MIT Libraries | DSpace@MIT

MIT Open Access Articles

Feature-Binding Errors After Eye Movements and Shifts of Attention

Massachusetts Institute of Technology

Feature-binding errors after eye movements and shifts of attention

Julie D. Golomb[1,*], **Zara E. L'Heureux**[2], and **Nancy Kanwisher**[2]

[1]Dept of Psychology, The Ohio State University

[2]McGovern Institute for Brain Research, Massachusetts Institute of Technology

Abstract

When we move our eyes, the eye-centered (retinotopic) locations of objects must be updated to maintain world-centered (spatiotopic) stability. Here we demonstrate that the attentional updating process temporarily distorts our fundamental ability to bind object locations with their features. Subjects were simultaneously presented with four colors after a saccade – one in a pre-cued spatiotopic target location – and instructed to report the target's color using a colorwheel. Subjects' reports were systematically shifted in color space, toward the color of the distractor in the retinotopic location of the cue. Probabilistic modeling exposed both crude "swapping" errors and subtler "feature mixing" (as if the retinotopic color had blended into the spatiotopic percept). Additional experiments conducted without saccades revealed that the two types of errors stem from different attentional mechanisms (attention shifting vs splitting). Feature mixing not only reflects a new perceptual phenomenon, but provides novel insight into how attention is remapped across saccades.

Keywords

saccade; retinotopic; spatiotopic; remapping; eye-centered; illusory conjunction

Introduction

Our sensory systems are constantly bombarded with information, and we cannot process everything. Eye movements and spatial attention are two fundamental ways our visual systems filter our complex environment. Eye movements are particularly interesting both in terms of their frequency (multiple times each second), and the additional challenge the movements introduce for spatial stability. With each movement, the images hitting our retinas change dramatically. How can we attend to a world-centered (spatiotopic) location, when the underlying visual representations are coded in retinotopic (eye-centered) coordinates (Cohen & Andersen, 2002; Gardner, Merriam, Movshon, & Heeger, 2008;

Golomb, Chun, & Mazer, 2008; Golomb, Nguyen-Phuc, Mazer, McCarthy, & Chun, 2010), even at higher stages of processing (Golomb & Kanwisher, 2012a; but see Crespi et al., 2011; d'Avossa et al., 2007)? The brain may solve this problem in part by predictive "remapping" (Duhamel, Colby, & Goldberg, 1992), that is, by updating receptive fields – or spatial pointers (Cavanagh, Hunt, Afraz, & Rolfs, 2010) – with each saccade, sometimes even before it is executed. However, visual stability requires that we not only update spatial locations, but that the updated spatial information is correctly bound to features in the environment.

We have recently argued that updating of spatial attention across saccades entails two distinct processes: a rapid (sometimes anticipatory) remapping to the new location (Rolfs, Jonikaitis, Deubel, & Cavanagh, 2011), and a slower process of extinguishing the previous representation (Golomb et al., 2008; Golomb, Marino, Chun, & Mazer, 2011; Golomb et al., 2010). Here we predicted that if the attentional updating process is not complete by the end of the saccade, such that both representations are temporarily active at the same time (the newly remapped location and the not-yet-extinguished previous location), we might be susceptible to errors beyond spatial misperception. Might we even find a mixing of *features* at these two locations? Furthermore, might such mixing be found not only when eye movements occur, but whenever two attentional traces are active at the same time?

We used a continuous report paradigm (Wilken & Ma, 2004; Zhang & Luck, 2008) where subjects were presented with an array of four colored stimuli, and were instructed to report the color of a designated stimulus by clicking the appropriate place on a colorwheel (Figure 1). The target location was cued before the saccade, but all four colors were presented simultaneously after the saccade – thus this task is not about trans-saccadic integration of color (integrating features from the same location at two points in time: Hunt & Cavanagh, 2011; Wittenberg, Bremmer, & Wachtler, 2008), but the ability to correctly bind features to their locations (associating a single color with a single location: Treisman, 1996). While previous studies have reported peri-saccadic errors involving spatio-temporal mislocalization (Burr, Ross, Binda, & Morrone, 2010; Ross, Morrone, & Burr, 1997) or general perceptual impairments (Latour, 1962; Ross, Morrone, Goldberg, & Burr, 2001), to our knowledge the current study is the first to investigate distortions of feature binding following a saccade. Furthermore, we predict a novel, specific disruption of binding: after a saccade, the presence of a retinotopic distractor (but not a distractor at a "control" location) will systematically distort perception at the spatiotopic location, via either erroneous "swapping" of retinotopic and spatiotopic features, or perhaps even "feature mixing", producing a blended percept.

In the experiments below we test this hypothesis that distortions in feature binding are present immediately following a saccade when attention is to be maintained at a spatiotopic location (Expt 1), and compare it to other scenarios involving potentially ambiguous attentional states: when attention is maintained at a retinotopic location across a saccade (Expt 2), and when attention is shifted (Expt 3) or split (Expt 4) across two locations in the absence of a saccade.

Materials and Methods

Subjects

Sixteen subjects (8 female; mean age 27.2) participated in Experiment 1, and 9 subjects participated in Experiment 2; three subjects participated in both experiments. Twelve subjects and 18 subjects participated in Experiments 3 and 4. Additional subjects were excluded for not successfully performing the task (>50% probability of random guessing on no-saccade trials, γ parameter from Model A). See supplemental methods for additional details on subjects and exclusions.

