

Complexity and specificity of experimentally induced expectations in motion perception

Nikos Gekas

IANC, School of Informatics, University of Edinburgh
Edinburgh, UK



Matthew Chalk

IANC, School of Informatics, University of Edinburgh
Edinburgh, UK



Aaron R. Seitz

UC Riverside, Department of Psychology
Riverside, CA, USA



Peggy Seriès

IANC, School of Informatics, University of Edinburgh
Edinburgh, UK



Our perceptions are fundamentally altered by our expectations, a.k.a. 'priors' about the world. In previous statistical learning experiments (Chalk, Seitz, & Seriès, 2010), we investigated how such priors are formed by presenting subjects with white low contrast moving dots on a blank screen, and using a bimodal distribution of motion directions such that two directions were more frequently presented than the others. We found that human observers quickly and automatically developed expectations for the most frequently presented directions of motion. Here, we examine the specificity of these expectations. Can one learn simultaneously to expect different motion directions for dots of different colors? We interleaved moving dot displays of two different colors, either red or green, with different motion direction distributions. When one distribution was bimodal while the other was uniform, we found that subjects learned a single bimodal prior for the two stimuli. On the contrary, when both distributions were similarly structured, we found evidence for the formation of two distinct priors, which influenced significantly the subjects' behavior when no stimulus was presented. Our results can be modeled using a Bayesian framework and discussed in terms of a sub-optimality of the statistical learning process under some conditions.

Keywords: expectation, motion perception, specificity, Bayesian, psychophysics

Introduction

A large body of work has led to the idea that human performance can be described as being probabilistically (or 'Bayesian') optimal, from cue integration (Hillis, Watt, Landy, & Banks, 2004), multisensory integration (Ernst & Banks, 2002), and sensorimotor learning (Kording & Wolpert, 2004), to high-level cognitive decisions (Trommershauser, Maloney, & Landy, 2008; Griffiths & Tenenbaum, 2006; Chater, Tenenbaum, & Yuille, 2006). It is found that humans can successfully learn about the regularities of particular scenes and the uncertainty related to them, and use this knowledge to guide perception and action. However, the scope of this optimality is controversial, in particular for tasks of high complexity. There is indeed evidence that there are strong limits on the types of statistical regularities that sensory systems can automatically detected (Fiser & Aslin, 2001, 2002). A common opinion is thus that only approximations of Bayesian inference could be implemented in the brain (Fiser, Berkes, Orbán, & Lengyel, 2010). However, in general, the constraints imposed on these approximations are still unclear.

Most experiments have looked at simple statistical learning tasks, involving regularities in one dimension. How humans learn probability distributions integrating multiple features simultaneously is still unclear, but there is indication that statistics about more than one variable can be learned jointly. Seitz et al. (2007), for example, showed that it is possible for stimulus associations to form across modalities, when there are statistical contingencies. They used a novel audio-visual procedure where subjects were passively exposed to a rapid serial presentation of audio-visual pairings and were later tested to report about the degree of familiarity of the stimuli. They found that subjects acquired knowledge of visual-visual, audio-audio, and audio-visual stimulus associations and that the learning of these types of associations occurred in an independent manner. Looking at vision alone, Turk-Browne et al. (2008) investigated how subjects learn regularities involving both shape and color. They showed that visual statistical learning could be both object-based and feature-based in regards to how feature dimensions co-vary. When shapes co-varied perfectly with a particular color, statistical learning

was object-based: the expression of learning required the presence of both shapes and colors at test. However, when shape and color were partially decoupled during learning, subjects showed robust statistical learning for each feature separately.

Another related question that has been poorly explored is that of specificity or transfer of the learning that is acquired. Is statistical learning very specific to the learned objects, similar to what is found in perceptual learning (Sasaki, Nanez, & Watanabe, 2009), or do human spontaneously transfer the acquired knowledge to similar objects? Two studies have directly investigated transfer of statistical learning. Adams et al. (2004) have shown that the so-called ‘light-from-above prior’ can be temporarily modified with active visual-haptic experience and that the resultant adaptation generalizes to a different visual task. Using statistical learning of sequences of shapes, Turk-Browne and Scholl (2009) explored transfer across space and time and across temporal order. They found that learning of statistically defined temporal sequences could be expressed in static spatial configurations, and that learning of statistically defined spatial configurations facilitated detection performance in temporal streams.

In previous work (Chalk, Seitz, & Seriès, 2010), we investigated whether expectations can be acquired through fast statistical learning, and how these expectations affect the perception of simple stimuli. Subjects were presented with a field of coherently moving white dots, shown at low contrast and had to report about the direction of motion of the dots (estimation task), as well as whether they had perceived the dots or not (detection task). Unknown to the subjects, two motion directions (-32° and 32°) were presented more frequently than the other directions. After only minutes of exposure to these stimuli, subjects developed expectations for the frequent directions, and this affected their behavior in two ways. First, subjects tended to perceive motion directions as being more similar to the expected directions than they really were (attractive bias). Second, even when there was no stimulus presented on the screen, subjects tended to ‘hallucinate’ motion particularly in the expected directions. This learning was implicit: when asked about the stimulus distribution after the experiment, subjects indicated no conscious knowledge that some directions had been presented more frequently than others. Chalk et al. constructed a simple Bayesian model that provided a good fit of subjects’ estimation biases. The modeling results suggest that subjects combined their sensory evidence with a learned prior distribution of expected motion directions in a probabilistically optimal manner. The model also provided correct predictions for the behavior when no stimulus was presented.

Here, we use a variation of this experimental paradigm to investigate the complexity and specificity of the priors that can be learned. Instead of using one color for the dots and one distribution for the motion directions, we now use two colors for the dots (green and red) corresponding to two different distributions of motion directions. We ask whether participants are able to form distinct expectations for the two different colors, or whether subjects form only a color nonspecific prior. Color information is mostly processed through the ventral stream, while motion information primarily through the dorsal stream. Our hypothesis is that a sub-optimal integration of the two information streams may provide insight on the mechanisms underlying the formation of priors and their constraints.

Experiment 1

Methods

Participants

Twenty-two naive participants with normal color acuity were recruited from the University of Edinburgh. All participants gave informed written consent in accordance with the University Of Edinburgh, School Of Informatics Ethics Panel and the Declaration of Helsinki, and received monetary compensation.

Stimuli

The motion stimuli consisted of a field of colored dots, either red or green, with density of 2 dots/deg², moving coherently at a speed of 9°/sec within a circular annulus, which had a minimum and a maximum diameter of 2.2° and 7° respectively. They were generated using the Matlab programming language with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997), and displayed on a Mitsubishi DiamondPro 750SB monitor running with resolution of 1024 x 768 at 100 Hz. Participants viewed the display in a darkened room at a viewing distance of 70cm. The display luminance was calibrated and linearized with a Cambridge Research Systems Colorimeter separately for each color. The background luminance was set to 5 cd/m².

Procedure

At the beginning of each trial, a central white dot (0.15° diameter) was presented as a fixation point (Figure 1a). Then, the field of dots was presented, in an annulus around the fixation point, along with a gray bar, which projected out from

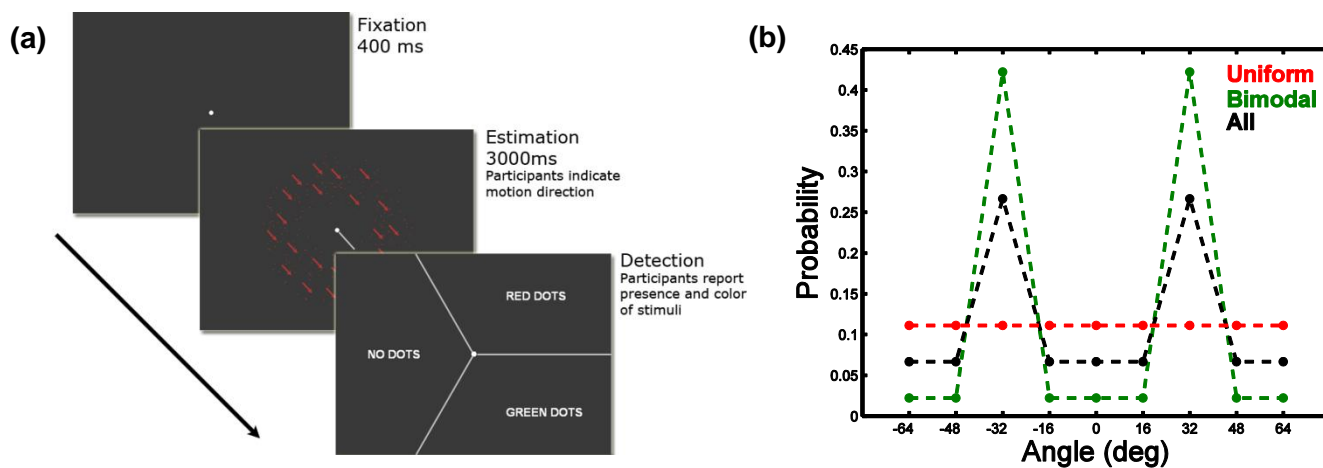


Figure 1. (a) Experimental procedure. Participants were presented with a fixation point, followed by the motion stimulus and the response bar, which they rotated to indicate their perceived direction of motion. After a period of 3000ms or the press of the mouse button, the screen was cleared and divided into three separate sections, and participants clicked in the appropriate section to indicate their choice. (b) Probability distributions of motion directions for Experiment 1. In the uniform condition, all directions were equally presented, while in the bimodal condition, two directions, 32° away from the central direction, were presented in a larger number of trials than other directions. The combined distribution was identical to the one used in Chalk et al. (2010). The central motion direction was randomized for each participant.

the central dot. The initial angle of the bar was randomized for each trial. Participants reported the direction of motion by rotating the bar using a mouse (estimation task). The display cleared when either the participant clicked on the mouse button, validating her choice, or when a period of 3000 ms passed.

