

The Impact of Interference on Short-Term Memory for Visual Orientation

Rosanne L. Rademaker
Maastricht University

Ilona M. Bloem
Maastricht University and Boston University

Peter De Weerd and Alexander T. Sack
Maastricht University

Visual short-term memory serves as an efficient buffer for maintaining no longer directly accessible information. How robust are visual memories against interference? Memory for simple visual features has proven vulnerable to distractors containing conflicting information along the relevant stimulus dimension, leading to the idea that interacting feature-specific channels at an early stage of visual processing support memory for simple visual features. Here we showed that memory for a single randomly orientated grating was susceptible to interference from a to-be-ignored distractor grating presented midway through a 3-s delay period. Memory for the initially presented orientation became noisier when it differed from the distractor orientation, and response distributions were shifted *toward* the distractor orientation (by $\sim 3^\circ$). Interestingly, when the distractor was rendered task-relevant by making it a second memory target, memory for both retained orientations showed reduced reliability as a function of increased orientation differences between them. However, the degree to which responses to the first grating shifted toward the orientation of the task-relevant second grating was much reduced. Finally, using a dichoptic display, we demonstrated that these systematic biases caused by a consciously perceived distractor disappeared once the distractor was presented outside of participants' awareness. Together, our results show that visual short-term memory for orientation can be systematically biased by interfering information that is consciously perceived.

Keywords: visual short-term memory, orientation, memory bias, interference, memory masking

Supplemental materials: <http://dx.doi.org/10.1037/xhp0000110.supp>

Visual perception is a highly complex process aimed at making sense of a dynamic external world by constructing a coherent visual percept from rapidly changing retinal images. To keep visual information online in order to perform necessary computations, the brain needs to maintain this information after it can no longer be sensed directly, and inoculate it against interference from other inputs. Visual “short-term” or “working” memory serves as a highly efficient buffer that temporarily stores sensory informa-

tion for future use. Research aiming to better characterize visual working memory often focuses on one particularly captivating feature of the system, which is its profound capacity limitation, composed of about three to four items (Bays & Husain, 2008; Fournie, Suchow, & Alvarez, 2012; Fukuda, Awh, & Vogel, 2010; Luck & Vogel, 1997; Palmer, 1990; van den Berg, Shin, Chou, George, & Ma, 2012; Wilken & Ma, 2004; Zhang & Luck, 2008). However, an emphasis on *quantity* is often confounded with limitations equally applicable to encoding and perception in general (Gazzaley & Nobre, 2012; Palmer, 1990), and does little to inform memory *quality*. What happens to the quality of visual memories once they have been well and truly transferred into internal representations?

Traditionally, research into memory quality falls under the label “short-term memory”, and a core question concerns the extent to which new information from the eyes has the potency to interfere with information already in memory. One way to investigate this is by keeping the stimuli at the encoding stage constant, as well as keeping set size within the confines of supposed (cognitive) capacity limitations. Doing precisely this, early psychophysical work into visual short-term memory has shown that memories are not immune against interference from other stimuli: When people remembered a single spatial frequency (Bennett & Cortese, 1996; Magnussen, Greenlee, Asplund, & Dyrnes, 1991; Nemes, Whitaker, Heron, & McKeefry, 2011), direction of motion (Magnussen & Greenlee, 1992; McKeefry, Burton, & Vakrou, 2007; Pasternak

This article was published Online First August 10, 2015.

Rosanne L. Rademaker, Cognitive Neuroscience Department, Maastricht University; Ilona M. Bloem, Cognitive Neuroscience Department, Maastricht University, and Department of Psychological and Brain Sciences, and the Center for Computational Neuroscience and Neural Technology, Boston University; Peter De Weerd and Alexander T. Sack, Cognitive Neuroscience Department, Maastricht University.

We thank Sam Ling for contributing the foundation to the mouse-probe code, forgetting about this contribution, and for generally being awesome; Caroline Benjamins for technical assistance; Ruben van Bergen and Pascal Mamassian for helpful discussions; Jan Brascamp for his input regarding the design of Experiment 4; and Daryl Fournie for feedback on the MemToolbox.

Correspondence concerning this article should be addressed to Rosanne L. Rademaker, Cognitive Neuroscience Department, Maastricht University, Oxfordlaan 55, 6229 EV Maastricht, The Netherlands. E-mail: rosanne.rademaker@maastrichtuniversity.nl

& Zaksas, 2003), or color (Nemes, Parry, Whitaker, & McKeefry, 2012; Nilsson & Nelson, 1981), a second irrelevant stimulus shown during the retention interval degraded performance on a delayed discrimination task. Interference effects (also referred to as “memory masking”) have also been found when a to-be-ignored stimulus was presented shortly before the memory task (Lalonde & Chaudhuri, 2002), and in other modalities such as short-term memory for pitch (Deutsch, 1970, 1973).

Visual memory for simple features proved vulnerable only to distractors containing conflicting information along the relevant stimulus dimension (Magnussen et al., 1991; Magnussen & Greenlee, 1992; Magnussen & Greenlee, 1999). For example, in the case of a remembered spatial frequency, this meant that an irrelevant but different spatial frequency affected thresholds, whereas an irrelevant orientation change—without a change in spatial frequency—did not (Lalonde & Chaudhuri, 2002; Magnussen et al., 1991; Nemes et al., 2011). The fact that interference acts independently on low-level features implicates mid-level visual areas beyond primary visual area V1 as a locus of interference (Magnussen, 2000; Magnussen & Greenlee, 1999). This idea is further supported by findings demonstrating that interference obeyed size constancy (Bennett & Cortese, 1996), was location invariant (Nemes et al., 2011, 2012; Phillips, 1974), and still evident under free-viewing conditions (Magnussen et al., 1991). Recent neuroimaging work confirms the likely locus of interference at a middle level of the visual hierarchy, strongly implicating V4 in particular (Sneve, Sreenivasan, Alnæs, Endestad, & Magnussen, 2015).

These findings have spawned the idea that memories are stored in narrowly tuned feature-specific channels in visual cortex, where inhibitory cross-channel interactions are responsible for the psychophysically observed distractor effects (Magnussen, 2000; Magnussen & Greenlee, 1992, 1999; Nemes et al., 2011). In this view, information loss is due to inhibition between different memory stores maintaining conflicting information of a shared visual feature. This idea aligns with the observation that an irrelevant stimulus did not impact performance when it matched the memorized stimulus on the task relevant feature, and that the most prevalent disruption occurred when a distractor differed from the memory target by one octave of more (in the case of spatial frequency), or by about twice the velocity (in the case of direction of motion) (Lalonde & Chaudhuri, 2002; Magnussen et al., 1991; Magnussen & Greenlee, 1992).

More recently it was discovered that the deleterious effects of a distractor on memory did not result from a drop in memory fidelity, but from an attraction of representations in memory toward the distractor (Huang & Sekuler, 2010a). When participants viewed two subsequently presented gratings of different spatial frequencies, having to report only one of them, it was found that the reported spatial frequency was biased toward the spatial frequency of the unreported grating, whereas the variability of report remained unchanged (Huang & Sekuler, 2010a). Earlier work relying on delayed discrimination tasks had been unable to uncover memory attraction, as such tasks index memory quality or fidelity by a single threshold measure. Instead, this study employed a method of adjustment procedure, allowing a measure of response variability as well as a measure of the response mean (or central tendency). Critically, attraction was stronger when two spatial frequencies were task relevant compared with when one of

the two was irrelevant and could be ignored (Dubé, Zhou, Kahana, & Sekuler, 2014; Huang & Sekuler, 2010a).

Based on these findings a modified version of the channel interaction account emerged, in which a second stimulus exerts its influence at a visual stage of processing via (incomplete) perceptual averaging. Such averaging presumably occurs in a population of spatial frequency selective mechanisms. In this view, the degree to which two stimuli are averaged depends on their respective weights, and selective attention might alter these weights such that the influence of irrelevant information can be mostly filtered out (Dubé et al., 2014; Huang & Sekuler, 2010a). A number of key predictions follow from this account: The more a distractor differs from an item in memory, the more strongly it will attract memory due to averaging. Attraction should be reduced for irrelevant and unattended distractors by means of attentional filtering. It is furthermore implied that the influence of the second grating will always be one of attraction, as illustrated by a magnet metaphor (Huang & Sekuler, 2010a).

Although attraction has been shown in the context of spatial frequency (Dubé et al., 2014; Huang & Sekuler, 2010a; Nemes et al., 2011) and color (Nemes et al., 2012), it remains unclear whether and how a distractor might bias memory representations of orientation. Both short- and long-term memory for orientation, draw upon visual cortical regions (Bosch, Jehee, Fernández, & Doeller, 2014; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009; Sneve, Alnæs, Endestad, Greenlee, & Magnussen, 2012), making it likely that orientation, like other low-level features, is susceptible to interference. Such susceptibility should come as no surprise, considering that interfering visual information necessarily enters the same sensory regions as those responsible for memory maintenance.

We designed our study to investigate memory interference for orientation under a variety of circumstances, testing the channel interaction (Magnussen, 2000) and perceptual averaging (Dubé et al., 2014; Huang & Sekuler, 2010a) theories proposed by previous work. First, it is currently not known whether memory biases other than attraction exist. Specifically, changes in memory variability or “noise” due to a distractor have not been demonstrated before, but an increase in noise logically follows if one adopts the channel interaction account (Magnussen, 2000). Conversely, theories such as optimal cue integration (Ernst & Banks, 2002) predict a noise reduction, since integration of a distractor into a target representation would presumably cause a decrease in variance. Another open question is whether or not attraction might depend on the range of memory target and distractor differences used within a single experiment. We explored whether attraction might become stronger, returns to baseline, or scales with a larger range of target–distractor differences (Experiments 1 and 2), as the current evidence is conflicting (Huang & Sekuler, 2010a; Magnussen et al., 1991; Magnussen & Greenlee, 1992; Nemes et al., 2011, 2012). We furthermore contrasted interference from both task relevant and irrelevant distractors (Experiments 1 and 2 vs. Experiment 3). A perceptual averaging mechanism (Dubé et al., 2014; Huang & Sekuler, 2010a) assumes less interference from task irrelevant information, as attention acts to largely filter out a distractor’s influence. Conversely, the channel interaction account (Magnussen, 2000) is agnostic to the relevance of the interfering information. How well these accounts hold for orientation memory remains to be seen. Finally, to address why irrelevant information

would be integrated in the first place we looked at whether mere bottom-up information that is not processed consciously would be sufficient to interfere with information in memory (Experiment 4).

