



Research

Cite this article: Zamboni E, Ledgeway T, McGraw PV, Schluppeck D. 2016 Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proc. R. Soc. B* **283**: 20160263. <http://dx.doi.org/10.1098/rsob.2016.0263>

Received: 4 February 2016
Accepted: 24 May 2016

Subject Areas:

cognition, neuroscience

Keywords:

vision, sensorimotor bias, perceptual decisions, psychophysics, motion

Author for correspondence:

Denis Schluppeck
e-mail: denis.schluppeck@nottingham.ac.uk

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rsob.2016.0263> or via <http://rsob.royalsocietypublishing.org>.

Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception

Elisa Zamboni, Timothy Ledgeway, Paul V. McGraw and Denis Schluppeck

School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, UK

TL, 0000-0002-2000-2420; PVM, 0000-0001-9484-8560; DS, 0000-0002-0634-7713

Visual perception is strongly influenced by contextual information. A good example is reference repulsion, where subjective reports about the direction of motion of a stimulus are significantly biased by the presence of an explicit reference. These *perceptual biases* could arise early, during sensory encoding, or alternatively, they may reflect *decision-related processes* occurring relatively late in the task sequence. To separate these two competing possibilities, we asked (human) subjects to perform a fine *motion-discrimination* task and then estimate the direction of motion in the presence or absence of an oriented reference line. When subjects performed the discrimination task with the reference, but subsequently estimated motion direction in its absence, direction estimates were unbiased. However, when subjects viewed the same stimuli but performed the estimation task only, with the orientation of the reference line jittered on every trial, the directions estimated by subjects were biased and yoked to the orientation of the shifted reference line. These results show that judgements made relative to a reference are subject to *late, decision-related biases*. A model in which information about motion is integrated with that of an explicit reference cue, resulting in a late, decision-related re-weighting of the sensory representation, can account for these results.

1. Introduction

Traditionally, perceptual errors (or misperceptions) have been explained as either arising from limitations in the sensory apparatus that transform stimuli into a neural representation or from higher-level cognitive biases reflecting our prior knowledge about the world. It is also known that the *responses* obtained from subjects about their perceptions can be influenced by factors such as the history of preceding responses, prior expectations and reward contingencies.

Several previous studies have demonstrated systematic biases away from a reference when subjects estimate the direction of motion of a stimulus, a phenomenon termed *reference repulsion*. Other types of repulsive effects can also be observed in judgements between two transparently moving stimuli [1]. While these effects have been framed in terms of inhibitory interactions between directionally selective neurons as a low-level, *sensory phenomenon* (see [1–4]), it has also been suggested they could result from higher-level cognitive effects [5].

In a study by Jazayeri & Movshon [6], subjects performed a fine-discrimination task on a moving random-dot stimulus in the presence of an oriented reference line that served as a discrimination boundary. They had to judge whether the motion was in a direction more clockwise (CW) or counterclockwise (CCW) than the reference. Subjects received feedback on whether their response was correct. In 30% of the trials, this feedback was withheld and they had to estimate the perceived direction on that trial using a manual matching task. Estimates of motion direction were consistently biased away from the reference following this fine discrimination task.

Interestingly, the reported bias in direction depends on the reliability of the motion signal: the lower the reliability (as determined by the coherence—the proportion of local stimulus elements that move in the same direction) of the motion stimulus, the larger the bias. A model that re-weights sensory information away from the reference direction could quantitatively account for these biases. The

weighting function in this model captures the shape of a mechanism tuned optimally for the fine-discrimination task [7,8], suggesting the re-weighting of information during the fine discrimination task as a potential cause for the observed perceptual biases. However, the biases in the reported motion directions could also have occurred later in the trial [5]. The oriented reference was present throughout the trial, from the time of stimulus presentation all the way to the response interval, and it is therefore not possible to determine at what point during the trial these biases actually arose.

Thus, the well-known biases observed in fine-discrimination tasks could arise at the time of stimulus presentation (during initial encoding by the sensory apparatus), during subsequent decoding of that information for processing by higher-level visual areas, or during the decision process and formation of the subject's response. To address this issue, we designed an experiment in which we manipulated the properties of the orientated reference line during the estimation phase. In the first experiment, following a fine direction-discrimination task, we removed the oriented reference completely during a subsequent direction estimation task. In the second experiment, subjects performed an estimation task following passive viewing of a moving stimulus—however, in this case, we systematically manipulated the orientation of the reference just before the onset of the estimation task. If the observed biases are owing to a process that operates at the time of stimulus presentation, encoding by the sensory apparatus, or even decoding from the activity in early visual processing, these manipulations should have no effect. However, we found that manipulating information about the oriented reference relatively late in the trial (just before subjects give their estimate) affects these biases, suggesting that the biases arise at a later stage than previously thought.

2. Material and methods

(a) Participants

Seven subjects participated in the study (six female; age 23.4 ± 2.8 years, mean \pm s.d.). Of this group, five completed each of the two experiments. One subject was one of the authors, and the remaining subjects were naive to the purpose of the experiments. All had normal or corrected-to-normal vision and gave written, informed consent in accordance with regulations by the Ethics Committee of the School of Psychology, University of Nottingham.

(b) Apparatus and stimuli

Visual stimuli were presented on a CRT monitor (Mitsubishi Diamond Plus 73) at a resolution of 1024×768 pixels, a refresh rate of 85 Hz and at a viewing distance of 57 cm. Stimuli were generated using the MGL toolbox (available from <http://gru.stanford.edu/mgl>) in MATLAB (The MathWorks, Natick, MA) on Apple Mac OS X.

A white fixation spot (diameter 0.5°) was displayed at the centre of the screen and remained throughout each trial. Motion stimuli were random-dot-kinematograms (RDKs) composed of a field of moving 'white' (147 cd m^{-2}) dots on a 'grey' (38 cd m^{-2}) background. Each dot was 0.12° in diameter, moving within a circular aperture 5° in diameter. RDKs were computed at different coherence levels (4%, 7%, 13% and 25%), selected randomly at the beginning of each trial) such that *signals* dots were moving at a speed of 4° s^{-1} in a predefined direction. The remaining *noise* dots were randomly re-plotted within the circular aperture.