After the estimation response, there was a 200ms delay before the detection response screen was shown. The screen was divided into 3 equal parts by 3 lines originating from the center of the screen and ending at the borders of the screen, resulting in an upper right, a lower right, and a left compartment, which displayed the text RED DOTS, GREEN DOTS, and NO DOTS respectively. Participants moved a cursor to the appropriate compartment on the screen to indicate their choice, and validated with a mouse click. The cursor then flashed green or red for a correct or an incorrect choice respectively. After the detection task was complete, the screen was cleared, and the next trial began after a 400 ms delay. The participants were presented with block feedback on the estimation task every 20 trials: a message was shown on the screen indicating their average performance in terms of their estimation error in the previous 20 trials.

Design

The experiment consisted of two sessions, lasting around 1 hour each, with each session comprising 765 trials. The 2 sessions were taken over successive days. The stimuli were presented at four different contrast levels and 2 colors (red and green), all randomly interleaved. There were 225 trials for each color where the contrast was determined using 3/1 staircases on detection performance (one staircase for each color). Additionally, there were 45 trials for each color, where the stimuli had the highest contrast level (high contrast), and 225 trials where no stimulus was presented (zero contrast). For the two staircased contrast levels, the direction of motion could be 0° , $\pm 16^\circ$, $\pm 32^\circ$, $\pm 48^\circ$, or $\pm 64^\circ$, with respect to a central reference angle. At the highest contrast level, the dots were given random directions, even outside the nine directions mentioned, in order to emphasize that the estimation task is continuous rather than a categorical decision out of 9 alternative directions.

Depending on their color, the dots could move according to one of 2 different direction distributions: uniform, or bimodal. In the uniform condition, all 9 directions were equally likely to be presented in each session or 25 trials per direction per session, whereas in the bimodal condition, there were 95 trials per session with direction -32° and $+32^\circ$, and 5 trials per session for each of the other directions of motion (Figure 1b). The combined distribution was identical to the one used by Chalk et al. (2010) with the difference that the total number of trials was slightly less. The distributions of the two colors were counter-balanced between participants in order to avoid any biases caused by color selection or color sensitivity.

Data analysis

As detailed above, the presented directions were symmetrical around a central motion angle, randomly chosen for each participant. This symmetry allowed us to average the results for stimuli moving to either side of the central motion direction. ‘Unfolded versions’ of the data can be found in the Supplementary materials (Supplementary figures 2 & 6).

The first 200 trials from each session were excluded from the analysis in order to allow the staircases to reach stable contrast levels (Supplementary figure 1). Participants’ performance was evaluated after their first session to eliminate those that could not perform adequately in the estimation task. Out of the original 22 participants, 4 failed to have a mean absolute estimation error less than 30° in the highest contrast trials, and they were not asked back for a second session. The remaining 18 participants passed that criteria for 2 sessions, and their data were analyzed. Responses to the highest contrast stimuli were disregarded from the analysis, and they were used only as a performance benchmark.

Participants were significantly better than chance in reporting the color of the stimulus even at low contrast (82% correct on average). However, they still made an error about color in a significant number of trials. Because those cases were ambiguous in terms of our hypothesis regarding color-specific expectations (i.e. whether participants would be using expectations corresponding to the stimulus color or to the reported color), these trials were removed from our analysis.

In the estimation task, the variance of participants’ direction estimates was large. Similarly to Chalk et al. (2010), we hypothesized that this was due to the fact that in some trials participants made completely random estimates, thus increasing the variance of motion direction estimates quite substantially. To account for this, we fitted the estimation responses of each participant to the distribution $(1 - \alpha) \cdot V(\mu, \kappa) + \alpha/2\pi$, where α is the proportion of trials where the participant makes random estimates, and $V(\mu, \kappa)$ is a von Mises (circular normal) distribution with mean μ and width $1/\kappa$, given by: $V(\mu, \kappa) = \exp[\kappa \cdot \cos(\theta - \mu)/(2\pi \cdot I_0(\kappa))]$. The parameters were chosen by maximizing the likelihood of generating the data from the distribution. Participants’ estimation mean and standard deviation were taken as the circular mean and standard deviation of the von Mises distribution. The use of this approach allows for more consistent and significantly smaller variances across participants, motion directions, and contrasts, than merely averaging over trials, without compromising the qualitative aspect of the results.

In trials where no stimulus was presented, we aggregated all participants’ responses and fit the population estimation performance using a linear combination of 6 circular normal distributions. The mean (peak) and variance (width) of each distribution were fit by maximizing the log likelihood of generating the data from the combined distribution. The number of distributions were chosen so as to avoid underfitting and overfitting the data, but was kept constant for all sets of data in order to avoid biasing the results towards a certain combined distribution (unimodal, bimodal, or trimodal).

For 11 participants, the motion direction of the red dots followed the uniform distribution and the directions of green dots the bimodal one, whereas for 7 participants this was reversed. There was no significant effect of color on the estimation bias, given a particular distribution ($p = 0.37$ and $p = 0.52$ for the uniform and bimodal distributions respectively, one-way ANOVA), so the data was combined across all participants. Additionally, there was no significant interaction between experimental session and motion direction on the bias and standard deviation ($p = 0.65$ and $p = 0.46$ respectively, four-way within-subjects ANOVA), so the data across sessions was combined as well.

Results

Motion direction estimates when a stimulus was presented

We asked whether the more frequently presented directions of the bimodal distribution would bias the participants’ estimation performance and whether that bias would be present only in the bimodal condition or in both conditions. We first analyzed the estimation responses obtained by averaging over all trials irrespective of the color condition. These results showed a distinctive pattern very similar to the one reported by Chalk et al. (Figure 2a) and as expected if participants are biased towards perceiving motion directions as being more similar to the most frequent directions than they really are. Estimates of the central motion direction were largely unbiased, whereas estimates of $\pm 16^\circ$ and $\pm 48^\circ$ were positively and negatively biased respectively. This indicates an attractive bias, which, however, is not centered around $\pm 32^\circ$ (the most frequent directions), but slightly shifted (to around $\pm 40^\circ$). Finally, estimation responses for stimuli moving at $\pm 64^\circ$, which lies at the largest distance from the most frequent directions, were significantly biased towards more central directions.

There was a significant effect of motion direction on the estimation bias ($p < 0.001$, two-way within-subjects ANOVA between motion direction and subjects), and the estimation bias of participants at $\pm 16^\circ$ and $\pm 48^\circ$ was significantly larger and smaller respectively than the bias at $\pm 32^\circ$ ($p = 0.043$ and $p = 0.001$, signed rank test). This verified that participants made estimates that were closer to the most frequently presented directions than the actual directions of the stimulus.

We next explored whether the bias towards the most frequently presented motion directions was found in both uniform and bimodal conditions, or whether participants behaved differently for the two conditions. The results showed that the estimation bias for the uniform condition closely resembles the bias for the bimodal condition (Figure 2b). Figure 2b




Figure 2. Effect of expectations on estimation biases and standard deviations for all trials. Results are averaged over all participants and error bars show within-subject standard error. The vertical dashed line corresponds to the two most frequently presented motion directions ($\pm 32^\circ$). (a) Participants' mean estimation bias for Experiment 1 averaged over all trials is plotted against presented motion direction. (b) Participants' mean estimation biases for Experiment 1 separated for the uniform and bimodal conditions are plotted against presented motion direction. Inset: Expected bias if participants use a prior that approximates the stimulus statistics. Here: The prior has modes centered on the most frequent directions, and the widths of the prior distributions and of the sensory likelihoods correspond to the average values across subjects. (c) The standard deviation in participants' estimation for Experiment 1 averaged over all trials is plotted against presented motion directions.

(inset) shows the estimation biases predicted by an “ideal” observer that has learned the true statistics of the stimuli. There was a significant effect of motion direction on the estimation bias for both conditions ($p < 0.001$ and $p < 0.001$ for the uniform and bimodal respectively, two-way within-subjects ANOVA), and direct comparison of the estimation biases between the two conditions showed that there was no significant difference between them ($p = 0.1$, three-way within-subject ANOVA between motion direction, color condition and subjects). Additionally, for the uniform condition, the estimation bias of participants at $\pm 16^\circ$ was significantly larger than the bias at $\pm 32^\circ$ ($p = 0.012$, signed rank test), and the bias at $\pm 48^\circ$ was significantly smaller than the bias at $\pm 32^\circ$ ($p = 0.047$, signed rank test). This suggests that participants tended to perceive motion direction as being more similar to $\pm 32^\circ$ than it really was, independently of the color of the dots. There appears to be large positive bias in the bimodal condition at 0° . However, the number of trials is very small for this condition and bootstrap analysis indicates that the difference between biases at 0° between the two conditions is in fact not significant (Supplementary figure 3a).

In accordance with Chalk et al., the standard deviation at the most frequently presented directions was lower than at other directions (Figure 2c). Overall, there was a significant effect of motion direction on the estimation standard deviation ($p = 0.017$, two-way within-subjects ANOVA between standard deviation and subjects). The standard deviations showed no significant difference between the two conditions ($p = 0.08$, three-way within-subjects ANOVA between standard deviation, color condition, and subjects).

Motion direction estimates when no stimulus was presented

On average, participants reported seeing a stimulus in $8.53\% \pm 3.95\%$ of the trials when no stimulus was presented. When participants reported detecting a stimulus, they reported the color of the uniform condition in 46.4% of the trials and the colors of the bimodal condition in 53.6% of the trials. We divided all participants' estimation responses in three sets; trials where they reported detecting stimulus of the uniform condition, trials where they reported detecting a stimulus of the bimodal condition, and trials where they reported detecting no stimulus. We fit a linear combination of 6 circular normal distributions to each data set and for all the combined data. We show the data along with the fitted distributions for the sets where participants reported detecting a stimulus of each condition (Figure 3a & Figure 3b). Detailed descriptions of the distributions can be found in the Supplementary materials. In order to compare statistics of the distributions, we used a Bootstrapping analysis. The data for each condition were re-sampled with replacement and distributions were fit to each data set. The process was repeated 10,000 times, and 95% percentile confidence intervals were calculated for each condition. A direct comparison of the folded distributions can be seen in Figure 3c.