To answer these questions it was vital to parse more general performance changes into independent contributions. By using the method of adjustment we were able to construct error distributions, the shape of which can disclose fundamental mechanisms behind changes in memory performance (Ma, Husain, & Bays, 2014). For example, studies quantifying working memory limitations routinely rely on error distributions to infer information about whether or not an item is in memory (Wilken & Ma, 2004; Zhang & Luck, 2008), whether people might inadvertently report the wrong item (Bays, Catalao, & Husain, 2009), or whether memories are variable from trial-to-trial and item-to-item (Fougnie et al., 2012; van den Berg et al., 2012). Here we exploit this methodology to investigate memory quality, fitting a circular Gaussian (or “von Mises”) distribution to retrieve the noisiness of memory (indexed by the standard deviation, or SD) and the distribution mean (μ), allowing a more comprehensive insight into the dynamics underlying memory interference. Additionally, because we expected memory attraction to shift the distribution mean, we compared the fits from a von Mises with and without a parameter describing the distribution mean. Such a comparison provides us with additional information, namely, instead of examining whether a shift exists, it assesses whether assuming such a shift helps explain the data.

Here we followed the psychophysical short-term memory tradition of presenting only a single stimulus that is then translated into a high fidelity memory, and looked at its robustness against interference. We showed that short-term memory for orientation is not immune against interference, and systematic biases emerge when a distractor orientation is presented that differs from an orientation held in memory. These biases required awareness and consisted of increased memory noise, which has not been demonstrated before, and a shift in the response distribution toward the distractor orientation or “attraction.” The range of target–distractor differences did not impact these biases, whereas attention reduced memory attraction, even leading to occasional instances of memory repulsion. These findings are not predicted by previous theories, and require a revision of current models on memory quality and interference.

Experiment 1

Method

Participants. Eight healthy volunteers (6 female) between the ages of 21 and 29 ($M = 25.38$; $SE = 1.12$) participated in Experiment 1. With the exception of two of the authors (Rosanne L. Rademaker and Ilona M. Bloem), participants received monetary reimbursement for their time and were naive to the purpose of the study. Participants in all of the experiments described in this paper had normal or corrected-to-normal vision, and provided informed consent. All experiments took place under the approval of the standing ethical committee of the Psychology and Neuroscience Department at Maastricht University.

Stimuli. All experimental stimuli were viewed in a dark room on a luminance-calibrated CRT monitor with $1,280 \times 1,024$ resolution and 60-Hz refresh rate. Visual stimuli were generated using MATLAB 7.5.0 (R2007b) and the Psychophysics Toolbox (Brain-

ard, 1997; Pelli, 1997) on a PC running Windows XP. Stimuli consisted of centrally presented oriented gratings with a spatial frequency of $2\text{ c}/^\circ$, and a diameter of 3° of visual angle. Gratings were presented at 20% Michelson contrast with added jitter (randomly selected from a uniform distribution with a range spanning $\pm 10\%$ contrast), within a wide Gaussian envelope ($SD = 2.5^\circ$) on a uniform gray background that shared the same mean luminance of $40.8\text{ cd}/\text{m}^2$. Grating phase was randomized between 0 and 2π . The test stimulus used to obtain participants’ responses was a mouse-probe consisting of the centrally presented white bull’s eye fixation (0.5° of visual angle in diameter) and an interrupted white line, of which each segment was 0.025° wide and 0.125° long. The two line segments were spaced 3° apart to ensure that their visual field position was nonoverlapping with that of the previously presented grating. By moving the mouse, the interrupted white line rotated about the fixation bull’s eye, allowing participants to replicate the orientation in memory by method-of-adjustment. Participants were seated at a viewing distance of 57 cm, and a chinrest assisted in maintaining head stability. Participants were instructed to maintain steady fixation throughout all experimental trials.

Procedure. Throughout all the experiments described here, the general outline of the task was the same (see Figure 1 for reference). First, a target grating with a randomly chosen orientation between 1° and 180° was presented for 200 ms, and participants remembered the orientation of this grating. After the retention-interval a test (mouse-probe) was presented at an initially random orientation also between 1° and 180° , and participants rotated this dial to match the orientation in memory. Once a participant was satisfied with the response, a left mouse-click allowed them to continue to the next trial. Precision of replication-performance was stressed throughout all experiments described in this article, and there were never any time constraints for participants’ responses.

We first established a baseline performance for this particular method-of-adjustment probe when a single orientation was memorized. In two separate blocks of 200 trials each, participants remembered a randomly oriented grating for 1, 3, 6, or 12 s




Figure 1. Trial sequence, Experiment 1. Participants viewed a randomly orientated grating for 200 ms at the start of each trial. They retained the orientation in memory over a 3-s interval while fixating a white bull’s eye presented against a mean-gray background. Midway through the retention interval, a second, to-be-ignored grating was presented for 200 ms. The orientation of this to-be-ignored distractor grating could be the same as the orientation in memory, or it could be rotated 2° , 4° , 7° , or 15° counter-clockwise or clockwise relative to the orientation in memory (randomly interleaved). After the retention interval participants were presented with a test stimulus, which they could rotate by using the computer mouse to match the orientation in memory as precisely as possible. When satisfied with their response, participants clicked the mouse and continued to the next trial 1 s later.

(randomly interleaved), after which they rotated the dial to match the orientation in memory as closely as possible. The trials with a retention interval of 3 s were subsequently used as the “no-distractor” baseline for the rest of the experiment.

For the main part of the experiment a second distractor grating was introduced, and presented for 200 ms halfway through a fixed 3-s retention interval (see Figure 1). The orientation of the distractor could be one of several orientations that were fixed relative to the target. These relative orientations were sampled in a Gaussian ($\delta = 25$) fashion, resulting in denser sampling around the target orientation. The distractor orientation could differ -15° , -7° , -4° , -2° , 0° , 2° , 4° , 7° , or 15° from the randomly chosen target orientation. Thus, the distractor grating could appear at an orientation that was counter-clockwise, the same, or clockwise relative to the target (randomly interleaved). Participants were told that the second grating was completely irrelevant to the task, and were instructed to ignore it. This part of the experiment consisted of 900 trials in total, divided over five blocks (~ 22 min per block). Participants were allowed a short practice before the start of the experiment. Half of them started with the two baseline blocks, and the other half started with the main experiment.

Analyses. For each condition of interest, a distribution of response errors was obtained by calculating the difference between target and response (reported orientation minus target orientation). Memory accuracy was the average (absolute) orientation-error in each condition. In order to look beyond simple accuracies, and to take the entire response distribution into account, we also estimated relevant characteristics from these response distributions by fitting a von Mises function (circular analog of a normal distribution) to the response distributions for the experiments described in this article. A von Mises describes the data in terms of the mean (μ) and circular variance (SD).

Data analysis were performed in MATLAB using custom functions as well as functions provided by the MemToolbox (Suchow, Brady, Fougine, & Alvarez, 2013), and the Circular Statistics Toolbox (Brens, 2009). Here, we used maximum likelihood estimation to obtain estimates for each parameter value (on which we performed regular repeated-measures statistics). Additionally, we used the Bayesian information criterion (BIC; Schwarz, 1978) to compare models with and without a distribution mean as a free parameter.

Results and Discussion

We tested whether an irrelevant distractor presented during retention could systematically affect memory of a single target by parametrically varying the orientation of the distractor relative to the orientation of the target. A within-subject analysis of variance (ANOVA) revealed that the absolute response error was not the same across the relative orientation differences between the target and distractor gratings (Figure 2A), $F(8, 56) = 2.286$, $p = .034$. This trend was quadratic, $F(1, 7) = 10.559$, $p = .014$, indicating that larger relative orientation differences between target and distractor led to bigger performance decrements, compared with smaller relative differences. This quadratic effect is reflected by the characteristic “V-shape” in the data. Additional post hoc tests (paired t tests) did not reveal a difference between the no-distractor baseline and any of the target–distractor conditions (all $ps > .149$), implying that the quadratic effect is a mixture of improved mem-

ory performance at small, and impaired memory performance at large target–distractor differences.

To examine why performance suffered when the distractor orientation differed more from the memorized target orientation, we fit a von Mises to estimate the degree of variability in the report of the target (see Figure 2B). There was a strong trend implicating that memory noise varied as a function of target–distractor difference, $F(8, 56) = 2.072$, $p = .054$; this trend was quadratic, $F(1, 7) = 5.214$, $p = .056$, implying noisier memory representations when the target and distractor orientations differed more. Post hoc t tests showed no deviations from baseline (all $ps > .14$).

In addition to a marginally noisier memory, larger orientation errors arose because of a shift of the response-error distribution’s mean (Figure 2C), $F(8, 56) = 12.404$, $p < .001$. This effect was linear, $F(1, 7) = 22.475$, $p = .002$, demonstrating that when the distractor orientation was rotated relative to the orientation in memory, participants’ responses were shifted in the direction of the rotation. For a schematic overview of this “attraction” effect, also see the General Discussion. Post hoc t tests demonstrated that the distribution mean was significantly different from the distribution mean during baseline trials when the target and distractor differed by -15° , -4° , and 7° ($p = .01$, $p = .013$, and $p = .007$, respectively).

We compared the fits from a von Mises with and without a parameter describing the distribution mean (Figure 2D). BIC difference values (“no bias” minus “with bias”) of >0 indicated that a von Mises with bias term better fits the data than a von Mises without bias. Figure 2D shows that a von Mises with a bias term best described our data (testing Δ BIC against 0), $F(1, 7) = 5.722$, $p = .048$, although we did not find statistical support for the idea that it does so more for larger relative orientation distances between the target and distractor (testing whether Δ BIC differs for different target–distractor conditions), $F(8, 56) = 1.154$, $p = .343$.

Experiment 2

Method

Participants. Participants in Experiment 2 were eight volunteers (seven female), ages 21 to 33 years old ($M = 24.75$; $SE = 1.57$), four of whom had already participated in Experiment 1. Participants were naive to the purpose of the study and were reimbursed for their time (with the exception of two of the authors, Rosanne L. Rademaker and Ilona M. Bloem).

Stimuli and procedure. Experiment 2 was virtually identical to Experiment 1, with two minor exceptions: First, the range of possible distractor orientations around the target was wider, spanning 90° in total. The distractor orientation could differ from the target orientation by -45° , -30° , -15° , -7° , 0° , 7° , 15° , 30° , or 45° and these conditions were presented in a randomly interleaved fashion. Second, instead of measuring a baseline in separate blocks as was done in Experiment 1, here we randomly interleaved trials without a distractor during retention. In total we collected 1,000 trials divided over five blocks (~ 25 min per block).