Experimental Setup

Stimuli were generated using the Psychtoolbox extension (Brainard, 1997) for Matlab and presented on a 21" flat-screen CRT monitor. Subjects were seated at a chinrest 64cm from the monitor. Eye position was monitored using ISCAN (Experiments 1–2) and Eyelink 1000 (Experiments 3–4) eye-tracking systems recording pupil and corneal reflection. The monitors were color calibrated with a Minolta CS-100 colorimeter.

Experiment 1: Maintain attention across saccade (spatiotopic) task (Figure 1)

—Each trial began with a white fixation dot presented at one of four locations on the screen (arranged as the corners of an 8.7°x8.7° square). Once subjects were accurately fixating for 1sec (determined by real-time eye-tracking), a spatial pre-cue (black 2°x2° square) was presented for 500ms. After another 1sec fixation period, on half of trials the fixation dot jumped to a horizontally or vertically adjacent position. On these "saccade" trials, subjects had to immediately move their eyes to the new location. On the other half of trials ("no-saccade" trials), the fixation dot remained at the original location and subjects held fixation for an equivalent amount of time based on average saccadic latency from a prior study (~350ms). Both the location of the cue (any of 5 locations on the screen: center-center, center-top, center-bottom, left-center, or right-center) and the presence/direction of the saccade were randomized and unpredictable.

After a delay of either 50ms or 500ms from the time of successful saccade completion, an array of four differently colored squares appeared at equidistant locations around fixation (7.4° eccentricity). One of these squares occupied the same spatiotopic (absolute) location as the pre-cue – subjects were instructed to report the color of the square that appeared at this cued location. The colored squares appeared for 50ms, followed by 200ms of masks (colored with a random color value at each pixel location, covering each of the four stimulus locations). A large colorwheel (diameter 16.4°) was then presented in the center of the screen – at a random rotation – and subjects clicked with the mouse to report the color of the cued location. They were then given feedback showing them the correct color.

On "saccade" trials, one of the four colored stimuli appeared at the spatiotopic location of the cue – this was the color subjects were supposed to report. Another stimulus occupied the same retinotopic location as the cue, and the two remaining stimuli occupied the mirror-symmetric control locations. On "no-saccade" trials, the cued location was both spatiotopic and retinotopic, and the other three stimuli were all considered control locations. The color

at the cued (spatiotopic) location was chosen randomly on each trial from 180 possible colors (evenly distributed along a circle in CIE L*a*b* color space, according to the parameters in Zhang & Luck, 2008). The colors of the remaining stimuli were chosen so that the retinotopic and equidistant control stimuli were equally different from the spatiotopic color, but in opposite directions (90° clockwise or counterclockwise along the colorwheel, with direction randomly varying from trial to trial), and the stimulus at the diagonal location was set at 180° away in color space.

At any point in the trial, if the subject's eye position deviated more than 2° from the correct fixation location, or if saccadic latency was greater than 600ms, the trial was immediately aborted and repeated later in the block.

Experiment 2: Maintain attention across saccade (retinotopic) task—The task was identical to Experiment 1, except subjects reported the color of the stimulus at the retinotopic (not spatiotopic) location of the cue after the saccade.

Experiment 3: Shift attention task—Subjects remained fixated on a central fixation dot (Figure 5a). One of four peripheral locations was cued for 250ms. On half of the trials ("shift" trials), after a 1sec delay a second cue appeared in a different location for 50ms. After either a 50ms or 500ms delay, the array of four colored squares appeared for 50ms, followed by the masks and colorwheel as in Experiment 1. The task was to report the color that appeared at the most recent location of the cue; i.e., subjects had to shift attention from the original location to the final location. On "no-shift" trials, a second cue never appeared and subjects simply reported the color at the initially cued location. Inclusion of "no-shift" trials ensured that subjects had to attend to the first cue and could not simply wait for the second.

Experiment 4: Split attention task—Subjects fixated on a central fixation dot, and two of the four stimulus locations were simultaneously cued (Figure 5b). Subjects were instructed to attend to both locations (share/split attention). After 1sec the four colors appeared, followed by the masks. When the colorwheel appeared, a post-cue was presented indicating which of the locations to report. The post-cued location was always one of the two pre-cued locations, but it was unpredictable which one.

Analyses

The location on the colorwheel where subjects clicked on each trial was recorded and converted into a difference score in degrees of visual angle, with the correct spatiotopic color at 0° difference, and retinotopic and control distractors at 90° and −90°, respectively. (Experiment 2: retinotopic color at 0°, spatiotopic at 90°; Experiment 3: color at final cued location 0°, original cued location 90°; Experiment 4: color at post-cued location 0°, other attended location 90°.) The mean of the distribution was calculated separately for each subject and condition, and 2-tailed t-tests were run to test if the means were significantly different from zero.