The results suggest that participants were strongly biased to report motion in the most frequently presented directions when no stimulus was presented but they reported detecting a stimulus. A possible explanation for this behavior could have been that participants automatically moved the estimation bar towards one of the two most frequently presented directions, irrespective of their response in the detection task. However, such a response bias could be ruled out: participants

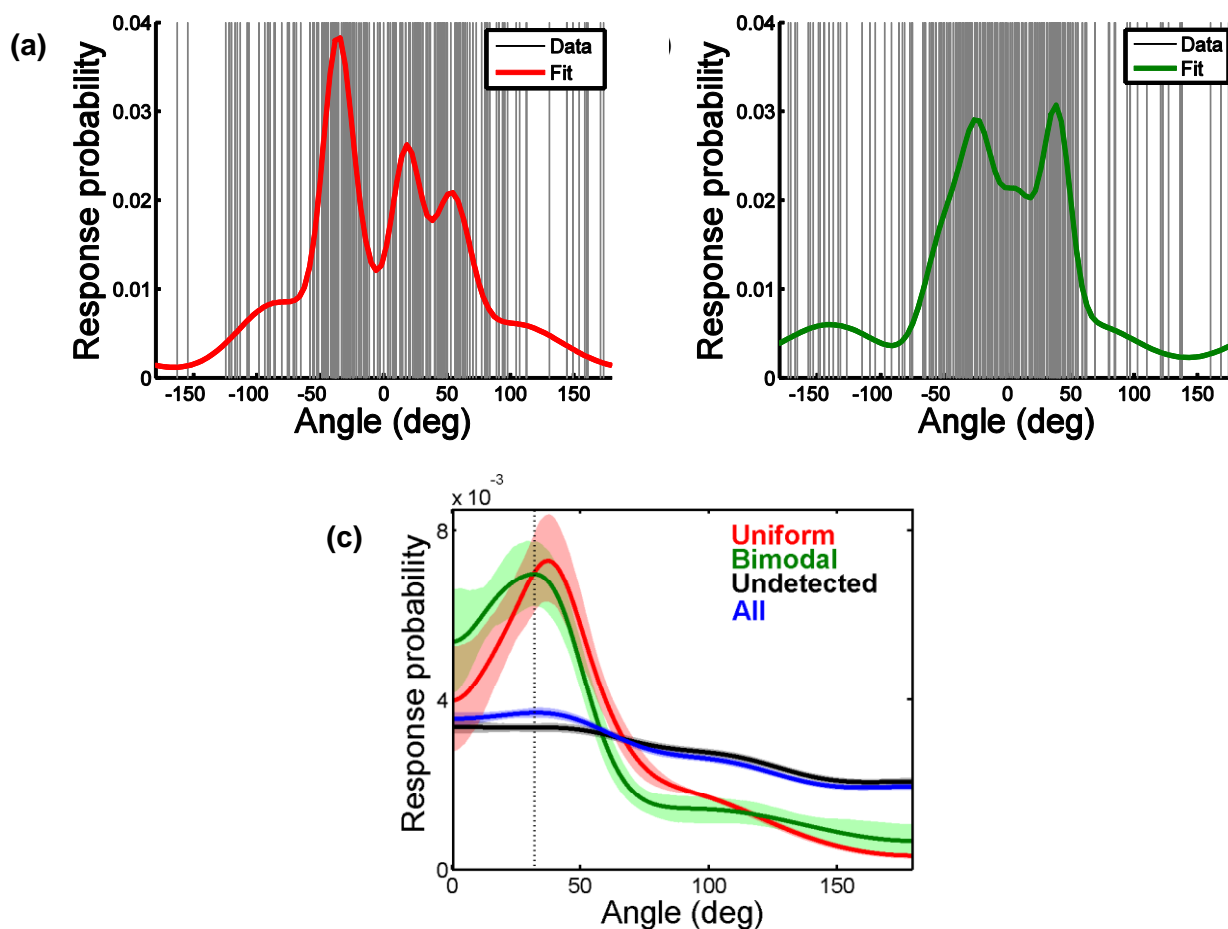


Figure 3. Participants' responses in trials where no stimulus was presented but they reported detecting a stimulus of the uniform condition (a) or bimodal condition (b) and fitted distributions. The vertical grey lines correspond to all the data points (estimation response), pooled across participants. (c) Fitted response distributions for trials where participants reported the color of the uniform distribution (red), the color of the bimodal distribution (green), trials where they did not detect a stimulus (black), and for all trials (blue). Data points from either side of the central motion direction have been averaged together. The vertical dotted line corresponds to the two most frequently presented motion directions ($\pm 32^\circ$). Error bars show 95% confidence intervals.

were not significantly more likely to move the estimation bar close to the frequent directions on trials that they reported seeing no stimulus. There was no significant difference between the participants' estimation behavior when reporting different colors, and the peaks of the combined distributions for both conditions were close to the frequently presented directions ($\pm 32^\circ$) but shifted slightly outwards, closer to $\pm 36^\circ$, to a greater degree for the uniform condition. These results largely replicate those of Chalk et al (2010); however, seem to suggest that the learned perceptual biases may be largely color invariant.

Experiment 2

The results of Experiment 1 suggested that participants may have ignored color information. Results showed that participants learned to expect the most frequently presented directions irrespective of color and used those expectations to guide perception for both conditions. In other words, participants may have learned only the combined distribution of the stimuli, which was bimodal. Alternatively, participants may have transferred the structure from the bimodal distribution to the unstructured (uniform) distribution. To disentangle between these alternatives, we designed a complementary experiment in which the combined distribution is uniform, and the two different distributions have comparable structural properties.

Methods

Participants

Twenty-three naive participants with normal color acuity were recruited from the University of Edinburgh. All participants in the study gave informed written consent in accordance with the University Of Edinburgh, School Of Informatics Ethics Panel and the declaration of Helsinki, and they received monetary compensation.

Stimuli & Procedure

The stimuli and experimental procedure were identical to those described in Experiment 1.

Design

The design of the experiment was identical to that described in Experiment 1, except for the motion-direction distributions of the stimuli. These distributions were chosen so that the combined distribution over both colors was uniform. The first distribution was bimodal, similar to the one used in Experiment 1, while the second distribution was complimentary of the first one in which the peaks were shifted by 32° (Figure 4). For convenience, we refer to the latter distribution as ‘trimodal’. The trimodal distribution had 40 trials per session for directions -64° , 0° , and $+64^\circ$, 4 trials per session for directions -32° and $+32^\circ$, and 24 trials per session for each of the other directions. The bimodal distribution had 44 trials per session for directions -32° and $+32^\circ$, 8 trials per session for directions -64° , 0° , and $+64^\circ$, and 24 trials per session for each of the other directions. The distributions of the two colors were counter-balanced between participants in order to avoid any biases caused by color selection or color sensitivity.

Data analysis

We performed the same analysis as for Experiment 1. Twenty-three observers participated in the first session of Experiment 2. Two of them were excluded from the second session, as their estimation error in high contrast trials was greater than 30° (our pre-hoc criterion), and three of them ignored the estimation task completely. The remaining eighteen participants returned for the second session, and their data were used in the results.

For 9 participants, the motion direction of the red dots followed the trimodal distribution and the directions of the green dots the bimodal one, whereas for 9 other participants this was reversed. There was no significant effect of the color presented on the estimation bias ($p = 0.73$ and $p = 0.6$ for the trimodal and bimodal distributions respectively, one-way ANOVA), so the data was combined across all participants. Additionally, there was no significant interaction between experimental session and motion direction on the bias and standard deviation ($p = 0.4$ and $p = 0.55$ respectively, four-way within-subjects ANOVA), so the data across sessions was combined as well.

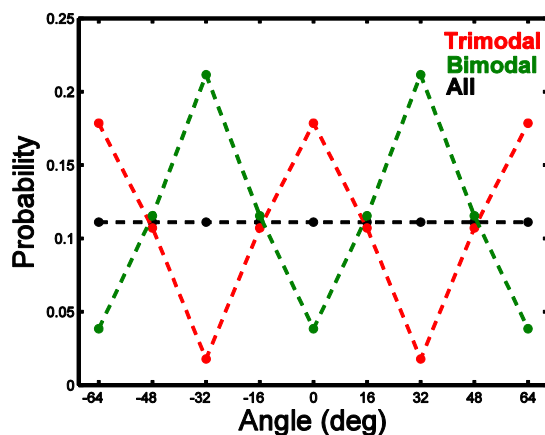


Figure 4. Probability distributions of presented motion directions for Experiment 2. In the trimodal condition, three directions, the central direction along with directions 64° away from it, were presented in a larger number of trials than other directions, while in the bimodal condition there were two directions, 32° away from the central direction more frequently presented. The combined distribution was uniform. The distributions of the two colors were counter-balanced between participants. The central motion direction was randomized for each participant.

Results

Motion direction estimates when a stimulus was presented

When a stimulus was presented and the data was pooled across color conditions, we found that, on average, participants' motion estimation performances did not vary significantly with motion direction (Figure 5a) ($p = 0.29$, two-way within-subjects ANOVA between motion direction and subjects). This was not unexpected; the combined distribution of the stimuli being uniform, potential biases in the estimation of each color condition might have cancelled each other out when averaged. There was no significant effect of motion direction on the estimation bias for both conditions ($p = 0.12$ and $p = 0.15$ for the trimodal and bimodal respectively, two-way within-subjects ANOVA), but there was a significant difference between the estimation biases for the two color conditions ($p = 0.046$, three-way within-subjects ANOVA between motion direction, color condition, and subjects, Figure 5b). However, these biases were weaker than in Experiment 1. The largest difference between the conditions was at $\pm 48^\circ$, where, on average, estimates were positively biased for the trimodal condition and slightly negatively biased (or unbiased) for the bimodal condition. Additionally, at $\pm 64^\circ$, estimates were largely unbiased for the trimodal condition but negatively biased for the bimodal condition. In contrast, at 0° (resp. $\pm 16^\circ$), participants' estimates were negatively biased (resp. unbiased) for both conditions. Figure 5b (inset) shows the estimation biases predicted by an "ideal" observer that has learned the true statistics of the stimulus. These results suggest that the participants' motion-direction estimates were approximately biased towards the most frequent directions for each color condition for outwards angles (i.e. 32° and 64° respectively – with an outwards shift for the bimodal condition) but dominated by an attraction towards the central direction for small angles, independently of the color condition.

