CrossMark

Divided spatial attention and feature-mixing errors

Julie D. Golomb [1]

Abstract Spatial attention is thought to play a critical role in feature binding. However, often multiple objects or locations are of interest in our environment, and we need to shift or split attention between them. Recent evidence has demonstrated that shifting and splitting spatial attention results in different types of feature-binding errors. In particular, when two locations are simultaneously sharing attentional resources, subjects are susceptible to feature-mixing errors; that is, they tend to report a color that is a subtle blend of the target color and the color at the other attended location. The present study was designed to test whether these feature-mixing errors are influenced by target–distractor similarity. Subjects were cued to split attention across two different spatial locations, and were subsequently presented with an array of colored stimuli, followed by a postcue indicating which color to report. Target–distractor similarity was manipulated by varying the distance in color space between the two attended stimuli. Probabilistic modeling in all cases revealed shifts in the response distribution consistent with feature-mixing errors; however, the patterns differed considerably across target-distractor color distances. With large differences in color, the findings replicated the mixing result, but with small color differences, repulsion was instead observed, with the reported target color shifted *away* from the other attended color.

✉ Julie D. Golomb
golomb.9@osu.edu

[1] Department of Psychology, The Ohio State University, 201 Lazenby Hall 1827 Neil Avenue, Columbus, OH 43210, USA

Spatial attention has long been thought to play a role in feature binding. In Treisman's classic feature integration theory, spatial attention is described as the "glue" that binds object features together (Treisman, 1988, 1998; Treisman & Gelade, 1980). The idea is that spatial attention helps solve the "binding problem," such that features falling within the same window of attention are grouped together into a coherent object, and features outside this focus are excluded. When attention is diverted, binding errors can occur, resulting in "illusory conjunctions" (Treisman & Schmidt, 1982)—for example, if a subject were to view a blue circle and a red square, but report seeing a blue square.

If spatial attention is important for accurate feature binding, then what happens to feature binding when we need to shift or split attention? In the real world, multiple objects or locations are often of interest in the environment. Golomb, L'Heureux, and Kanwisher (2014) recently reported that a unique pattern of feature-binding errors can occur under circumstances with unstable or ambiguous spatial attention. In the few hundred milliseconds following a saccadic eye movement—when spatial attention has not fully remapped—Golomb et al. reported both "swapping" errors (misreporting a distractor color in the display, similar to illusory conjunctions) and "mixing" errors (reporting a color blend between the target color and the retinotopic distractor color). These effects were not specific to eye movements, but could be found under other, more general attentional scenarios as well: "Swapping" errors were found when spatial attention needed to be rapidly shifted from one location to another, whereas "mixing" errors occurred when two spatial locations were simultaneously attended.

The present article explores these feature-mixing errors in more detail, asking whether the mixing effect is sensitive not only to the distribution of spatial attention, but also to the properties of the features themselves. Specifically, target–

distractor similarity was manipulated, by varying the distance in color space between the items. Target–distractor similarity could be expected to influence feature-mixing errors in a few different ways. On the one hand, distractors that are more similar to the target might make the target less distinctive and increase errors, as has been observed in visual search (Duncan & Humphreys, 1989), multiple-object tracking (Makovski & Jiang, 2009), working memory (Shapiro & Miller, 2011), and illusory conjunction (Donk, 1999) tasks. On the other hand, greater similarity might decrease errors and improve precision for the target, as in change detection tasks (Lin & Luck, 2009).

Alternatively, since in the present task the target was cued and defined spatially, similarity in feature space might have less of a fundamental influence on the type or amount of errors; rather, the mean of the response distribution (i.e., the reported target color) might simply shift accordingly based on the target–distractor distance. In other words, if the distractor is similar in color to the target, the shift would be small, but if the distractor is farther in color space from the target, the shift would be larger, reflecting a mix between two very different colors. Finally, however, a different sort of prediction could be made: When very similar colors share the focus of attention, they might actually inhibit each other, resulting in repulsion away from the distractor, rather than blending toward it. In a 2009 article, Johnson, Spencer, Luck, and Schöner described a dynamic neural field model of visual working memory, which included implications for target–distractor similarity; when two working memory items are similar in feature space (e.g., color), they interact in a strongly inhibitory fashion, which makes it easier to detect subsequent changes at test.

In the present task, these possibilities were explored by adopting the same paradigm as in the split-attention experiment of Golomb et al. (2014, Exp. 4). 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 color wheel (Fig. 1). Before stimulus presentation, subjects were cued to attend to two of the four spatial locations. The to-be-reported color was always in one of the attended locations, and was indicated with a postcue during the response period. In the original Golomb et al. study, the adjacent distractor colors were always quite different in color from the target, located ±90° along the color wheel in color space. In the present experiment, distance in color space was varied systematically, by testing distractors both closer to and farther from the target color. The analyses investigate whether target–distractor similarity alters the probability and nature of feature-binding errors, by examining the distribution of responses and using probabilistic mixture modeling to evaluate how the reported target color was influenced by the color of a distractor sharing the focus of attention.