Results and Discussion

Here we tested whether larger target–distractor differences would result in even larger shifts of the response distribution

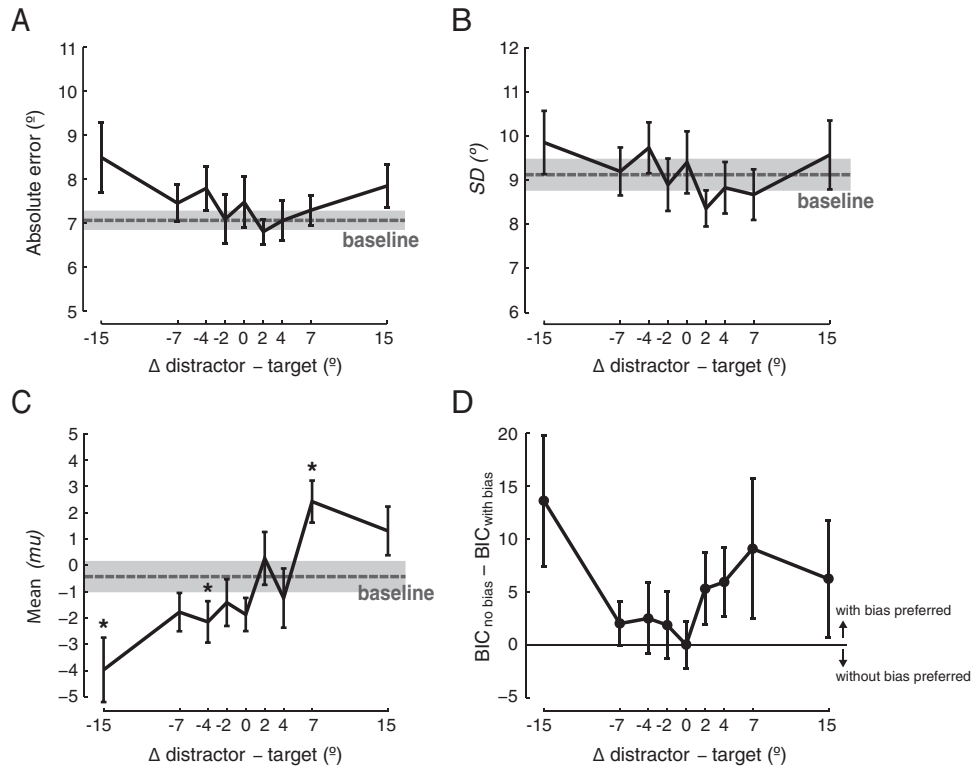


Figure 2. Results, Experiment 1. Panel A: When the distractor orientation differed more from the remembered orientation (be it counterclockwise or clockwise), the absolute error went up relative to when the distractor orientation differed less from the remembered orientation. Panel B: The loss of accuracy at larger relative orientation differences between the target and distractor grating could be in part explained by a similar trend in memory precision (increase in *SD*). Panel C: The increase in the absolute error could furthermore be attributed to a shift in the entire response distribution *toward* the orientation of the distractor grating. This shift was more prominent on trials where the distractor differed more from the target, as indicated by a linear trend. Panel D: A von Mises with a bias term, modeling the distribution mean (*μ*) in addition to the distribution variability (*SD*), fit the data better than a von Mises without a bias term. This advantage did not become significantly stronger at larger relative orientation differences between the target and irrelevant distractor. Group-averaged data are plotted in solid black lines with error bars representing ± 1 *SEM*; gray dashed lines and shaded regions indicate the mean (± 1 *SEM*) on baseline trials during which no distractor was presented during the 3-s retention interval—obtained during separate experimental blocks. Asterisks indicate a significant difference from baseline trials. BIC = Bayesian information criterion.

toward the distractor orientation, or might alternatively return to baseline, by expanding the range of relative orientation differences with respect to the range used in Experiment 1. Replicating our previous findings, Figure 3A shows that with the wider range of orientation differences, the absolute response error differed at various relative target–distractor orientation differences, $F(8, 56) = 3.899$, $p = .001$. As in Experiment 1, this effect was V-shaped, or quadratic, $F(1, 7) = 7.966$, $p = .026$, indicating bigger performance decrements when the relative orientation difference between target and distractor was larger, compared with when it was smaller. Post hoc paired *t* tests show that this was primarily due to the contrast between relatively small (i.e., 0° and 7° differences) compared with relatively large (i.e., 15°, 30°, and 45° differences) target–distractor differences (p values between .001 and .308; $p = .09$ on average).

A von Mises was fit to examine the respective contributions of memory variability (Figure 3B) and mean response (Figure 3C). Memory noise differed across the various target–distractor con-

ditions, $F(8, 56) = 3.872$, $p = .001$, and a quadratic trend indicated that memory became noisier as the target and distractor differed more, compared with when they differed less, $F(1, 7) = 6.539$, $p = .038$. Post hoc paired *t* tests showed that contrasting relatively small (i.e., 0° and 7°) with relatively large (i.e., 15°, 30°, and 45°) target–distractor differences generally accounts for this V-shape (p values between .016 and .48; $p = .121$ on average). Finally, memory noise during trials with a distractor did not differ from no-distractor baseline trials, $F(1, 7) = 0.050$, $p = .829$, implying that the V-shaped effect reflected a mixture of enhanced precision at smaller target–distractor differences, and reduced precision at larger target–distractor differences.

The mean response was shifted for various target–distractor difference conditions, $F(8, 56) = 10.589$, $p < .001$. This linear effect, $F(1, 7) = 20.421$, $p = .003$, indicated that, as in Experiment 1, the distractor orientation attracted the representation of a single orientation held in memory.




Figure 3. Results, Experiment 2. Panel A: Participants made larger errors replicating an orientation in memory when the orientation of an irrelevant distractor grating differed more from the memorized orientation. Panel B: Memory was noisier for the remembered orientation (larger *SD*) at larger relative target–distractor orientations, compared with less noisy memory (smaller *SD*) at smaller relative target–distractor orientations. Panel C: The error-response distribution shifted *toward* the orientation of the distractor grating; this shift was larger when the difference in orientation between the two was larger. Panel D: Despite a shift in the mean of the response distribution, our data were no better described by a von Mises that included this shift parameter, nor was there any evidence suggesting that a von Mises including a bias term fit the data better at larger relative orientation differences. Group-averaged data are plotted in solid black lines with error bars representing ± 1 *SEM*; Gray dashed lines and shaded regions indicate the no-distractor baseline mean ± 1 *SEM* derived from randomly interleaved trials. Asterisks indicate conditions for which there was a significant difference with the no-distractor baseline. BIC = Bayesian information criterion.

Figure 3D shows that including a bias term to the von Mises did not result in significantly better data fits than not including such a bias term (testing Δ BIC against 0), $F(1, 7) = 3.146$, $p = .119$. Despite a shift in responses toward the distractor orientation (Figure 3C), adding a bias parameter did not help describe our data better at larger relative orientation differences, $F(8, 56) = 0.890$, $p = .531$.

We did not find that the wider range of orientation differences presented here changed the attraction signature found in Experiment 1. First, we did not observe that attraction returned to baseline at a 45° target–distractor difference. Second, there was no change in the magnitude of the distribution shifts between the first two experiments, $t(14) = 0.329$, $p = .747$: In Experiment 1 the distribution of response errors shifted from an average $\mu = -3.98^\circ$ (when the distractor was rotated 15° clockwise relative to the memory target) to an average $\mu = 2.424^\circ$ (when the distractor was rotated 7° relative to the memory target), resulting in a maximum observed shift of 6.4°. In Experiment 2 the maximum shift observed was 5.822° (from -2.852° to 2.971° in the -45° and 45° difference conditions, respectively).

In addition, we found that memory noise did not differ between the two experiments—repeated-measures ANOVA with a between-subjects factor: $F(4, 56) = 0.843$, $p = .504$ —when comparing the conditions both had in common (-15° , -7° , 0° , 7° , and 15° target–distractor differences). We furthermore calculated the maximum impact that the distractors had on precision (largest *SD* minus smallest *SD*), and found that the maximum impact in Experiment 1 (of 1.5° between target–distractor conditions of -15° and 2°) did not differ from the maximum impact in Experiment 2 (of 2.32° between target–distractor conditions of 0° and 45° difference), $t(14) = 0.654$, $p = .524$.

Experiment 3

Method

Participants. Eight participants (5 female) took part in Experiment 3 (ages between 21 and 34 years, $M = 26$; $SE = 1.64$), including four volunteers who had previously participated in Experiments 1 and 2, and one who had previously participated in

Experiment 2 only. With the exception of two of the authors (Rosanne L. Rademaker and Ilona M. Bloem), participants were reimbursed for their time and naive to the purpose of the experiment.

Stimuli and procedure. Stimuli and procedures in Experiment 3 (see Figure 4) were identical to those from Experiment 2, apart from three exceptions. We reduced the possible orientation differences between the two gratings to -45° , -15° , 0° , 15° , and 45° . Furthermore, participants were no longer instructed to ignore the second stimulus, but instead had to attend and remember it. The two stimuli had an equal probability of being cued once the test display was presented, with the cue being a black number ('1' or '2') displayed at fixation for as long as the test stimulus was on the screen. Third, because both stimuli were always task relevant we did not obtain a baseline during which one of the two stimuli was not present. In total, participants performed 1,000 trials divided over five blocks (~24 min each).

Results and Discussion

Here we investigated what would happen to the V-shaped and attraction effects found in the previous two experiments, when the second grating was task relevant. We also looked at how memory for the second orientation was affected by the first, and whether its representation would be prone to the same biases. Replicating our previous experiments, we found that participants' responses were more erroneous as the difference between the first and second grating was larger (Figure 5A), $F(4, 28) = 12.881$, $p < .001$, with a quadratic trend, $F(1, 7) = 43.567$, $p < .001$. This effect was due solely to the difference between conditions during which the target and distractor did not, and did share the same orientation (i.e., Δ nonprobed minus probed of 0° vs. all other Δ nonprobed minus probed). Specifically, when the target and distractor differed by 45° the errors were larger than when the target and distractor differed by 0° (all paired t test $p < .024$), and a similar trend was found comparing the conditions with 15° versus 0° differences (p values between .011 and .066). None of the conditions where the target and distractor differed from one another revealed any differences in the absolute error (all $ps > .284$). In other words, the V-shaped effect leveled off once the target and distractor started to differ. Finally, accuracy was better for the second grating compared with the first, $F(1, 7) = 6.05$, $p = .043$.