The distribution of responses was also fit with probabilistic models (Bays, Catalao, & Husain, 2009; Zhang & Luck, 2008) accounting for various sources of error:

A. Basic "mixture model" combining a circular Gaussian (von Mises) probability density function (pdf) and a uniform guessing component:

$$p(\theta)=(1-\gamma)\phi_{\mu,\kappa}+\gamma\left(\frac{1}{2\pi}\right), \quad \text{(A)}$$

where θ is the difference in radians between the reported and target color values, γ is the proportion of trials in which the subject responds at random, and φ is the von Mises distribution with mean μ and concentration κ (standard deviation = $1/\kappa$).

B. Model combining three Gaussian distributions, centered on the spatiotopic target, retinotopic distractor, and control distractor color values, respectively.

$$p(\theta)=(1-\beta-\delta)\phi_{0,\kappa_0}+\beta\,\phi_{\pi,\kappa_\pi}+\delta\,\phi_{-\pi,\kappa_{-\pi}}, \quad \text{(B)}$$

where the means of the von Mises distributions (φ) are fixed at 0 (spatiotopic target), π (retinotopic distractor) and $-\pi$ (control distractor). β is the probability of misreporting the retinotopic color value, δ is the probability of misreporting the control color value, and the parameters κ_0, κ_π, and $\kappa_{-\pi}$ (standard deviation = $1/\kappa$) vary independently.

C. Combination model allowing for both a shift in spatiotopic percept and a misreport of retinotopic distractor colors, plus guessing.

$$p(\theta)=(1-\beta-\gamma)\phi_{\mu,\kappa}+\beta\,\phi_{\pi,\kappa_\pi}+\gamma\left(\frac{1}{2\pi}\right), \quad \text{(C)}$$

where γ is the probability of random guessing, β is the probability of misreporting the retinotopic color value (defined by a von Mises distribution with a fixed mean of π and flexible κ_π), and μ and κ are the mean and concentration of the primary von Mises distribution.

Maximum likelihood estimates of the parameters μ, κ, κ_0, $\kappa_{-\pi}$, κ_π, γ, β, and δ were obtained separately for each subject and condition using Matlab's fminsearch optimization procedure (Nelder & Mead, 1965). A range of initial parameter values were tested to ensure that global minima were reached.

Results

In Experiment 1, on "no-saccade" trials response distributions were centered on the correct spatiotopic color value (Figure 2), with means not significantly deviating from 0° (t's<1, p's>0.38, for early and later delays). However, when the stimuli were presented 50ms after completion of a saccade, the distribution was subtly but significantly shifted in color space in the direction of the retinotopic distractor color (shift=7.7°, t(15)=3.12, p=0.007). This shift was only present immediately after the saccade: when stimuli were presented 500ms after the saccade, the distribution was again centered around 0° (t(15)=0.22, p=0.828); a

pairwise comparison confirmed that the retinotopic bias was significantly greater in the post-saccade early than post-saccade later condition (t(15)=2.25, p=0.040).

In Experiment 2, subjects performed the same task, but reported the retinotopic color. Strikingly, there was no influence of the spatiotopic distractor on the retinotopic percept (Figure 2B,D). None of the means were significantly deviated from 0° -- in fact, at the critical post-saccade early delay, the mean color reported was only 0.29° different from the true retinotopic color (t(8)=0.372, p=0.720), and there was no significant difference between post-saccade conditions (t(8)=1.19, p=0.270). A between-groups comparison revealed a significant difference between the retinotopic bias in Experiment 1 and the spatiotopic bias in Experiment 2 (t(17.8)=2.87, p=0.010, equal variances not assumed; Linear Mixed Model: F(1,23)=4.838, p=0.038).

These two experiments reveal a highly selective new form of perceptual interference: systematic color misperception is induced following a saccade, driven only by a retinotopic distractor (equidistant control distractors do not alter the spatiotopic percept, nor do spatiotopic distractors alter the retinotopic percept), and only for a brief period of time. What is the source of this interference? To evaluate contributions of different sources of error, we fit the data with probabilistic "mixture models" (Figure 3).

A standard "mixture model" (Zhang & Luck, 2008) assumes that performance can be characterized as a mixture of trials in which the subject successfully perceived the stimulus (with some Gaussian deviation around the correct response, such that standard deviation reflects the resolution of the representation), and some trials in which the subject randomly guessed (uniform distribution). Mixture models can also test another important source of error: the probability of misreporting (swapping) one of the distractor colors instead of the target color (Bays et al., 2009). We thus asked whether the reported retinotopic bias was driven by a shift in the mean of the Gaussian and/or an increase in the probability of retinotopic swapping. Given the large number of free parameters in the models, we first tested each of these effects in isolation, and then tested a combined model including both types of error.

In the first model (single-Gaussian plus guessing), standard deviation did not significantly differ across delays (t's<1 and p's>0.6 for both tasks). The probability of random guessing was slightly higher in the post-saccade early condition in both tasks (though only significant in Expt 1: t(15)=2.30, p=0.036; Expt 2: t(8)=1.28, p=0.236). Critically, however, the mean of the distribution representing the "successful" trials was significantly shifted in the postsaccade early condition of Experiment 1, in the direction of the retinotopic distractor color (t(15)=2.97, p=0.01).