There was a significant effect of motion direction on the estimation standard deviation (Figure 5c) ($p < 0.001$, two-way within-subjects ANOVA between standard deviation and subjects). The highest values were at $\pm 16^\circ$, while the lowest were at $\pm 64^\circ$. On average, stimuli closer to the central direction produced larger standard deviations than those further away. There was no significant difference between the standard deviations for the two color conditions ($p = 0.23$, three-way within-subjects ANOVA between standard deviation, color condition, and subjects).

Motion direction estimates when no stimulus was presented

On average, participants reported seeing a stimulus in $10.34\% \pm 9.1\%$ of trials when no stimulus was presented. Participants reported seeing the color corresponding to the trimodal condition in 49.9% of the trials, and that corresponding to the bimodal condition in 50.1% of the trials. We divided all participants' estimation responses into three sets; trials where they reported detecting stimulus of the trimodal condition, trials where they reported detecting stimulus of the bimodal condition, and trials where they reported detecting no stimulus. We fit a linear combination of 6 circular normal distributions to the estimation performances for each data set and for all the combined data. Figure 6 presents the data along with the fitted distributions for the sets where participants reported detecting a stimulus of each condition (Figure 6a & Figure 6b). Detailed descriptions of the distributions can be found in the Supplementary materials. As in Experiment 1, we used a




Figure 5. Effect of expectations on estimation biases and standard deviations for all trials. Results are averaged over all participants and error bars show within-subject standard error. The red and green vertical dashed lines corresponds to the most frequently presented motion directions for the red and green conditions (0° and $\pm 64^\circ$) and $\pm 32^\circ$ respectively. (a) Participants' mean estimation bias for Experiment 2 is plotted against presented motion direction. (b) Participants' mean estimation biases for Experiment 2 separated for the trimodal and bimodal conditions are plotted against presented motion direction. Inset: Expected bias if participants use a prior that approximates the stimulus statistics. Here: The prior has modes centered on the most frequent directions, and the widths of the prior distributions and of the sensory likelihoods correspond to the average values across subjects (c) The standard deviation in participants' estimation for Experiment 2 is plotted against presented motion direction.

Bootstrapping analysis to compare the statistics of the distributions. A direct comparison of the folded distributions can be seen in Figure 6c.

The results suggest that participants were strongly biased to report motion around $\pm 54^\circ$ when no stimulus was presented but they reported detecting a stimulus. As in Experiment 1, participants were not significantly more likely to move the estimation bar close to those directions on trials for which they reported seeing no stimulus. Here though, there are differences between the participants' estimates depending on the color they report. The trimodal and bimodal shapes of the distributions resemble the stimulus distributions for the 2 conditions: It was significantly more likely for participants to report the color of the trimodal condition at 0° and nonsignificantly at $\pm 64^\circ$, while it was significantly more likely for them to report the color of the bimodal condition at $\pm 32^\circ$. This suggests that when no stimulus was present, but participants reported perceiving one, they were biased to make direction estimates consistent with the most frequent directions of the color condition they reported. However, the peaks of the fitted distributions were not exactly centered at the frequent directions but shifted ($\pm 59^\circ$ for the trimodal and $\pm 44^\circ$ for the bimodal condition), consistent with the estimation biases when a stimulus was presented. The probability of reporting stimulus of the trimodal condition at 0° was lower than the probability at $\pm 64^\circ$.

We also assessed the tendency for participants to make estimates close to the most frequent directions, relative to other directions, by multiplying the probability of estimation within 8° of 0° and of $\pm 64^\circ$ for the trimodal distribution and of $\pm 32^\circ$ for the bimodal distribution by the total number of 16° bins. This probability ratio p_{rel} would be equal to 1 if estimation was equally likely between the most frequently presented directions and other 16° bins. It is possible to investigate

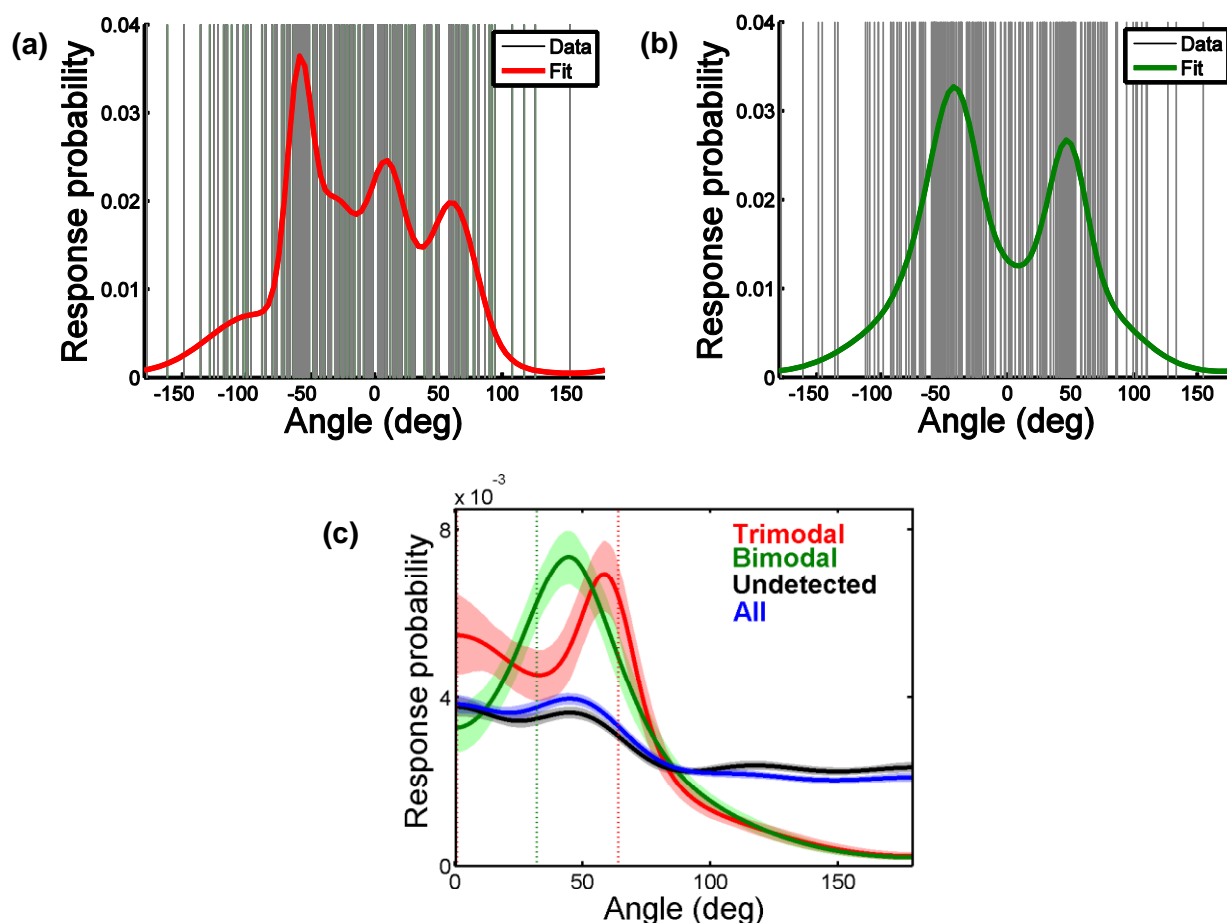


Figure 6. Participants' estimation responses and fitted distribution for trials where no stimulus was presented but they reported detecting a stimulus of the trimodal condition (a) or of the bimodal condition (b). The vertical grey lines correspond to all the data points (estimation response), pooled across participants. (c) Fitted response distributions for trials where participants' reported the color of the trimodal distribution (red), the color of the bimodal distribution (green), where they did not detect a stimulus (black), and for all trials (blue). Data points from either side of the central motion direction have been averaged together. The red and green vertical dotted lines corresponds to the most frequently presented motion directions for the trimodal and bimodal distributions (0° and $\pm 64^\circ$) and ($\pm 32^\circ$) respectively. Error bars show 95% confidence intervals.

how quickly these biases developed by calculating the probability ratio for individual participants every 100 trials for both sessions (including all responses up to that point) (Supplementary figure 9). For the bimodal condition, the median value of p_{rel} was significantly larger than 1 at the most frequently presented directions of that distribution ($\pm 32^\circ$) after only 200 trials of the first session. On the other hand, for the trimodal condition it took approximately 400 and 900 trials for the probability ratio to become significantly larger than 1 for the most frequent presented direction of that distribution (0° and $\pm 64^\circ$ respectively), suggesting that it may have taken longer to learn the trimodal distribution. Also, the probability ratios for the most frequently presented directions of the opposite distribution ($\pm 32^\circ$ for the trimodal and 0° and $\pm 64^\circ$ for the bimodal) were never significantly larger than 1.

These results suggest that color information was used when learning the direction distributions and support the postulate that the lack of color specificity found in Experiment 1 was due to a transfer of information between the colors rather than an inability for the system to use the color information when present. To better understand these results we applied several models to the data.

Modeling

In Chalk et al (2010), we used two classes of models to understand how participants' expectations may be combined with the presented stimulus to produce the observed behavior. The first class assumed that participants developed response strategies unrelated to perceptual changes, while the second class assumed that participants used a Bayesian strategy in which they combined a learned prior of the stimulus statistics with their sensory evidence in a probabilistic way.

The models were assessed with the use of a metric called the 'Bayesian information criterion' (BIC), which is defined as $BIC = -2 \cdot \ln(L) + k \cdot \ln(n)$, where L is the likelihood of generating the experimental data from the model, k is the number of parameters in the model, and n is the number of data points available. The first term quantifies the error between the data and the model predictions, while the second term penalizes increasing model complexity, and the model with the lower value of BIC should be preferred when comparing two models (Schwarz, 1978). The Bayesian model was found to exhibit significantly smaller BIC values than all other models, and produced fits for the estimation bias and the standard deviation that were at least on par with the first class of models, despite having fewer free parameters. This suggests that a Bayesian strategy was the best description of the participants' behavior.