Memory variability (or SD ; Figure 5B) mirrored the effects found in the absolute errors, indicating noisier memory when the two orientations differed more compared with when they differed less, $F(4, 28) = 10.547$, $p < .001$, with a quadratic trend, $F(1, 7) = 44.893$, $p < .001$, and noisier memory for the orientation that was presented first, $F(1, 7) = 46.413$, $p < .001$. Post hoc t tests showed that when target–distractor differences existed (i.e., by 15° or 45°) memory was noisier than when target and distractor were of the same orientation (i.e., 0° difference; all $ps < .058$). Comparing all conditions where target–distractor differences existed yielded only one instance where the larger target–distractor difference resulted in more memory variability than the smaller difference (for Grating 2, comparing the -15° and 45° conditions: $p = .011$), whereas all other comparisons did not (all $ps > .069$).

Memory for the first orientation was attracted toward the (now relevant) second orientation (solid black line in Figure 5C), $F(4, 28) = 2.79$, $p = .046$; linear trend, $F(1, 7) = 5.711$, $p = .048$. However, the maximum distribution shift of 2.056° (defined as the largest clockwise shift minus the largest counterclockwise shift) was much reduced compared with Experiment 1, $t(14) = 3.4$, $p = .004$, and Experiment 2, $t(14) = 2.503$, $p = .025$, where the maximum shifts were 6.4° and 5.822° , respectively. Distribution means for the second target (dashed black lines in Figure 5C) show a different shift from those of the first target, $F(4, 28) = 4.647$, $p = .005$. In fact, there was a trend indicating that memory for the second target was shifted away from the orientation of the first target, $F(4, 28) = 1.808$, $p = .155$.

The distribution of responses to the first target was equally well described by a von Mises with and without bias term (testing ΔBIC against 0), $F(1, 7) = 4.244$, $p = .078$, and this did not change across the various orientation differences between the two memory items, $F(4, 28) = 1.018$, $p = .415$. However, there was a clear benefit to include a bias term for the second target (testing ΔBIC against 0), $F(1, 7) = 5.200$, $p = .057$, which was not the same at the various relative orientation differences between the two memory items, $F(4, 28) = 4.936$, $p = .004$. Adding a bias parameter improved the fit for conditions during which the two targets differed in orientation, compared with when they did not differ (paired t tests of $\Delta 0^\circ$ against -45° , -15° , 15° , and 45° difference conditions: $p = .016$, $p = .031$, $p = .036$, and $p = .069$, respectively).

Adding a bias parameter helped fit the data for the second target (Figure 5D), whereas a tentative distribution shift away from the first orientation (Figure 5C) did not reach significance. Exploring individual participant biases to the second target (see the online supplemental material for Figure S1) uncovered that the majority of participants ($N = 5$) showed "repulsion" with reports shifted away from the orientation of the first target (by 9.13° on average; $SE = 1.934^\circ$), and the others ($N = 3$) showed an attraction (of 4.59° on average; $SE = 1.133^\circ$). Thus, all participants had a certain degree of bias in their responses to the second target, but the direction of that bias differed between individuals.

Why might performance be better for the second grating compared with the first? One possibility is that the task on the second grating was easier: once the first grating was presented and its orientation was known, the observer was inadvertently provided with information about the orientation of the second grating. Because the second grating could only have a limited number of orientations relative to the first, participants could have (implicitly)




Figure 4. Trial sequence, Experiment 3. Participants remembered two gratings, both presented for 200 ms each. The first was randomly orientated and retained for 3 s. The second had an orientation that was rotated -45° , -15° , 0° , 15° , or 45° relative to the first grating (randomly interleaved) and was retained for 1,400 ms. A central number cue (not veridically depicted here in the interest of legibility) indicated which grating to report, and a test stimulus could be rotated by using a mouse to match the cued orientation in memory as precisely as possible.

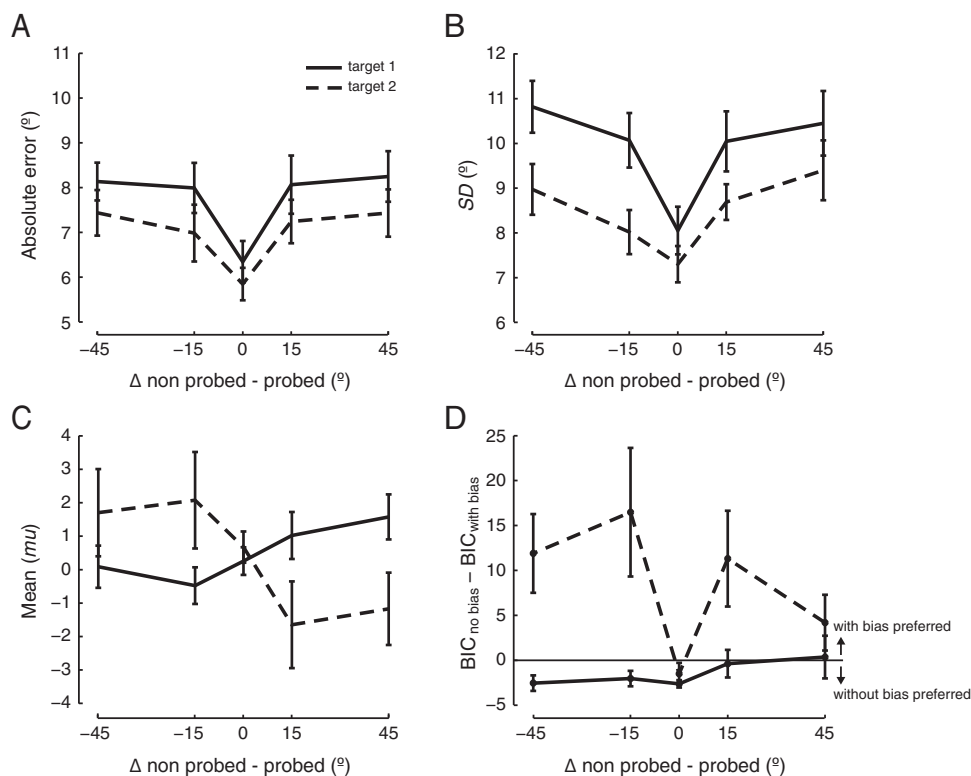


Figure 5. Results, Experiment 3. Panel A: Replication errors were larger when two (remembered) orientations differed from one another. Moreover, performance was better for the orientation that was presented last. Panel B: The effects in Panel A were mirrored by the variability of a von Mises that was fit to the distribution of error-responses. Panel C: The error distribution for responses to the first target shifted *toward* the orientation of the second target, but this attraction was smaller than the attraction in previous experiments where the second orientation was ignored. The error distribution for responses to the second target appeared to be shifted *away* from the orientation of the first target, but this effect did not reach statistical significance. Panel D: Responses to Target 1 were equally well described by a von Mises with and without bias term. Responses to Target 2 were better fit when a bias term was included, and this benefit was more apparent when the two remembered orientations differed from one another. Group-averaged data for Target 1 and 2 is plotted in solid black and dashed black lines, respectively. Error bars represent ± 1 SEM. BIC = Bayesian information criterion.

learned this relationship over time. However, this explanation is unlikely, as we found no evidence that memory for the second grating improved as participants were exposed to more experimental trials (see the online supplemental material for Figure S2). Alternatively, the benefit for the second target could have reflected the shorter retention duration over which it was remembered. From the baseline data obtained as part of the first experiment, we observed a 0.426° ($SE = 0.218^{\circ}$) accuracy reduction when information had to be remembered for 2 more seconds (comparing a 1- and 3-s retention interval). In this experiment the difference in retention intervals between the two targets was slightly shorter, (1.6 s), although the loss of accuracy was slightly larger ($M = 0.768^{\circ}$; $SE = 0.312^{\circ}$). Although it is difficult to directly compare data from different participants in different experiments, it is conceivable that temporal decay played a role in the performance differences between the two targets. Finally, there is evidence from other studies suggesting that the last item of a series of sequentially presented—or fixated—stimuli is assigned more resources than previous items, making its representation in memory more precise in a way that cannot be explained by temporal decay (Bays &

Husain, 2008; Gorgoraptis, Catalao, Bays, & Husain, 2011; Warden & Miller, 2007; Zokaei, Gorgoraptis, Bahrami, Bays, & Husain, 2011). Such an account could also explain the benefit enjoyed by the second grating in this experiment—the last in the series of two to-be-remembered stimuli.

Experiment 4

Method

Participants. For Experiment 4 we tested 10 participants (ages between 23 and 31 years, $M = 24.3$; $SE = 0.79$; 5 female). All 10 participants were naive to the purpose of the experiment, though one of them participated in all of the previous experiments reported here, and two others had previously participated in Experiment 3 only. Participants all received monetary reimbursement for their time.

Stimuli and procedure. Unique to Experiment 4 was that stimuli were viewed dichoptically through a mirror-stereoscope, presenting each of the two eyes with separate and independent

information. Displaying stimuli dichoptically allowed us to capitalize on a phenomenon called binocular rivalry. This entails that two different images shown one to each eye will rival with each other, resulting in only one of the two reaching awareness while the other is suppressed.

Participants' eye-dominance was determined before the start of the experiment by a procedure that matched the relative strength of two rivaling images by adjusting their respective contrasts (adapted from: Pearson, Clifford, & Tong, 2008; Pearson, Rademaker, & Tong, 2011). Six out of the 10 participants tested here were left-eye dominant.

In general, the outline of the task in this experiment (Figure 6A) was identical to that from Experiments 1 and 2, with any events occurring during the retention interval being ignored. Stimuli were identical to those used in prior experiments with two exceptions: The first was that in addition to a second distractor grating (oriented -45° , -15° , 0° , 15° , or 45° relative to the target) also a circular 3°-diameter full-contrast checkerboard could be presented midway through the retention interval. The second exception was that, while the target was still presented at 20% Michelson contrast, the distractor was presented at 10% Michelson contrast (both $\pm 10\%$ random uniform jitter) resulting in the distractor grating always having lower contrast than the target. This was done to increase the chance that the distractor would be suppressed when a checkerboard was simultaneously presented to the other eye. To further maximize the chances that the distractor grating would be rendered invisible when presented alongside a checkerboard, we stacked the odds in favor of the checkerboard by always presenting it to the dominant eye—where it was more likely to be consciously perceived. The distractor grating was ipso facto presented to the nondominant eye where it was more likely to be suppressed. This

also meant that the target grating was always presented to the nondominant eye (as was the mouse probe).

Five types of events (Figure 6B), lasting 200 ms, could occur midway through the retention interval. One, a distractor grating could be presented to the nondominant eye (same eye as target). Two, a distractor grating could be presented to the dominant eye (different eye from target). Three, a distractor grating could be presented to the nondominant eye (same eye as target) while a checkerboard was simultaneously presented to the dominant eye (different eye from target). Four, a checkerboard could be presented to the dominant eye (different eye from target). And five, no intervening stimulus could be presented midway through the retention interval (true baseline).