Patterns were updated on every alternate frame (42.5 Hz). The average dot density was $40 \text{ dots deg}^{-2} \text{ s}^{-1}$. A 'black' (0.5 cd m^{-2}) segment reference line (length, 0.5° ; width, 0.15°), starting 3.5° from the fixation point, served as a discrimination boundary (figure 1*a,b*). Stimulus parameters were chosen to closely match those used in previous experiments [6].

(i) Procedure

Each trial started with the presentation of the fixation spot at the centre of the screen for 1 s, followed by the onset of the discrimination boundary (black line segment) at a randomly chosen orientation. After 0.5 s, an RDK with a particular coherence level was displayed for 1 s, moving in a direction randomly selected to be within $\pm 24.5^\circ$ of the boundary. As the RDK disappeared, the fixation spot turned orange, cueing the subjects to perform a fine-discrimination task on the direction of motion. Subjects had to press one of two keys on a standard keyboard, indicating whether the dots were moving CW or CCW with respect to the discrimination boundary. A maximum of 2 s were allowed for the response. On 70% of trials, as soon as the key was pressed, the subject received feedback on the trial (the fixation spot turned green for a correct response, red otherwise). On the remaining 30% of trials, feedback was withheld and a black circular ring (3.6° diameter, centred on the fixation spot) was presented immediately after the discrimination task, cueing the subjects to manually estimate the direction in which they saw the dots moving by extending a line 'tethered' to the fixation spot. As soon as the line hit the ring, the task was completed. A maximum of 10 s were allowed for these estimation responses (figure 1*c*).

Subjects performed a minimum of 5000 trials over 100 experimental runs. Prior to engaging in the experimental trials, subjects underwent extensive practice to familiarize themselves with the psychophysical tasks. First, each participant performed 400 trials (over eight runs) at suprathreshold coherence levels (200 trials at a coherence of 50%, then a further 200 at 25%) with a discrimination boundary that was always at vertical. Next, in another 400 practice trials, at the same coherence levels, the discrimination boundary could take on any value between 0° and 360° . Finally, in another 400 practice trials, the coherence levels were adjusted to those used in the experiment.

(ii) Experiment 1: reference present or absent at estimation

Subjects performed the task in two situations that were similar, but different in one important aspect. In one set of trials ($n = 2500$), the reference line was present throughout the duration of each trial. In the other set of trials ($n = 2500$), the discrimination boundary was displayed only for the duration of the discrimination task. Thus, the black line segment was not presented at the time of the estimation task (figure 1*d*).

(iii) Experiment 2: manipulating consistency of the reference information

Subjects were presented with the same stimuli as in the first experiment with the same timing and procedure, with the following modifications: (i) subjects were instructed to passively view the stimuli and not to perform the fine-discrimination task (hence also, no feedback was given), (ii) subjects manually estimated the direction of motion in all the trials and *crucially* (iii) unbeknownst to the subjects the orientation of the reference line was manipulated just before the onset of the estimation task. In particular, the reference could either remain consistent with that at the start of the trial, or it could vary by $\pm 6^\circ$ from the initial orientation. In order to decrease the likelihood of subjects perceiving this change in orientation, a brief (100 ms) blank interval between the presentation of the RDKs and that of the circular ring was added. Thus, for this experiment, the task consisted

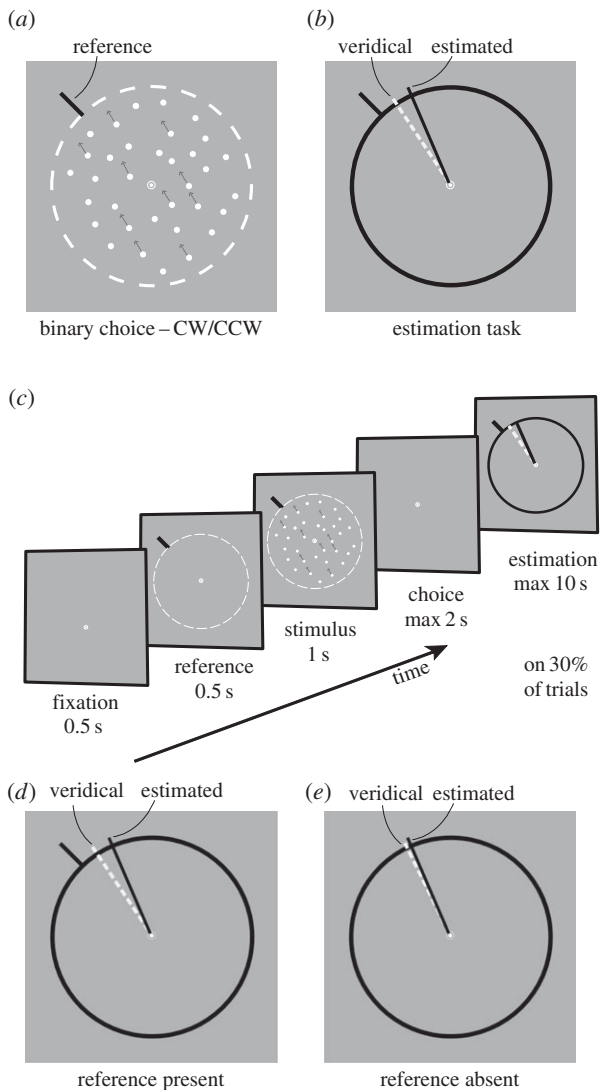


Figure 1. Stimuli and task. (a) Subjects were presented with a random-dot-kinematogram (RDK) in a circular aperture (dashed line for illustration purpose only). They had to indicate whether the stimulus direction was more CW or CCW than a reference line in a two-alternative-forced-choice task. (b) On a proportion of trials (30%), subjects also reproduced the perceived direction of motion in an adjustment task by extending a line (shown in black) from the fixation point towards a circle delineating the stimulus aperture. For each trial, we recorded the estimated and veridical directions of motion (illustrated by dashed line). (c) Timing of stimuli. A reference (orientation randomly chosen for each trial) was presented for 0.5 s. The RDK stimulus was shown for 1 s (coherence either 4%, 7%, 13% or 25%; direction $\pm 24.5^\circ$ from the reference). Signal dots moved with a speed of 4° s^{-1} . Noise dots were replotted at random locations within the stimulus aperture on every other screen refresh. After another 0.5 s delay, subjects reported whether the stimulus direction was more clockwise or counterclockwise than the reference, followed by feedback. In a subset of trials, no feedback was given, and subjects performed the estimation task instead. (d,e) Importantly, subjects performed the estimation task either (d) in the presence or (e) in the absence of the reference line.

purely in *estimating* the direction of motion of the coherently moving dots. Five subjects performed this experiment, each performing a minimum of 2500 trials. It is worth noting that although there were three conditions (-6° , 0° and $+6^\circ$ shift in the reference line), because subjects gave an estimate on all trials (not just 30%), the number of direction estimates was approximately the same as in Experiment 1.

