



Turning the world around: Patterns in saccade direction vary with picture orientation

Tom Foulsham^{a,*}, Alan Kingstone^b, Geoffrey Underwood^a

^aSchool of Psychology, University of Nottingham, University Park, Nottingham, NG7 2RD, UK

^bDepartment of Psychology, University of British Columbia, Vancouver, BC, Canada

ARTICLE INFO

Article history:

Received 3 December 2007

Received in revised form 21 May 2008

Keywords:

Eye movements
Top-down guidance
Bottom-up guidance
Scene perception
Animal classification
Human

ABSTRACT

The eye movements made by viewers of natural images often feature a predominance of horizontal saccades. Can this behaviour be explained by the distribution of saliency around the horizon, low-level oculomotor factors, top-down control or laboratory artefacts? Two experiments explored this bias by recording saccades whilst subjects viewed photographs rotated to varying extents, but within a constant square frame. The findings show that the dominant saccade direction follows the orientation of the scene, though this pattern varies in interiors and during recognition of previously seen pictures. This demonstrates that a horizon bias is robust and affected by both the distribution of features and more global representations of the scene layout.

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

There is a wealth of evidence showing that the places where people fixate are not randomly distributed throughout a scene or picture (Buswell, 1935; Mackworth & Morandi, 1967; Yarbus, 1967). These fixations, and the saccadic eye movements between them, position the part of the retina with the highest resolution (the fovea) on certain points in an image, leading to preferential processing compared to non-fixated or peripheral regions. For this to be efficient the locations selected must be informative with regard to perceiving the parts of the scene and accomplishing any task goals. Rather than being uniformly distributed through a scene (which would lead to an even coverage of all regions), fixations are indeed targeted at points which are “informative” (Mackworth & Morandi, 1967); objects inconsistent with scene expectations (Loftus & Mackworth, 1978; Underwood & Foulsham, 2006); and regions with high local contrast (Reinagel & Zador, 1999) or entropy (Raj, Geisler, Frazor, & Bovik, 2005). These observations have led to formalised models which aim to predict fixation locations. A particular class of these, saliency map models, combine information on low-level contrast within various neurobiologically plausible features (Itti & Koch, 2000; Itti & Koch, 2001; Koch & Ullman, 1985). These models perform better than chance in some circumstances (Foulsham & Underwood, 2008; Parkhurst, Law, & Niebur, 2002), although their limitations have

also been frequently pointed out (Henderson, Brockmole, Castelano, & Mack, 2007; Tatler, Baddeley, & Gilchrist, 2005). Other models take into account top-down goal (such as the known identity of a target object; Navalpakkam & Itti, 2005; Rao, Zelinsky, Hayhoe, & Ballard, 2002) and abstract, “gist” information (Torralba, 2003; Torralba, Oliva, Castelano, & Henderson, 2006).

A common assumption in bottom-up models is that, prior to any saliency computation, all possible eye movements are equally likely. Given that visual acuity decreases rapidly with eccentricity, treating all retinotopic locations as equally likely to be fixated is problematic, and some researchers have addressed this (Parkhurst et al., 2002; Vincent, Troscianko, & Gilchrist, 2007). In other work it has been argued that systematic biases in which part of a display are fixated, in particular a central bias, should also be considered (Tatler et al., 2005). It is not always clear, however, whether there is a tendency to fixate centrally independent of the distribution of salient features, or whether fixations are biased towards the centre because salient objects are often located there. For example, the horizon in landscape photographs often provides a high contrast edge, which might attract attention in a bottom-up fashion.

In this paper we investigate a related bias found in picture viewing: asymmetry in saccade direction. Several authors have reported that there are many more horizontal (leftwards or rightwards) than vertical saccades, and that there are even fewer oblique angle saccades. This is a general observation aside from the length of the saccade or its starting point, and the pattern has been seen in a variety of tasks and stimuli (Brandt, 1945; Crundall & Underwood, 1998; Gilchrist & Harvey, 2006). Why do people

* Corresponding author.

E-mail address: lpxtf@psychology.nottingham.ac.uk (T. Foulsham).

routinely move their eyes in this way when viewing scenes? There are several possibilities.

First, in what we will call the *oculomotor* explanation, the distribution of saccade directions might be due to dominance of the muscle or neural apparatus which preferentially triggers horizontal shifts of the eyes, regardless of the stimulus being viewed. Although most physiological research has concentrated on horizontal saccades there is some evidence that vertical and oblique saccades are slower (Collewyn, Erkelens, & Steinman, 1988) and exhibit more curvature (Becker & Jurgens, 1990) than horizontal saccades (see Becker, 1991 for a review of saccade dynamics). This would support a low-level bias for horizontal saccades, and such a bias would be expected across different stimuli.

Second, an *image-characteristics* explanation could explain this bias in terms of the distribution of salient features in pictures and natural scenes. The horizon often features in outdoor scenes and this is normally marked by a high contrast edge between dark ground and lighter sky. It might also be the case that the semantically important objects in the scene (people, cars, buildings) are found near this horizon. Photographs of the natural environment are usually composed with these objects in the centre (in fact a beginners' photography heuristic suggests that the horizon should be around two-thirds of the way up the picture). This non-uniform distribution of salient features has been identified as a confound in studies which show a correlation between saliency and fixations (Tatler et al., 2005). If people are reflexively drawn to regions of high contrast or high saliency in the periphery (as suggested by saliency map models), and if these regions tend to be positioned horizontally from each other, then this could cause a predominance of horizontal saccades. It has also been noted that natural and manmade scenes tend to have more horizontally and vertically oriented contours than oblique ones (Coppola, Purves, McCoy, & Purves, 1998), which could be an image-based determinant of saccade direction.

Alternatively, a *learned* account could predict a horizontal bias based on our experience with pictures and the environment. By this account, horizontal saccades are not favoured automatically by neurophysiology or caused by relatively automatic orienting to salient objects on the midline but rather learned over time and initiated top-down. Following multiple experiences with scenes where important information (both visually salient and semantically interesting) is located on the horizon we learn to move our eyes in this way, in order to maximize the details observed in the fewest saccades. This learning might be subject to cultural and experiential differences, for example, in terms of reading habits. Consistent with this Abed (1991) reported differences in scanning direction between Western, Middle-Eastern and East Asian participants looking at simple dot patterns. While Western readers made more shifts moving from left to right, Arabic readers were more likely to show the opposite pattern. East Asian participants showed a 1:1 ratio of horizontal to vertical saccades unlike the 2:1 ratio seen in other readers. In a different learned domain Western drivers showed a ratio which was closer to 4:1, though interestingly there was no difference between experts and novices despite the former presumably having more experience with the layout of the road (Crundall & Underwood, 1998; Underwood, Chapman, Bowden, & Crundall, 2002).

A more specific instance of this would be a *learned layout* explanation. In this more flexible account, basic cues about the layout of the scene might influence the likelihood of moving along each axis. For example it is known that the overall gist of a picture (for example whether it is outdoors or indoors) can be acquired very rapidly (Biederman, Rabinowitz, Glass, & Stacy, 1974; Kirchner & Thorpe, 2006; Potter, 1976; Potter, Staub, & O'Connor, 2004). Coarse information gathered from the first glimpse might also include simple knowledge about the location of the horizon or the overall structure. In a series of experiments, Sanocki (2003) showed that briefly

shown scenes can prime spatial layout, and that this priming affects subsequent perception. This supports the proposition that scenic layout, and not just scene category, is represented following an initial glimpse. This knowledge could affect which way the eyes are likely to move.

A final possibility is that the biases in saccade direction are *display-specific*, an artefact of laboratory-based eye tracking studies that present scenes on a computer monitor. These monitors are normally wider than they are high, and pictures are often presented in the landscape orientation filling the screen. Thus it may be more efficient to move horizontally than vertically in order to cover the whole area. Experiments often cue attention with a fixation cross in the centre of the screen at the start of a trial. In this case there is more information to the left and right of fixation than above or below, and this continues to be the case with an asymmetric display. In addition, older studies often suffered from large tracking errors, which tended to be greater in vertical saccades than in horizontal ones (Yee et al., 1985). In the real world, biases in saccade direction might be different, although Crundall and Underwood (1998) found the bias occurred with real roads. Of course the human field-of-view with two eyes is also asymmetric, spanning up to 180 degrees horizontally and only around 90 degrees vertically, and this is extended further with head movements which move further horizontally than vertically. These head movements affect the way the eyes move in the real world (Pelz, Hayhoe, & Loeber, 2001; Smeets, Hayhoe, & Ballard, 1996).

The experiments reported here investigate the distribution of saccade directions in an attempt to distinguish between these explanations. The accounts summarised above are not mutually exclusive and several of them might combine in natural vision. For example when combined with a central starting position in the laboratory bottom-up information might drive saccade direction. The environment in which humans find themselves is likely to have shaped our physiology in terms of field-of-view and head and eye musculature. Studies of non-human animals demonstrate that the environment, and in particular the prominence of the horizon in the natural habitat, affects the organisation of the retina (Hughes, 1977). How flexible is the human tendency to move their eyes in a certain way, and when in viewing does this tendency emerge?

To answer these questions in Experiments 1 and 2 we rotated natural images from the horizontal whilst recording the eye movements made in a simple scene understanding task. We controlled various elements of the display in order to remove artefacts of the laboratory set-up. If the pattern of saccade directions is due to oculomotor factors or long-term learning then it should be insensitive to trial-by-trial variations in scene orientation. On the other hand, if the pattern changes with scene rotation the bias must arise from changes in the distribution of salient features or early recognition of the scene layout. A special case concerns scenes which are rotated 180 degrees and are therefore upside-down. Inverted scenes will preserve any clustering of features around the horizon but scene inversion might disrupt gist acquisition and scene recognition. In this case it will be informative to discover whether the saccades resemble those in normal, correctly oriented scenes.

2. Experiment 1

2.1. Method

2.1.1. Participants

Thirteen student volunteers from the University of Nottingham with normal or corrected-to-normal vision took part for payment. All were naïve to the purpose of the experiment and gave their full, informed consent to participate.

2.1.2. Stimuli and design

Experiment 1 used landscape photographs that all had a visible horizon and thus tended to have a large contrast boundary running horizontally through the middle third of the picture. Forty colour photographs of landscapes and outdoor scenes were chosen from a commercially available CD-ROM or taken with a high resolution digital camera. To create the normally oriented set these pictures were then cropped into square images of 768 by 768 pixels around their centre (see Fig. 1a). Making all the images square helped to control for any effect a rectangular display might have on saccade direction, and gave a consistent frame of reference across all trials. The original full-size images were then rotated by 45, 90, 135 and 180 degrees using photographic manipulation software. Cropping these into squares gave the rotated stimuli sets (see Fig. 1b–e). Preparing the stimuli in this way ensured that the borders of the picture were square in each case, whilst the horizon was rotated. The final stimuli were composed of five rotation sets (0, 45, 90, 135 and 180 degrees) with eight pictures in each set. As this experiment was mainly concerned with the axis on which the picture was aligned, the pictures in the 45, 90 and 135 degrees sets contained pictures that were rotated both clockwise and anti-clockwise from the horizontal (so that pictures in the 45 degree set were equally likely to contain an orientation aligned at +45 degrees and +225 degrees).