In the second class of models, we estimated the probability of misreporting one of the neighboring distractors instead of the correct target color. A misreport of the retinotopic or control colors would result in additional peaks in the distribution at 90° and −90°, respectively. (Although distributions were aligned for figures and analyses with retinotopic set at 90°, in the actual task retinotopic was equally likely to be colored 90° or −90° different from the spatiotopic color.) The probability of misreporting one of the distractor

colors was relatively low in most conditions (Fig. 3); however, the probability of misreporting the retinotopic color doubled in the post-saccade early condition, and was significantly greater than control misreport at this delay (t(15)=2.35, p=0.033).

Fitting the data with a combination model (Figure 4A–B) revealed significant effects of both sources of retinotopic interference: The model captured the increased probability of retinotopic misreport, but even after accounting for these trials, the primary distribution was still significantly shifted toward the retinotopic color value at the post-saccade early delay (t(15)=2.52, p=0.024). In other words, even on trials when subjects think they are reporting the correct spatiotopic color, the retinotopic color is unconsciously bleeding into the spatiotopic percept.

To test the dual sources of error in a different way, in Figure 4C we binned the raw data as a function of absolute distance from the correct spatiotopic color value, to directly compare the retinotopic and control halves of the distribution. Critically, there were two distinct portions where the curves diverged. Responses in the retinotopic direction were more common than in the control direction in the bin centered at 90° (t(15)=2.17, p=0.046), consistent with misreport/swapping. But retinotopic influence was also significant at much smaller deviations in colorspace (bins centered at 26°, 39°, 51°: t's>2.62, all p's<0.02), supporting the perceptual mixing hypothesis.

Finally, in Experiments 3 and 4 we tested whether these two sources of error are only caused by eye movements, or if they would be present in other tasks involving the shifting or splitting of attention (Figure 5) without saccades. Experiment 3 was designed to induce shifting of spatial attention from one location to another, without involving saccades. We found clear evidence for swapping errors at the early delay, with subjects more likely to erroneously report the color at the original cue location than at the equidistant control (bins centered at 64°, 77°, and 90°: t(11)=2.24, p=0.046; t(11)=2.79, p=0.018; and t(11)=3.14, p=0.0094, respectively; model fits for probability of cue1 vs control misreport: t(11)=2.18, p=0.052). However, we did not find evidence for mixing errors in this context (t(11)=0.38, p=0.713). On the other hand, in Experiment 4, when attention had to be shared between two locations, the pattern of responses was consistent only with mixing errors, with subjects more likely to make subtle errors in the direction of the other cued color than the control color (bins centered at 26° and 39°: t(17)=3.05, p=0.007 and t(17)=2.15, p=0.046, respectively; mean of combination model distribution shifted toward other cued color: t(17)=2.31, p=0.034). Swapping errors were not seen in this context (pMisreport other cue vs control: t(17)=0.54, p=0.594). A significant between-group interaction with Experiment x Tail (other cue vs control) x Bin(1:14) confirmed the difference in error patterns in the shifting vs splitting contexts (F(4.2,116.8)=3.79, p=0.005, equal variances not assumed).

Discussion

This report documents a new perceptual/attentional phenomenon: the systematic distortion of color perception caused by residual retinotopic interference. Our primary goal was to better understand how attention remaps across saccades, and whether this process affects feature binding. We discovered a pattern of binding errors – a systematic bias – that not only

carries important implications for stability across saccades, but also sheds light on attentional mechanisms in general.

The distortion we found after a saccade was highly spatially and temporally specific. It occurred only for a brief period after each eye movement, temporally overlapping with the "retinotopic attentional trace" (Golomb et al., 2008, 2010). It was driven by the presence of a distractor color in the retinotopic – but not equidistant control – location. And finally, it was asymmetric: the spatiotopic task was susceptible to retinotopic interference, but not vice versa, which is particularly notable given that spatiotopic coordinates are the more ecologically relevant and intuitive coordinate system.

Critically, this perceptual distortion arose from two distinct types of errors: a "swapping" of features (Treisman & Schmidt, 1982), and a "mixing", or blending, in feature space between features from two different locations. Perceptual blending has been reported in other contexts; e.g., an object's features may be biased by and/or averaged with the features of other objects in the display or in memory (Brady & Alvarez, 2011; Hsieh & Tse, 2009; Huang & Sekuler, 2010), but here it is particularly notable how specific – and distinct from swapping – the mixing is. Our additional experiments conducted without saccades suggest that these two types of errors stem from different attentional mechanisms. Swapping errors were found when subjects shifted the locus of attention from one location to another and misreported the color at the previous location, as if attention had not had time to update on those trials. Mixing errors, on the other hand, were found when subjects simultaneously attended to two locations (but were only tested on one). These data suggest that swapping errors stem from incomplete updating, whereas mixing errors occur when two locations simultaneously share attentional resources.

This set of findings has potential implications for attentional updating in a wide range of contexts, as future studies may explore more fully. In terms of remapping of attention across saccades, it follows from the retinotopic attentional trace that delayed spatial updating could cause feature errors after a saccade, but we could not predict whether these errors would be swapping or mixing. The fact that both types of feature errors occur immediately following a saccade indicates that not only does attention take time to update following each saccade, but – crucially – at some point during the remapping process, attention is simultaneously selecting two different locations. (Interestingly, unlike the attention shifting and splitting contexts, in the saccade context, subjects aren't explicitly attending to two different locations; the task is to maintain attention at a single spatiotopic location, making the binding errors even more remarkable.)