(c) Data analysis

Psychometric curves were obtained from the results of the direction of motion-discrimination task in Experiment 1. The proportion of times a subject reported a CW response was plotted as a function of the difference (in degrees) between the orientation of the discrimination boundary and the direction of motion of the coherent dots. We used nonlinear least-squares (*fminsearch* in Matlab; Nelder–Mead algorithm [9]) to estimate the two parameters (μ , σ) of the best-fitting cumulative Gaussian distribution. The parameter σ is a measure of a subject's discrimination threshold in the task.

Data from the direction estimation task were used to compare the presented (true) direction of motion with the perceived (estimated) one. To allow the combination of data across trials, we considered all data relative to the reference orientation for each trial (i.e. we subtracted the reference orientation from the true and estimated directions of motion). The range of true directions falls within $\pm 24.5^\circ$ of the reference. For each true stimulus direction, we then constructed histogram representations of the estimated directions. Figure 2a,b shows histograms from one representative subject for trials in the ambiguous case: 0° stimuli, moving exactly in the direction of the reference. To allow direct comparison with previous studies [6,10], we also visualized the joint histogram as density plots (figure 2c,d). Biases in the estimated motion direction (figure 2a,c) lead to a distinct pattern in these plots, quite distinct from those that arise if subjects' responses are noisy, but veridical (figure 2b,c).

The pattern of biased responses across all trials was well captured by the *dip*-statistic [11], a descriptive statistic that quantifies deviations from unimodality. For the two experimental conditions in Experiment 1 (reference present versus absent during estimation trials), and each subject and coherence level separately, we computed the *dip*-statistic across the estimated directions. We combined data across all (true) stimulus directions, which corresponds to taking the marginal histogram as shown in figure 3a. A large *dip*-statistic provides evidence for deviations from a unimodal distribution [11]. To assess the probability that a particular, observed value of the *dip*-statistic was owing to chance, we calculated 1000 bootstrap estimates of the *dip*-statistic by sampling, with replacement, from the uniform random distribution [12].

(d) Decoding model

We implemented a version of the Jazayeri & Movshon [6] decoding model, comprising two main components, to fit the data. The first is a sensory representation of the moving stimulus, modelled as a Gaussian probability density function, centred on the true direction of motion and whose variance is derived from the discrimination psychometric functions. In particular, the amount of spread of the representation changes with the coherence level: as the proportion of dots moving coherently decreases, the variance of the sensory representation increases. The second component of the model is a weighting profile, which serves to boost the responses of directionally tuned mechanisms that are optimal for the preceding fine-discrimination task: neurons whose tuning functions peak slightly to the sides of the decision boundary are more informative for such tasks than those tuned to directions very close to or far away from the boundary [7,8].

In order to account for the manipulations applied to the task in the second experiment, we modified the model. Because no direction discrimination was performed (i.e. subjects passively viewed the stimuli without making a judgement), the variance of the internal representation of the moving stimulus was also parametrized (instead of being derived from the discrimination psychometric function). The effect of shifting the reference line during the trial was estimated by, effectively, allowing the weighting function to slide along the x-axis (direction of motion), until the best fit to the data was obtained. This parameter thus returned the