2.1.3. Apparatus

Pictures were displayed on a 19 inch monitor with a refresh rate of 125 Hz. The square images were centred on a white background, which filled the screen resolution of 1024×768 pixels. At a fixed viewing distance of 60 cm this subtended 34×27 degrees of visual angle. The eyes were tracked using the Eyelink II (SR Research), a head-mounted, video-based system that sampled the pupil image at 500 Hz. Fixations and saccades were identified by the Eyelink system on the basis of the manufacturer's displacement, velocity and acceleration thresholds, which were 0.10° , $30^\circ/s$ and $8000^\circ/s^2$, respectively. A validation procedure before the experiment selected the eye that provided the best spatial resolution for tracking, and this gave a mean error of less than 0.5° for all subjects. Spatial resolution was high for both horizontal and vertical movements (error was less than 0.5° in each direction), and where this wasn't the case the calibration was repeated until this level of accuracy was achieved. A chin rest was used and participants were instructed not to move their head during trials. A head camera on the eye tracker also monitored head movements but these occurred very infrequently. When these occurred participants were

reminded to keep their head still, the trial number was recorded, and data from this trial was excluded from the results. Responses were entered using a gamepad.

2.1.4. Procedure

Participants were first calibrated in the eye tracker using a 9-dot grid, and this was repeated several times to minimize error. A short practice example was shown prior to the experimental trials. Each trial (see Fig. 2) began with a drift correct dot. This was a small circular marker, displayed using the Eyelink calibration software, which re-aligned the calibration and also had the effect of forcing the participant to start scanning at a particular place in the image. The participant was required to press a button whilst fixating this location and the Eyelink software shifted the calibration to compensate for any small changes due to head movements or slippage of the eye tracker. Large deviations resulted in a new calibration. In order to avoid artefactual effects on saccade direction of a central starting point, the drift correct location was varied from trial to trial and appeared in one of the four corners of the display (approximately 4 degrees from the edge of the monitor) and just outside the boundary of the square stimulus. Each starting point was used equally within each rotation set but was otherwise random. After fixation of this point was confirmed the picture was displayed for 2.5 s.

Participants were instructed to take in as much information as possible, and each picture was followed by a written sentence presented on the screen. Participants were required to verify the truth of the sentence in terms of the picture they had just inspected, and they did this by pressing one of two keys on a keypad to indicate true or false. All sentences were active declaratives referring to the identity or location of objects or scene features (see Fig. 2 for example). Each rotation set was associated with an equal number of true and false sentences. The trial ended with the participant's response, and then the next trial began. All forty stimuli were presented in a randomised order that was unique for each subject. Eye movements were recorded while the picture was presented, and image onset and offset times were also written to the data file to ensure that no eye movements from reading the sentence (which would largely be horizontal) were included.

2.2. Results

Participants were quite accurate at the sentence verification, showing that they were able to do the task (mean proportion

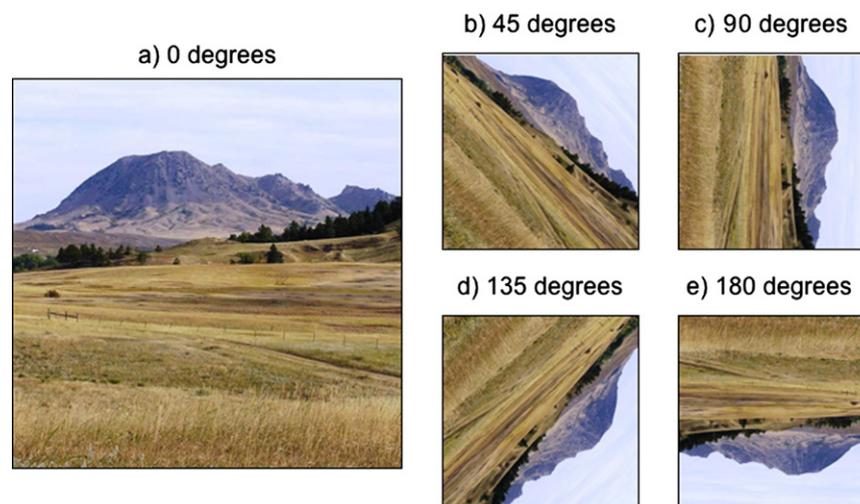


Fig. 1. An example stimulus from Experiment 1 (a). The final rotation sets also contained images at four different rotations (b–e).

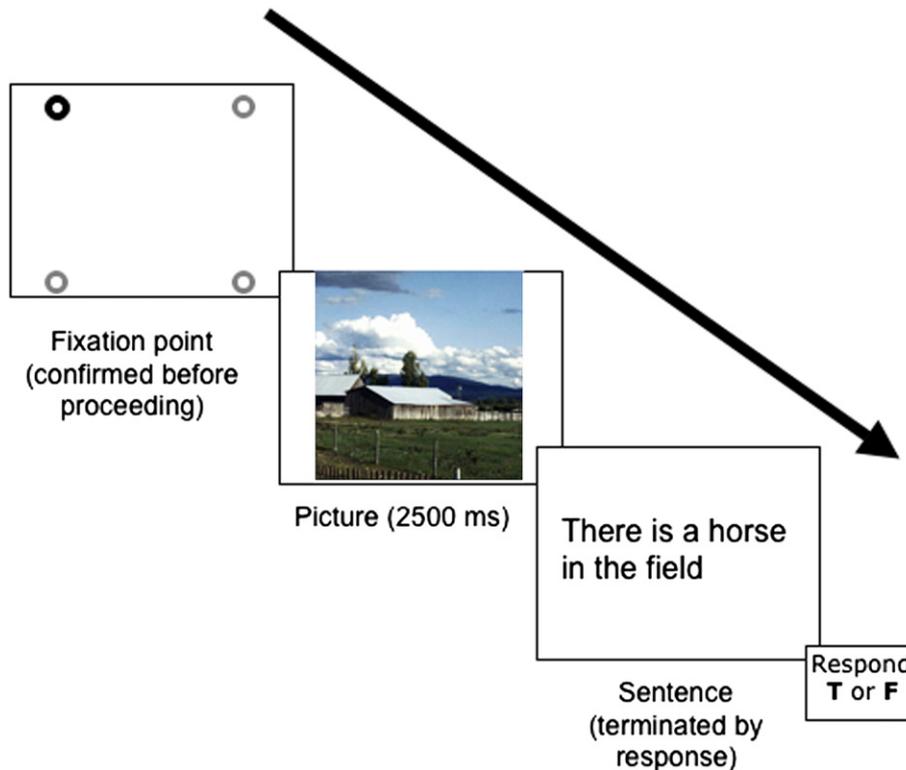


Fig. 2. The procedure for one trial. Fixation started at one of four points (grey circles). The picture was then shown for a fixed period before a sentence verification checked picture understanding.

correct 0.823). The remainder of the results aim to characterise the saccades made whilst encoding the pictures for the sentence task.

2.2.1. Saccade direction

The angular direction of all saccades made whilst inspecting the pictures was recorded. The first saccade in each trial started at the experimenter-controlled drift correct location and as such it was examined separately. Inspection of these saccades and their landing points showed that in the vast majority of cases the direction and amplitude was determined by the start point, with a general tendency for the saccade to move to the centre of the screen (see Fig. 3). As a result all further analyses looked only at subsequent saccades ($N = 4320$).

To describe the overall pattern of saccade directions associated with each picture orientation, the following procedure was followed. First the small number of saccades shorter than 1 degree were removed, in order to exclude readjustive and microsaccades. This removed less than 5% of the data. All possible directions were then divided into 36 bins of 10 degrees each. To remain consistent with the labels for picture orientation, these bins were numbered clockwise from the horizontal, starting at 0 for saccades that lay between 350 and 0 degrees (i.e., almost exactly leftwards). The proportion of saccades in each bin was then plotted in a polar plot. These plots are shown in Fig. 4 separately for each picture orientation (a–e). Looking first at the normally oriented condition (a), it is clear that there is a strong horizontal bias, with more than twice as many saccades in the 0 and 180 degree bins than in the 90 and 270 degree bins. There are also more saccades in the vertical than in the oblique. The peak of saccades on the horizontal axis is particularly pronounced for rightwards (180 degree) saccades. Fig. 4b–d show clearly that this pattern changes with the orientation of the picture, so that more saccades are made in the axis in which the horizon of the picture lies; saccade direction is effected by picture orientation. Fig. 4f combines all orientation conditions by rotating

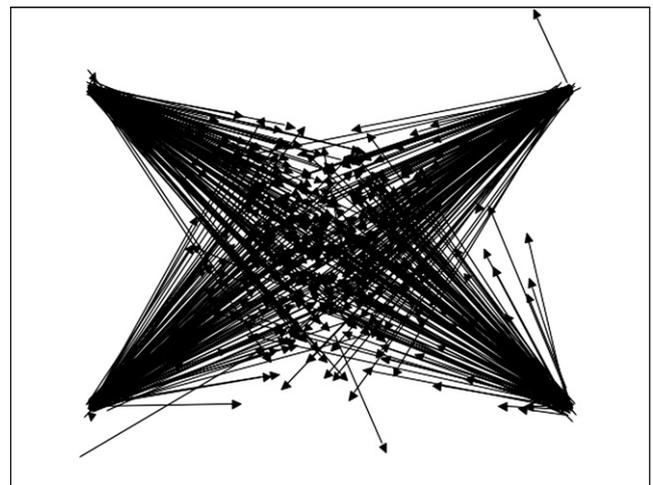


Fig. 3. The first saccade in each trial, across all participants in Experiment 1. Trials started in one of four locations, but the first saccade was almost always made towards the centre.

the direction bins so that the charts horizontal (0 degrees) is aligned with the original orientation of the picture, showing a relative horizontal bias across conditions.

In order to make some clearer statistical comparisons, the same data were divided into four bins according to the closest axis (horizontal, 45/225 degrees, vertical and 135/315 degrees). The bins were defined using all eight directions (four cardinal and four oblique) $\pm 22.5^\circ$, so that, for example, the 45/225 degrees bin contained all saccades greater than or equal to 22.5° and less than 67.5° , along with all those greater than or equal to 202.5° and less than 247.5° . Whilst this loses any asymmetries in terms of left or rightwards saccades, it allows for a more straightforward analysis.