Materials and method

Subjects

Twenty subjects (14 female, six male; mean age = 19.4 years) participated in the experiment. Four additional subjects were excluded for not successfully performing the task (>50 % probability of random guessing, according to the γ parameter from Model A; see Golomb et al., 2014).

Experimental setup

Stimuli were generated using the Psychophysics Toolbox extension (Brainard, 1997) for MATLAB and presented on a 21-in. flat-screen CRT monitor. Subjects were seated at a chinrest 61 cm from the monitor, and their eye position was monitored using an EyeLink 1000 eyetracking system recording pupil and corneal reflection. The monitor was color calibrated with a Minolta CS-100 colorimeter.

Task

The task (Fig. 1) was the same as in Golomb et al.'s (2014) Experiment 4. Each trial began with a white fixation dot presented at the center of the screen. Once subjects had accurately fixated for 1.5 s (determined by real-time eyetracking), two of the four stimulus locations (adjacent horizontal or vertical locations) were simultaneously precued. The stimulus locations were 2° × 2° squares located to the upper left, upper right, lower left, and lower right of fixation (7.4° eccentricity). The precues were black square outlines presented for 200 ms. Subjects were instructed to attend to both precued locations (i.e., to share or split attention). After another 1.5-s fixation period, an array of four differently colored squares appeared at the stimulus locations. The colored squares appeared for 50 ms, followed by a 200-ms mask (colored with a random color value at each pixel location, covering each of the four stimulus locations).

A large color wheel (diameter 16.4°) was then presented at the center of the screen at a random rotation. When the color wheel appeared, a postcue was also presented, indicating to subjects to report the color that had appeared at that location. The postcued location was always one of the two precued locations, but it was unpredictable which one. Subjects clicked with the mouse to report the color of the stimulus at the postcued location. They were then given feedback showing them the correct color. At any point in the trial, if the subject's eye position deviated more than 2° from the fixation location, the trial was immediately aborted and repeated later in the block.

Stimulus colors were chosen as follows: The color at the postcued (target) location was chosen randomly on each trial from 180 possible colors (evenly distributed along a circle in
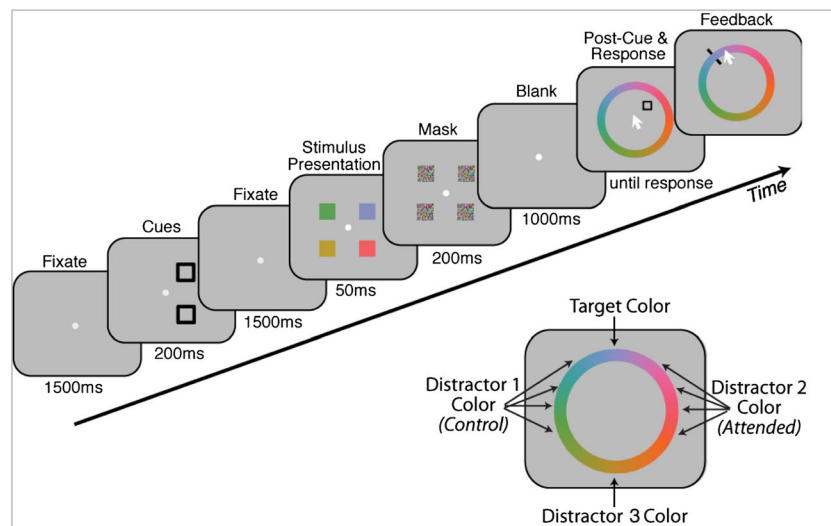
Fig. 1 Task. While subjects fixated the fixation dot, two spatial precues briefly appeared, in adjacent horizontal or vertical positions. Subjects were instructed to attend to both locations (split attention). An array of four colored stimuli was then presented briefly and masked. Only after the stimuli disappeared were subjects given a postcue instructing which of the two attended stimuli was the target. A large color wheel (randomly rotated) was presented at the center of the screen, and subjects used the mouse to report the target color. (Inset) The similarity in color space between the target and distractor colors was manipulated: The attended distractor color could be either 30°, 60°, 90°, or 120° different from the target color. The control distractor was always the same distance in color space from the target as the attended distractor, but in the opposite direction

CIE L*a*b* color space, according to the parameters in Zhang & Luck, 2008, and Golomb et al., 2014). The color at the other attended location was chosen to be 30°, 60°, 90°, or 120° different in color space from the target color (clockwise or counterclockwise along the color wheel, with the direction and magnitude varying pseudorandomly from trial to trial). The color at the equidistant control location was chosen to have the same-magnitude color difference, but in the opposite direction. The stimulus at the fourth location was always set 180° away in color space.