Note that the first and second event types were perceptually identical, since participants could not resolve which eye the stimuli originated from when viewing them through a mirror-stereoscope. These events thus replicated the perceptual experiences from participants in Experiments 1 and 2. Note that the third and fourth event types were perceptually identical because the checkerboard was designed to always dominate perception, which meant that even if a distractor grating was presented simultaneously to the other eye, only the checkerboard was consciously perceived. To ensure the latter was true, we included a "visibility check" after all trials during which a checkerboard was presented. We asked participants to indicate on a 4-point scale what they had seen (1 = *only the checkerboard*; 2 = *maybe something else*; 3 = *something else but unclear what*; 4 = *something else and its orientation*). Results from the visibility check showed that participants were indeed unable to perceive the distractor grating when presented dichoptically alongside the checkerboard. A rating of "4" was never given, and a rating of "3" was given only once throughout the entire experiment (this trial was removed and repeated




Figure 6. Trial sequence, Experiment 4. Panel A: Participants remembered a randomly orientated grating presented to the nondominant eye, ignored any events occurring during the retention phase, and replicated the orientation in memory 3 s later by rotating a test stimulus. Panel B: Five types of 200-ms events could occur during the retention phase: A distractor grating was visibly presented to the same eye or to the other eye, relative to the eye that the target was presented to. These first two event types were perceptually identical. Thirdly, a distractor grating was presented to the same (nondominant) eye and rendered invisible by a full-contrast checkerboard presented simultaneously to the other (dominant) eye (see also the example trial in Panel A). The fourth event type was perceptually identical to the third (participants only perceived the checkerboard), although only a checkerboard was presented to the dominant eye. The fifth event type was a true no-distractor baseline, with nothing presented during the retention phase. The distractor gratings in event types 1, 2, and 3 were rotated -45° , -15° , 0° , 15° , or 45° relative to the memorized grating (randomly interleaved). The color coding corresponds to that in later figures for the purpose of convenience and comparability. nondom = viewing of stimuli with the nondominant eye; dom = viewing of stimuli with the dominant eye. See the online article for the color version of this figure.

later on during the run in which it had occurred). A rating of “2” was given on <7% of trials (mean across participants 1.1% with $SE = 0.613\%$), and this was independent of whether or not there was actually an irrelevant grating presented to the other eye, $t(9) = 0.0007$, $p = .998$. In total, participants performed 1,700 trials divided over 10 blocks (~20 min each).

Results and Discussion

This experiment investigated the role of awareness: Is memory still biased by a distractor that people do not consciously perceive? A second question concerned binocular convergence: If biases persist when target and the distractor are presented to different eyes, interference occurs at a level of the visual hierarchy where information from the two eyes has been combined. Figure 7A shows that the V-shaped effect for the absolute (errors was replicated when participants consciously perceived the distractor grating: We observed quadratic effects when the distractor was presented to the same eye (yellow/light gray line), $F(1, 9) = 12.451$, $p = .006$ —and to the other eye (orange/medium gray line), $F(1, 9) = 12.239$, $p = .007$ —relative to the eye presented with the memory target, and it made no difference to which eye the distractor was presented (compare yellow/light gray and orange/medium gray lines), $F(1, 9) = 0.896$, $p = .368$. Paired t tests showed that the quadratic effect arose from the comparison between the conditions where the target and distractor are the same (i.e., differ 0°) with all other conditions (i.e., target–distractor differences of 15° and 45° , all $ps < .057$) and not from comparisons between conditions where target–distractor differences existed (all $ps > .12$). Additionally, when a distractor was consciously perceived, performance suffered relative to the no-distractor baseline—same eye: $F(1, 9) = 11.733$, $p = .008$; other eye: $F(1, 9) = 19.117$, $p = .002$ —but only when the target and distractor differed in orientation—same eye: $F(4, 36) = 5.415$, $p = .002$; other eye: $F(4, 36) = 3.971$, $p = .009$.

When the distractor grating was rendered invisible, no evidence of systematic changes in accuracy emerged across conditions (red/dark gray line), $F(4, 36) = 0.450$, $p = .772$, and participants’ responses were no different from when a checkerboard was presented by itself (compare red/dark gray and dashed lines), $F(1, 9) = 0.045$, $p = .836$.

Estimates of the von Mises SD (Figure 7B) followed the absolute errors made by participants: Quadratic trends indicated that memory was noisier at larger target–distractor orientation differences, but only if the distractor grating was visibly presented to the same eye, $F(1, 9) = 18.358$, $p = .002$, or the other eye, $F(1, 9) = 5.025$, $p = .052$, relative to the eye viewing the target. This was independent of eye-of-origin (compare yellow/light gray and orange/medium gray lines representing distractors to the same and other eye respectively), $F(1, 9) = 0.728$, $p = .416$. Again, the V-shape, or quadratic effect, could be accounted for by comparing 0° difference trials with 15° and 45° difference trials (paired t test p values between .001 and .075), and not because there were any changes between conditions where the target and distractor did not have the same orientation (all $ps > .231$). When the distractor was invisible, no systematic effects of memory variability were observed, $F(4, 36) = 0.136$, $p = .968$, and memory variability was no different from trials on which a checkerboard was presented alone (compare red/dark gray and dashed lines), $F(1, 9) = 0.078$, $p = .786$.

We replicated the shifts of the mean response when the distractor was visible—same eye: $F(4, 36) = 9.97$, $p < .001$, and other eye: $F(4, 36) = 8.387$, $p < .001$ —but not when the distractor was invisible: $F(4, 36) = 0.914$, $p = .466$. The attraction toward visible distractors was independent of the eye to which they were presented (compare yellow/light gray and orange/medium gray lines), $F(1, 9) = 0.016$, $p = .902$.

In terms of magnitude, the maximum distribution shifts were 2.598° and 2.547° for visible distractors presented to the same and

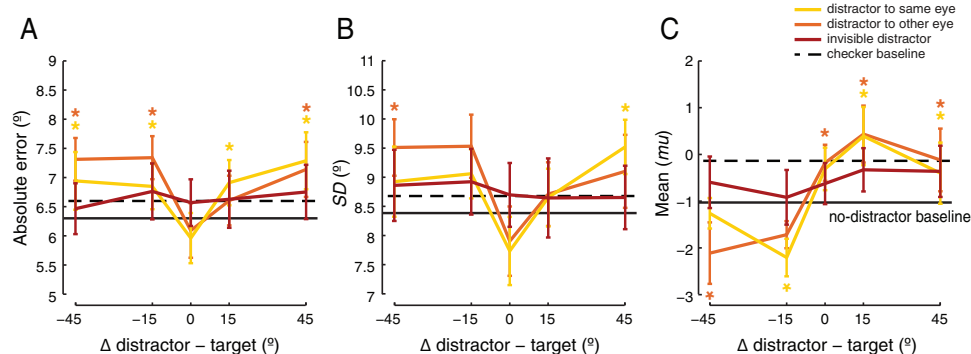


Figure 7. Results, Experiment 4. Panel A: Participants made larger errors when a distractor orientation differed from a memorized orientation, but only when the distractor was consciously perceived, in which case it did not matter if it was presented to the same eye as the memory target or not. Panel B: Memory was noisier when a consciously perceived but irrelevant distractor orientation differed from the orientation in memory. No such noise differences were observed when the distractor was rendered invisible by simultaneous presentation of a checkerboard stimulus to the other eye. Panel C: Only when a distractor orientation was consciously perceived (irrespective of the eye to which it is presented) did the response distribution shift toward the orientation of the distractor grating. Error bars represent $\pm 1 SEM$. Asterisks indicate conditions for which there was a significant difference with the no-distractor baseline (colors/gray-shades correspond to eye-viewing conditions). See the online article for the color version of this figure.

other eye, respectively. Compared with the previous experiments, this was similar in magnitude to the shift found in Experiment 3 (of 2.056°) where the second grating was task-relevant—same eye: $t(16) = 0.636$, $p = .534$; other eye: $t(16) = 0.681$, $p = .505$ —whereas it was smaller than shift-magnitudes found in Experiments 1 and 2 (of 6.4° and 5.822° , respectively), where the second grating was irrelevant (all $ps < .031$). During this experiment participants were expected to ignore all events occurring during the retention interval, but it is likely that some deployment of attention occurred nonetheless because participants were effectively performing a second task in parallel to the primary memory task. Namely, they had to conclude trials where a checkerboard was presented with a perceptual report (the visibility check—see Method), which required a certain degree of attention toward stimulus events taking place during retention. This could explain why the shift-magnitudes found here more closely resembled those from Experiment 3 (where the second grating was attended), than to those from Experiments 1 and 2 (where the second grating was ignored).

General Discussion

The experiments described here revealed systematic biases for a single orientation in memory that arose in response to interfering information (for a schematic overview, see Figure 8). First of all, when a second (to-be-ignored) distractor grating was presented during retention—its orientation parametrically varied relative to the remembered orientation—larger differences between the memorized and distractor orientations led to noisier memory compared with smaller differences between the two orientations. Additionally, the orientation represented in memory was attracted toward the distractor orientation, biasing the correct answer by $\sim 3^\circ$. Increasing the range of target–distractor differences did not affect the magnitude of these two biases, while other factors such as attention and awareness did have an influence. Memory attraction was reduced when the interfering information was made task relevant, and memory biases were completely abolished when participants did not consciously perceive the distractor.