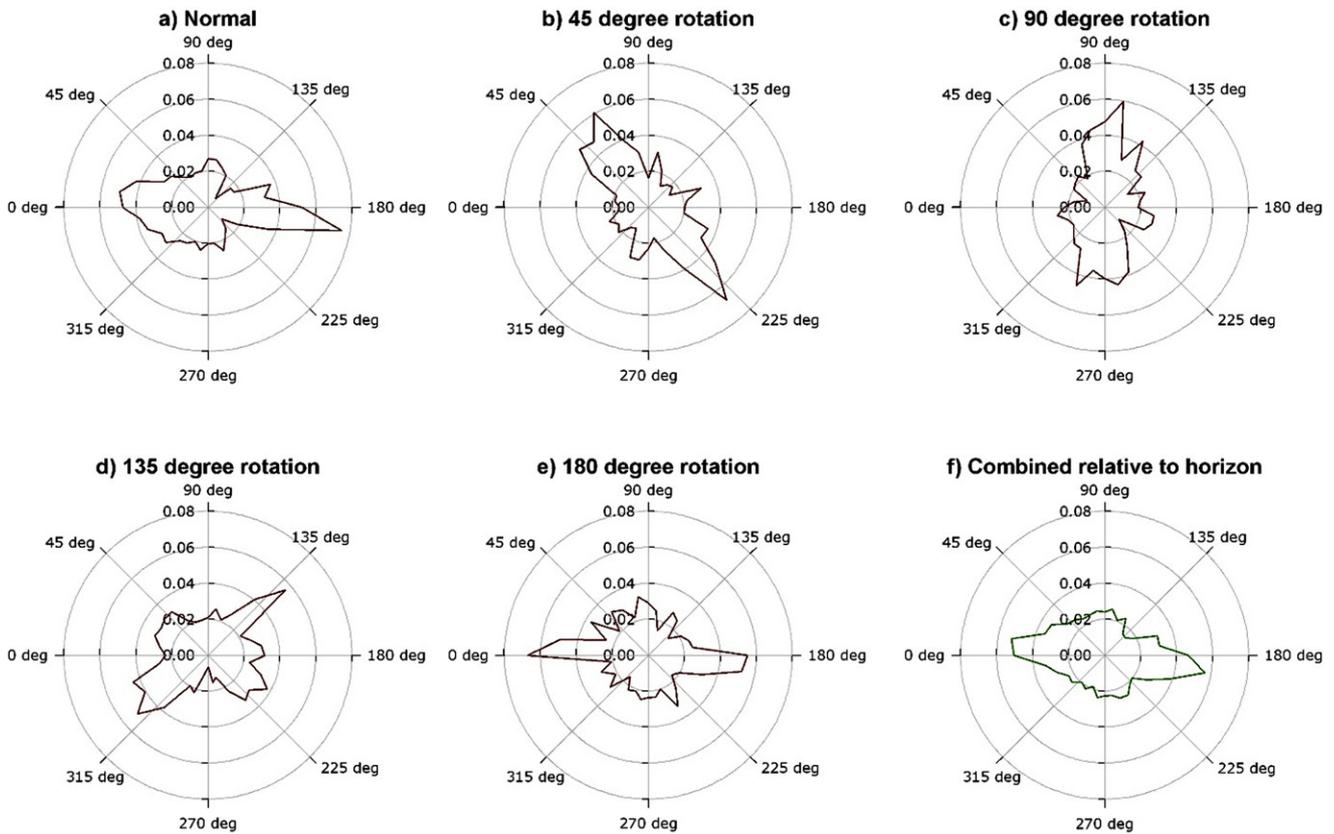


Fig. 4. Radial histograms for all saccades (excluding the first). Each plot shows the proportion of saccades (y axis) in each of 36 direction bins. The majority of saccades occur in the axis of the horizon.

We confirmed that most of the variation in the distributions is symmetrical by comparing the frequency of leftward (<90° or >270°) and rightward (>90° and <270°) saccades. This is equivalent to comparing the left and right sides of Fig. 4f, and there was no difference ($t_{12} < 1$) and therefore no evidence of any asymmetry. The frequency of saccades within each axis was computed, for each subject, in each orientation condition. Fig. 5 shows these data as a proportion of the total number of saccades made in each condition. The data were compared using two-way (4 directions by 5 picture orientations) repeated-measures ANOVA. The oblique picture orientations contained pictures rotated both clockwise and anti-clockwise. Planned *t*-test comparisons then compared the frequency of saccades in the same axis as the picture orientation with the mean of the other three axes.

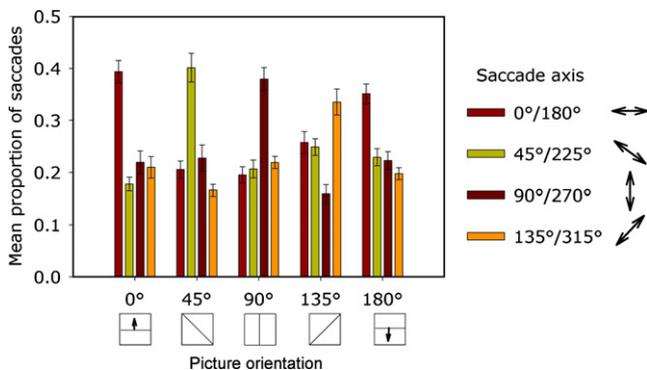


Fig. 5. Mean proportion of saccades within each saccade axis, as a function of picture orientation. Error bars indicate plus/minus one standard error of the mean across participants.

Across picture conditions there was a reliable effect of direction on saccade frequency ($F_{3,36} = 3.97$, $MSE = 36.2$, $p < 0.05$), with fewer saccades being made in the 135 degree axis than in the horizontal axis (post hoc *t*-test, $t_{12} = 3.58$, $p < 0.005$) or the 45 degree axis ($t_{12} = 2.91$, $p < 0.05$). No other comparisons were reliable. There was no main effect of picture orientation ($F_{3,36} = 2.11$, $p = .094$) indicating that roughly the same number of saccades was made regardless of how the picture was rotated. Most importantly, the change in the complete direction distribution is shown in this analysis by an interaction between orientation and axis ($F_{12,144} = 21.25$, $MSE = 20.65$, $p < 0.001$). It can be seen from Fig. 5 that this interaction is due to the greatest proportion of saccades in any one condition being made close to the horizon's axis. This was confirmed with planned comparisons which compared the mean frequency of saccades in the axis where the picture's horizon was located with the mean frequency of saccades elsewhere, collapsed across the remaining levels. In all cases the comparison was highly reliable (at 0 degree orientation, $t_{12} = 6.61$, $p < 0.001$; at 45 degree orientation, $t_{12} = 5.45$, $p < 0.001$; at 90 degree orientation, $t_{12} = 5.72$, $p < 0.001$; at 135 degree orientation, $t_{12} = 3.31$, $p < 0.01$; at 180 degree orientation, $t_{12} = 5.41$, $p < 0.001$). Thus whichever way the picture was oriented there were more saccades in the axes corresponding to the picture's (horizontal) orientation than elsewhere. Of particular interest is the comparison between pictures oriented normally (0 degrees) and those inverted (180 degrees) and Fig. 5 shows that the distribution of saccades in the four axes is highly similar between the two.

2.2.2. Saccade direction over time

How early does the orientation of the image begin to effect the eye movement direction? To investigate this we looked at the frequency of saccades in each axis as a function of ordinal saccade

number. As previously the first saccade was removed and to remain consistent the following saccades are numbered from 2 to 6. For statistical analysis we pooled data across the picture orientation conditions by rotating each saccade population so that the orientation of the picture was aligned with the horizontal. The four axes can then be thought of as relative to the dominant (horizon) axis. A two-way repeated-measures ANOVA was then possible with axis and ordinal saccade number (from 2nd to 6th) as factors. The data for this analysis is shown in Fig. 6. Across the five saccades the orientation bias is shown by a main effect of axis ($F_{3,36} = 33.38$, $MSE = 13.55$, $p < 0.001$). Pairwise comparisons showed a predominance of 0° saccades (versus 45°, $t_{12} = 6.70$, $p < .001$; versus 90°, $t_{12} = 5.84$, $p < .001$; versus 135°, $t_{12} = 7.96$, $p < .001$). There were no other reliable differences. There was also an interaction showing that the asymmetry in saccade direction varied with ordinal saccade number ($F_{12,144} = 23.55$, $MSE = 9.92$, $p < 0.01$). Planned comparisons showed that there were more saccades made in the 0° axis than those in other directions (averaged across levels) and that this was the case for all saccades except the second (on 2nd saccade, $t_{12} = 1.47$; 3rd, $t_{12} = 5.13$, $p < 0.001$; 4th, $t_{12} = 3.51$, $p < 0.005$; 5th, $t_{12} = 4.52$, $p = 0.001$; 6th, $t_{12} = 5.78$, $p < 0.001$). Thus the pattern in saccade directions emerges relatively early on the second free saccade.

2.2.3. Saccadic amplitude

A further question concerns the amplitude of saccades in each direction. The stimuli used here were square so did not require larger saccades in any particular direction. Are there also asymmetries in saccade length that vary according to picture orientation? Fig. 7 shows that there are, with mean saccade amplitudes showing a similar pattern to that of the saccade directions in Fig. 5. ANOVA

showed a marginally significant main effect of picture orientation ($F_{4,48} = 2.47$, $MSE = 1.54$, $p = 0.057$) and there was no effect of axis ($F_{3,36} = 2.29$, $MSE = 2.26$, $p = .095$). There was an interaction of axis and picture orientation ($F_{12,144} = 8.59$, $MSE = 1.81$, $p < 0.001$). Overall, saccades within the picture's original horizontal axis were larger than those within the other directions, although this effect was not as large as that seen with saccade frequency. As previously, planned comparisons quantified this, and in all cases saccades were longer in this dominant orientation than the mean of the other directions (at 0 degree orientation, $t_{12} = 5.66$, $p < 0.001$; at 45 degree orientation, $t_{12} = 9.56$, $p < 0.001$; at 90 degree orientation, $t_{12} = 3.29$, $p < 0.01$; at 135 degree orientation, $t_{12} = 2.39$, $p < 0.05$; at 180 degree orientation, $t_{12} = 2.38$, $p < 0.05$). Looking at Fig. 7, the trend is less clear when the picture was oriented at 90 degrees, and in this case vertical and horizontal saccades had a similar mean amplitude. The effect is largest in the 45 degree orientation, although this may be because, due to the square frame, an oblique horizon was in fact slightly longer.

2.3. Conclusions from Experiment 1

It is clear from these findings that there is a strong systematic tendency for saccades to occur along the axis of the natural horizon. As this tendency changes when the picture is rotated, an inflexible oculomotor account that fully explains the bias in terms of asymmetries in muscle control can be discounted. This does not contradict findings that horizontal saccades are faster or easier to make (Becker, 1991), and they may be more common in natural behaviour. However, we can suppress horizontal eye movements and make a larger proportion in other directions if the stimulus is oriented in a different way. The fact that the distribution of saccade directions changed on a trial-by-trial basis suggests that horizontal saccades are not just habitual patterns that have been learned and cannot be altered. The square image frame and scattered starting positions means that a horizontal bias is not just due to the laboratory artefacts of a rectangular display (although the bias may act in conjunction with a default strategy of moving towards the centre of an image). There was also an interesting converging result in the data for saccadic amplitude, showing that people make saccades in the horizon axis that are on average longer than those made in other directions.