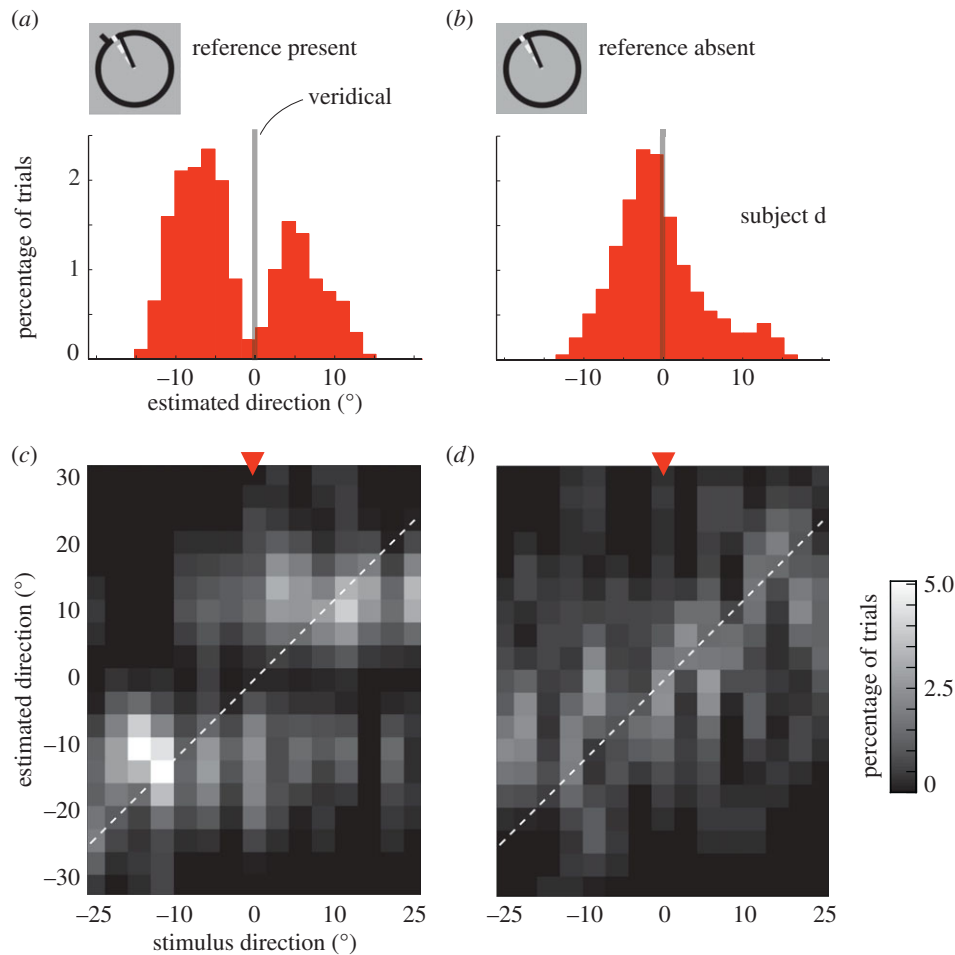


Figure 2. Behavioural results, reference present or absent. Estimated motion directions on the 30% of trials in which subjects performed the estimation task, following fine discrimination. (a) Example histogram for one stimulus direction and coherence (subject d, 0° , coherence 4%). When an explicit reference is present during the estimation period, estimated directions are shifted away from the veridical motion direction (grey line), resulting in a bimodal distribution. (b) However, in trials that are otherwise identical, when the reference is taken away while subjects estimate the perceived direction, the bias disappears resulting in a unimodal distribution. (c,d) The joint histograms show estimated directions for all (true) stimulus directions. Data in (c), obtained when the reference line was present, show an idiosyncratic pattern of responses. There is repulsion from the reference: few stimuli are estimated as moving at 0° (the direction of the reference). There is also some evidence of attraction to the reference for stimuli moving in directions more than about $\pm 10^\circ$ from the reference. When the reference was absent during estimation, the repulsive and attractive effects disappear and most estimates fall along the diagonal, indicating veridical responses. (Online version in colour.)

bias, in degrees, compared with the angular position of the reference at the start of the trial. Importantly, when considered together with the task sequence in our experiments, this second change to the model implies a fundamentally different process from that described by the original model. This is because it ties the repulsive biases to the position of the reference at the time of estimation and crucially not at the time of stimulus presentation.

The specific shape of the weighting function (and the fact that it is bimodal) captures both repulsion away from the reference line (CW, CCW), as well as the influence of the implicit boundary imposed by the stimulus range. Possible choices for such a function are the gamma distribution or any other parametric curve that has a similar shape.

3. Results

(a) Experiment 1: reference present or absent at estimation

All observers were able to perform the discrimination task at all coherence levels tested, and the precision of the discrimination performance was lawfully related to the amount of coherent signal present in the stimulus. The slopes of the psychometric

functions, obtained by fitting data with a cumulative Gaussian distribution function, become steeper as the coherence of the RDK increases (see electronic supplementary material, figure S1). In addition, discrimination was more precise as the coherent signal dots moved farther away from the discrimination boundary than when the motion was closer to the boundary. Moreover, we found no systematic direction bias (shifts in the point of subjective equality) for the discrimination data.

However, data from the estimation task revealed an interesting pattern. Assuming a veridical subject, i.e. one that reports precisely the direction of motion presented in the task, one would expect the distribution of estimates for a given direction to be normally distributed around the true direction of motion. However, this is not what we found. Estimated directions of motion were characterized by a bimodal distribution, indicating a repulsive effect consistent with the discrimination judgement: when the observer's choice in the discrimination task was CW, the estimated direction of motion was CW—shifted with respect to the discrimination boundary. This effect was particularly marked for trials in which the signal dots moved close to the discrimination boundary (figure 2a,c), and less evident as the coherent dots moved farther from the boundary.

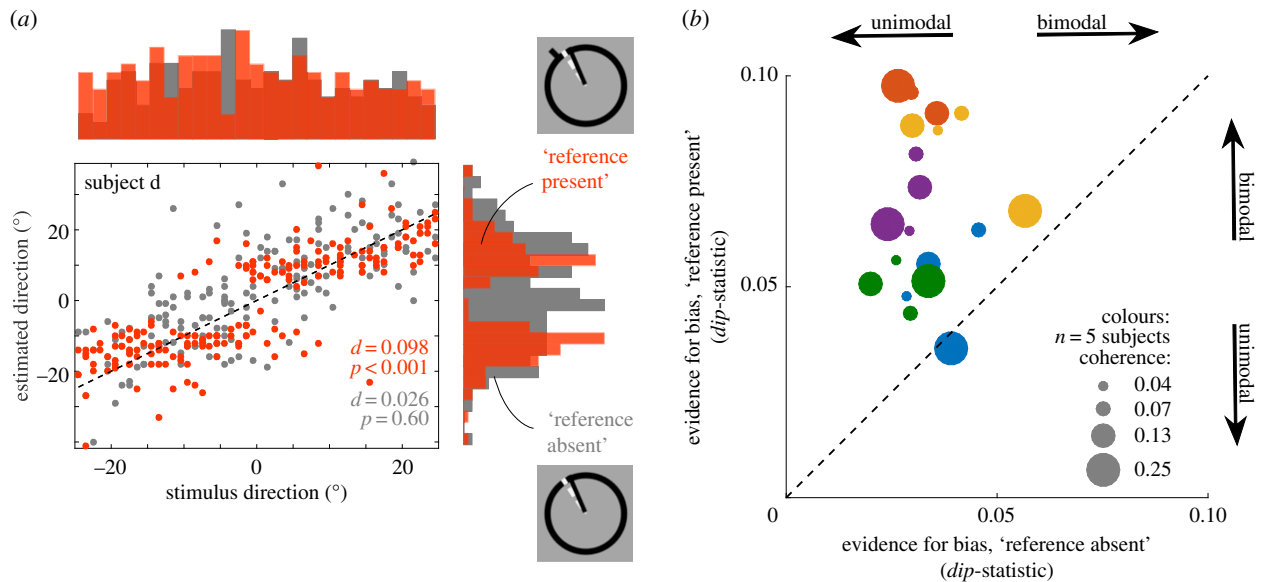


Figure 3. Behavioural results, across observers. Scatter plot of estimated directions as a function of (true) stimulus direction with respect to the reference for one example subject and condition (subject d, coherence 25%). Data for the ‘reference present’ condition are shown in red, those for the ‘reference absent’ condition in grey. The repulsive bias from the reference at stimulus directions around 0° can clearly be seen for the ‘reference present’ data and can be appreciated in the marginal histogram of estimated directions (red, reference present). For the data shown in grey (reference absent), the bias is not apparent. To assess deviations from unimodality in the marginal histograms, we computed Hartigan’s *dip*-statistic (d) and tested for statistical significance by a bootstrapping procedure (see S2c for details). For the data shown in (a) the *dip*-statistic indicates deviation from a unimodal distribution for the ‘reference present’ data ($d = 0.098$, $p < 0.01$), but not the ‘reference absent’ data ($d = 0.026$, $p = 0.60$). (b) Summary of psychophysical bias plots. The scatter plot shows *dip*-statistic for the ‘reference present’ (ordinate) and ‘reference absent’ (abscissa) data. Colours indicate data from different subjects ($n = 5$). Symbol sizes indicate data from four different levels of motion coherence used in the experiment (4–25%). For all subjects and all coherence levels (with one exception), the bias plots are unimodal when the reference was absent, bimodal when the reference was present (all individually, statistically significant at $p < 0.01$).