These data support the hypothesis that the remapping of attention entails two temporally overlapping stages: updating to the spatiotopic location and disengaging from the previous retinotopic location. We have previously raised the idea of a two-stage remapping process; e.g., based on evidence that there is a point in time where both retinotopic and spatiotopic locations are facilitated, but not the locations in between (Golomb et al., 2011). However, these prior results could be caused by two independent processes/phases (i.e., a "turning on" of the new location that occurs before the "turning off" of the previous location), or by a single-stage remapping process that occurs with variable latency (such that on some trials

attention has already updated to the spatiotopic location, and on others it is still stuck at the retinotopic location). If the latter were true, then in the current study we should expect a mixture of fast remapping trials, where subjects correctly report the color at the spatiotopic location (within some normally distributed variance), and slower remapping trials, where we might see "swapping" errors if attention were still stuck at the retinotopic location. In other words, in a one-stage model, on any given trial attention should either be still stuck at the retinotopic location, or already updated to the spatiotopic location, but not both. The existence of "mixing" errors suggests that there is a period of time where both locations are still active. Thus, even after spatial pointers have been updated to the correct spatiotopic location, lingering facilitation at the retinotopic location means that retinotopic distractors can continue to interfere with perception.

These data converge to paint a picture where retinotopic representations are the "native language" of the visual system, and although spatial pointers or receptive fields can shift to the updated location in anticipation of a saccade (Cavanagh et al., 2010; Duhamel et al., 1992; Rolfs, et al, 2011), lingering processing at the previously attended retinotopic location can carry costs for stability even after the saccade is completed. A system in which retinotopic representations serve as the default but can be converted into other reference frames on demand allows for flexible and neurally efficient representations (Cohen & Andersen, 2002), but can also carry costs for behavior, such as a loss of spatial precision with each update (Golomb & Kanwisher, 2012b). The current study reveals that these potential costs are not limited to the encoding of spatial locations, but affect the binding of features to those locations, as features from two different locations may be simultaneously bound to the same object.

The ability to maintain or remap spatial attention is an important aspect of visual stability, and our report is a striking example of how the perceptual world is not nearly as stable as it feels. Understanding the mechanisms – and errors – of attention across eye movements is crucial, as saccades are arguably the most frequent shifts of attention we make during daily life (2–3 per second). Importantly, it is not just location information that is disrupted by a saccade; object features can also be distorted, and in ways more complicated than simple location swapping. Such perceptual instabilities could have important consequences for real-world visual processing, where multiple objects are often simultaneously present in the environment.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

Bays PM, Catalao RF, Husain M. The precision of visual working memory is set by allocation of a shared resource. J Vis. 2009; 9(10):7.1–11. [PubMed: 19810788]

Brady TF, Alvarez GA. Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. Psychol Sci. 2011; 22(3):384–92. [PubMed: 21296808]

Brainard DH. The psychophysics toolbox. Spat Vis. 1997; 10(4):433–436. [PubMed: 9176952]

Burr DC, Ross J, Binda P, Morrone MC. Saccades compress space, time and number. Trends Cogn Sci. 2010; 14(12):528–33. [PubMed: 20971675]

Cavanagh P, Hunt AR, Afraz A, Rolfs M. Visual stability based on remapping of attention pointers. Trends Cogn Sci. 2010; 14(4):147–53. [PubMed: 20189870]

Cohen YE, Andersen RA. A common reference frame for movement plans in the posterior parietal cortex. Nat Rev Neurosci. 2002; 3(7):553–62. [PubMed: 12094211]

Crespi S, Biagi L, d'Avossa G, Burr DC, Tosetti M, Morrone MC. Spatiotopic coding of BOLD signal in human visual cortex depends on spatial attention. Plos ONE. 2011; 6(7):e21661. [PubMed: 21750720]

d'Avossa G, Tosetti M, Crespi S, Biagi L, Burr DC, Morrone MC. Spatiotopic selectivity of BOLD responses to visual motion in human area MT. Nat Neurosci. 2007; 10(2):249–55. [PubMed: 17195842]

Duhamel JR, Colby CL, Goldberg ME. The updating of the representation of visual space in parietal cortex by intended eye movements. Science. 1992; 255(5040):90–2. [PubMed: 1553535]

Gardner JL, Merriam EP, Movshon JA, Heeger DJ. Maps of visual space in human occipital cortex are retinotopic, not spatiotopic. J Neurosci. 2008; 28(15):3988–99. [PubMed: 18400898]

Golomb JD, Kanwisher N. Higher level visual cortex represents retinotopic, not spatiotopic, object location. Cereb Cortex. 2012; 22:2794–2810. [PubMed: 22190434]

Golomb JD, Kanwisher N. Retinotopic memory is more precise than spatiotopic memory. Proc Natl Acad Sci U S A. 2012; 109(5):1796–1801. [PubMed: 22307648]

Golomb JD, Chun MM, Mazer JA. The native coordinate system of spatial attention is retinotopic. J Neurosci. 2008; 28(42):10654–62. [PubMed: 18923041]

Golomb JD, Marino AC, Chun MM, Mazer JA. Attention doesn't slide: Spatiotopic updating after eye movements instantiates a new, discrete attentional locus. Atten Percept Psychophys. 2011; 73(1): 7–14. [PubMed: 21258903]

Golomb JD, Nguyen-Phuc AY, Mazer JA, McCarthy G, Chun MM. Attentional facilitation throughout human visual cortex lingers in retinotopic coordinates after eye movements. J Neurosci. 2010; 30(31):10493–506. [PubMed: 20685992]

Hsieh PJ, Tse PU. Feature mixing rather than feature replacement during perceptual filling-in. Vision Res. 2009; 49(4):439–50. [PubMed: 19135078]

Huang J, Sekuler R. Distortions in recall from visual memory: Two classes of attractors at work. J Vis. 2010; 10(2):24.1–27.