Two particularly interesting questions emerge from the findings. First, could the pattern of saccades arise as a consequence of the distribution of bottom-up information in the world? It was especially true in the landscape stimuli used here that edges and other visual features tended to be clustered around the horizon. If these features attract eye movements, as suggested by saliency map models, then horizontal saccades are expected purely because that is where conspicuous information is located. This information was also likely to be the most useful to perceive and remember in order to perform the sentence verification task. However, an alternative explanation might rely on the preattentive recognition of gist or layout. By this account people perceive the orientation of a picture very quickly, and they use this with their knowledge of horizons to determine where information is likely to occur. One way of studying this is by asking when the bias for horizontal saccades occurs. It appears that the preference for saccades in the same axis as the horizon occurs early, although not immediately. It has been shown that gist information becomes available very early (Potter, 1976), and that it can influence initial eye movements (Castelhano & Henderson, 2007), although the current data did not find an effect on the first two saccades. It has also been argued that the influence of salience is greatest nearer the start of viewing a picture (Parkhurst et al., 2002). Experiment 2 compares two different types of scenes to see if the presence of a clear horizon can explain the scanning pattern.

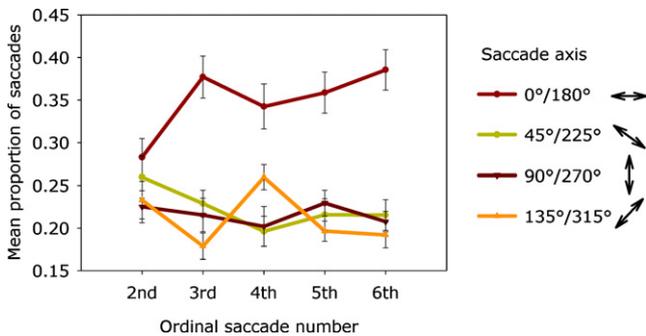


Fig. 6. Mean proportion of saccades in each direction (with standard error bars) as a function of ordinal saccade number. Data is collapsed across orientation conditions, with saccade axis being relative to the original horizon.

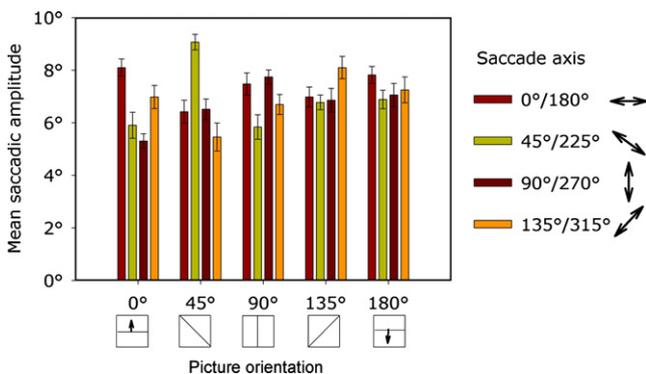


Fig. 7. Mean saccadic amplitude (with error bars indicating the standard error of the mean) as a function of direction and picture orientation.

A second consideration is whether the biases in viewing have a deleterious effect on processing of the scene. With this in mind participants in Experiment 2 were tested later to see how well they had encoded the rotated images. There has been considerable interest in whether the eye movements made when encoding an image affect those made when viewing it again (Althoff & Cohen, 1999; Foulsham & Underwood, 2008; Noton & Stark, 1971). For this reason Experiment 2 explored whether there are any systematic effects of saccade direction at encoding on eye movements made whilst viewing the same images in a memory test. Do prior orientation and the resulting pattern of saccade directions have an impact on later re-viewing and recognition?

3. Experiment 2

In Experiment 2 landscapes were compared with interior scenes. As interiors do not feature a natural horizon, a purely bottom-up account should lead to a reduced bias for saccades that follow the horizon. An explanation that relies only on early recognition of scene layout need not distinguish between landscapes and interiors and so would predict a similar bias in both cases.

In addition, participants were later given an unexpected memory test to see whether recognition memory is affected by prior orientation and the resulting scanning strategy. This task also gives the opportunity of looking at saccade biases at recognition, in addition to those during an encoding task. Will the demands of this task alter the saccade bias?

3.1. Method

3.1.1. Participants

Twelve participants took part who had not been tested previously. All participants were from the University of British Columbia, had normal or corrected-to-normal vision, and completed standard consent procedures. The participants took part for course credit.

3.1.2. Apparatus

The eye tracking set-up was very similar to that used in Experiment 1. The same make and model of eye tracker was used (Eye-link II), and a validation procedure showed that it was recording at a high spatial resolution of at least 0.5° .

3.1.3. Stimuli and design

The same 40 landscapes were used here as in Experiment 1. In addition, the same number of interior photographs was also used. These were colour photographs at the same resolution as the landscapes, and they were prepared in the same way to give a selection of different rotations. Only half of the resulting 80 stimuli were presented with sentences in the first part of the task, whilst the other half were presented at the normal orientation in the recognition phase.

In order to describe the distribution of features in the two types of stimuli they were analysed using two different methods. The Itti and Koch (2000) saliency algorithm was used to generate saliency maps for each stimulus. The source code for this model is freely available (<http://ilab.usc.edu/toolkit/>) and it defines saliency as the contrast between points in the image and their surround, based on three feature channels: intensity, colour and orientation. These features are extracted at various spatial frequencies, the different scales are combined in a centre-surround fashion, and the channels are summed to give an overall saliency map. The saliency map thus represents the places where features stand out from their background. An average saliency map was created, both for the land-

scapes and the interiors, by simply adding the saliency (pixel intensity) values for each point to those for the same spatial location in all other images and then dividing by the number of images. These maps weight each type of feature equally. The resulting average maps are shown in Fig. 8a, and it is apparent that the salient points in the landscapes tend to be near the horizon, whilst saliency is more distributed in interiors. To test this we divided the images into three horizontal bands and compared the average saliency in the central third with that in the top and bottom bands. Looking at the landscape stimuli, the mean saliency was higher in the central third of the picture than in the outer two-thirds (paired t-test across stimuli, $t_{39} = 4.4$, $p < 0.001$). However, in interiors saliency was not reliably greater in the centre than elsewhere ($t_{39} = 1.3$, $p = 0.18$).

If the horizon is a potent cue for saccade orientation, then horizontally oriented edges might be particularly important. Coppola et al. (1998) reported that natural scenes tend to have a predominance of oriented contours at the cardinal, rather than the oblique orientations. To see if this was the case in our stimuli, we analysed the edge content of the landscapes and interiors using the same method as Coppola et al. (1998). Images were converted to greyscale and analysed using a simple Sobel filter and the software MATLAB. Convolution with a three pixel square kernel gave the local gradient at each point in the image for horizontal and vertical directions. Combining these outputs resulted in an estimate of the orientation and magnitude of contours at each point in the image (see Coppola et al., 1998, for further details). Fig. 8b plots the summed magnitude for each orientation: the frequency of each direction weighted by the magnitude of the gradient. The results, which are similar to those in Coppola et al. (1998), show that the distribution of edges is biased towards horizontal and vertical orientations, with relatively fewer oblique contours. It is interesting to note that while both types of picture contain many horizontal lines, interiors are more likely to have vertical edges. Specifically, the summed magnitude of near ($\pm 22.5^\circ$) horizontal orientations was greater than that in near vertical orientations in landscapes ($t_{39} = 5.0$, $p < 0.001$), but in interiors both horizontal and vertical orientations were just as strongly represented ($t_{39} = 1.7$, $p = 0.10$).

Investigating the features of landscapes and interiors in this way allows us to make some general observations about the two types of stimuli. In landscapes, points that stand out from their surround tend to be clustered near the horizontal midline, and there are more horizontal than vertical edges and fewer still oblique contours. In interiors, however, salient features are more distributed. There is no pronounced horizon and there are more vertical edges, presumably because of the presence of walls and the manmade edges of objects. If eye movements are following these features their direction should be differently distributed in the two types of scene.

3.1.4. Procedure

The first part of this experiment was the same as that in Experiment 1 (see Fig. 9). The encoding stimuli were presented in a random order, and each one was followed by a sentence verification test. Landscapes and interiors were intermixed. This order was chosen as opposed to blocking by picture type, so as to avoid participants becoming accustomed to one type of picture and adopting a less dynamic strategy. As previously, participants were instructed to pay attention to the stimuli so that they could verify the sentence as accurately as possible.

Participants were then given an intervening task consisting of viewing fractal images, a task which took approximately 15 min. Once this task was complete, a surprise memory test was presented with the images from the sentence verification test. Participants had no reason to suspect this would occur. Following a drift correct fixation marker in the centre of the screen, each test

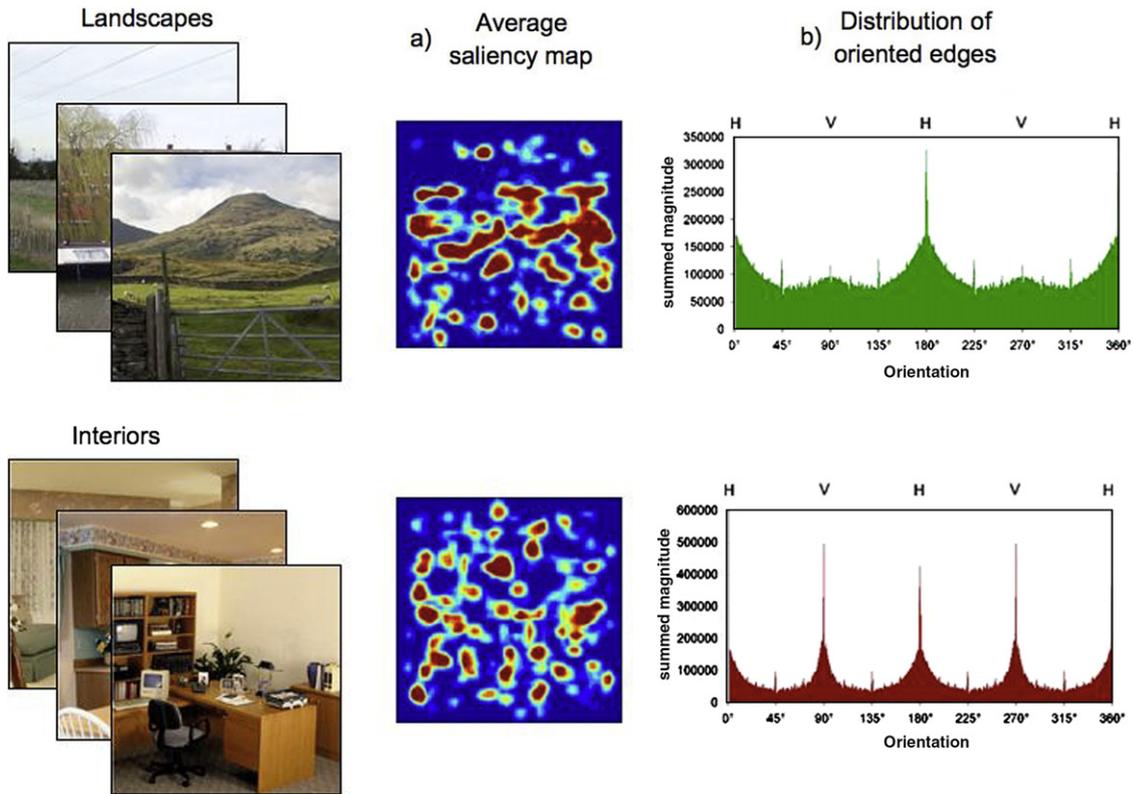


Fig. 8. Analysis of features in the different types of stimuli in Experiment 2. Maps show average saliency (a, with brighter areas indicating higher saliency) and the distribution of oriented contours (b). Plots in (b) show the mean summed magnitude across all the pictures in the set, reflecting the frequency and gradient intensity of contours at each orientation. The full range of orientations is shown, although symmetrical orientations (0/180, 90/270, etc) are equivalent and feature only because the edge filter distinguished between gradients from dark to light and those from light to dark. The images contained an abundance of horizontal (H) and vertical (V) edges.