Here, we implemented both the simple Bayesian model and the 'response strategy' models, and we again found that the Bayesian model was able to fit the data accurately and exhibited significantly better BIC values than the other models (Supplementary Figure 5). Next, we evaluated several extended versions of the simple Bayesian model that took into account the statistical information of the two colored conditions and compared them to the simple model. We will briefly describe the Bayesian models before reporting their performances. A more detailed description of the simple Bayesian model can be found in (Chalk, Seitz, & Seriès, 2010).

The simple Bayesian model

The simple Bayesian model assumed that participants combined their sensory evidence with a learned prior of the stimulus directions in a probabilistic manner. Participants were assumed to make noisy observations (θ_{obs}) of the stimulus motion direction (θ), with a probability $p_l(\theta_{obs}|\theta) = V(\theta, \kappa_l)$, where $V(\theta, \kappa_l)$ is a circular normal distribution with width $1/\kappa_l$. The posterior probability that the stimulus is moving in a particular direction θ , using Bayes' rule, is given by multiplying the likelihood function $p_l(\theta_{obs}|\theta)$ with the prior probability $p_{prior}(\theta)$:

$$p(\theta|\theta_{obs}) \propto p_{prior}(\theta) \cdot p_l(\theta|\theta_{obs}), \quad (1)$$

It was hypothesized that participants could not access the 'true' prior, $p_{prior}(\theta)$, so they learned an approximation of this distribution, $p_{exp}(\theta)$. This approximation was defined as the sum of two circular normal distributions, each with width determined by $1/\kappa_{exp}$, and centered on motion directions $-\theta_{exp}$ and θ_{exp} respectively:

$$p_{exp}(\theta) = \frac{1}{2} [V(-\theta_{exp}, \kappa_{exp}) + V(\theta_{exp}, \kappa_{exp})], \quad (2)$$

Participants were assumed to make perceptual estimates of motion direction θ_{perc} by choosing the mean of the posterior distribution:

$$\theta_{perc} = \frac{1}{Z} \int \theta \cdot p_{exp}(\theta) \cdot p_l(\theta_{obs}|\theta) \cdot d\theta, \quad (3)$$














Model Name	Description	Free parameters	Example priors	
			Both conditions	
1Bimodal	A single Bimodal prior is applied in both conditions.	4		
Uni+Bi	A Uniform and a Bimodal prior are combined in a single prior applied in both conditions.	5		+  = 
Gaus+Bi	A Gaussian and a Bimodal prior are combined in a single prior applied in both conditions.	7		+  = 
			Uniform condition	Bimodal condition
Split_UniBi	A Uniform prior is applied in the uniform condition, while a Bimodal prior is applied in the bimodal condition.	4		
Split_GausBi	A Gaussian prior is applied in the uniform condition, while a Bimodal prior is applied in the bimodal condition.	6		
Split_2Bimodal	Two distinct Bimodal priors are applied in the uniform and bimodal conditions.	6		

Table 1. Names and descriptions of the models proposed to describe the participants' behavior in Experiment 1 are provided along with the number of free parameters required by each model, and example priors.

where Z is a normalization constant. Finally, it was hypothesized that there is a certain amount of noise associated with moving the mouse to indicate the direction the stimulus is moving, and that the participants make completely random estimates in a fraction of trials α . The estimation response θ_{est} given the perceptual estimate θ_{perc} is then:

$$p(\theta_{est}|\theta_{perc}) = \frac{1}{2\pi} [(1 - \alpha) \cdot V(\theta_{perc}, \kappa_m) + \alpha], \quad (4)$$

where the magnitude of the motor noise is determined by $1/\kappa_m$. We assumed that the perceptual uncertainty at the highest contrast was close to zero ($1/\kappa_l \sim 0$). So, by substituting $\theta_{exp} = \theta$ and using Equation 4 we fit participants' estimation distributions at high contrast in order to approximate the width of the motor noise ($1/\kappa_m$) for each participant for all models.

In total, the free parameters that were fitted to the estimation data for each participant were the center and width of the expected distribution (θ_{exp} and $1/\kappa_{exp}$, respectively), the width of the participants' sensory likelihood ($1/\kappa_l$), and the fraction of trials where they made completely random estimation (α).

Experiment 1

Extended models

In the current experiment, the simple Bayesian model that assumes a single bimodal prior corresponding to the combined distribution of the stimuli is a non-optimal model, as it is blind to stimulus color. We will refer to this model as '1Bimodal'.

Five different model variations of the simple model were proposed (Table 1). The first 2 variations, 'Uni+Bi' and 'Gaus+Bi' also assume that participants form a unique prior, used in all trials irrespective of the color of the dots. 'Uni+Bi' assumes that participants learned a linear combination of the two different motion distributions used to construct the stimuli, a uniform and a bimodal distribution. These distributions are combined to create the final prior $p_{prior}(\theta)$ in the following way:

$$p_{prior}(\theta) = c_{pr} \cdot p_{bimodalr}(\theta) + (1 - c_{pr}) \cdot p_{uniform} \quad (5)$$

where c_{pr} is a free parameter fitted for each participant, $p_{uniform}$ is a uniform distribution identical to the distribution of the uniform stimuli (Figure 1b), and $p_{bimodal}(\theta)$ is equal to equation 2. The model had a total of 5 free parameters (θ_{exp} , κ_{exp} , κ_l , α , and c_{pr}).

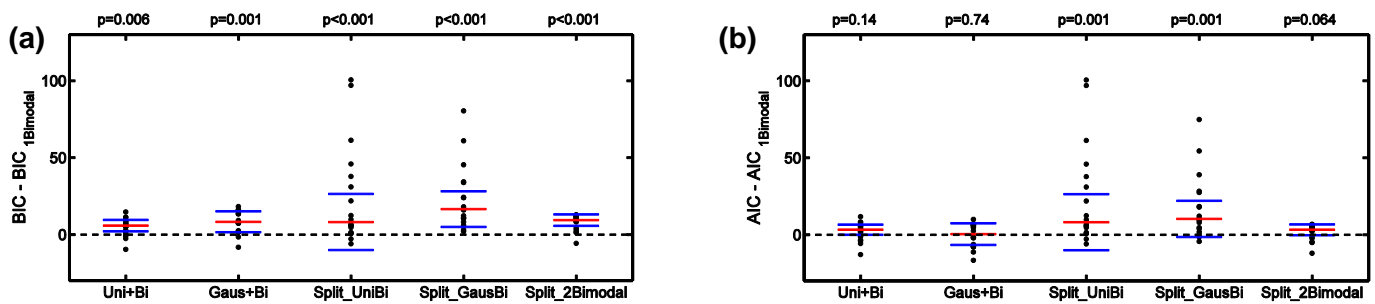


Figure 7. Model comparison. The (a) Bayesian information criterion (BIC) and (b) the Akaike Information Criterion (AIC) values of each model subtracted by the AIC and BIC values of the ‘1Bimodal’ model are plotted for each participant (black dots), along with median values (red lines), and the 25th and 75th percentiles (blue lines). p -values indicate whether the median was significantly different from zero for each model (signed rank test).

The model ‘Gaus+Bi’ assumed that participants developed a Gaussian prior instead of a uniform one, which was of the form $p_{\text{gaussian}}(\theta) = V(-\theta_u, \kappa_u)$. This model was inspired by data inspection showing that participants tend to exhibit an attraction towards the central direction. This model required two additional free parameters (θ_u, κ_u).

The other 3 variations correspond to the hypothesis that participants may form distinct priors for the 2 color conditions. The data was split between the two color conditions, which were fit separately. The model ‘Split_UniBi’ assumed that participants correctly learned a uniform prior for the uniform condition, and a bimodal prior for the bimodal condition. This model corresponds to the optimal observer model. The model ‘Split_GausBi’ assumed that participants learned a Gaussian prior for the uniform condition, and a bimodal prior for the bimodal condition. The form of the Gaussian prior was the same as in ‘Gaus+Bi’, so it required two additional free parameters. Finally, the model ‘Split_2Bimodal’ assumed that participants learned two different bimodal priors for each condition. This model required two additional free parameters for the center and width of the additional bimodal prior.

Model evaluation

All five extended models were implemented and assessed. We compared all extended models with the simple model, ‘1Bimodal’, that assumed participants only learned the combined, ‘color-blind’, distribution of the motion directions. The BIC values obtained for each model were subtracted by the BIC value obtained for the ‘1Bimodal’ model (Figure 7a). The results show that all BIC values were significantly greater than that obtained with the ‘1Bimodal’ model ($p = 0.006$ for ‘Uni+Bi’, $p = 0.001$ for ‘Gaus+Bi’, and $p < 0.001$ for ‘Split_UniBi’, ‘Split_GausBi’, and ‘Split_2Bimodal’, signed rank test). The single prior models performed better than the ‘Split’ models, and the bimodal prior dominated over the uniform and Gaussian priors with averaged c_{pr} values of 0.87 ± 0.14 and 0.8 ± 0.22 respectively. The best performing ‘Uni+Bi’ model was still significantly worse than that of the simple ‘1Bimodal’ model.

Among the ‘Split’ models, the BIC values exhibited by the theoretically optimal ‘Split_UniBi’ model were not signif-

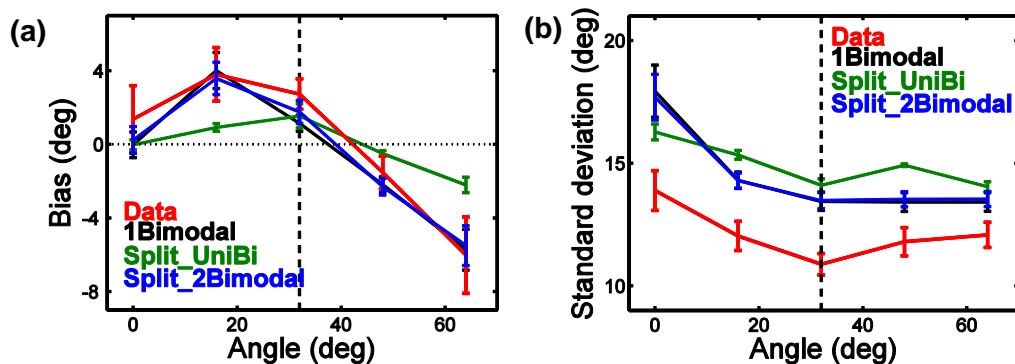


Figure 8. Predicted averaged estimation biases (a) and standard deviations (b) for each model. Predictions for the ‘Split_UniBi’ (green), the ‘Split_2Bimodal’ (blue), and the ‘1Bimodal’ models (black) are plotted with the experimental data (red). Results are averaged over all participants and error bars represent within-subject standard error.

icantly better than the values of the other ‘Split’ models despite having 2 free parameters less than the other two models (Figure 7a). This strongly suggests that participants did not learn a uniform prior for the uniform condition. Moreover, the ‘Split_2Bimodal’ model exhibited significantly better values than the other two models ($p = 0.02$ and $p = 0.002$ compared to ‘Split_UniBi’ and ‘Split_GausBi’ respectively), which indicates that participants learned bimodal priors for both the uniform and bimodal conditions. However, the values were significantly worse compared to the values of the simple ‘1Bimodal’ model.