Hunt AR, Cavanagh P. Remapped visual masking. J Vis. 2011; 11(1):13. [PubMed: 21245278]

Latour PL. Visual threshold during eye movements. Vision Res. 1962; 2:261–262.

Nelder JA, Mead R. A simplex method for function minimization. The Computer Journal. 1965; 7(4): 308.

Rolfs M, Jonikaitis D, Deubel H, Cavanagh P. Predictive remapping of attention across eye movements. Nature Neuroscience. 2011; 14:252–256.

Ross J, Morrone MC, Burr DC. Compression of visual space before saccades. Nature. 1997; 386(6625):598–601. [PubMed: 9121581]

Ross J, Morrone MC, Goldberg ME, Burr DC. Changes in visual perception at the time of saccades. Trends Neurosci. 2001; 24(2):113–21. [PubMed: 11164942]

Treisman A. The binding problem. Curr Opin Neurobiol. 1996; 6(2):171–8. [PubMed: 8725958]

Treisman A, Schmidt H. Illusory conjunctions in the perception of objects. Cognit Psychol. 1982; 14(1):107–41. [PubMed: 7053925]

Wilken P, Ma WJ. A detection theory account of change detection. J Vis. 2004; 4(12):1120–35. [PubMed: 15669916]

Wittenberg M, Bremmer F, Wachtler T. Perceptual evidence for saccadic updating of color stimuli. J Vis. 2008; 8(14):9.1–9. [PubMed: 19146310]