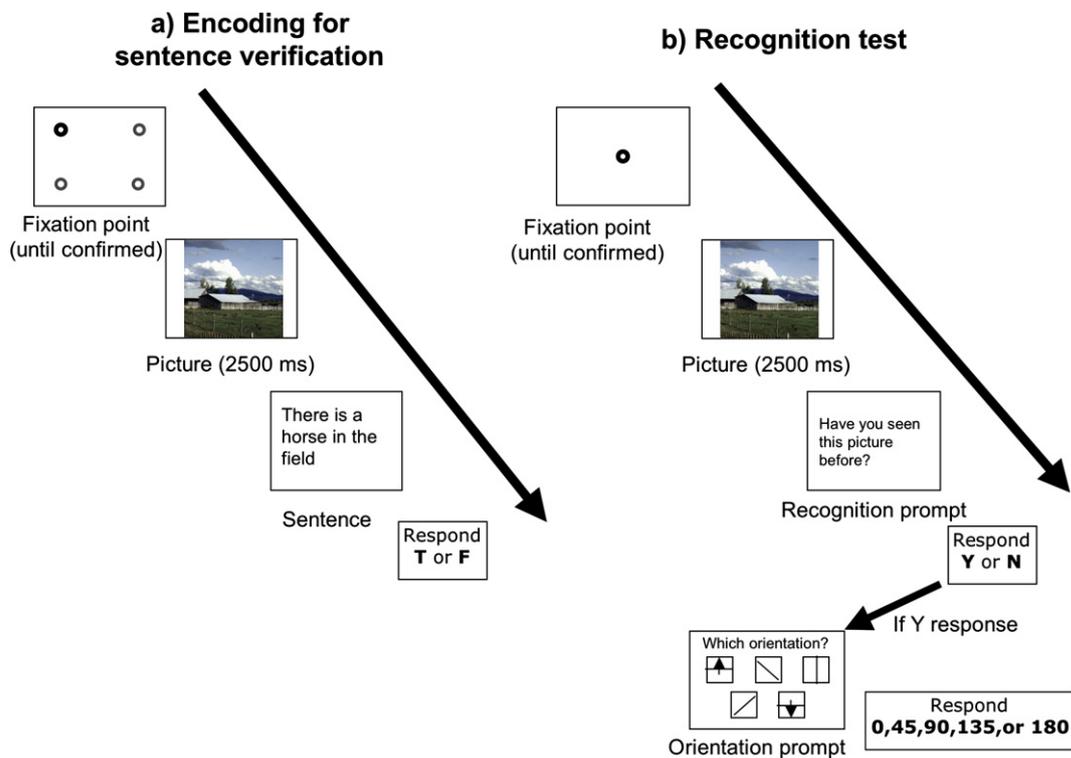


Fig. 9. The procedure in Experiment 2. The first part was a picture-sentence verification test as in Experiment 1 (a). In the second part, following an intervening task, there was a surprise test of recognition for the images and their previous orientation (b).

picture, half of which had been seen previously, was presented for 2500 ms. This duration was the same as used in the encoding

phase, allowing a similar number of saccades to be compared. All trials in this part of the experiment began in the centre of the

screen, unlike picture viewing in the encoding phase. This meant that (in)congruency in starting location between first and second viewing was equal across all orientations. All pictures were presented at the normal, non-rotated orientation. Following picture offset, a response screen asked subjects to respond with one of two keys whether the picture had been seen previously. One of the points of interest in this experiment was whether previous orientation affected responses at recognition, and so it was also useful to know whether participants remembered at which orientation pictures had been presented. With this in mind, if participants responded that they had seen the picture before an additional screen asked them to indicate which way the picture had been rotated. There were five possible responses (0, 45, 90, 135 and 180 degrees) and these were tied to five different keys on the keyboard.

3.2. Results

As in Experiment 1, the proportion of correct answers in the sentence task was relatively high (a mean of 0.771) and these responses will not be considered further. The recognition test results, however, can tell us how well the pictures were processed in the general encoding-for-sentence-verification task. After looking at whether this encoding was affected by orientation the remaining analyses concentrate on whether a horizontal bias exists in both landscapes and interiors, and whether it changes when reviewing pictures for the second time.

3.2.1. Recognition test responses

How accurate were participants at recognising the previously rotated scenes? The mean proportion of hits (the proportion of images seen previously, at any orientation, which were recognised as such) is shown in the top half of Table 1, as a function of previous orientation and scene type. The false alarm rate was generally low ($M = 0.126$, $SE = 0.024$). There was an effect of orientation on recognition ($F_{4,44} = 17.76$, $MSE = 0.0372$, $p < 0.001$), but there was no effect of scene type ($F_{1,11} < 1$) and no interaction ($F_{4,44} = 1.82$, $MSE = 0.0435$, $p = 0.141$). Both landscapes and interiors were recognised equally well and showed a similar pattern of hits across the different rotations. Comparing between orientations, recognition was better for those pictures that had originally been presented normally (ie. at 0 degrees) when compared with all other levels (versus 45 degrees, $t_{11} = 7.70$, $p < 0.001$; versus 90 degrees, $t_{11} = 6.05$, $p < 0.001$; versus 135 degrees, $t_{11} = 3.85$, $p < 0.005$; versus 180 degrees, $t_{11} = 5.85$, $p < 0.001$). There were also some other reliable differences, namely that pictures shown previously at 45 degrees were recognised more poorly than those shown at 90 degrees ($t_{11} = 2.30$, $p < 0.05$) and that the 135 degree condition led to better performance than elsewhere (versus 45 degrees, $t_{11} = 4.65$, $p = 0.001$; versus 90 degrees, $t_{11} = 2.99$, $p < 0.05$; versus 180 degrees, $t_{11} = 4.27$, $p = 0.001$). Chance performance in this task was

50%, and a one-sample t-test compared each condition against this value. Whilst performance in the 0, 90 and 135 degrees rotated condition was better-than-chance that in the 45 and 180 degrees conditions was not (0 degrees, $t_{11} = 15.28$, $p < 0.001$; 45 degrees, $t_{11} < 1$; 90 degrees, $t_{11} = 2.56$, $p < 0.05$; 135 degrees, $t_{11} = 5.14$, $p < 0.001$; 180 degrees, $t_{11} < 1$).

An additional question was whether participants could remember at which rotation a picture had been seen. It is clear from the mean proportion of correct orientation responses (Table 1 bottom half) that performance on this task was poor for all but the pictures shown at 0 degrees. This was shown in an effect of orientation on accuracy ($F_{4,44} = 28.09$, $MSE = 0.057$, $p < 0.001$) such that 0 degrees led to higher accuracy than the other orientations. Orientation recognition accuracy was similar for both picture types ($F_{1,11} < 1$). Comparing between orientations, the 0 degrees condition led to much better performance than elsewhere (versus 45 degrees, $t_{11} = 6.70$, $p < 0.001$; versus 90 degrees, $t_{11} = 6.51$, $p < 0.001$; versus 135 degrees, $t_{11} = 6.98$, $p < 0.001$; versus 180 degrees, $t_{11} = 5.02$, $p < 0.001$). Of the remaining comparisons, only 180 degree (upside-down) presentations were significantly different from the other levels (versus 45 degrees, $t_{11} = 2.63$, $p < 0.05$; versus 90 degrees, $t_{11} = 3.97$, $p < 0.005$; versus 135 degrees, $t_{11} = 2.60$, $p < 0.05$). As there were five possible responses, a chance level of accuracy would be 20%, and only 0 degrees and 180 degrees conditions led to better-than-chance performance (0 degrees, $t_{11} = 7.57$, $p < 0.001$; 45 degrees, $t_{11} = 2.12$, $p = 0.057$; 90 degrees, $t_{11} = 1.34$, $p = 0.207$; 135 degrees, $t_{11} = 1.62$, $p = 0.133$; 180 degrees, $t_{11} = 3.36$, $p < 0.01$). There was also an interaction between scene type and orientation ($F_{4,44} = 3.44$, $MSE = 0.028$, $p = 0.016$) which was driven by a difference between accuracy to 180 degree pictures. Accuracy was similar between landscapes and interiors for all orientations (all $t_{11} < 1$) with the exception of the 180 degrees condition. These upside-down pictures were remembered as such more accurately in landscapes than interiors ($t_{11} = 2.44$, $p = 0.033$) and only above-chance in the former. Thus, although all items at test were presented at their normal orientation, the orientation at which they had been presented previously was important.

3.2.2. Saccade direction at encoding

In the first part of this experiment we were concerned with whether the distribution of saccade directions was affected by the orientation of the picture, and in particular whether this effect was different in landscapes and interiors. As previously the first saccade was excluded as it was assumed that it was generally directed to the centre of the screen and that it was therefore dependent on the (randomised) starting location. The plots in Fig. 10 show the proportion of saccades in each direction as a function of picture type. Fig. 10a and b show the distribution of saccade directions within the four axes, as a function of picture orientation,

Table 1
Mean and standard errors for responses to the memory test in Experiment 2

		Original picture orientation				
		0	45	90	135	180
<i>Proportion of hits</i>						
Landscapes	M	0.833	0.569	0.646	0.729	0.542
	SE	0.036	0.086	0.057	0.065	0.103
Interiors	M	0.958	0.396	0.583	0.785	0.563
	SE	0.028	0.084	0.071	0.061	0.045
<i>Proportion correctly recognised orientation</i>						
Landscapes	M	0.688	0.042	0.083	0.062	0.375
	SE	0.076	0.028	0.047	0.033	0.092
Interiors	M	0.667	0.069	0.083	0.146	0.125
	SE	0.094	0.048	0.036	0.048	0.038

and for landscapes and interiors, respectively. These data are broadly similar to those in Fig. 5. The variability is somewhat larger, probably due to a smaller number of experimental trials.

Two separate repeated-measures ANOVAs were computed to analyse the effects of picture orientation and saccade axis on saccade frequency in landscapes and interiors. As in Experiment 1 it was the interaction that was most important and we predicted that the most frequent direction would change with the orientation of the picture. In landscapes, there was no main effect of picture orientation ($F_{4,44} = 1.29, MSE = 1.16, p = 0.29$) but a reliable effect of saccade axis ($F_{3,33} = 8.69, MSE = 17.42, p < 0.005$). Regardless of picture orientation there were slightly more saccades in the vertical axis than in other directions (versus horizontal saccades, $t_{11} = 3.21$; versus 45°, $t_{11} = 2.88$; versus 135°, $t_{11} = 3.43$, all $p < 0.05$). As predicted, the interaction was significant ($F_{12, 132} = 13.62, MSE = 8.60, p < 0.001$) and planned t -tests compared the frequency of saccades in the axis parallel to the horizon with the mean of the other three directions. This comparison was reliable in all picture conditions except 180 degrees rotation. For instance, when the landscape's horizon was tilted at 45 degrees there were more saccades in the 45/135 degree bin than elsewhere, and this was true for the direction parallel to the horizon in each of the first four orientation conditions (0 degree rotation, $t_{11} = 2.23, p < 0.05$; 45 degrees, $t_{11} = 7.58, p < 0.001$; 90 degrees, $t_{11} = 5.63, p < 0.001$; 135 degrees, $t_{11} = 3.74, p < 0.005$). In upside-down pictures this comparison was not reliable ($t_{11} < 1$). Whilst the trend in the 0 and 180 degree sets is not as clear-cut as in Experiment 1 (due to a large number of vertical saccades) the overall pattern is the same: the dominant saccade direction shifts with the horizon.