In addition to the *BIC*, the models were evaluated with the Akaike Information Criterion (*AIC*), which is defined as $AIC = 2 \cdot k - 2 \cdot \ln(L) + (2 \cdot k \cdot (k + 1)/n - k - 1)$, where L is the likelihood of generating the experimental data from the model, k is the number of parameters in the model, and n is the number of data points available. The *AIC* penalizes the number of parameters less strongly than the *BIC*. Even so, the extended models perform worse than the ‘1Bimodal’ model, non-significantly for the single prior models, and significantly for the ‘Split’ models (Figure 7b).

While most of the extended models had more free parameters than the ‘1Bimodal’ model, they did not produce significantly better fits for the participants’ estimation biases and standard deviations (Figure 8a & 8b). The ‘Split_2Bimodal’ model had almost identical results to the ‘1Bimodal’ model ($p = 0.45$, three-way within-subjects ANOVA between motion direction, model and subjects), further suggesting that participants did not form two independent priors for the two color conditions. Additionally, while the ‘Split_UniBi’ predictions were not significantly different compared to ‘1Bimodal’ predictions ($p = 0.14$, three-way within-subjects ANOVA), they exhibited a larger mean absolute error (3.25° compared to 1.87° for the ‘1Bimodal’ model, and 1.85° for the ‘Split_2Bimodal’).

Experiment 2

The participants’ behavior in Experiment 2 was more difficult to quantify, as the estimation performances varied greatly on a participant-by-participant basis. We first assessed whether the data could be described using a simple model called ‘Uniform’ that assumed that participants only learned the combined statistics of the stimuli. The prior distribution was defined as a ‘uniform’ distribution identical to the combined distribution over both colors, and the model required only 2 free parameters (κ_l and α). The model’s predictions on the estimation biases differed significantly from participants’ estimation biases with a mean absolute error of 6.02° (Figure 9a), suggesting that participants formed non-uniform priors.

We next implemented models with more complex priors, assuming that the same prior was used for both color condi-

Model Name	Description	Free parameters	Example priors	
Uniform	A single Uniform prior is applied in both conditions.	2	Both conditions	
2Circ	A single prior equal to the sum of 2 pairs of circular normal distributions is applied in both conditions.	6		
3Circ	A single prior equal to the sum of 3 pairs of circular normal distributions is applied in both conditions.	8		
Split_TriBi	A Trimodal prior is applied in the trimodal condition, while a Bimodal prior is applied in the bimodal condition.	6	Trimodal condition	Bimodal condition
Split_2Circ	Two distinct priors for each condition of the stimuli equal to the sums of 2 pairs of circular normal distributions.	10		
Split_3Circ	Two distinct priors for each condition of the stimuli equal to the sums of 3 pairs of circular normal distributions.	14		

Table 2. Names and descriptions of the models proposed to describe the participants’ behavior in Experiment 2 are provided along with the number of free parameters required by each model, and example priors.

tion (Table 2); In the ‘2Circ’ model the prior is defined as the sum of two pairs of circular normal distributions, and in the ‘3Circ’ model the prior is defined as the sum of three pairs of circular normal distributions. The former model requires 6 free parameters (θ_{exp} , κ_{exp} , θ_{2exp} , κ_{2exp} , κ_l , and α) and the later 8 (adding θ_{3exp} , κ_{3exp}). The predictions of these models were significantly more accurate than the predictions of the ‘Uniform’ model with mean absolute errors of 5.89° and 5.76° respectively (Figure 9a). The prior distributions predicted by the models differed extensively between participants. The standard deviation predicted by the models was larger than the experimental results (Figure 9c). However, the qualitative trend for the standard deviation to decrease away from the central direction displayed by all models (but the uniform) was consistent with the data.

We next tried models that assumed distinct priors depending on the color condition. The model ‘Split_TriBi’ assumes that participants correctly developed a trimodal prior for the trimodal condition and a bimodal prior for the bimodal condition. The center of the trimodal distribution was fixed at 0° , while the other peaks were fit to the data. The ‘Split_TriBi’ model requires 6 free parameters (θ_{tri} , κ_{tri} , θ_{bim} , κ_{bim} , κ_l , and α), and corresponds to a model of the optimal observer. The models ‘Split_2Circ’ and ‘Split_3Circ’ are similar to 2Circ and 3Circ defined above but now with two distinct priors for each condition, requiring 10 and 14 free parameters respectively. As can be expected from the models’ increased complexity, the estimation biases predicted by the ‘Split’ models were closer to the experimental results with mean absolute errors of 5.72° , 5.66° , and 5.67° respectively. The models predict different biases for the two color conditions (Figures 9b₁ & 9b₂). The ‘Split_TriBi’ model provides very accurate predictions for both conditions at $\pm 48^\circ$ and at $\pm 64^\circ$ but fails at $\pm 16^\circ$. This suggests that the participants’ estimation performances were more weakly biased towards the central direction for the trimodal condition than expected and that this attractive bias possibly transferred to the bimodal condition. The predicted standard deviations of these models did not differ much from those predicted by the single prior models (Figure 9d).

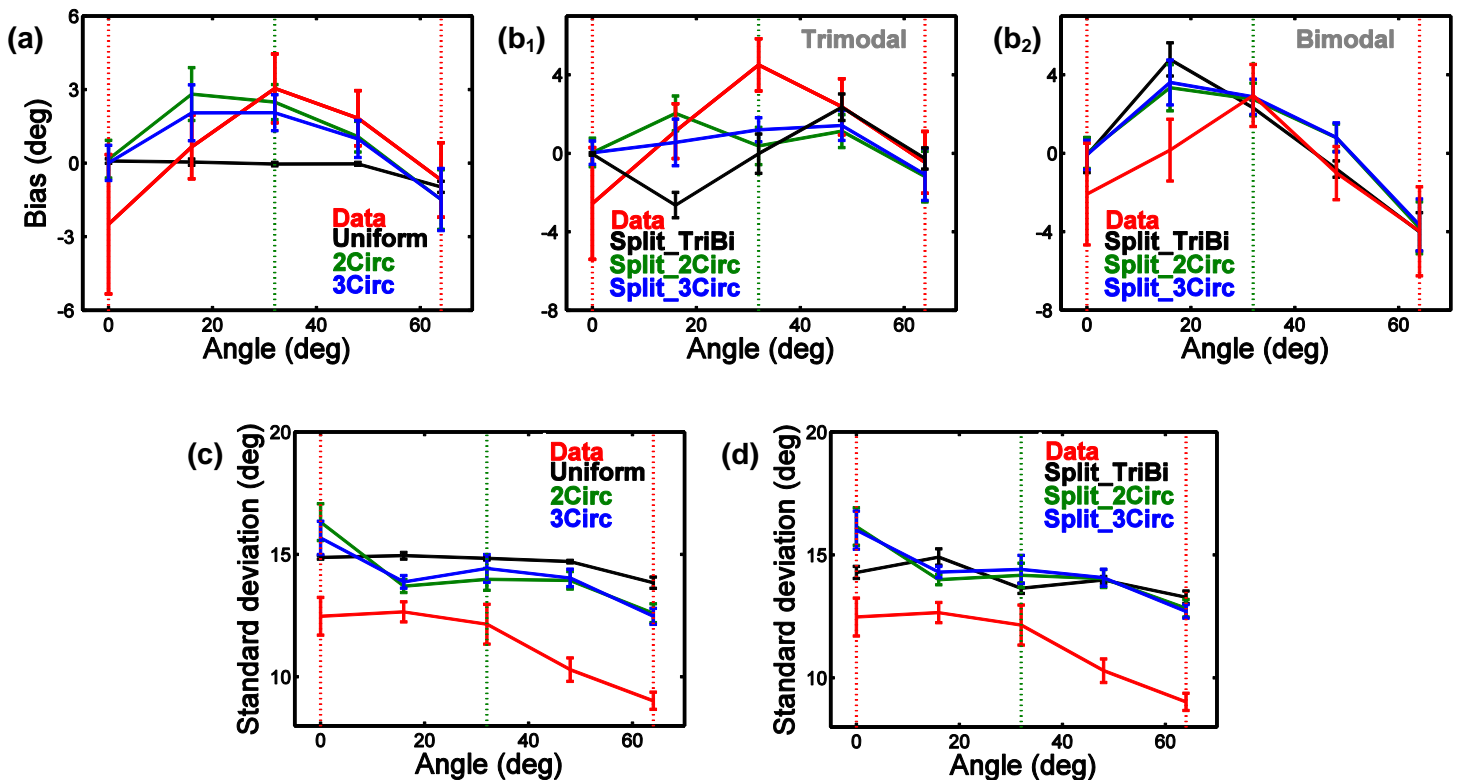


Figure 9. (a) Comparison between the predicted averaged estimation biases for the ‘3Circ’ (blue), the ‘2Circ’ (green), and the ‘Uniform’ models (black) and the experimental data (red). (b) Predicted averaged estimation biases for the ‘Split_3Circ’ (blue), the ‘Split_2Circ’ (green), and the ‘Split_TriBi’ models (black) and experimental data (red) of the trimodal condition (b₁) and of the bimodal condition (b₂). (c) Predicted averaged estimation standard deviations for the ‘3Circ’ (blue), the ‘2Circ’ (green), and the ‘Uniform’ models (black) and the experimental data (red). (d) Predicted averaged estimation standard deviations for the ‘Split_3Circ’ (blue), the ‘Split_2Circ’ (green), and the ‘Split_TriBi’ models (black) and the experimental data (red). Results are averaged over all participants and error areas show within-subject standard error. The red and green vertical dotted lines corresponds to the most frequently presented motion directions for the trimodal and bimodal distributions (0° and $\pm 64^\circ$) and ($\pm 32^\circ$) respectively.