Zhang W, Luck SJ. Discrete fixed-resolution representations in visual working memory. Nature. 2008; 453(7192):233–5. [PubMed: 18385672]
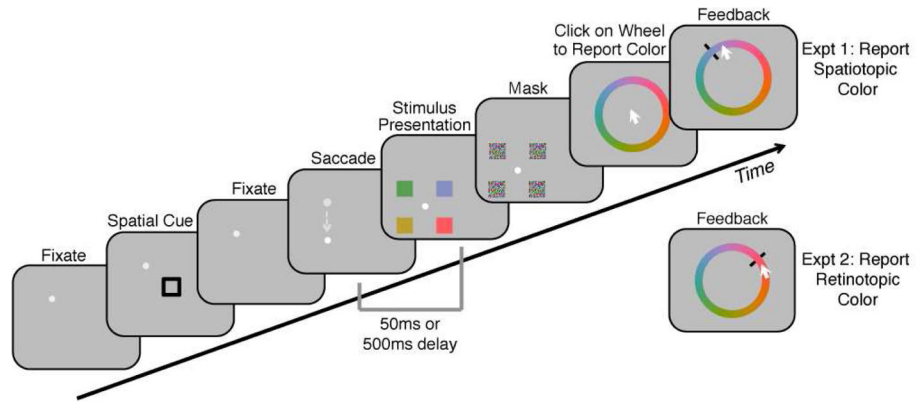
Figure 1. Task (Experiments 1–2: "Maintain attention across saccade")
While subjects fixated the fixation dot, a spatial pre-cue briefly appeared. Subjects were instructed to report the color of whichever stimulus subsequently appeared at the cued location, in either spatiotopic (Expt 1) or retinotopic (Expt 2) coordinates. On "saccade" trials, as depicted, the fixation dot moved between cue and stimulus presentation, and subjects needed to accurately saccade to the new fixation location. At either 50ms or 500ms after completion of the saccade (determined by real-time eye-tracking), an array of 4 colored stimuli appeared for 50ms, followed by a mask array. A large color wheel (randomly rotated) was then presented at the center of the screen, and subjects moved the mouse to report the appropriate target color. On "no-saccade" trials, the fixation dot never moved, and the array of stimuli was presented around the current fixation location. Note that the example pictured is just one of several combinations of saccade and cue configurations, including both horizontal and vertical saccades.
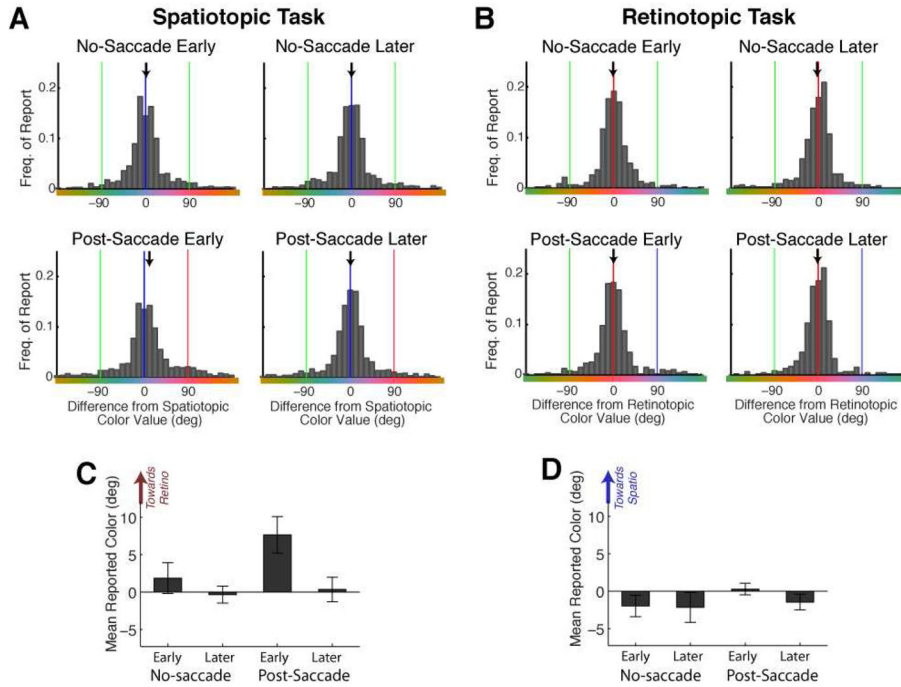
Figure 2. Distribution of Responses
A–B, Response histograms (combined across subjects) are shown for each condition and task. **A,** Spatiotopic Task (Expt 1). Data are plotted as differences in color value relative to the correct spatiotopic color. Difference scores were calculated by aligning all trials such that the spatiotopic color was defined as 0°, and the retinotopic distractor color was +90°. Note, however, that the actual retinotopic color could have been distanced 90° in either direction along the color wheel – the color strip shown here is just for illustrative purposes. **B,** Retinotopic task (Expt 2). Data are plotted as differences in color value relative to the correct retinotopic color (0°), aligned so that spatiotopic color is at +90°. Spatiotopic, retinotopic, and control color values are indicated with blue, red, and green lines, respectively. The black arrows indicate the mean of the distribution. **C–D,** Mean reported color value is re-plotted for each condition and task. Error bars are SEM. N=16 (Spatiotopic Task); N=9 (Retinotopic Task).
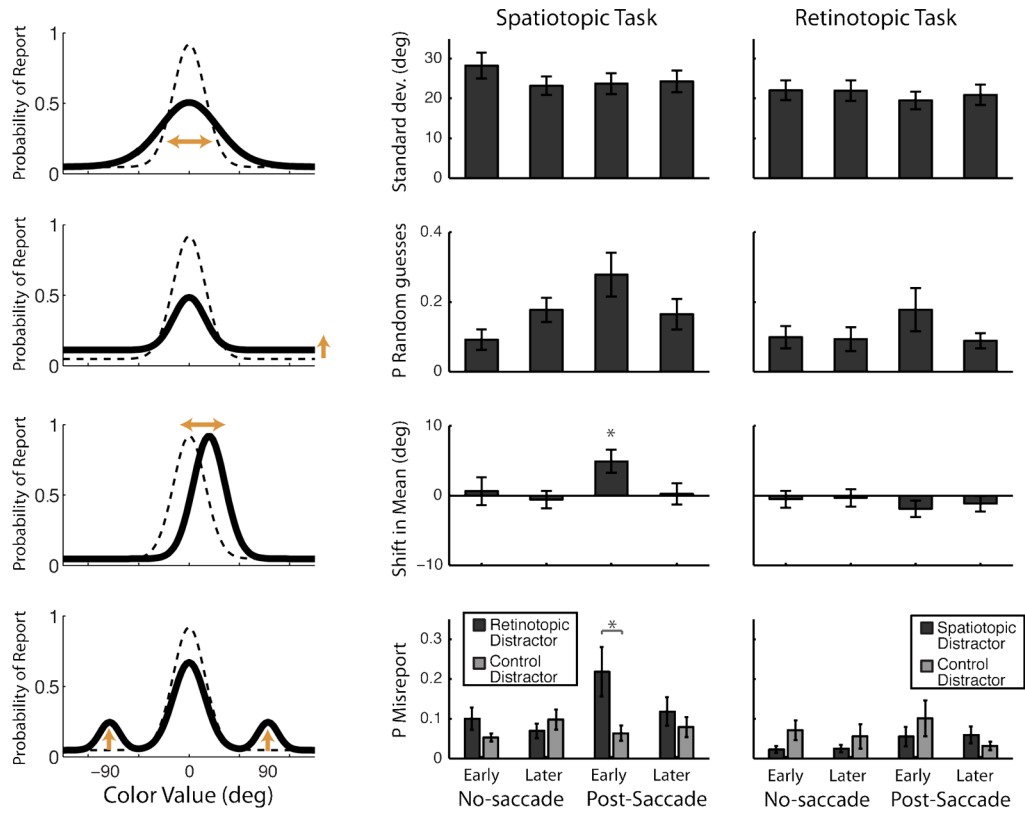
Figure 3. Model fits