Looking now at the interiors (Fig. 10b) the change in distributions is qualitatively similar to that in landscapes. There was not a reliable main effect of picture orientation ($F_{4,44} = 2.26, MSE = 1.15, p = 0.078$) or saccade axis ($F_{3,33} < 1$), but again there was an interaction ($F_{12,132} = 10.95, MSE = 6.91, p < 0.001$). The predicted pattern—that the modal direction in each condition would be that parallel to the picture's original horizontal—was observed in all cases except the 45 degree rotation (where there were more saccades in the 135/315 degree bin). The planned comparisons confirmed this (0 degree rotation, $t_{11} = 6.25, p < 0.001$; 45 degrees, $t_{11} = 1.19, p = 0.258$; 90 degrees, $t_{11} = 4.15, p < 0.005$; 135 degrees, $t_{11} = 2.61, p < 0.05$; 180 degrees, $t_{11} = 2.20, p = 0.05$).

To facilitate a comparison between the two types of scene Fig. 10c and d plot the data from all orientation conditions, rotated so that the saccades counted as horizontal (0/180 degrees) are those made in the same direction as the picture's normal horizontal, and all other directions are measured relative to this. While the expected bias for horizontal saccades is clear in landscapes (Fig. 10c), in interior pictures (Fig. 10d) the shape of the plot is somewhat different. To characterise this statistically, we divided the arc into four symmetrical axes (0, 45, 90 and 135) and looked at the frequency of saccades in each axis relative to the original orientation as a function of picture type. These data were analysed by a two-way repeated-measures ANOVA with axis (0, 45, 90 and 135 degrees from the horizon) and picture type (landscape or interior) as factors. Although picture type was not reliable ($F_{1,11} < 1$), there was a main effect of relative axis ($F_{3,33} = 40.3, MSE = 62.6, p < 0.001$). Across both types of picture there were more saccades in the axis of the original horizon (0 degrees) than elsewhere

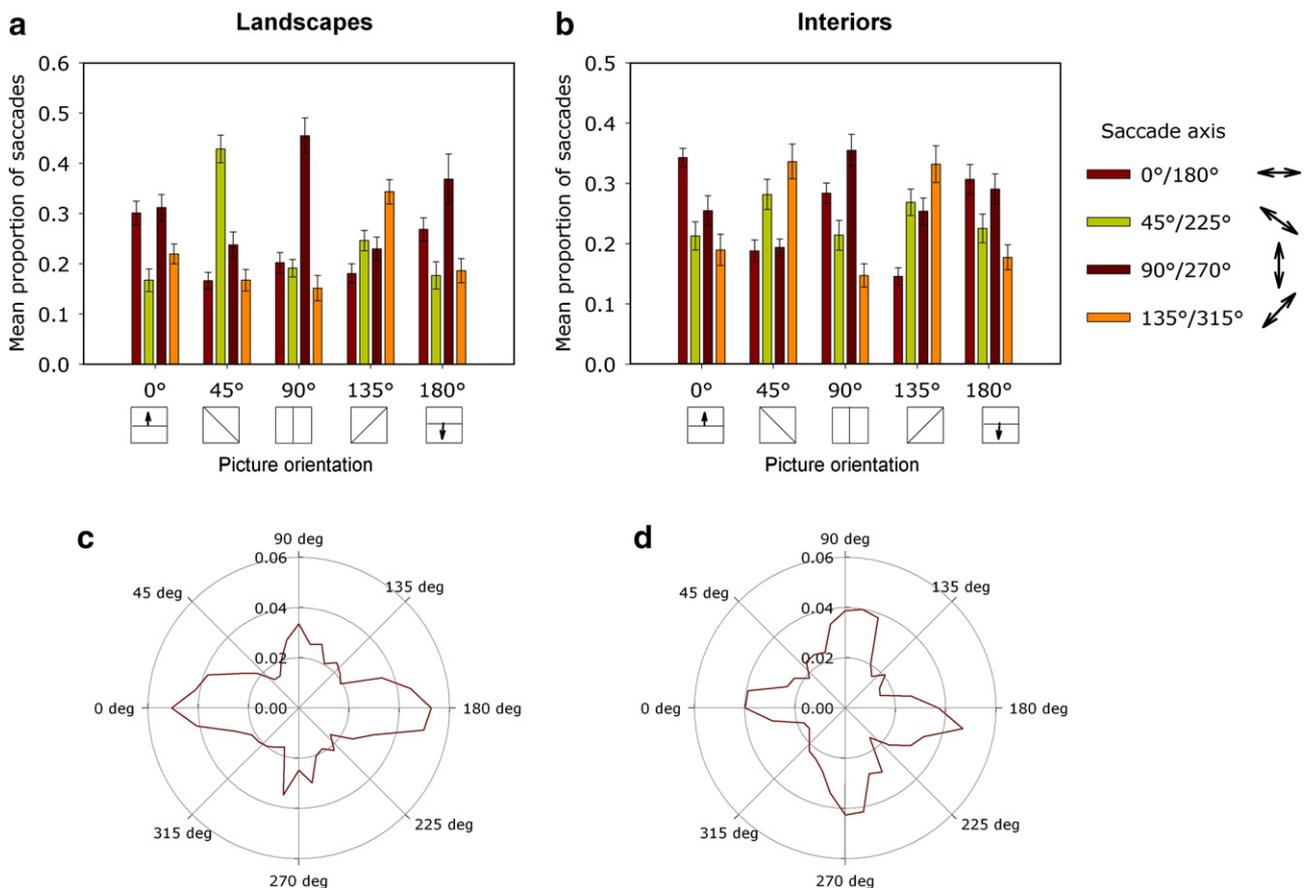


Fig. 10. Mean proportion of saccades in each direction in landscapes (left panels) and interiors (right panels) during the encoding task. Panels (a) and (b) show the mean proportion of saccades in each axis as a function of picture orientation and with error bars indicating plus/minus one standard error across participants. In panels (c) and (d) data are shown for the full range of directions and collapsed across orientation conditions, with saccade axis being relative to the picture's original horizontal.

(planned t-tests; versus 45 degrees, $t_{11} = 16.23$, $p < 0.001$; versus 90 degrees, $t_{11} = 3.33$, $p < 0.01$; versus 135 degrees, $t_{11} = 17.00$, $p < 0.001$). There were also more 90 degree saccades than oblique angle saccades (versus 45 degrees, $t_{11} = 3.58$, $p < 0.005$; versus 135 degrees, $t_{11} = 4.23$, $p = 0.001$), though the frequency of saccades in the two different oblique axes did not differ.

Interestingly there was also an interaction, although this was only marginally significant ($F_{3,33} = 2.86$, $MSE = 38.6$, $p = 0.052$). It is clear from the figures that this interaction is due to a larger proportion of “vertical” saccades (that is, those perpendicular to the original horizontal) whilst viewing interiors. To confirm this, the frequency of 90 degree saccades was compared to the frequency of saccades in other directions, and this was computed separately for each type of picture. In landscapes, there were more vertical than oblique saccades (versus 45 degrees, $t_{11} = 2.24$, $p < 0.05$; versus 135 degrees, $t_{11} = 2.69$, marginal at $p = 0.046$) but, as in Experiment 1, there were far fewer vertical than horizontal eye movements ($t_{11} = 3.57$, $p < 0.005$). In interiors there were also more vertical than oblique saccades (versus 45 degrees, $t_{11} = 4.38$, $p = 0.001$; versus 135 degrees, $t_{11} = 4.39$, $p = 0.001$). Crucially, there was no difference between the frequency of vertical and horizontal saccades in interiors ($t_{11} = 1.82$, $p = 0.095$).

To summarise the saccade directions at encoding, the following conclusions can be drawn. First, the distribution of these directions remains broadly constant in the picture reference frame; as in Experiment 1, and in the majority of picture orientations, the most common saccades were those in the plane of the original horizontal. Second, this effect was subtly different in interiors, which showed a relatively higher frequency of saccades in the picture’s original vertical axis.

3.2.3. Saccade direction at test

How was saccade direction affected by the recognition task? In this part of the experiment all stimuli were presented at the normal orientation, so we would expect the standard bias for horizontal saccades as seen in the zero degree condition in Experiment 1. Saccade direction plots are shown in Fig. 11 for old, previously seen pictures as a function of original orientation (a–e), and for novel, unseen pictures (f). These plots, and the subsequent analyses, include the first saccade as in this part of Experiment 2 viewing began in the centre so direction was unconstrained. Any differences in the eye movements made whilst viewing pictures at test cannot be because of their orientation and must be due to recognising or reprocessing them in some way. There are some noticeable differences. Although most plots show a horizontal bias this is not as clear as in previous analyses and in those pictures previously shown at 180 degrees there is large vertical bias. However, there appears to be no systematic effect of prior orientation; it is not the case that orientation at encoding carries over to effect viewing when presented the correct way round at test. To ascertain any non-specific effect of previous exposure, the saccade distribution associated with old pictures shown in the encoding phase was compared to that from novel items in the recognition test. Both types of items were equally likely to be interiors or landscapes.