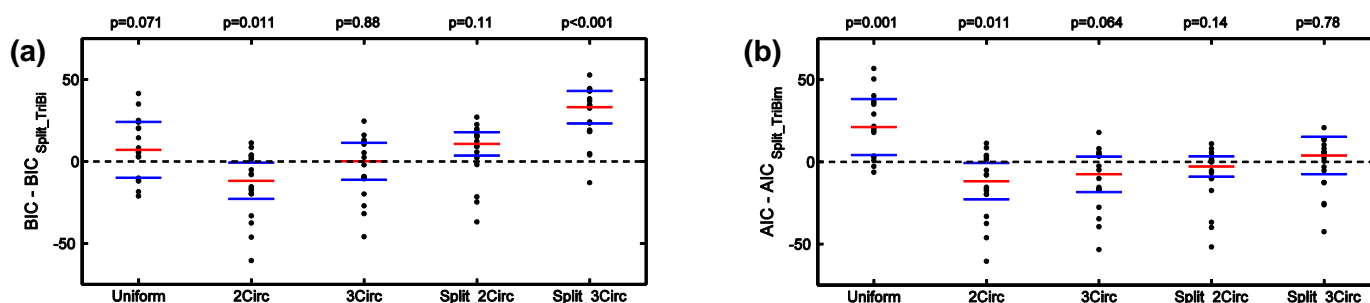


Figure 10. Model comparison. The (a) Bayesian Information Criterion (BIC) and (b) the Akaike Information Criterion (AIC) values of each model subtracted by the AIC and BIC values of the ‘Split_TriBi’ model are plotted for each participant (black dots), along with median values (red lines), and the 25th and 75th percentiles (blue lines). p -values indicate whether the median was significantly different from zero for each model (signed rank test).

We computed the AIC and BIC values to evaluate the models’ goodness of fit for the estimation performances. The simple ‘Uniform’ model performed significantly worse than all the other models, so we selected the optimal ‘Split_TriBi’ model as the reference for comparison. The BIC results (Figure 10a) showed that the ‘2Circ’ model was the only one that performed significantly better than the ‘Split_TriBi’ model, whereas the ‘Split_3Circ’ model performed significantly worse. In terms of AIC , the ‘Uniform’ model was now significantly worse, while again only the ‘2Circ’ model was significantly better (Figure 10b). The fact that more complex models provided significantly better results suggests that participants learned complicated priors even though the combined distribution was uniform. Moreover, it is interesting that while, on average, the ‘Split_TriBi’ model was significantly worse than ‘2Circ’, it performed better for one third of the participants. This suggests that, for at least some participants, estimation performances were better explained by assuming distinct priors for the two color conditions.

We compared the centers of the circular normal distributions predicted by the models ‘2Circ’, ‘3Circ’, ‘Split_2Circ’, and ‘Split_3Circ’ with the most frequently presented directions by measuring the minimum absolute difference between the model θ_{exp} values and 0° , 32° , and 64° . We found that the θ_{exp} values were distributed in a similar way across all models. We calculated the ratios of θ_{exp} values based on their proximity to each direction for each of the 4 models; the percentage of θ_{exp} values (averaged over all models) which fall closest to one of the most frequent directions (than to the other two directions) were $46.4\% \pm 5.2\%$ for $\pm 64^\circ$, $37\% \pm 4.5\%$ for $\pm 32^\circ$ and $16.6\% \pm 4.7\%$ for 0° . The average minimum absolute difference was $8.45^\circ \pm 2.6^\circ$ from the frequently presented directions (averaged over all models). This suggests that, on average, participants do learn a distribution with peaks located around the most frequent directions. However, it seems that the representation of the central direction is suppressed, compared to the other directions ($\pm 32^\circ$, and $\pm 64^\circ$).

The performance of the ‘2Circ’ model suggests that the results are best explained by a model that assumes a single prior. However, the performance of the optimal ‘Split_TriBi’ model indicated that at least 6 participants developed a distinct prior for each stimulus condition, and even though the differences between these priors are sometimes subtle, the best-fitting peaks of the distributions are close to the frequently presented directions of the corresponding conditions. The ‘Split_TriBi’ model (but not ‘2Circ’) can also qualitatively explain participants’ behavior when there was no stimulus presented. This suggests that at least some participants may have learned approximations of the ‘true’ priors.

Discussion

The results of Experiment 1 showed that participants quickly developed expectations for the most frequently presented directions over all trials, irrespective of the color of the dots. They exhibited estimation biases towards those directions similarly for both the uniform and bimodal color conditions. On trials where no stimulus was presented but participants reported seeing a stimulus, they were strongly biased to make estimates in the most frequently presented motion directions regardless of the color reported. Participants’ estimation behavior in Experiment 1 was described successfully by a probabilistic model, which assumed that they used a non-optimal Bayesian strategy, which combined their sensory evidence with a unique learned prior of the combined stimulus statistics, applied to both color conditions in a probabilistic way. The model could accurately predict the participants’ behavior when no stimulus was presented, and performed better than ‘response strategy’ models, which assumed that participants developed response strategies unrelated to perceptual changes.

Several variations of the Bayesian model were proposed that took into account the color information of the stimuli presented. The models assumed that participants formed a single prior by linearly combining two different distributions or that they used two different priors to estimate the motion direction of the two color conditions. While the new models provided slightly more accurate fits to the data, they performed significantly worse in terms of the *BIC* and *AIC* criteria compared to the simple Bayesian model. This suggests that participants developed a single bimodal prior that was used for all stimuli independent of color.

In Experiment 2, participants' estimation performances on trials where a stimulus was presented were significantly different between the two color conditions, but the induced biases were comparably weaker than and not as clear as in Experiment 1. Participants' estimation performances on trials where no stimulus was presented but where participants reported seeing a stimulus, were significantly different depending on the color they reported, and there were indications that participants increasingly perceived the most frequently presented directions of the color condition they reported as the sessions progressed. Participants' estimation behavior in Experiment 2 was described effectively by a model that assumed that a distinct prior was learned for each color condition. For a number of participants, the model was able to provide significantly better fits to the experimental data than a simple model that assumed that participants learned only the combined statistics of the stimuli, and had smaller *AIC* and *BIC*, despite having a larger number of free parameters. Moreover, the prior distributions predicted by the model for each color condition provided a possible explanation for participants' behavior in trials where no stimulus was presented.

Optimality

In general, it is not easy to distinguish between biases that occur at the perceptual or decision-making level (Schneider & Komlos, 2008). However, the modeling results in the current work showed that participants combined their expectations with sensory observations, not by following different strategies on each trial but by using both in a probabilistic way. At the end of the second session in both Experiments, we questioned the participants as to whether they had identified directions that were presented more frequently and whether they had noticed a difference between the directions of the two color conditions. In Experiment 1, the majority (13 out of 18 participants) reported that there were equal number of stimuli moving in all directions, and even more participants (15 out of 18) reported that they did not notice any difference between the motion directions of the two colored stimuli. In Experiment 2, 12 out of 18 participants thought that there were an equal number of stimuli moving in all directions, and 14 out of 18 thought that the red and green dots had the same distribution. The participants who reported noticing one or two directions being more frequent (5 out of 18 in Experiment 1 and 6 out of 18 in Experiment 2) were only able to give an approximate estimate of the quadrant of motion directions. The participants who reported noticing a difference between the two distributions (3 out of 18 in Experiment 1 and 4 out of 18 in Experiment 2) failed to report the correct frequent directions of each distribution, and their estimation performance was not in accordance with their conscious estimates.

In both experiments and when stimulus was and was not presented, there is evidence that, on average, the priors used by the participants are slightly shifted compared to the stimulus distribution. A similar shift was also present in Chalk et al. (2010). That participants would learn only an approximation of the stimulus distribution with slight variability in its peaks was not unexpected. However, it is possible that this shift reflects more than inter-individual variability. In this case, we can only speculate about its potential origin. A possible explanation is that it would reflect a perceptual repulsion from the main reference direction, which in this case is the central direction (Raubert & Treue, 1998), or an apparent repulsion between the two modes of the motion distributions (Treue, Hol, & Rauber, 2000).

The Bayesian model in Chalk et al. (2010) showed that participants behaved similarly to an optimal Bayesian observer. Here, in Experiment 1, participants failed to behave optimally, as they did not incorporate the information provided by the color condition to make more accurate estimates of the motion directions of the stimuli. This suboptimal behavior could have resulted from either ignoring the color of the dots in the estimation task (since it was not directly relevant to this task) and learning only the overall combined distribution of the motion directions, or from learning only the distribution of the bimodal condition and using it also for the uniform condition. The results of Experiment 1 could not disentangle between these two explanations. Experiment 2, on the contrary, was designed so that those two explanations would lead to different predictions. It showed that participants didn't ignore color information but tried to learn the statistical properties of the stimuli for each condition.

That only one motion distribution was learned in Experiment 1 is intriguing. A possible explanation is that neural mechanisms might exist that restrict the formation of multiple priors in order to reduce extensive plasticity in sensory cortex, or that force competition between different sensory adaptations. In Experiment 1, learning of the bimodal prior seemed to have dominated over the uniform prior. This might be related to a compromise between plasticity and performance cost for the task at hand. Indeed, one can argue that, in Experiment 1, the benefit of learning about the uniform dis-

tribution is not as large as that of learning about the bimodal distribution or about each of the two distributions in Experiment 2, in terms of the detection task performances.