Maximum likelihood estimate fits for different parameters of the "mixture models". Model A (rows 1–3) combines a circular Gaussian (von Mises) distribution with a uniform (random guessing) distribution. Best-fit model parameters for standard deviation, proportion of random guesses, and shift in mean are shown for spatiotopic and retinotopic tasks at right. Left panels illustrate each error source with cartoon models: the dotted line represents the baseline "no-saccade" distribution, and the thick black line depicts possible ways the distribution could change after a saccade. The final row depicts Model B, which accounts for errors caused by misreporting one of the distractor colors instead of the target color. Models were fit separately for each subject and then parameter values were averaged across subjects. Error bars are SEM. N=16 (Spatiotopic Task); N=9 (Retinotopic Task).
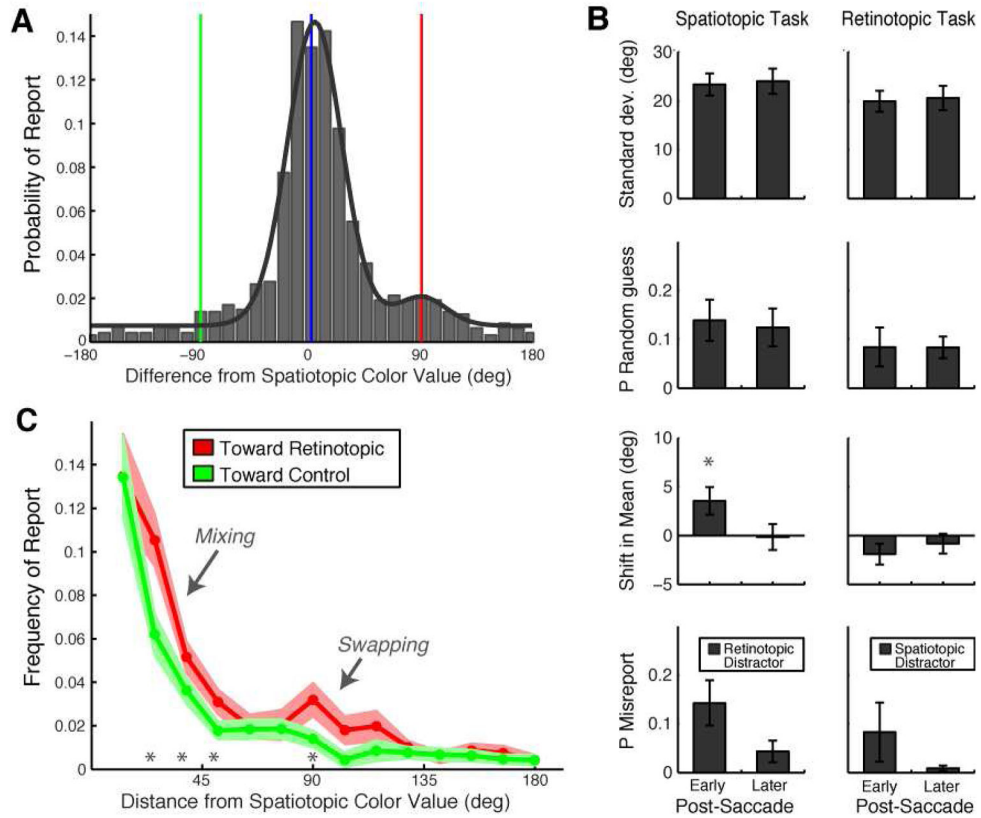
Figure 4. Two sources of retinotopic error

A, Data from the spatiotopic task, post-saccade early delay fit with the combination model (Model C). The combination model includes a von Mises distribution with flexible mean and standard deviation, a second "misreport" von Mises distribution centered on the retinotopic distractor, and a uniform (guessing) distribution. **B,** Best-fit parameters reflecting each source of error. Data are only shown for post-saccade conditions, since the misreport distribution is specifically defined for the retinotopic distractor (or spatiotopic in Expt 2). Error bars are SEM. N=16 (Spatiotopic Task); N=9 (Retinotopic Task). **C,** The raw data for the spatiotopic task, post-saccade early delay are binned and plotted as a function of absolute distance from spatiotopic color value. I.e., the two halves of the histogram are folded over one another for comparison. The retinotopic and control probabilities significantly deviate along two portions of the curve, reflecting the co-existence of "blending" and "misreport" effects. Light red and green shaded areas indicate the SEM around the retinotopic and control probabilities; asterisks indicate bins where the conditions significantly differed ($p < 0.05$). N=16.

Figure 5. Non-saccade attention experiments
A, Experiment 3: "Shift attention". As in Experiments 1–2, while subjects fixated, a spatial pre-cue briefly appeared. Subjects were instructed to attend to the cued location to report the color of whichever stimulus subsequently appeared there. On "shift" trials, a second cue appeared before stimulus presentation, and subjects needed to shift their attention and report the color at the final cued location. On no-shift trials (not shown), subjects only received the first cue. **B,** Experiment 4: "Split attention". Subjects initially received two simultaneous cues and were told to attend to both locations. Only after the colors had appeared were they given a post-cue instructing which of the two colors to report. Plots below depict the tails of the histograms for each task, as in Figure 4c, and the parameter estimates for the models, as in Figure 3. N=18 (Expt 3), N=12 (Expt 4).