As previously the data were organised into the four major axes relative to the horizon and comparisons were performed on the frequency of saccades in each axis (Fig. 12). The horizontal bias resulted in an effect of axis on saccade frequency ($F_{3,33} = 5.04$, $MSE = 1164.5$, $p < 0.01$). Across old and new items more saccades were made in the horizontal than in the oblique axes (planned t-tests; versus 45 degrees, $t_{11} = 4.79$, $p = 0.001$; versus 135 degrees,

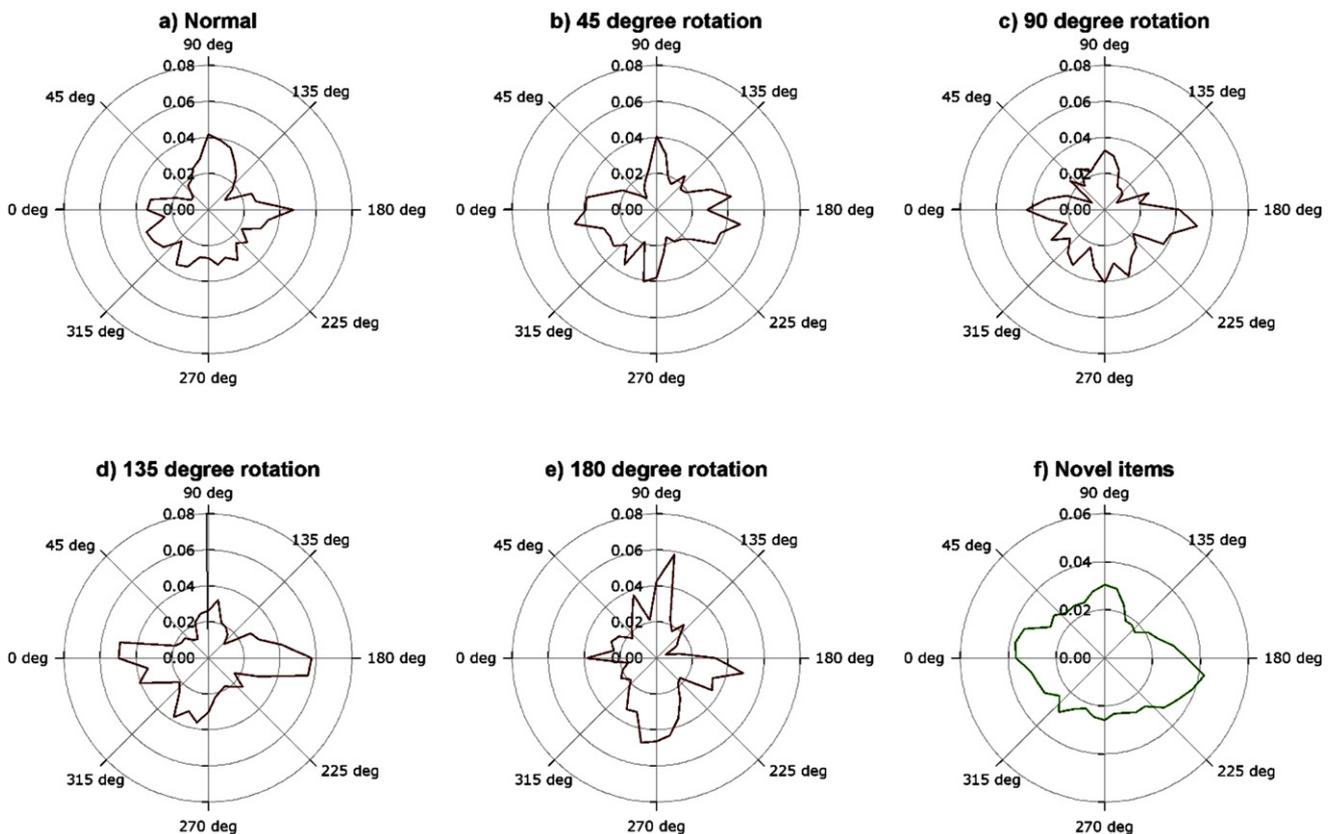


Fig. 11. Saccade direction distributions for pictures viewed during the recognition test. All pictures were shown at the normal orientation, but some had previously been shown at various rotations (old pictures; a–e). New pictures had not been seen previously (f).

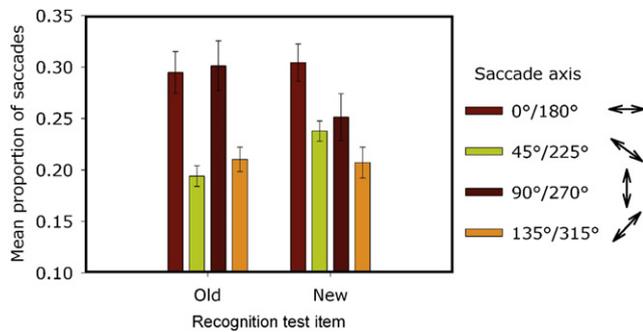


Fig. 12. Mean proportion of saccades in the four axes, for old and new items when viewed in the recognition test.

$t_{11} = 3.62$, $p < 0.005$). There were also more vertical than 135 degree saccades ($t_{11} = 2.34$, $p < 0.05$). However, there was no difference between the frequency of horizontal and vertical saccades, and no other comparisons reached significance. There was no main effect of exposure on frequency ($F_{1,11} < 1$), but there was an interaction indicating that the distribution of saccades across different directions varied with exposure ($F_{3,33} = 19.52$, $MSE = 52.0$, $p < 0.001$). It is clear from Fig. 12 that the difference between the saccades made in old and new pictures is that there are more vertical saccades in the former. Looking at paired comparisons between axes in old and new pictures the only reliable differences are between vertical and oblique saccades in old, previously seen pictures; in these items there were more 90 degree saccades (than 45 degrees, $t_{11} = 3.27$, $p < 0.01$; than 135 degrees, $t_{11} = 3.07$, $p < 0.05$). No other comparisons were significant; although the pattern in new pictures resembles that seen elsewhere none of the levels were reliably different. As a precaution we repeated this analysis excluding the first saccade, so as to bring the results in line with those from the encoding phase. The pattern was unchanged and there were reliably more vertical saccades in old trials than new trials ($t_{11} = 4.10$, $p < 0.005$).

Why did participants make more vertical saccades in previously seen images? If these saccades were associated with making a manual response (due to looking down at the keypad for example) then they would also be seen in new trials, but they were not. However, it seems possible that they may have come about when observers recognised an image, perhaps because of the impending orientation recognition test. If the predominance of vertical saccades was due to this response then we would also expect it in cases where pictures were incorrectly recognised (false alarms). To test this, the mean proportion of all saccades that were in the vertical axis was compared between old trials that were correctly responded to (hits) and those that led to an error (false alarms; FAs). There was a significant difference ($t_{11} = 3.56$, $p < 0.005$). A greater proportion of saccades were made vertically when a hit occurred ($M = 0.289$, $SE = 0.024$) than when an FA was made ($M = 0.228$, $SE = 0.015$).

4. General discussion

Experiments 1 and 2 measured several thousand saccades across people and pictures and found a robust bias for saccades in the scene horizontal. We are not aware of any other published results that explore saccade direction in natural images. When square pictures were presented at the normal orientation there were many more saccades made in the leftwards and rightwards directions than in the vertical or oblique directions. This supports previous reports, dating back to Brandt (1945). When the image was rotated, this pattern changed so that the most frequent direction for a saccade was parallel to the original orientation of the pic-

ture, suggesting a sensitivity to the content of the image that can override the effects of rotation. Saccadic amplitude also varied systematically with direction (Experiment 1); on average larger saccades were made in the axis of the horizon.

4.1. Explaining biases in saccade direction

With regard to the possible explanations for a horizontal bias from the introduction, several conclusions can be drawn. First, it is unlikely that the bias can be fully accounted for by laboratory artefacts. The predominance of horizontal saccades is not caused solely by rectangular stimuli. It would be interesting to look at saccades in pictures that are higher than they are wide (e.g., portraits), or in images with a circular frame, though the current findings suggest that horizontal saccades would still be more common. Even with an earth-fixed frame of reference (a non-square monitor and the room behind) the degree to which the dominant eye movement direction rotated was highly systematic. Given the relatively difficult task there is no reason to think that participants were paying attention to cues outside the image, but the results suggest that if these cues were masked, removing any stable frame of reference, the effects of rotation would be even more pronounced. Outside the confines of a monitor, a fixed head position and a two-dimensional, discrete image, biases in saccade direction might change. Alternatively, in combination with head and trunk movements that tend to move horizontally, a horizontal saccade bias might be increased. This is a strong impetus for further research in more realistic settings. Although starting position was varied, the vast majority of trials began with a saccade into the centre of the image and thus it is possible that this central bias contributes to a predominance of horizontal saccades.

Both of our studies demonstrated that a strong oculomotor explanation should also be discounted. People can easily make saccades in the vertical and oblique axes if this is the way in which the picture is oriented. Similarly, although our environment may tend to be laid out horizontally, and our experience with this and other fixed situations (such as reading) may affect our propensity to move in any particular direction, we can alter this within one or two self-initiated eye movements on a scene. There were no explicit instructions to alter scanning behaviour, and in Experiment 2 later memory for this orientation was poor, suggesting that it was not explicitly encoded.

So why do people move their eyes horizontally (or in parallel with the horizon of a rotated picture)? Two remaining possibilities from the introduction can be considered. The distribution of image features, such as edges or salient points, might guide attention in a bottom-up fashion. If these features were oriented or clustered along a horizontal axis of the picture then their distribution would change as the picture was rotated and this might account for the predominance of horizontal saccades. It may also have been that the task (verifying a sentence concerning objects and other details) biased participants to look at features that were distributed in this way. It has been shown elsewhere that fixation locations are dependent on expectations of target location in natural scenes (Neider & Zelinsky, 2006); modifying the task might change the importance of the horizon and the frequency of horizontal saccades. In Experiment 2 interior photographs were contrasted with landscapes, to see whether the more distributed features in the former led to less of a horizontal bias. This was indeed the case, with the ratio of horizontal to vertical saccades changing from around 2:1 in landscapes to approach 1:1 in interiors. This is the pattern that would be predicted based on the analysis of edge content in the two types of picture, which showed relatively more vertical edges in interiors. Coppola et al. (1998) reported that cardinal orientations were over-represented in the natural world in comparison to oblique contours, and they linked this to the sensitivity

shown by humans and other animals to different orientations. It is possible, therefore, that the reason people make fewer oblique saccades is due to the edges in the environment. Experiment 2 was not designed to manipulate edge content, but the results suggest an interesting link between contour magnitude and saccade direction. The landscapes used had more horizontal than vertical edges, and in most cases they led to a predominance of horizontal saccades. In interiors, where vertical edges were also prevalent, there was less of a horizontal bias and vertical saccades were common. This supports a bottom-up, image-characteristics account of saccade direction control. At its strongest this account suggests that each subsequent saccade is targeted at the next most salient image feature (as in saliency map models), and that due to the distribution of saliency in landscapes these saccades tend to be moving between horizontally aligned points, whereas in interiors they are more spread out.

A different explanation might posit the early recognition of scene type and layout. This fast perception of “gist” has been widely reported (Biederman et al., 1974; Potter, 1976; Potter, Staub, & O’Connor, 2004) and can be predicted based on the analysis of low spatial frequency image statistics (Torralba, 2003). For example, if participants gleaned enough information from the first fixation to realise the type (landscape or interior) and orientation of the image, this information might activate stored representations of where objects and interesting features occur in this class of image. This information could then guide the eyes. Without a clear definition of the features that make up gist it is difficult to conclusively test their effect on eye movement direction. Two other results from the present research are relevant for distinguishing between bottom-up control and early gist acquisition. First, when saccade direction was inspected as a function of time since picture onset, the disparity between the frequency of horizontal saccades and those made in other directions increased from the second saccade (where there was no significant difference) to the third saccade and beyond. Parkhurst et al. (2002) suggest that the influence of visual saliency declines over time, and if points on a saliency map are selected and then inhibited based on a winner-take-all system then to some degree later fixations should be made to less salient regions. Therefore, if horizontal saccades were due to the distribution of saliency, we would expect the bias to be greatest on the first free saccade and to decrease over multiple saccades, but this was not the case. It seems likely that gist and layout information build up over several fixations, so this could explain the increase in saccade bias following the first fixation on the picture. A second point of interest is the 180 degree rotation condition, where the image was completely inverted. Some researchers have reported that inverting a scene disrupts the acquisition of gist, so one might expect a different pattern of eye movements, even though the distribution of features relative to the horizontal axis will be the same as in the normally oriented picture. In fact, saccades in the 180 degree condition showed an equally strong horizontal bias.