The pattern of responses was markedly different for trials in which the discrimination boundary was absent when subjects had to estimate the perceived direction of motion of the RDKs. Given that subjects were asked to estimate the direction in which they perceived the coherent motion earlier in the trials, we would not expect the absence of the reference line at this late stage to have an effect. However, the distribution of estimates differed from the bimodal one observed in the other trials: estimated directions of motion were centred on the true signal dot direction of the stimulus (figure 2*b,d*). The repulsive effect observed when the reference was present while estimating was not apparent, showing that the presence of the discrimination boundary during the estimation task plays an important role.

(b) Experiment 2: manipulating consistency of the reference information

No overt discrimination task was performed in this experiment. Data from the estimation task were divided according to the shift of the discrimination boundary (i.e. -6° , 0° , $+6^\circ$ with respect to the angular position of the boundary at the beginning of the trial; see figure 4*a*) and distributions of estimates are shown in figure 4*b*. When the reference line remained consistent throughout the trial, the same bias observed in Experiment 1 (reference present) emerged. Interestingly, when the boundary was shifted by 6° CW from its original location, the distribution of estimates shifted with it, presenting the same bias: repulsion from the reference line at time of estimation towards the boundary location at the start of the trial. Equally, the distribution of estimates shifted in a consistent direction for the CCW manipulation.

In order to fit these data, we used a maximum-likelihood procedure to fit both the original Jazayeri & Movshon [6] model and our modified model, described previously, to data from each participant. Because the models are nested, this allowed us to perform a likelihood ratio test, taking into account any advantages owing to the increase in free parameters. For all the diagnostic conditions across all subjects, our modified model outperformed the original one.

The estimated mean directions (± 1 s.e.m.) for the -6° , 0° and $+6^\circ$ conditions were -3.48° ($\pm 0.054^\circ$), -0.37° ($\pm 0.045^\circ$) and 3.49° ($\pm 0.069^\circ$). A repeated-measures two-way ANOVA showed a significant main effect of shift condition ($F_{2,8} = 58.731$, $p < 0.0001$), but no significant main effect of coherence level ($F_{3,12} = 0.776$, $p = 0.5294$). The interaction between shift condition and coherence level was not significant ($F_{6,24}$, $p = 0.8154$). Thus, the bias observed was driven by the position of the reference line during the estimation task but was independent of stimulus uncertainty (coherence; figure 5).

4. Discussion

Perceptual estimates can be biased by contextual cues. Explanations of perceptual illusions are often framed by how the early sensory apparatus transforms stimuli into a neural representation and how this neural representation is further processed. For example, *Mach bands*, the well-known optical illusion, can be explained by the particularities of stimulus encoding by mechanisms that have a centre-surround organization [13]. Many other biases or misperceptions are often explained by higher-level cognitive biases reflecting prior knowledge or assumptions about the world [14]. It is thought that the ‘hollow face’ illusion, for example, is owing to the