The distractor colors were designed to be symmetric in color space around the target color so that the attended distractor and control distractor would be equidistant in both feature and physical space. The fourth color was included to balance the array, so that the target location would not be predictable. However, in the 30°, 60°, and 120° conditions, there was a possibility that subjects technically could have predicted which color would be the target, on the basis of the symmetric distribution in color space. This possibility seems quite unlikely in practice, though, since the extremely short, masked presentation times would have made it almost impossible for subjects to perceive all four colors in enough detail to figure out which one was the target, especially when the colors and similarity distances varied randomly from trial to trial.

Analyses

The location on the color wheel where subjects clicked on each trial was recorded and converted into a difference score in degrees of angle, with the correct target color having a 0°

difference. Responses in the direction of the other attended location color were aligned to be positive differences, with the control color direction being negative. The mean of the distribution was calculated separately for each subject and condition, and submitted to within-subjects analyses of variance (ANOVAs) and t tests.

The distribution of responses was then fit with probabilistic mixture models (Bays, Catalao, & Husain, 2009; Golomb et al., 2014; Zhang & Luck, 2008) accounting for various sources of error. Several variations of models were tested that included parameters for single Gaussian (target color) distributions, multiple Gaussians ("swapping"/misreport of distractor colors), and uniform guessing components. For each of these models, log-likelihood and AIC (Akaike information criterion: Akaike, 1974) values were calculated for each subject. Lower AIC values indicate a better model fit after taking into account the number of parameters in the model. On the basis of the AIC comparison, the subsequent analysis focused on two variations of the model. The lowest AICs were found for a simple model with a single Gaussian (with flexible mean) plus noise (Model A below). In addition, data are presented from the best-fitting variation of a more complex model that included parameters for "swapping" errors from the distractor colors (Model B below).

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

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

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

B. Full mixture model, allowing for both a shift in mean and misreport of the distractor colors, plus guessing.

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

where γ is the probability of random guessing, β is the probability of misreporting the other attended color value (defined by a von Mises distribution with a fixed mean centered on the attended distractor color value), δ is the probability of misreporting the control color value (defined by a von Mises distribution with a fixed mean centered on the control distractor color value), μ is the mean of the primary von Mises distribution, and κ is the concentration of the distributions.

Maximum-likelihood estimates of the parameters μ, κ, γ, β, and δ were obtained separately for each subject and condition using the MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013) and MATLAB's fminsearch optimization procedure (Nelder & Mead, 1965). A range of initial parameter values were tested to ensure that global minima were reached.

Results

For each condition, response histograms were generated (Fig. 2A), plotting responses on each trial in terms of the difference in color value between reported colors and correct target colors. To assess the effects of attended versus control distractors on these distributions, the data were analyzed in several ways:

First, the mean reported color was calculated for each subject and condition (Fig. 2C). A value of 0 would indicate no systematic deviation from the correct color; positive values indicate a greater tendency to report colors closer to the other attended color ("attraction"), and negative values indicate a greater tendency to report colors in the opposite direction (closer to the control color; "repulsion"). A one-way ANOVA revealed a significant main effect of target–distractor similarity [$F(3, 57) = 14.88$, $p < .001$, $\eta_p^2 = .44$], as well as a significant linear contrast [$F(1, 19) = 60.56$, $p < .001$, $\eta_p^2 = .76$]. When distractor colors were very different from the target color (90° or 120°), there was attraction toward the other attended color, as in Golomb et al. (2014). But when the distractor colors were similar to the target color, subjects were actually more likely to report colors shifted in the opposite
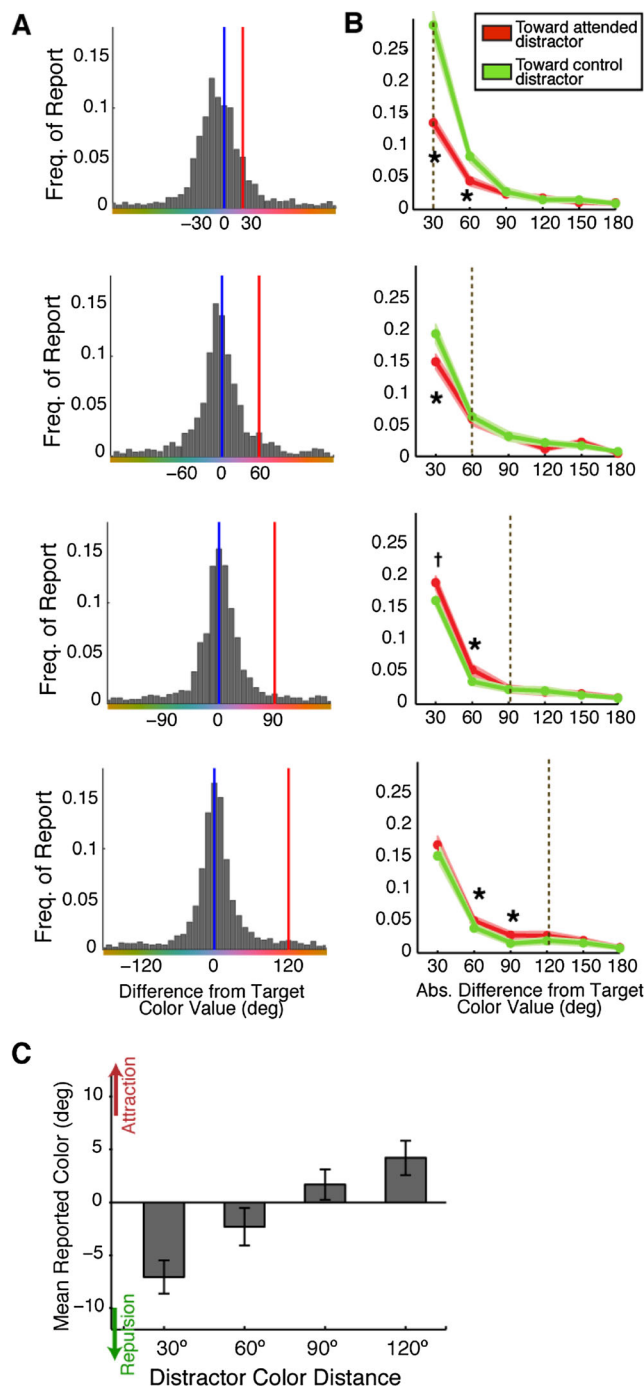
direction, as if there were repulsion—in feature space—away from the attended distractor.

