, \varphi_1)$ and $f_2 = (\lambda_2, \varphi_2)$ with longitudes λ_1, λ_2 and latitudes φ_1, φ_2 .

The *saccade amplitude* between fixations f_1 and f_2 is defined by the great-circle distance

$$\Delta\sigma = r \cdot \cos^{-1}(\sin \varphi_1 \sin \varphi_2 + \cos \varphi_1 \cos \varphi_2 \cos |\Delta\lambda|),$$

where $\Delta\lambda = \lambda_2 - \lambda_1$ and r is the sphere radius.

The *saccade direction* α_1 starting at fixation f_1 is defined by the azimuth of the great circle at f_1 , is defined by

$$\alpha_1 = \tan^{-1}(\cos \varphi_2 \sin \Delta\lambda / \cos \varphi_1 \sin \varphi_2 - \sin \varphi_1 \cos \varphi_2 \cos \Delta\lambda)$$

Appendix 2: Gaze-Head Lag Analysis

Give are two sequences, a sequence $G = g_1 g_2 \dots g_n$ of gaze positions and a sequence $H = h_1 h_2 \dots h_n$ of head positions. The calculation of the head position sequence H is described in section on gaze - head lag.

The minimum distance method is based on calculating all $g_i - h_j$ distances within a fixed temporal window of size $[-k, k]$. We used $k = 8$ for our analyses. For each gaze direction g_i and a set of head positions h_j with $j = i-k, i-(k-1), \dots, i-1, i, i+1, \dots, i+(k-1), i+k$, the pair $g_i - h_j$ with minimum distance is determined. If $i < j$, i.e., the h_j with minimum distance occurred before g_i , then it was concluded that head was leading gaze, otherwise if $i > j$, i.e., h_j with minimum distance occurred after g_i then it was concluded that head was lagging behind gaze.

References

- Batschelet, E. (1981). *Circular Statistics in Biology*. London: Academic Press.
- Birmingham, E., Bischof, W. F. and Kingstone, A. (2008). Gaze selection in complex social scenes. *Visual Cognition*, 16, 341-355.
- Birmingham, E., Bischof, W. F. and Kingstone, A. (2009). Saliency does not account for fixations to eyes within social scenes. *Vision Research*, 49, 2992-3000.
- Bizzi, E., Kalil, R. E., & Tagliasco, V. (1971). Eye-head coordination in monkeys: Evidence for centrally patterned organization. *Science*, 173, 452-454.
- Blignaut, P. (2009). Fixation identification: The optimum threshold for a dispersion algorithm. *Attention, Perception, & Psychophysics*, 71, (4), 881-895.
- Buswell, G. T. (1935). *How People Look at Pictures: A Study of the Psychology of Perception in Art*. Chicago, IL: University of Chicago Press.
- Coppola, D. M., Purves, H. R., McCoy, A. N., & Purves, D. (1998). The distribution of oriented contours in the real world. *Proceedings of the National Academy of Sciences of the USA*, 95(7), 4002-4006.
- Cristino, F., & Baddeley, R. (2009). The nature of visual representations involved in eye movements when walking down the street. *Visual Cognition*, 17(6-7), 880-903.
- Doshi, A., & Trivedi, M. M. (2012). Head and gaze dynamics during visual attention shifts in complex environments. *Journal of Vision*, 12(2), 1-15. doi:10.1167/12.2.9
- Fang, Y., Nakashima, R., Matsumiya, K., Kuriki, I., & Shioiri, S. (2015) Eye-head coordination for visual cognitive processing. *PLoS ONE*. doi:10.1371/journal.pone.0121035
- Foulsham, T., Chapman, C., Nasiopoulos, E., & Kingstone, A. (2014). Top-down and bottom-up aspects of active search in a real world environment. *Canadian Journal of Experimental Psychology*, 68, 8-19.
- Foulsham, T., Cheng, J. T., Tracy, J. L, Henrich, J., & Kingstone, A. (2010). Gaze allocation in a dynamic situation: Effects of social status and speaking. *Cognition*, 117, 319-331.

- Foulsham, T. & Kingstone, A. (2010). Asymmetries in the direction of saccades during perception of scenes and fractals: Effects of image type and image features. *Vision Research* 50, 779–795. doi:10.1016/j.visres.2010.01.019
- Foulsham, T. & Kingstone, A. (2017). Are fixations in static natural scenes a useful predictor of attention in the real world? *Canadian Journal of Experimental Psychology*, 71, 172-181.
- Foulsham, T., Kingstone, A., & Underwood, G. (2008). Turning the world around: Patterns in saccade direction vary with picture orientation. *Vision Research* 48, 1777–1790. doi:10.1016/j.visres.2008.05.018
- Foulsham, T., Teska, R., & Kingstone, A. (2011). Saccade control in natural images is shaped by the information visible at fixation: evidence from asymmetric gaze-contingent windows. *Attention, Perception & Psychophysics*, 73, 266-283.
- Foulsham, T., Walker, E., & Kingstone, A. (2011). The where, what and when of gaze allocation in the lab and the natural environment. *Vision Research*, 51, 1920-1931. doi:10.1016/j.visres.2011.07.002.
- Freedman, E. G. (2008). Coordination of the eyes and head during visual orienting. *Experimental Brain Research* 190, 369–387. doi:10.1007/s00221-008-1504-8
- Gilchrist, I. D., Brown, V., Findlay, J. M., & Clarke, M. P. (1998). Using the eye-movement system to control the head, *Proceedings of the Royal Society of London B*, 265, 1831–1836. doi: 10.1098/rspb.1998.0509.
- Hessels, R. S., Niehorster, D. C., Nyström, M., Andersson, R., & Hooge, I. T. C. (2018). Is the eye-movement field confused about fixations and saccades? A survey among 124 researchers. *Royal Society Open Science*, 5, 180502. doi:10.1098/rsos.180502
- Komogortsev, O. V., Gobert, D. V., Jayarathna, S., Koh, D., & Gowda, S. (2010). Standardization of automated analyses of oculomotor fixation and saccadic behaviors, *IEEE Transactions on Biomedical Engineering*, 57(11), 2635-2645.

- Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception & Performance*, 4(4), 565-572.
- Mackworth, N. H., & Morandi, A. J. (1967). The gaze selects informative details within pictures. *Perception & Psychophysics*, 2, 547-552.
- Mardia, K. V., & Jupp, P. E. (2000). *Directional Statistics*. Chichester: Wiley.
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, 42(1), 107-123.
- Pelz, J., Hayhoe, M., & Loeber, R. (2001). The coordination of eye, head, and hand movements in a natural task. *Experimental Brain Research*, 139(3), 266-277.
- Potter, M. C., Staub, A., & O'Connor, D. H. (2004). Pictorial and conceptual representation of glimpsed pictures. *Journal of Experimental Psychology – Human Perception and Performance*, 30(3), 478-489.
- Salvucci, D. D., & Goldberg, J. H. (2000). Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the EyeTracking Research and Applications Symposium* (pp. 71-78). New York: ACM Press.
- Schizo. (2010). Descent into Fractal Core and Fractal Matrix [Video]. Retrieved from <http://www.wearvr.com/apps/fractal-canyon-vr>
- Sitzmann, V., Serrano, A., Pavel, A., Agrawala, M., Gutiérrez, D., Masia, B., & Wetzstein, G. (2018). Saliency in VR: How do people explore virtual environments? *IEEE Transactions on Visualization and Computer Graphics*, 24(4), 1633-1642.
- Solman, G. J. F., Foulsham, T., & Kingstone, A. (2017). Eye and head movements are complementary in visual selection. *Royal Society Open Science* 4, 160569. doi:10.1098/rsos.160569
- Solman, G. J. F., & Kingstone, A. (2014). Balancing energetic and cognitive resources: Memory use during search depends on the orienting effector. *Cognition*, 132, 443-454. doi:10.1016/j.cognition.2014.05.005

- Tatler, B. W. (2007) The central fixation bias in scene viewing: selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision*, 7(14):4, 1-17.
- Tatler, B. W., Baddeley, R. J., & Gilchrist, I. D. (2005). Visual correlates of fixation selection: Effects of scale and time. *Vision Research*, 45, 643–659.
- Tatler, B. W., & Vincent, B. T. (2008). Systematic tendencies in scene viewing. *Journal of Eye Movement Research*, 2(2):5, 1-18.
- Technologies, U. (2015). *Unity - Manual: Unity Manual*. [online] Docs.unity3d.com. Retrieved from <http://docs.unity3d.com/Manual/index.html>
- 't Hart, B. M., Vockeroth, J., Schumann, F., Bartl, K., Schneider, E., König, P., & Einhäuser, W. (2009). Gaze allocation in natural stimuli: Comparing free exploration to head-fixed viewing conditions. *Visual Cognition*, 17(6-7), 1132-1158.
- 't Hart, B. M., & Einhäuser, W. (2012). Mind the step: complementary effects of an implicit task on eye and head movements in real-life gaze allocation. *Experimental Brain Research*, 223(2), 233-249.
- Torralba, A. (2003). Modeling global scene factors in attention. *Journal of the Optical Society of America A – Optics Image Science and Vision*, 20(7), 1407-1418.
- Underwood, G., & Foulsham, T. (2006). Visual saliency and semantic incongruency influence eye movements when inspecting pictures. *Quarterly Journal of Experimental Psychology*, 59(11), 1931-1949.
- Underwood, G., Foulsham, T., & Humphrey, K. (2009). Saliency and scan patterns in the inspection of real-world scenes: Eye movements during encoding and recognition. *Visual Cognition*, 17, 812–834.
- Unema, P. J. A., Pannasch, S., Joos, M., & Velichkovsky, B. M. (2005). Time course of information processing during scene perception: The relationship between saccade amplitude and fixation duration. *Visual Cognition*, 12(3), 473-494.
- Wikipedia (2019). *Equirectangular projection*. Retrieved from https://en.wikipedia.org/wiki/Equirectangular_projection

- Xiao, J., Hays, K., Ehinger, A., Oliva, A., & Torralba, A. (2010). SUN database: Large-scale scene recognition from abbey to zoo". *IEEE Conference on Computer Vision and Pattern Recognition*, June 2010. doi: 10.1109/CVPR.2010.5539970
- Yarbus, A. L (1967). *Eye movements and vision* (B. Haigh, Trans.). New York: Plenum Press. (Original work published 1965)