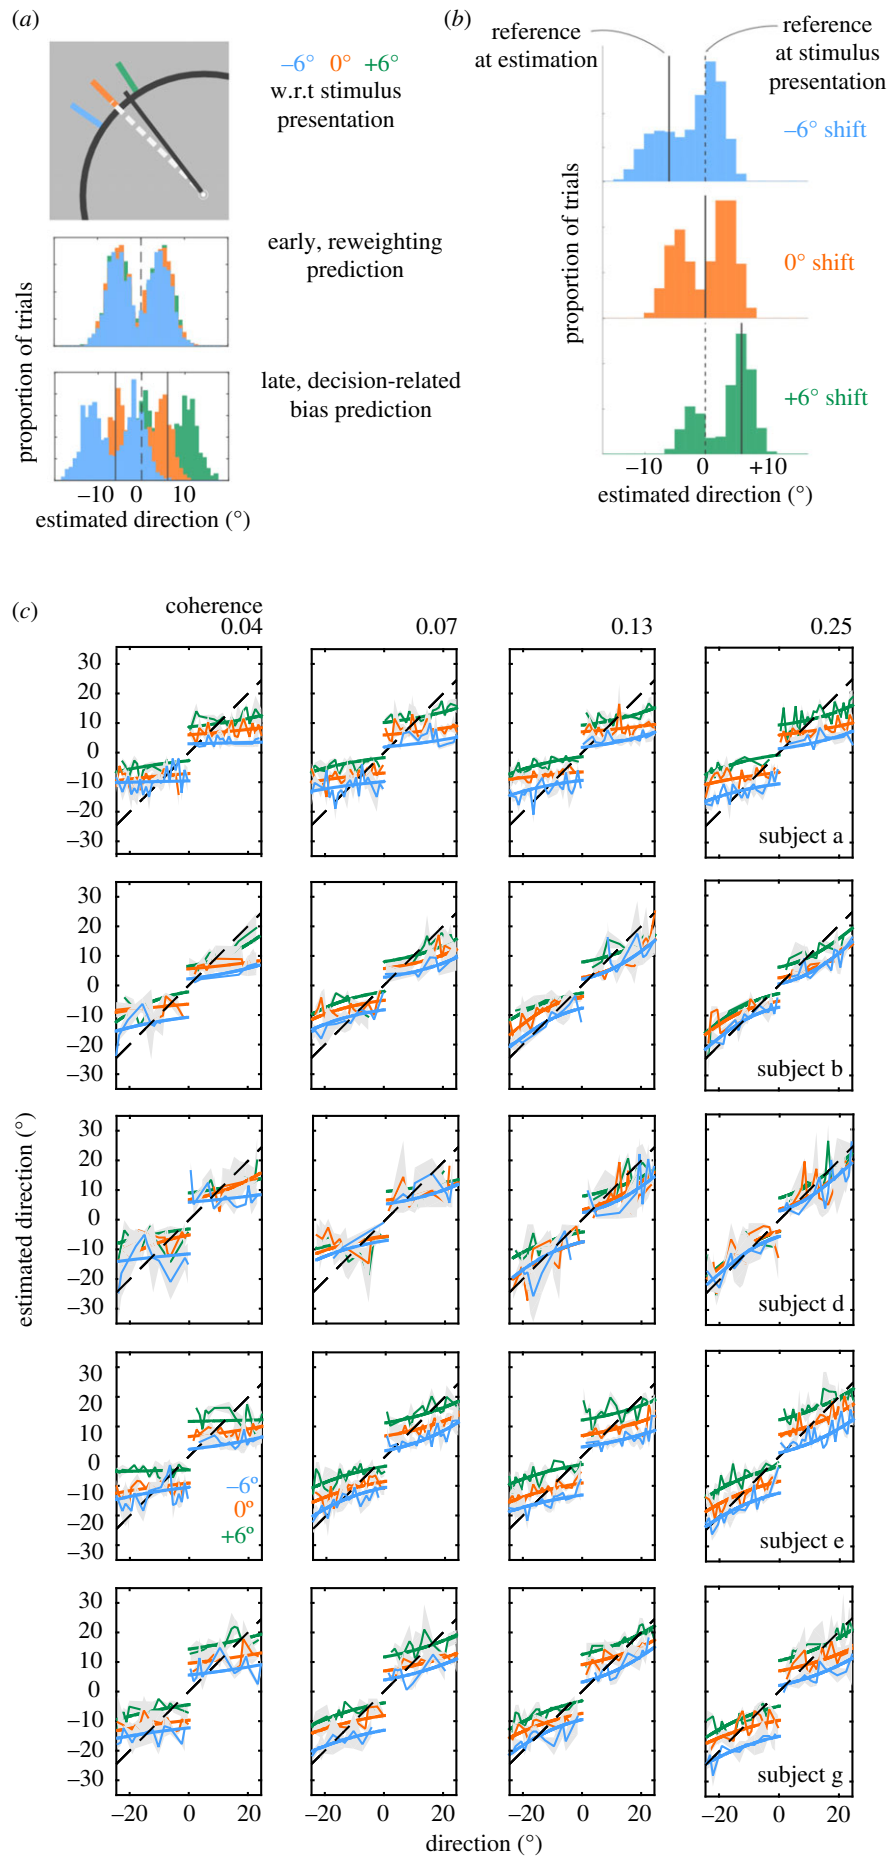


Figure 4. (Caption overleaf.)

Figure 4. (*Overleaf.*) Behavioural results, reference shifted. Estimated directions are affected by a late (just before the estimation task) shift in the location of the reference with respect to the stimulus. (a) Model predictions using simulated data for illustration. Following the RDK presentation, the reference line was shifted by either -6° , 0° , or $+6^\circ$ with respect to the beginning of the trial. The first part of each trial was identical for all three conditions (one for each reference shift): subjects passively viewed the moving stimulus without performing any task. If there is re-weighting of sensory evidence during this period of stimulus presentation and decoding, then the distribution of estimates should be the same for all three conditions (histograms overlap). If, however, the bias is due to a late mechanism, then the relative shifts in the reference between the first and second part of each trial should result in a commensurate shift in the estimated directions (histograms are separated). (b) Histograms of estimated directions for repeated presentations of a 0° stimulus. Colours indicate the different shifts in the reference: -6° , 0° and $+6^\circ$ (blue, orange, green). The locations of the reference at stimulus presentation (0° , black dashed line) and during the estimation task (-6° , 0° , $+6^\circ$, black solid line) are indicated. The distributions are clearly anchored to the location of the shifting reference line (illustrative data from one subject (d) at one coherence level). (c) Mean direction estimates as a function of the true stimulus direction. Five rows show data from different subjects. The four panels in each row show data for four different levels of stimulus coherence (4%, 7%, 13% and 25%). Colours indicate data for -6° (blue), 0° (orange) and $+6^\circ$ (green) shifts in the reference. Thin lines and shaded areas, mean ± 1 s.e.m. Thick lines, model fits of the late, decision-related bias model. (Online version in colour.)