There were additional noteworthy findings. In Experiment 1, saccades in the axis of the horizon also had, on average, a larger amplitude. Due to the square dimensions of the image, oblique directions had a longer plane than cardinal directions in which to move, and so we should be cautious about the effect this may have had on some rotations. However, even in the 0, 90 and 180 degree conditions larger saccades were made in the plane of the horizon, despite the presence of a longer oblique axis running in a different direction. Saccadic amplitude could be taken as an indication of the degree of peripheral processing; greater processing of peripheral regions allow more distant saccade targets to be selected, leading to larger saccades. If this is the case then it suggests an asymmetry in the way processing of information away from fixation takes place, perhaps with covert attention spreading further along the perceived horizon than in the direction perpendicular to it.

4.2. Saccade direction and recognition

Experiment 2 looked at participants’ later, incidental memory for pictures that had been rotated. Subsequent recognition was much better for pictures that had been shown the correct way up. Thus, even though people could adjust their scanning patterns in line with the rotation of the picture encoding may have been better when not rotated. Recognition of a picture’s previous orientation was rather poor for almost all the conditions. This suggests that although the pictures were encoded into memory (as shown by above-chance old/new recognition), orientation was not remembered. Thus if an oriented representation of the scene is formed early in viewing to guide saccade direction, this representation is not maintained and available for later retrieval. The memory results were consistent for both landscapes and interiors. While these results are interesting they should be treated with some caution. Investigating memory was not one of the main aims of the work presented here, and there are several issues that might have affected the results. First, the cropping of rotated images meant that some of the correctly oriented test images contained slightly different information from when they were presented initially. The differences were small and peripheral, but it is possible that they may have led to more errors for rotated pictures. Second, further research is needed to unravel the effects of orientation at encoding and at test. In Experiment 2, the recognition advantage for non-rotated images might be due either to better encoding, or to the congruency between encoding and test orientation. The starting fixation location used at encoding and test, and the congruency between these, was controlled in the present study but this might also have an effect on memory. We are pursuing these memory effects elsewhere, but the remainder of the discussion will concentrate on the eye movement data.

How did re-exposure to pictures in the test phase of a memory test effect scanning? Scanpath theory suggests that eye movement sequences are stored along with the features of an image, and recapitulated when that image is seen again (Noton & Stark, 1971). If this were the case then we might expect carry-over effects of the dominant scanning direction on the eye movements made at test. However, there was no systematic effect of prior orientation on the direction of saccades made when pictures were viewed for the second time. Although the present study did not look at scanpaths (chains of multiple, sequential saccades) this finding, along with others in the literature, suggests that the predictions of scanpath theory are too strong (Foulsham & Underwood, 2008; see also Henderson, 2003). Recent work by Althoff and Cohen (1999) has examined a “reprocessing effect” for pictures of scenes and faces, whereby there are differences in the eye movements made when a stimulus has been seen before compared with when it is novel, even in the absence of explicit recognition. The analysis of saccade direction in old versus new pictures at test gives an indication that a similar effect is happening here. Old pictures could be distinguished from those that had not been seen on the basis of the proportion of vertical saccades. Correct recognition of old pictures was associated with more vertical saccades, and this was not found in trials leading to a false alarm. It is unclear what caused this difference, although it may have been related to the requirement to make further orientation recognition responses. This observation also suggests that a top-down, task-driven factor can modify and reduce the tendency to make horizontal saccades.

4.3. Implications for models of eye guidance

What are the implications of the saccade pattern discussed for models that aim to predict eye movements in natural scenes? Bottom-up models, such as the saliency map model of Itti and Koch (2000), have been adapted to take into account the space-variant

sampling of the retina (Vincent et al., 2007). However, locations situated in the horizontal, vertical and oblique directions are equally likely to become saccade targets in this model. The data presented here suggest that this is not the case, and that across a range of images horizontal saccades are more likely, and that oblique movements are rare. Saliency-based models might produce better predictions if they incorporated a saccade generator that took into account what is known about saccade dynamics and direction distributions. Of course, if the distribution of salient features is asymmetric then this pattern might emerge naturally, and the comparison with interior scenes, which have more distributed features and show less of a horizontal bias, supports this account. On the other hand several findings in the present research suggest that this bias varies according to gist and top-down goals—it forms after several fixations and is affected by previous viewings. Perhaps a more realistic framework can therefore be provided by the contextual guidance model of Torralba et al. (2006). In this model local saliency is computed in parallel with the extraction of global features which can provide gist and layout information. This information provides contextual priors to bias the saliency map to certain locations, and if it included rough knowledge of scene orientation and the likely location of important features this would produce saccade asymmetries. The model was designed to predict real world visual search for a known target object, so it would need to be generalised to account for the encoding task used here.

5. Conclusion

In sum, this research provides a novel way of exploring eye movements in natural scenes by looking at the distribution of saccade directions. The bias for saccades parallel to the horizon is robust even in square photographs. When these photographs are rotated participants are able to adjust the modal direction of their saccades quickly. Models of eye movements need to be able to predict this observation based on either bottom-up feature distributions or, as seems likely, in concert with higher level knowledge of scene layout.

Acknowledgments

T.F. was supported by a Universitas 21 travel scholarship to work in A.K.'s laboratory at UBC. We are also grateful to the EPSRC for project award EP/E006329/1 to G.U., and to two anonymous reviewers for their very helpful comments.

References

- Abed, F. (1991). Cultural influences on visual scanning patterns. *Journal of Cross-Cultural Psychology*, 22(4), 525–534.
- Althoff, R. R., & Cohen, N. J. (1999). Eye-movement-based memory effect: A reprocessing effect in face perception. *Journal of Experimental Psychology—Learning Memory and Cognition*, 25(4), 997–1010.
- Becker, W. (1991). Saccades. In R. H. S. Carpenter (Ed.), *Eye movements* (pp. 95–137). London: Macmillan Press.
- Becker, W., & Jurgens, R. (1990). Human oblique saccades—Quantitative analysis of the relation between horizontal and vertical components. *Vision Research*, 30(6), 893–920.
- Biederman, I., Rabinowitz, J. C., Glass, A., & Stacy, E. (1974). On the information extracted from a glance at a scene. *Journal of Experimental Psychology—General*, 103, 560–597.
- Brandt, H. F. (1945). *The psychology of seeing*. The Philosophical Library: Springer.
- Buswell, G. T. (1935). *How people look at pictures: A study of the psychology of perception in art*. Chicago: University of Chicago Press.
- Castelano, M. S., & Henderson, J. M. (2007). Initial scene representations facilitate eye movement guidance in visual search. *Journal of Experimental Psychology—Human Perception and Performance*, 33(4), 753–763.
- Collewijn, H., Erkelens, C. J., & Steinman, R. M. (1988). Binocular coordination of human vertical saccadic eye-movements. *Journal of Physiology—London*, 404, 183–197.
- 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.
- Crundall, D. E., & Underwood, G. (1998). Effects of experience and processing demands on visual information acquisition in drivers. *Ergonomics*, 41(4), 448–458.
- Foulsham, T., & Underwood, G. (2008). What can saliency models predict about eye movements? Spatial and sequential aspects of fixations during encoding and recognition. *Journal of Vision*, 8(62), 1–17.
- Gilchrist, I. D., & Harvey, M. (2006). Evidence for a systematic component within scan paths in visual search. *Visual Cognition*, 14(4–8), 704–715.
- Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, 7(11), 498–504.
- Henderson, J. M., Brockmole, J. R., Castelano, M. S., & Mack, M. L. (2007). Visual saliency does not account for eye movements during visual search in real-world scenes. In R. van Gompel, M. Fischer, W. Murray, & R. W. Hill (Eds.), *Eye movements: A window on mind and brain* (pp. 537–562). Amsterdam: Elsevier.
- Hughes, A. (1977). The topography of vision in mammals of contrasting lifestyles. In F. Crescitelli (Ed.), *Handbook of sensory physiology* (Vol. VII, pp. 614–642). Berlin: Springer.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40(10–12), 1489–1506.
- Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3), 194–203.
- Kirchner, H., & Thorpe, S. J. (2006). Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision Research*, 46(11), 1762–1776.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4, 219–227.
- Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology—Human Perception and Performance*, 4(4), 565–572.
- Mackworth, N. H., & Morandi, A. J. (1967). The gaze selects informative details within pictures. *Perception & Psychophysics*, 2, 547–552.
- Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention. *Vision Research*, 45(2), 205–231.
- Neider, M. B., & Zelinsky, G. J. (2006). Scene context guides eye movements during visual search. *Vision Research*, 46(5), 614–621.
- Noton, D., & Stark, L. (1971). Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Research*, 11(9), 929.
- 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. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology—Learning Memory and Cognition*, 2, 509–522.
- 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.
- Raj, R., Geisler, W. S., Frazor, R. A., & Bovik, A. C. (2005). Contrast statistics for foveated visual systems: Fixation selection by minimizing contrast entropy. *Journal of the Optical Society of America A—Optics Image Science and Vision*, 22(10), 2039–2049.
- Rao, R. P. N., Zelinsky, G. J., Hayhoe, M., & Ballard, D. (2002). Eye movements in iconic visual search. *Vision Research*(42), 1447–1463.
- Sanocki, T. (2003). Representation and perception of scenic layout. *Cognitive Psychology*, 47, 43–86.
- Reinagel, P., & Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Network—Computation in Neural Systems*, 10(4), 341–350.
- Smeets, J. B. J., Hayhoe, M. M., & Ballard, D. H. (1996). Goal-directed arm movements change eye-head coordination. *Experimental Brain Research*, 109(3), 434–440.
- Tatler, B. W., Baddeley, R. J., & Gilchrist, I. D. (2005). Visual correlates of fixation selection: Effects of scale and time. *Vision Research*, 45(5), 643–659.
- 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.
- Torralba, A., Oliva, A., Castelano, M. S., & Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological Review*, 113(4), 766–786.
- Underwood, G., Chapman, P., Bowden, K., & Crundall, D. (2002). Visual search while driving: Skill and awareness during inspection of the scene. *Transportation Research, Part F*(5), 87–97.
- Underwood, G., & Foulsham, T. (2006). Visual saliency and semantic incongruity influence eye movements when inspecting pictures. *Quarterly Journal of Experimental Psychology*, 59(11), 1931–1949.
- Vincent, B. T., Troscianko, T., & Gilchrist, I. D. (2007). Investigating a space-variant weighted saliency account of visual selection. *Vision Research*, 47(13), 1809–1820.
- Yarbus, A. L. (1967). *Eye movements and vision*. New York: Plenum.
- Yee, R. D., Schiller, V. L., Lim, V., Baloh, F. G., Baloh, R. W., & Honrubia, V. (1985). Velocities of vertical saccades with different eye-movement recording methods. *Investigative Ophthalmology & Visual Science*, 26(7), 938–944.