Memory interference by means of attraction has been fairly well established (Dubé et al., 2014; Huang & Sekuler, 2010a; Nemes et al., 2011, 2012). Here, we revealed changes in memory noise due to a distractor orientation presented during retention, accounting for an additional source of memory error that has not been previously demonstrated. Do these noise changes affect memory for better or for worse? When plotting memory noise (SD) against the target–distractor differences examined throughout our four experiments, we consistently found a quadratic effect. This V-shaped effect could comprise of increased memory precision when target–distractor differences were small (a change for the better), decreased memory precision when target–distractor differences were large (a change for the worse), or a mixture of the two. Predominantly, memory precision on trials where a distractor was present (of any orientation) did not differ from precision during a no-distractor baseline, implying a mixture of enhanced and impaired precision at smaller and larger target–distractor differences, respectively. Furthermore, changes in memory noise leveled off once target–distractor differences became sufficiently large. For example, no additional changes were observed between target–

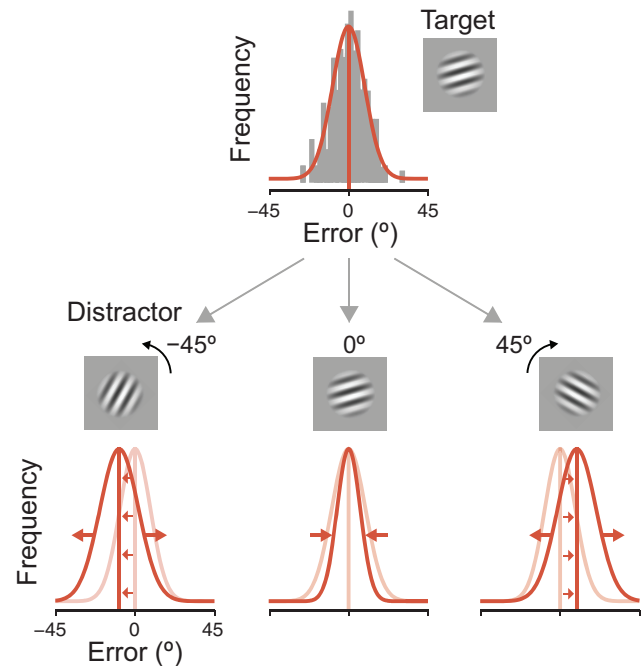


Figure 8. General overview of the main findings. When a single orientation was maintained in memory (top row), a second orientation presented during retention (middle row of distractor gratings) affected memory performance because, first, memory was noisier when the target and distractor differed, compared with when the target and distractor were the same. This is illustrated by changes in the width of the error distribution in response to the target—the original distribution without a distractor present in depicted in solid orange/gray in the top row and lighter orange/gray in the bottom row. In the presence of a distractor, response distributions showed a mixture of narrowing when target and distractor were the same, (bottom row, middle column, solid orange/gray), and broadening when target and distractor differed, (bottom row, outer two columns, solid orange/gray). Second, memory performance was affected because the orientation represented in memory was attracted toward the distractor orientation. This is illustrated by the distribution shift (bottom row, outer two columns, solid orange/gray) in the direction of the distractor orientation by $\sim 3^\circ$. For illustrative purposes, we have schematized and exaggerated the effects on memory noise and distribution means, and only present -45° and 45° difference conditions. See the online article for the color version of this figure.

distractor differences of 15° and 45° , where the degree of memory noise neither increased nor returned to baseline.

Increasing the range of orientation differences between the target and distractor did not impact the magnitude of attraction and noise biases to which memory is susceptible (Experiment 2). This refutes a perceptual averaging account (Dubé et al., 2014; Huang & Sekuler, 2010a) whereby larger target–distractor differences should have resulted in a stronger attraction. Attraction was also not found to return to baseline, which should necessarily occur once a target and distractor differ by 90° —no direction can be inferred when the two are orthogonal. Previous studies did observe that attraction returned to baseline once target–distractor differences fell outside of the bandwidth often assumed of early sensory processing (Nemes et al., 2011, 2012; Van der Stigchel, Merten, Meeter, & Theeuwes, 2007). For orientation, early sensory tuning

bandwidths are estimates around $40^{\circ}\sim 52^{\circ}$ (Albright, 1984; De Valois, Yund, & Hepler, 1982), implying that our target–distractor range with a maximum of 45° might have been just shy of revealing such a return to baseline. That said, there was no indication of this in Experiment 2 (Figure 3C). Instead we found that both biases in attraction, as well as precision, scaled to the range of target–distractor differences presented within a single experiment.

When the interfering information was made task relevant (Experiment 3), the extent to which the first orientation was attracted toward the second was much reduced compared with when information from the second grating could be completely discarded (as in Experiments 1 and 2). Additionally, we found that memory for orientation showed attraction only when the first (but not the second) of two task relevant orientations was probed. These findings directly contradict a comparable study where memory for spatial frequency was biased strongly toward a nonprobed spatial frequency irrespective of which of the two gratings was probed, and this attraction was reduced when the nonprobed item could be ignored (Huang & Sekuler, 2010a). Thus, an attentional manipulation rendering a distractor either task relevant or irrelevant yielded clearly opposite findings depending on whether memory for orientation or for spatial frequency was tested. Memory for these two features might be somehow fundamentally different, or subtle dissimilarities in methodology could be at fault. Regardless, it is clear that a perceptual averaging account of memory interference (Dubé et al., 2014; Huang & Sekuler, 2010a) did not hold for our data. For one, if the degree to which two stimuli are averaged depends on their respective weights, which are in turn determined by mechanisms of selective attention, stronger attraction toward an ignored distractor is unlikely (Experiments 1 and 2 vs. Experiment 3). Furthermore, perceptual averaging can only explain memory attraction, which is inconsistent with the majority of participants in Experiment 3 demonstrating a substantial repulsion.

Repulsion was expressed by the second (task-relevant) orientation being rotated in a direction opposite from the first. One possible explanation is that the second grating was already perceived as repelled away from the first and thus remembered as such (Scocchia, Cicchini, & Triesch, 2013). Another way to look at repulsion is by realizing that when two successive stimuli have different orientations there could be an implied rotation, or “movement.” For example, if the second grating was clockwise relative to the first, a “clockwise rotation” was implied. Memory for moving targets is biased in accordance with expectations people have about the physical characteristics of real-life moving objects, such as momentum, gravity, or friction (Hubbard, 1995a, 1995b). In our example this means a shift of both representations in the direction of the implied clockwise rotation: For the first grating this would look like “attraction” toward the second (clockwise) grating. For the second grating, the rotation would be clockwise relative to its own actual orientation, making it look like a “repulsion” relative to the first grating. Nonetheless, the fact that a subset of participants showed attraction rather than repulsion under the same circumstances indicates that individuals used different strategies, and that the direction of the bias exerted by a second stimulus is not fixed. To some extent, memory biases could be decisional rather than perceptual in nature (Park, Rademaker, & Tong, 2014), which contradicts theories assuming interference

occurring within the visual hierarchy (Huang & Sekuler, 2010a; Magnussen, 2000).

Memory biases were abolished when participants did not consciously perceive the interfering distractor orientation, implying that awareness of the distractor was necessary for it to influence representations held in memory (Experiment 4). This directly contradicts earlier work where a subliminal distractor negatively affected memory accuracy (Silvanto & Soto, 2012). The absence of distractor interactions during rivalrous stimulus presentation furthermore suggests a locus of interference at a relatively late stage of visual processing—while suppressed information can still be observed in V1 (Maier et al., 2008), this signal peters out along the visual hierarchy (Blake & Logothetis, 2002). When the distractor orientation was consciously perceived, the eye to which the distractor orientation was presented (relative to the target) did not impact short-term memory biases, implying a brain locus beyond the point of binocular convergence (Experiment 4). The memory biases observed in the present study are unlikely to arise at early sensory levels, which is supported by the existence of binocular cells in V1 (Parker, 2007; Poggio, Motter, Squatrito, & Trotter, 1985) in combination with our finding that interference occurred on binocularly combined signals that required awareness in a rivalrous setting.

One alternative explanation for the biases reported here might be that on some portion of trials participants reported the wrong orientation. As the memorized and distractor gratings differed more, such mistakes would lead to increases in fitted noise and larger apparent distribution shifts toward the distractor orientation. This is unlikely to be the case; first of all, our task was trivially easy, making it improbable that errors due to misreporting would occur frequently, if at all. The error distributions (see the online supplemental material for Figure S3) confirm this. Second, errors due to misreporting would amplify memory biases as the range of target–distractor differences increased, which was not found to be the case. Finally, we directly compared data fits from a von Mises to fits from a bimodal model that described the data in terms of a mean, circular variance, and probability of responses to the wrong orientation. We only included conditions where the target and distractor had different orientations, and where the second grating was visibly presented. A von Mises fit our data better in all experiments—Experiment 1: $F(1, 7) = 897.384, p < .001$; Experiment 2: $F(1, 7) = 21.30, p = .002$; Experiment 3: $F(1, 7) = 11.714, p = .011$; Experiment 4: $F(1, 9) = 3.68, p = .087$ —further ruling out the possibility of faulty reports by our participants.

Despite explicit instructions to ignore the second grating, its orientation nevertheless affected participants’ responses to the target orientation. This can be explained neither by weighted averaging and magnet metaphors, nor by misreporting of the stimulus. An alternative way of interpreting our findings is in the context of a Bayesian framework, which explains perception by integrating sensory information (or likelihood) with expectations about the world (prior). Within this framework, perceptual biases can emerge when people integrate prior information about a stimulus in a Bayesian fashion, following certain mathematical rules. Prior information can be derived from repeated long-term exposure to natural image statistics, such as the overrepresentation of horizontal and vertical information in natural scenes (Girshick, Landy, & Simoncelli, 2011) or default illumination and observer viewpoints being from above (Gerardin, Kourtzi, & Mamassian,

2010; Mamassian & Goutcher, 2001; Mamassian & Landy, 1998). Such priors might be implemented as early as primary sensory cortex, interacting with sensory evidence at very early stages of visual processing (Kok, Brouwer, van Gerven, & de Lange, 2013; Vintch & Gardner, 2014). Notably, priors are not necessarily static, and can be generated or updated in an experimental setting within a relatively short time (Brady, Konkle, & Alvarez, 2009; Chopin & Mamassian, 2012; Körding & Wolpert, 2004; Turk-Browne, Scholl, Chun, & Johnson, 2009).

In our experiment, the second grating had a limited number of possible orientations relative to the first, which meant it retroactively provided participants with information about the remembered orientation. Once learned, such statistical information could have been used as a prior and integrated into participants' responses. When trying to infer the memorized orientation, statistical information introduced by a distractor could have been more or less equally informative irrespective of the range of all possible distractors spanned around the target orientation: Although a larger range of distractor orientations around the target can be seen as a chance for stronger biases, it also implies a larger variance on the prior, and these two factors might cancel out. Potentially then, such integration could explain why we found that memory biases scaled with different target–distractor ranges in Experiments 1 and 2. However, we did not find support for the idea that repeated exposures to the distractor orientations in our study resulted in statistical learning: When we calculated memory attraction on a trial-by-trial basis, distractor effects were observed immediately, and did not evolve over time (see the online supplemental material for Figure S4).

Perhaps the integration of irrelevant feature information into short-term memory is an unavoidable feature of the system (Marshall & Bays, 2013). Such involuntary integration could deplete memory resources and reduce precision, which could explain our finding of larger variability when target and distractor orientations differed more. Alternatively, integration of information can be viewed as a common strategy, used because in many contexts it actually improves the accuracy of report (Huttenlocher, Hedges, & Duncan, 1991; Huttenlocher, Hedges, & Vevea, 2000). Memory interference could also be related to serial dependence in human vision, which assumes that both past and present inputs are used to inform perception. For example, a memorized spatial frequency (Huang & Sekuler, 2010a; Huang & Sekuler, 2010b), location (Simmering, Schutte, & Spencer, 2008) or perceived orientation (Fischer & Whitney, 2014) is attracted toward information from the recent past. Memory contents could be affected in a similar manner, with a distractor serving as a past influence on a current memory representation.