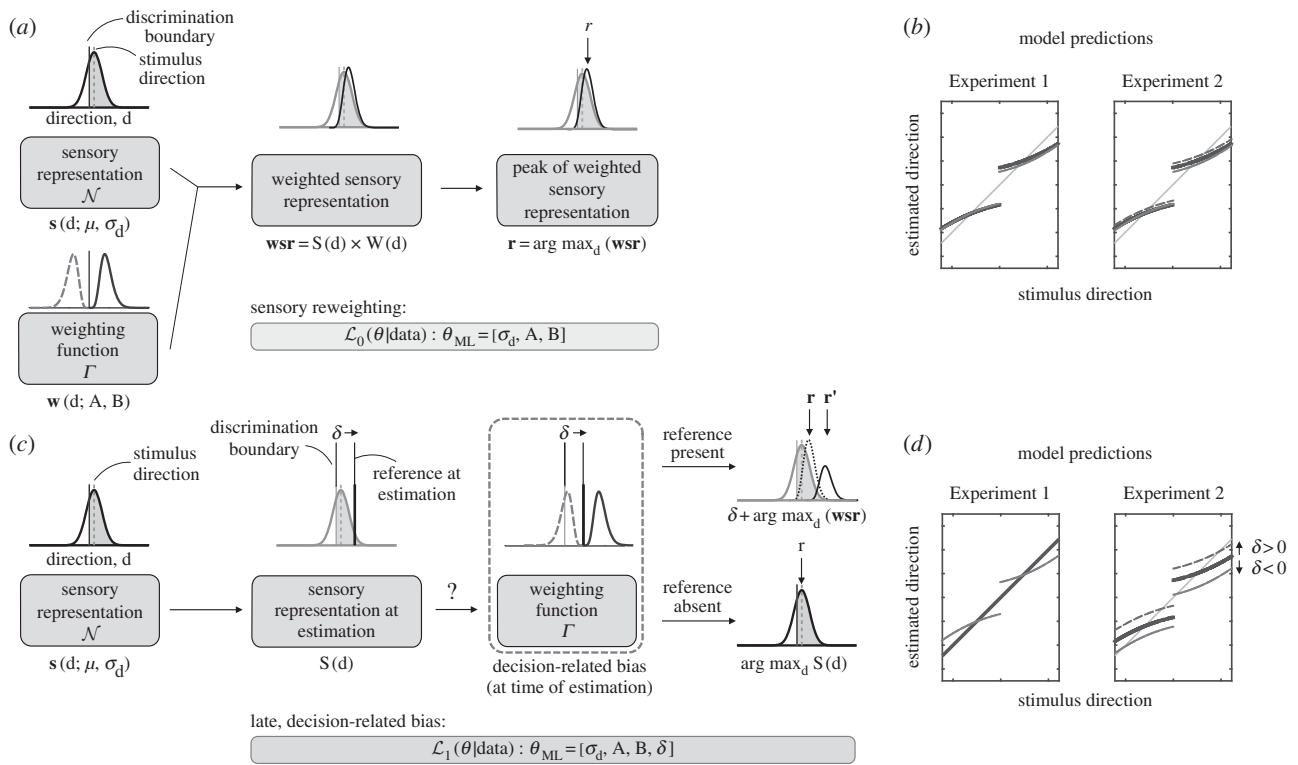


Figure 5. Sensory re-weighting and late, decision-related bias models. (a) Outline of the sensory re-weighting model. The sensory representation of the moving stimulus is modelled as a Gaussian probability density function (\mathcal{N}) centred on the true direction of motion (μ) and variance σ_d , a free parameter for both models. The weighting function \mathbf{w} is modelled as a gamma probability function (Γ) governed by two parameters: a shape parameter (A) and a scale parameter (B). Both \mathbf{s} and \mathbf{w} are derived in the same way for both the sensory re-weighting and late, decision-related bias models. The sensory representation of the motion stimulus is multiplied by the weighting profile, resulting in the weighted sensory representation (\mathbf{wsr}). To fit these models to the data, we obtained the (Gaussian) maximum-likelihood estimates for σ_d , A, B (and δ , for the modified model). (b) Model predictions for sensory re-weighting model. For both situations in our Experiment 1, reference present or absent during the estimation task, the model predicts the same biased responses, as the re-weighting of sensory information is tied to discrimination boundary which is unchanged in both situations. For the same reason, the original model predicts the same responses for Experiment 2, where the position of the reference is systematically changed at the time of the estimation task. (c) Outline of the late, decision-related bias model. The early sensory representation \mathbf{s} remains unchanged. The re-weighting of the sensory information by a weighting function \mathbf{w} is dependent on the presence of an explicit reference during the estimation task. In addition, the re-weighting is relative to the position of one or more references at the time of estimation. The parameter δ can absorb differences between the position of the discrimination boundary (during stimulus presentation) and the reference during the manual estimation task and correctly predicts shifts in the responses (from \mathbf{r} to \mathbf{r}'). (d) Model predictions for the late, decision-related bias model. In the situations where a reference is present during the estimation task (and not shifted with respect to the discrimination boundary), the original and modified models make the same predictions (light grey lines, left panel). However, if no reference is present at the estimation stage, the model predicts veridical responses (dark grey line, left panel). Additionally, if the reference present during estimation is displaced relative to the decision boundary, the modified model predicts concomitant changes in the responses (dashed and solid light grey lines, right panel).

combining of *prior* information about the distribution of the direction of illumination (more likely from above than below) and geometry of certain shapes (faces are convex, not concave).

When subjects are presented with two transparently moving stimuli, such as two superimposed populations of moving dots, judgements in the relative direction of motion are overestimated, a phenomenon that has been

termed *motion repulsion* [1]. The smaller the angular difference between the two dot populations, the bigger this repulsive effect. Interestingly, judging the direction of motion of a stimulus against a static reference can lead to repulsion in the estimates of the motion direction [15]. While the motion repulsion has been framed in terms of inhibitory interactions between directionally selective neurons as a low-level,

sensory phenomenon (see [1–4]), it has been suggested that reference repulsion could arise owing to higher-level cognitive effects ([5]; see also electronic supplementary material, discussion in [6]). That interpretation suggests that subjects have a veridical representation of the sensory stimulus at encoding and decoding in intermediate representations and that the bias arises from cognitive influences that are subsequently applied to them (as an example, see [16]).

Closely related to this phenomenon, Jazayeri & Movshon [6] reported a new perceptual illusion resulting from specific strategies applied during decoding of sensory information. Subjects performed a fine-discrimination task on a motion stimulus, judging whether the motion was in a direction more CW or CCW than an oriented reference line. Subjects received feedback on 70% of the trials, whereas on the remaining 30% feedback was withheld, and they were asked to match a line in the direction they perceived the stimulus moving. Discrimination performance decreased as the reliability of the motion signal decreased: the lower the reliability (coherence) of the motion stimulus, the higher the variability in responses. This effect was also reflected in the estimated perceived directions of motion: the lower the coherence of the motion stimulus, the larger the bias in the estimated responses. The authors accounted for these biases by implementing a model that re-weights sensory information away from reference direction. In particular, the shape of the weighting function allows an optimization of the fine-discrimination task (see [7,8]). Jazayeri & Movshon [6], therefore, interpret these results as a re-weighting of sensory information during discrimination, which later in the task sequence emerges as a bias in the estimated percept.

However, with the particular design of the task, it was impossible to disentangle whether the biases originated at the level of sensory representation (thus early in the trial sequence), or later in the trial [5]. This is because the oriented reference was present throughout the trial, from the time of stimulus presentation all the way to the response interval.

To address this, we designed an experiment in which we manipulated the oriented reference during the estimation phase. In the first experiment, we removed the oriented reference completely during estimation, following a fine direction-discrimination task. In the second experiment, subjects performed an estimation task following passive viewing of a moving stimulus—however, in this case, we systematically manipulated the orientation of the reference just before the onset of the estimation task. If the observed biases are owing to a process that operates at the time of stimulus presentation, encoding by the sensory apparatus, or even decoding from the activity in early visual processing, these manipulations should have no effect.

As originally observed by Jazayeri & Movshon [6], we found that subjects are biased when estimating the perceived direction of motion of an RDK in the presence of an oriented reference.

However, when we removed the reference line during the estimation task, this bias disappeared. The findings from Experiment 1 therefore suggest that subjects have a veridical representation of motion direction that can be used during the estimation task. Furthermore, this indicates that (i) making a binary judgement on a stimulus feature is not sufficient for biasing its representation, and (ii) the presence of the reference line during the estimation task is necessary for the bias to occur.