Another reason for sub-optimality in these tasks might be found in the classically postulated separation of the motion and color pathways in the visual cortex. One may wonder whether there is a single representation of the joint color-motion statistics (and if so, where in the brain?), or whether there is a representation of the motion statistics, for e.g. in MT, that would be further integrated with color information. The relationship between the neural processing of color and motion information has been a debated issue in visual neuroscience. However, there is evidence that chromatic information influences the responses in MT, and that this cortical area is an important component of the neuronal substrate of color-based motion processing (Dobkins and Albright 1998, Thiele, Dobkins & Albright, 2001). It has also been shown that color signals in MT of the macaque monkey influence behavior in speed judgment tasks (Seidemann, Poirson, Wandell, & Newsome, 1999), and, similarly, that the color sensitivity of motion-selective cortex MT+ in humans follows the color sensitivity of psychophysical observers in making speed judgments (Wandell et al., 1999).

Attention

The experimental and modeling results suggested that participants' estimation performances did not vary depending on the color condition in Experiment 1, and only moderately so in Experiment 2. Participants identified the correct color in 82% and 79.6% of the trials on average for Experiments 1 and 2 respectively, which indicates that they paid attention to the color as well as the motion direction of the stimuli. Additionally, there was no preference for one color condition over the other as the errors on color were evenly balanced in both Experiments (51% false green and 49% false red with 13% standard deviation in Experiment 1 and 53% false green and 47% false red with 16% standard deviation in Experiment 2). While on average there was no color preference, some participants exhibited a tendency towards reporting one color more than the other. However, there was no indication that this preference was affecting their estimation behavior (i.e. reporting a particular color at a specific direction as a response strategy when uncertain about the stimulus).

Selective attention towards one color at a time could possibly explain the results in Experiment 2. Participants could have attended more to one color for a certain amount of time, and then switched to the other color and so on. This behavior should result in learning being spread between the two color conditions, and in participants forming priors that would be a combination of the statistical properties of the two distributions. In order to evaluate this hypothesis, we implemented two models that assumed participants developed such priors; a linear combination of the stimuli distributions, i.e. a trimodal and a bimodal distribution (Supplementary Figure 8). The weight of each distribution was a free parameter that could take any value from 0 to 1. The first model assumed a single prior and the second assumed two distinct priors for each color condition. The performances of both models were worse than the optimal Bayesian model in fitting the experimental data and in the *AIC* and *BIC* criteria. Moreover, closer examination of the participants' performances in terms of color reported showed that color errors and 'hallucinations' were evenly spread over time between the two colors, which is inconsistent with the selective attention hypothesis.

In the original paradigm of Chalk et al., participants were not required to identify any other characteristic of the stimulus than its direction of motion. Here, the fact that participants were additionally asked to report on stimulus color, can be argued to introduce a processing component associated with the ventral stream, while the original task was associated only with the dorsal stream. This might have affected how attentional processes were deployed, particularly if attention serves to combine information from the ventral and dorsal streams (Van Der Velde & De Kamps, 2001). This might explain the performance differences we observe with some aspects of the original study. In particular, in Experiment 1, while there was a clear estimation bias towards the frequently presented directions, the detection rates and reaction times did not show a marked difference for those directions as found in Chalk et al. On the contrary, in both Experiments, detection rates were higher near the central direction and decreased away from it (Supplementary Figure 4).

In future work, we plan to further investigate this issue, and its possible impact on learning, by manipulating, instead of motion and color, two stimulus dimensions that are both processed in the dorsal stream, such as motion direction and speed.

Neural encoding

A lot of research has been devoted to modeling how visual neurons could encode information about sensory stimuli in the form of probability distributions, from the population level (Ma, Beck, Latham, & Pouget, 2006) to a single neuron (Deneve, 2008). While many studies have shown that selective attention increases the sensitivity of neurons (Treue & Maunsell, 1996; Treue, 2001; Wolfe & Horowitz, 2004), direct evidence for neural encoding of the prior is still scarce (Summerfield & Koechlin, 2008). It is unclear whether learned priors are encoded directly by gain changes in sensory

neurons as has been observed with attention, and how these changes reach higher cortical areas to affect perceptual behavior. Furthermore, the current study raises the question of whether features that are processed by different cortical pathways integrate and affect perceptual biases, and whether multiple priors can be encoded simultaneously from different sensory features.

The experimental results showed that the brain is endowed with mechanisms which adapt efficiently to different statistical properties of sensory information, and which, given enough time, can learn increasingly complex statistical structures or close approximations. The two models that best described participants' behavior in the two Experiments differed extensively. However, in theory, there should exist a single unifying model that could explain both behaviors and their dependency on the stimulus distributions. The exact form of this model and how it relates to existing theories of how probability distributions are encoded in the brain will be a matter of further investigation.

Conclusions

We presented human participants with two independent motion distributions differentiated by color and we asked whether they will learn the frequently presented directions of each condition and whether knowledge about the statistical properties of the two distributions will transfer between conditions. When one distribution was uniform and the other bimodal, we found that participants learned the statistics of the combined (bimodal) distribution and used that knowledge in both color conditions. When one distribution was bimodal and the other was complimentary of the first one in which the peaks are shifted by 32° , so that the combined distribution was uniform, we found that participants tried to learn the statistics of each distribution but did not clearly apply that knowledge only to the appropriate condition.

Our findings suggest that it is possible to learn the joint statistics of the stimuli by using a feature like color as a distinguishing factor but only under specific conditions. Complexity does not seem to be a limiting factor *per se* as the stimuli of Experiment 2 were more complex than that of Experiment 1. Thus, while it should be harder to learn the statistical information in Experiment 2 than in Experiment 1, there is potentially more to gain in regards to detection performances by learning the statistical information of the more complex stimuli. Further work will be needed to clarify whether more sessions would facilitate learning and whether the transfer of statistical information between conditions in both experiments would decrease with more training, and more importantly, to identify the neural locus and mechanisms of such learning.

Acknowledgments

We would like to thank Dagmara Panas for her work on an earlier version of the experiment that helped us identify potential issues with the experimental paradigm, and how to best address them. NG and MC were supported by funding from the Engineering and Physical Sciences Research Council, the Biotechnology and Biological Sciences Research Council, and the Medical Research Council of Great Britain. NG, MC and PS were also supported by Marie Curie PIRSES-GA-2009-247543 for collaborative visits to UC Riverside. ARS was funded by NSF (BCS-1057625).

Commercial relationships: none.

Corresponding author: Peggy Seriès.

Email: pseries@inf.ed.ac.uk.

Address: IANC, School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh, EH8 9AB, UK.

References

- Adams, W., Graf, E., & Ernst, M. (2004). Experience can change the 'light-from-above' prior. *Nature Neuroscience*, 7(10), 1057–1058.
- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, 14(3), 257–262.
- Blackwell, D. (1979). *Theory of games and statistical decisions*. Dover Pubns.

- Brainard, D. (1997). The psychophysics toolbox. *Spatial vision*, 10(4), 433–436.
- Chalk, M., Seitz, A., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, 10(8).
- Chater, N., Tenenbaum, J., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291.
- Deneve, S. (2008). Bayesian spiking neurons i: inference. *Neural computation*, 20(1), 91–117.
- Dobkins, K., & Albright, T. (1998). The influence of chromatic information on visual motion processing in the primate visual system. *High-level motion processing: Computational, neurobiological, and psychophysical perspectives*, 53–94.
- Ernst, M., & Banks, M. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429–433.
- Fiser, J., & Aslin, R. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological science*, 12(6), 499.
- Fiser, J., & Aslin, R. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99(24), 15822.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3), 119–130.
- Griffiths, T., & Tenenbaum, J. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767.
- Hillis, J., Watt, S., Landy, M., & Banks, M. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12).
- Kording, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247.
- Ma, W., Beck, J., Latham, P., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11), 1432–1438.
- Pelli, D. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision*, 10(4), 437–442.
- Rauber, H. J. & Treue S., (1998). Reference repulsion when judging the direction of visual motion, *Perception*, 27, 393–402.
- Sasaki, Y., Nanez, J., & Watanabe, T. (2009). Advances in visual perceptual learning and plasticity. *Nature Reviews Neuroscience*, 11(1), 53–60.
- Schneider, K., & Komlos, M. (2008). Attention biases decisions but does not alter appearance. *Journal of Vision*, 8(15).
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 461–464.
- Seidemann, E., Poirson, A., Wandell, B., & Newsome, W. (1999). Color signals in area mt of the macaque monkey. *Neuron*, 24(4), 911–917.
- Seitz, A., Kim, R., Wassenhove, V. van, Shams, L., & Seitz, A. (2007). Simultaneous and independent acquisition of multisensory and unisensory associations. *Perception*, 36(10), 1445–1454.
- Summerfield, C., & Koechlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron*, 59(2), 336–347.
- Thiele, A., Dobkins, K., & Albright, T. (2001). Neural correlates of chromatic motion processing. *Neuron*, 32(2), 351–358.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24(5), 295–300.
- Treue, S., Hol, K. & Rauber, H. J. (2000). Seeing multiple directions of motion- physiology and psychophysics. *Nature Neuroscience*, 3:270–276.
- Treue, S., & Maunsell, J. (1996). Attentional modulation of visual motion processing in cortical areas mt and mst. *Nature*, 382(6591), 539.
- Trommershauser, J., Maloney, L., & Landy, M. (2008). Decision making, movement planning and statistical decision theory. *Trends in cognitive sciences*, 12(8), 291–297.
- Turk-Browne, N., Isola, P., Scholl, B., & Treat, T. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(2), 399.

- Turk-Browne, N., & Scholl, B. (2009). Flexible visual statistical learning: Transfer across space and time. *Journal of Experimental Psychology: Human Perception and Performance*, 35(1), 195.
- Van Der Velde, F., & De Kamps, M. (2001). From knowing what to knowing where: Modeling object-based attention with feedback disinhibition of activation. *Journal of Cognitive Neuroscience*, 13(4), 479–491.
- Wandell, B., Poirson, A., Newsome, W., Baseler, H., Boynton, G., Huk, A., et al. (1999). Color signals in human motion-selective cortex. *Neuron*, 24(4), 901–909.
- Wolfe, J., & Horowitz, T. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5(6), 495–501