There is evidence that also other contexts are able to systematically bias the contents of visual short-term memory. Ensemble statistics are an example, such as a reported bias toward the average size of items in memory (Alvarez, 2011; Brady & Alvarez, 2011), or a central tendency biasing memory toward the average or prototypical value of a stimulus set (Dubé et al., 2014; Freyd & Johnson, 1987; Huang & Sekuler, 2010a; Huttenlocher et al., 1991; Spencer & Hund, 2002; Wilken & Ma, 2004). Eye movements can bias memory in the direction of the saccade (Bays & Husain, 2008). A spatial reference frame can bias memory for location away from salient axes (Simmering et al., 2008; Spencer & Hund, 2002). Notably, memory for spatial location was found to

shift toward the location of a task irrelevant distractor (Van der Stigchel et al., 2007), which parallels our findings for orientation. This shift only occurred when the distractor was relatively close by, which is reminiscent of the idea that attraction only occurs within a limited bandwidth (Nemes et al., 2011, 2012; Sneve et al., 2015). When memory uncertainty is higher (e.g., if items have to be retained for longer durations, if contrast is low, if set size is large, etc.), people are thought to rely more strongly on such biases to postulate a response (Dubé et al., 2014; Vintch & Gardner, 2014).

Thus, many types of interactions can arise among memory representations, or between stored information and incoming perceptual input. The extent of these interactions depends strongly on the metric distance between stimulus features relative to each other. These so-called metric interactions offer another compelling interpretation for the memory distortions reported here. To explore this in more detail we borrowed from work on dynamic neural network models consisting of a layered, neurally plausible architecture. This class of model can capture a wide variety of behaviors by assuming the existence of perceptual, as well as short- or long-term memory model-layers that interact via local excitatory and lateral inhibitory interactions (Johnson, Spencer, Luck, & Schöner, 2009; Simmering et al., 2008; Wei, Wang, & Wang, 2012). Such connectivity achieves sustained activation during working memory, but also leads to metric interactions between memory items presented at different times (Simmering et al., 2008) different spatial locations (Wei et al., 2012), and between stored information and perceptual input (Johnson et al., 2009). While items are maintained in working memory in these models, processes such as merging (Wei et al., 2012) and biasing (Simmering et al., 2008) can occur among representations in memory. A tentative explanation of the biases observed in our study could be derived from combining these delay-period dynamics with consequences from newly arriving sensory input—shown to add a subtle peak of activation in a neural model's memory layer (Johnson et al., 2009). If a small peak of distractor-centered excitation merges with the maintained representation one would expect a memory shift toward the distractor, as well as an increase in variance.

This raises the question which brain mechanisms implement interactions between memory and sensory processes, and give rise to memory biases observed in the literature. A recent neuroimaging study demonstrated that when people remembered a single spatial frequency, and were presented with a distractor whose spatial frequency differed by $1 c^\circ$ (but not $2 c^\circ$), activity to the distractor was suppressed in visual areas V3 and V4 (Sneve et al., 2015). This aligns with predictions of a “Mexican hat” shaped interaction profile from neural network models (Johnson et al., 2009; Simmering et al., 2008). Critically, modulations in V4 influenced activity in earlier areas V1–V3, where a stronger suppression correlated with performance decrements. This finding also complements previous psychophysics (Magnussen, 2000) by alluding to a possible brain mechanism involving interference and top-down influences from V4 initiating behaviorally relevant changes in V1.

However, the study by Sneve et al., (2015) only revealed how memorized information affected processing of a distractor, without addressing what happened to the memory representation itself. A lot is already known about how high-level cognitive states, such as

memory, can alter the readiness or modulate responses to new sensory inputs through feedback activity, biasing sensory neurons and perception (Johnson et al., 2009; Lui & Pasternak, 2011; Mendoza, Schneiderman, Kaul, & Martinez-Trujillo, 2011; Mendoza-Halliday, Torres, & Martinez-Trujillo, 2014; Miller & Desimone, 1994; Miller, Li, & Desimone, 1991, 1993; Scocchia et al., 2013; Zaksas & Pasternak, 2006). Indeed, working memory contents can guide selective attention (Downing, 2000; Soto, Hodsoll, Rotshtein, & Humphreys, 2008), and increase cortical excitability (Cattaneo, Pisoni, Papagno, & Silvanto, 2011). Our findings clearly demonstrate that dynamic interactions between bottom-up sensory information and top-down cognitive states (such as the maintenance of visual memories) work in both directions, with newly perceived sensory information also biasing the processing of information that is actively maintained in memory. However, the mechanisms by which perception might bias representations in visual short-term memory are much less understood.

Although it remains elusive whether memory interference serves a functional purpose, studying interference provides a window into the computational processes supporting visual memory by exploring how such mechanisms inoculate memories against the myriad of new images entering the eyes in rapid succession. Against a backdrop of complementary imaging, physiology, and psychophysical work, the current findings help further our understanding of the dynamic interactions involved in the maintenance of visual memories, and inform us about the fundamental question of how memories are stored.

References

- Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology*, *52*, 1106–1130.
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, *15*, 122–131. <http://dx.doi.org/10.1016/j.tics.2011.01.003>
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*, 1–11.
- Bays, P. M., & Husain, M. (2008, August 8). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854. <http://dx.doi.org/10.1126/science.1158023>
- Bennett, P. J., & Cortese, F. (1996). Masking of spatial frequency in visual memory depends on distal, not retinal, frequency. *Vision Research*, *36*, 233–238. [http://dx.doi.org/10.1016/0042-6989\(95\)00085-E](http://dx.doi.org/10.1016/0042-6989(95)00085-E)
- Berens, P. (2009). CircStat: A MATLAB toolbox for circular statistics. *Journal of Statistical Software*, *31*, 1–21.
- Blake, R., & Logothetis, N. (2002). Visual competition. *Nature Reviews Neuroscience*, *3*, 13–21. <http://dx.doi.org/10.1038/nrn701>
- Bosch, S. E., Jehee, J. F. M., Fernández, G., & Doeller, C. F. (2014). Reinstatement of associative memories in early visual cortex is signaled by the hippocampus. *The Journal of Neuroscience*, *34*, 7493–7500. <http://dx.doi.org/10.1523/JNEUROSCI.0805-14.2014>
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, *22*, 384–392. <http://dx.doi.org/10.1177/0956797610397956>
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, *138*, 487–502. <http://dx.doi.org/10.1037/a0016797>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436. <http://dx.doi.org/10.1163/156856897X00357>
- Cattaneo, Z., Pisoni, A., Papagno, C., & Silvanto, J. (2011). Modulation of visual cortical excitability by working memory: Effect of luminance contrast of mental imagery. *Frontiers in Psychology*, *2*, 1–9. <http://dx.doi.org/10.3389/fpsyg.2011.00029>
- Chopin, A., & Mamassian, P. (2012). Predictive properties of visual adaptation. *Current Biology*, *22*, 622–626. <http://dx.doi.org/10.1016/j.cub.2012.02.021>
- Deutsch, D. (1970, June 26). Tones and numbers: Specificity of interference in immediate memory. *Science*, *168*, 1604–1605. <http://dx.doi.org/10.1126/science.168.3939.1604>
- Deutsch, D. (1973). Interference in memory between tones adjacent in the musical scale. *Journal of Experimental Psychology*, *100*, 228–231. <http://dx.doi.org/10.1037/h0035440>
- De Valois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, *22*, 531–544. [http://dx.doi.org/10.1016/0042-6989\(82\)90112-2](http://dx.doi.org/10.1016/0042-6989(82)90112-2)
- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*, 467–473. <http://dx.doi.org/10.1111/1467-9280.00290>
- Dubé, C., Zhou, F., Kahana, M. J., & Sekuler, R. (2014). Similarity-based distortion of visual short-term memory is due to perceptual averaging. *Vision Research*, *96*, 8–16. <http://dx.doi.org/10.1016/j.visres.2013.12.016>
- Ernst, M. O., & Banks, M. S. (2002, January 24). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433. <http://dx.doi.org/10.1038/415429a>
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*, 738–743.
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, *3*, 1229. <http://dx.doi.org/10.1038/ncomms2237>
- Freyd, J. J., & Johnson, J. Q. (1987). Probing the time course of representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 259–268. <http://dx.doi.org/10.1037/0278-7393.13.2.259>
- Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current Opinion in Neurobiology*, *20*, 177–182. <http://dx.doi.org/10.1016/j.conb.2010.03.005>
- Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*, 129–135. <http://dx.doi.org/10.1016/j.tics.2011.11.014>
- Gerardin, P., Kourtzi, Z., & Mamassian, P. (2010). Prior knowledge of illumination for 3D perception in the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *107*, 16309–16314. <http://dx.doi.org/10.1073/pnas.1006285107>
- Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*, 926–932. <http://dx.doi.org/10.1038/nn.2831>
- Gorgoraptis, N., Catalao, R. F. G., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *The Journal of Neuroscience*, *31*, 8502–8511. <http://dx.doi.org/10.1523/JNEUROSCI.0208-11.2011>
- Harrison, S. A., & Tong, F. (2009, April 2). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*, 632–635. <http://dx.doi.org/10.1038/nature07832>
- Huang, J., & Sekuler, R. (2010a). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision*, *10*, 1–27.
- Huang, J., & Sekuler, R. (2010b). Attention protects the fidelity of visual memory: Behavioral and electrophysiological evidence. *The Journal of Neuroscience*, *30*, 13461–13471. <http://dx.doi.org/10.1523/JNEUROSCI.2560-10.2010>
- Hubbard, T. L. (1995a). Cognitive representation of motion: Evidence for friction and gravity analogues. *Journal of Experimental Psychology*:

- Learning, Memory, and Cognition*, 21, 241–254. <http://dx.doi.org/10.1037/0278-7393.21.1.241>
- Hubbard, T. L. (1995b). Environmental invariants in the representation of motion: Implied dynamics and representational momentum, gravity, friction, and centripetal force. *Psychonomic Bulletin & Review*, 2, 322–338. <http://dx.doi.org/10.3758/BF03210971>
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, 98, 352–376. <http://dx.doi.org/10.1037/0033-295X.98.3.352>
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129, 220–241. <http://dx.doi.org/10.1037/0096-3445.129.2.220>
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science*, 20, 568–577. <http://dx.doi.org/10.1111/j.1467-9280.2009.02329.x>
- Kok, P., Brouwer, G. J., van Gerven, M. A. J., & de Lange, F. P. (2013). Prior expectations bias sensory representations in visual cortex. *The Journal of Neuroscience*, 33, 16275–16284. <http://dx.doi.org/10.1523/JNEUROSCI.0742-13.2013>
- Körding, K. P., & Wolpert, D. M. (2004, January 15). Bayesian integration in sensorimotor learning. *Nature*, 427, 244–247. <http://dx.doi.org/10.1038/nature02169>
- Lalonde, J., & Chaudhuri, A. (2002). Task-dependent transfer of perceptual to memory representations during delayed spatial frequency discrimination. *Vision Research*, 42, 1759–1769. [http://dx.doi.org/10.1016/S0042-6989\(02\)00111-6](http://dx.doi.org/10.1016/S0042-6989(02)00111-6)
- Luck, S. J., & Vogel, E. K. (1997, November 20). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281. <http://dx.doi.org/10.1038/36846>
- Lui, L. L., & Pasternak, T. (2011). Representation of comparison signals in cortical area MT during a delayed direction discrimination task. *Journal of Neurophysiology*, 106, 1260–1273. <http://dx.doi.org/10.1152/jn.00016.2011>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17, 347–356. <http://dx.doi.org/10.1038/nn.3655>
- Magnussen, S. (2000). Low-level memory processes in vision. *Trends in Neurosciences*, 23, 247–251. [http://dx.doi.org/10.1016/S0166-2236\(00\)01569-1](http://dx.doi.org/10.1016/S0166-2236(00)01569-1)
- Magnussen, S., & Greenlee, M. W. (1992). Retention and disruption of motion information in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 151–156. <http://dx.doi.org/10.1037/0278-7393.18.1.151>
- Magnussen, S., & Greenlee, M. W. (1999). The psychophysics of perceptual memory. *Psychological Research*, 62, 81–92. <http://dx.doi.org/10.1007/s004260050043>
- Magnussen, S., Greenlee, M. W., Asplund, R., & Dyrnes, S. (1991). Stimulus-specific mechanisms of visual short-term memory. *Vision Research*, 31, 1213–1219. [http://dx.doi.org/10.1016/0042-6989\(91\)90046-8](http://dx.doi.org/10.1016/0042-6989(91)90046-8)
- Maier, A., Wilke, M., Aura, C., Zhu, C., Ye, F. Q., & Leopold, D. A. (2008). Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nature Neuroscience*, 11, 1193–1200. <http://dx.doi.org/10.1038/nn.2173>
- Mamassian, P., & Goutcher, R. (2001). Prior knowledge on the illumination position. *Cognition*, 81, B1–B9. [http://dx.doi.org/10.1016/S0010-0277\(01\)00116-0](http://dx.doi.org/10.1016/S0010-0277(01)00116-0)
- Mamassian, P., & Landy, M. S. (1998). Observer biases in the 3D interpretation of line drawings. *Vision Research*, 38, 2817–2832. [http://dx.doi.org/10.1016/S0042-6989\(97\)00438-0](http://dx.doi.org/10.1016/S0042-6989(97)00438-0)
- Marshall, L., & Bays, P. M. (2013). Obligatory encoding of task-irrelevant features depletes working memory resources. *Journal of Vision*, 13, 1–13.
- McKeefry, D. J., Burton, M. P., & Vakrou, C. (2007). Speed selectivity in visual short term memory for motion. *Vision Research*, 47, 2418–2425. <http://dx.doi.org/10.1016/j.visres.2007.05.011>
- Mendoza, D., Schneiderman, M., Kaul, C., & Martinez-Trujillo, J. (2011). Combined effects of feature-based working memory and feature-based attention on the perception of visual motion direction. *Journal of Vision*, 11, 11, 1–15.
- Mendoza-Halliday, D., Torres, S., & Martinez-Trujillo, J. C. (2014). Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nature Neuroscience*, 17, 1255–1262. <http://dx.doi.org/10.1038/nn.3785>
- Miller, E. K., & Desimone, R. (1994, January 28). Parallel neuronal mechanisms for short-term memory. *Science*, 263, 520–522. <http://dx.doi.org/10.1126/science.8290960>
- Miller, E. K., Li, L., & Desimone, R. (1991, November 29). A neural mechanism for working and recognition memory in inferior temporal cortex. *Science*, 254, 1377–1379. <http://dx.doi.org/10.1126/science.1962197>
- Miller, E. K., Li, L., & Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *The Journal of Neuroscience*, 13, 1460–1478.
- Nemes, V. A., Parry, N. R., Whitaker, D., & McKeefry, D. J. (2012). The retention and disruption of color information in human short-term visual memory. *Journal of Vision*, 12, 1–14.
- Nemes, V. A., Whitaker, D., Heron, J., & McKeefry, D. J. (2011). Multiple spatial frequency channels in human visual perceptual memory. *Vision Research*, 51, 2331–2339. <http://dx.doi.org/10.1016/j.visres.2011.09.003>
- Nilsson, T. H., & Nelson, T. M. (1981). Delayed monochromatic hue matches indicate characteristics of visual memory. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 141–150. <http://dx.doi.org/10.1037/0096-1523.7.1.141>
- Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 332–350. <http://dx.doi.org/10.1037/0096-1523.16.2.332>
- Park, Y., Rademaker, R. L., & Tong, F. (2014). Both variations in perceptual sensitivity and decisional response bias contribute to visual working memory performance. *Journal of Vision*, 14, 1375. <http://dx.doi.org/10.1167/14.10.1375>
- Parker, A. J. (2007). Binocular depth perception and the cerebral cortex. *Nature Reviews Neuroscience*, 8, 379–391. <http://dx.doi.org/10.1038/nrn2131>
- Pasternak, T., & Zaksas, D. (2003). Stimulus specificity and temporal dynamics of working memory for visual motion. *Journal of Neurophysiology*, 90, 2757–2762. <http://dx.doi.org/10.1152/jn.00422.2003>
- Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, 12, 334–341. <http://dx.doi.org/10.1016/j.tics.2008.05.006>
- Pearson, J., Clifford, C. W. G., & Tong, F. (2008). The functional impact of mental imagery on conscious perception. *Current Biology*, 18, 982–986. <http://dx.doi.org/10.1016/j.cub.2008.05.048>
- Pearson, J., Rademaker, R. L., & Tong, F. (2011). Evaluating the mind's eye: The metacognition of visual imagery. *Psychological Science*, 22, 1535–1542. <http://dx.doi.org/10.1177/0956797611417134>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442. <http://dx.doi.org/10.1163/156856897X00366>
- Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, 16, 283–290. <http://dx.doi.org/10.3758/BF03203943>
- Poggio, G. F., Motter, B. C., Squatrito, S., & Trotter, Y. (1985). Responses of neurons in visual cortex (V1 and V2) of the alert macaque to dynamic

- random-dot stereograms. *Vision Research*, 25, 397–406. [http://dx.doi.org/10.1016/0042-6989\(85\)90065-3](http://dx.doi.org/10.1016/0042-6989(85)90065-3)
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464. <http://dx.doi.org/10.1214/aos/1176344136>
- Scocchia, L., Cicchini, G. M., & Triesch, J. (2013). What's "up"? Working memory contents can bias orientation processing. *Vision Research*, 78, 46–55. <http://dx.doi.org/10.1016/j.visres.2012.12.003>
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, 20, 207–214. <http://dx.doi.org/10.1111/j.1467-9280.2009.02276.x>
- Silvanto, J., & Soto, D. (2012). Causal evidence for subliminal percept-to-memory interference in early visual cortex. *NeuroImage*, 59, 840–845. <http://dx.doi.org/10.1016/j.neuroimage.2011.07.062>
- Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. *Brain Research*, 1202, 68–86. <http://dx.doi.org/10.1016/j.brainres.2007.06.081>
- Sneve, M. H., Alnæs, D., Endestad, T., Greenlee, M. W., & Magnussen, S. (2012). Visual short-term memory: Activity supporting encoding and maintenance in retinotopic visual cortex. *NeuroImage*, 63, 166–178. <http://dx.doi.org/10.1016/j.neuroimage.2012.06.053>
- Sneve, M. H., Sreenivasan, K. K., Alnæs, D., Endestad, T., & Magnussen, S. (2015). Short-term retention of visual information: Evidence in support of feature-based attention as an underlying mechanism. *Neuropsychologia*, 66, 1–9. <http://dx.doi.org/10.1016/j.neuropsychologia.2014.11.004>
- Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, 12, 342–348. <http://dx.doi.org/10.1016/j.tics.2008.05.007>
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, 131, 16–37. <http://dx.doi.org/10.1037/0096-3445.131.1.16>
- Suchow, J. W., Brady, T. F., Fougny, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, 13, 1–8.
- Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, 21, 1934–1945. <http://dx.doi.org/10.1162/jocn.2009.21131>
- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 8780–8785. <http://dx.doi.org/10.1073/pnas.1117465109>
- Van der Stigchel, S., Merten, H., Meeter, M., & Theeuwes, J. (2007). The effects of a task-irrelevant visual event on spatial working memory. *Psychonomic Bulletin & Review*, 14, 1066–1071. <http://dx.doi.org/10.3758/BF03193092>
- Vintch, B., & Gardner, J. L. (2014). Cortical correlates of human motion perception biases. *The Journal of Neuroscience*, 34, 2592–2604. <http://dx.doi.org/10.1523/JNEUROSCI.2809-13.2014>
- Warden, M. R., & Miller, E. K. (2007). The representation of multiple objects in prefrontal neuronal delay activity. *Cerebral Cortex*, 17, i41–i50.
- Wei, Z., Wang, X. J., & Wang, D. H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience*, 32, 11228–11240. <http://dx.doi.org/10.1523/JNEUROSCI.0735-12.2012>
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4, 1120–1135. <http://dx.doi.org/10.1167/4.12.11>
- Zaksas, D., & Pasternak, T. (2006). Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *The Journal of Neuroscience*, 26, 11726–11742. <http://dx.doi.org/10.1523/JNEUROSCI.3420-06.2006>
- Zhang, W., & Luck, S. J. (2008, May 8). Discrete fixed-resolution representations in visual working memory. *Nature*, 453, 233–235. <http://dx.doi.org/10.1038/nature06860>
- Zokaei, N., Gorgoraptis, N., Bahrami, B., Bays, P. M., & Husain, M. (2011). Precision of working memory for visual motion sequences and transparent motion surfaces. *Journal of Vision*, 11, 2, 1–18.

Received November 7, 2014

Revision received June 9, 2015

Accepted June 11, 2015 ■