In Experiment 2, subjects were not required to perform a fine-discrimination task, but rather had to estimate the direction of motion of the RDK on a trial-by-trial basis. Again, estimates of motion direction were biased away from the reference line during the task, regardless of its angular location at stimulus presentation. This therefore suggests that an explicit binary judgement is not necessary to generate a biased response. Here subjects passively viewed the stimuli, but did not make a forced choice. Additionally, these results highlight the importance of the reference line during the estimation task: its availability is fundamental for the biased responses to arise.

It would be interesting to see if the reported biases are dependent on the response modality and whether perceptual reports via, e.g. saccadic eye movements show a similar pattern. Several studies have compared psychophysical and oculomotor performance in direction discrimination [17,18], but the results are equivocal. In particular, although effects such as *reference repulsion* and the *oblique effect* are consistently observed across different psychophysical experiments (see [19] as an example), studies investigating these effects with eye movements disagree in their conclusions [20,21].

Taken together, our results suggest that the oriented reference provided during the estimation task is used as an anchor to which the estimates of perceived directions of motion are yoked. We speculate that integrating information about the encoded stimulus with that about the reference results in a late, decision-related, rather than early, re-weighting of the sensory representation.

Ethics. Informed consent was obtained for each subject in accordance with regulations by the Ethics Committee of the School of Psychology, University of Nottingham.

Data accessibility. A comprehensive archive of the data reported in this study is available at <http://dx.doi.org/10.5061/dryad.ms84h>. In addition, an implementation of the model and code for producing illustrations is accessible at <https://github.com/schluppeck/zamboni-2016>.

Authors' contributions. E.Z., T.L., P.V.M. and D.S. designed research; E.Z. and D.S. performed research and analysed data; E.Z., T.L., P.V.M. and D.S. wrote the paper.

Competing interests. We declare we have no competing interests.

Funding. E.Z. is supported by a University of Nottingham Vice Chancellor's award.

Acknowledgements. We thank Mehrdad Jazayeri and two anonymous reviewers for insightful comments on our results and members of the Visual Neuroscience group for discussions of the experiments.

References

- Hiris E, Blake R. 1996 Direction repulsion between components in motion transparency. *Vis. Neurosci.* **13**, 187–197.
- Marshak W, Sekuler R. 1979 Mutual repulsion between moving targets. *Science* **205**, 1399–1401. (doi:10.1126/science.472756)
- Kim J, Wilson HR. 1996 Direction repulsion between components in motion transparency. *Vision Res.* **36**, 1177–1187. (doi:10.1016/0042-6989(95)00153-0)
- Mather G, Moulden B. 1980 A simultaneous shift in apparent direction: further evidence for a 'distribution-shift' model of direction coding. *Q. J. Exp. Psychol.* **32**, 325–333. (doi:10.1080/14640748008401168)
- Rauber H-J, Treue S. 1998 Reference repulsion when judging the direction of visual motion. *Perception* **27**, 393–402. (doi:10.1068/p270393)

6. Jazayeri M, Movshon JA. 2007 A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* **446**, 912–915. (doi:10.1038/nature05739)
7. Regan D, Beverley KI. 1985 Postadaptation orientation discrimination. *J. Opt. Soc. Am. A* **2**, 147–155. (doi:10.1364/JOSAA.2.000147)
8. Hol K, Treue S. 2001 Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Res.* **41**, 685–689. (doi:10.1016/S0042-6989(00)00314-X)
9. Nelder JA, Mead R. 1965 A simplex method for function minimization. *Comput J.* **7**, 308–313. (doi:10.1093/comjnl/7.4.308)
10. Stocker A, Simoncelli EP. 2008 A Bayesian model of conditioned perception. In *Advances in Neural Information Processing Systems 20* (eds JC Platt, D Koller, Y Singer, ST Roweis), pp. 1409–1416. Curran Associates, Inc. (See <http://papers.nips.cc/paper/3369-a-bayesian-model-of-conditioned-perception.pdf>.)
11. Hartigan JA, Hartigan PM. 1985 The dip test of unimodality. *Ann. Stat.* **13**, 70–84. (doi:10.1214/aos/1176346577)
12. Hartigan PM, Hartigan PM. 1985 Algorithm AS 217: computation of the dip statistic to test for unimodality. *Appl. Stat.* **34**, 320–325. (doi:10.2307/2347485)
13. Lotto RB, Williams SM, Purves D. 1999 Mach bands as empirically derived associations. *Proc. Natl Acad. Sci. USA* **96**, 5245–5250. (doi:10.1073/pnas.96.9.5245)
14. Gregory RL. 1997 Knowledge in perception and illusion. *Phil. Trans. R. Soc. Lond. B* **352**, 1121–1127. (doi:10.1098/rstb.1997.0095)
15. Rauber HJ, Treue S. 1999 Revisiting motion repulsion: evidence for a general phenomenon? *Vision Res.* **39**, 3187–3196. (doi:10.1016/S0042-6989(99)00025-5)
16. Huttenlocher J, Hedges LV, Duncan S. 1991 Categories and particulars: prototype effects in estimating spatial location. *Psychol. Rev.* **98**, 352–376. (doi:10.1037/0033-295X.98.3.352)
17. Beutter BR, Stone LS. 1998 Human motion perception and smooth eye movements show similar directional biases for elongated apertures. *Vision Res.* **38**, 1273–1286. (doi:10.1016/S0042-6989(97)00276-9)
18. Watamaniuk SN, Heinen SJ. 1999 Human smooth pursuit direction discrimination. *Vision Res.* **39**, 59–70. (doi:10.1016/S0042-6989(98)00128-X)
19. Gros BL, Blake R, Hiris E. 1998 Anisotropies in visual motion perception: a fresh look. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **15**, 2003–2011. (doi:10.1364/JOSAA.15.002003)
20. Churchland AK, Gardner JL, Chou I, Priebe NJ, Lisberger SG. 2003 Directional anisotropies reveal a functional segregation of visual motion processing for perception and action. *Neuron* **37**, 1001–1011. (doi:10.1016/S0896-6273(03)00145-4)
21. Rottach KG, Zivotofsky AZ, Das VE, Averbuch-Heller L, Discenna AO, Poonyathalang A, Leigh RJ. 1996 Comparison of horizontal, vertical and diagonal smooth pursuit eye movements in normal human subjects. *Vision Res.* **36**, 2189–2195. (doi:10.1016/0042-6989(95)00302-9)