The mean response provides a rough measure of biased responses, but it does not tell us what types of errors subjects were making to produce this shift. Golomb et al. (2014) demonstrated that when attention is split between two simultaneous cues, as here, the errors are driven by color "mixing" (a blend or shift toward the distractor color), rather than by color "swapping" (misreporting the color of the distractor

◀ **Fig. 2** Distribution of responses. (A) Response histograms (combined across subjects) are shown for each condition; data are plotted as differences in color values relative to the correct target color. Difference scores were calculated by aligning all trials such that the target color was defined as 0° and the attended distractor color was in the positive direction (+30°, 60°, 90°, or 120°). Note, however, that the actual attended distractor color could have been located in either direction along the color wheel—the color strip shown here is just for illustrative purposes. Vertical lines indicate target and attended distractor color values. (B) Responses are binned and plotted as a function of absolute distance from the target color value; that is, the two halves of the histogram are folded over one another for comparison. The lightly shaded areas around the lines indicate SEMs; asterisks indicate bins in which the two curves differed significantly ($p < .05$; the cross indicates $p < .10$). (C) Mean reported color values plotted for each condition. Error bars are SEMs; $N = 20$

instead of the target). These two types of errors can be dissociated using probabilistic mixture modeling (below), but they can also be visualized by comparing the two tails of the histogram and errors made in the directions of the attended versus control distractors. If the two halves of a histogram are compared to each other (Fig. 2B), mixing errors should be apparent at close-to-intermediate distanced bins, whereas swapping errors should appear as a peak at the actual color of the distractor (dashed lines).

For the larger target–distractor differences (60°, 90°, and 120°), no swapping errors were obvious, but there were signs of subtler mixing errors. For the largest target–distractor difference (120° condition), responses were shifted more toward the attended distractor color than in the control direction in the bins centered at 60° and 90° [$t(19) = 2.26$, $p = .036$, $d = 0.51$, and $t(19) = 2.31$, $p = .032$, $d = 0.52$, respectively]. For the 90° target–distractor condition, the effect was significant in the bin centered at 60° [$t(19) = 2.24$, $p = .037$, $d = 0.50$], and marginally so in the bin centered at 30° [$t(19) = 1.92$, $p = .069$, $d = 0.43$]. In the 60° target–distractor condition, the mixing effect reversed: Responses were more commonly shifted toward the control distractor color in the 30° bin [$t(19) = -2.18$ $p = .042$, $d = -0.49$]. The final, 30° target–distractor condition is harder to interpret, because at this small difference, mixing and swapping errors are not well dissociated. Nonetheless, responses were significantly shifted away from the attended distractor color and toward the control distractor color in the 30° and 60° bins [$t(19) = -5.92$ $p < .001$, $d = -1.32$, and $t(19) = -3.51$, $p = .002$, $d = -0.79$, respectively].

Finally, to quantify these effects using probabilistic modeling, the data from each subject and condition were fit to two types of mixture models (see the Materials and method section). Figure 3 illustrates the results from the full mixture model (Model B) and the simple mixture model (Model A). Both models included a primary, target-centered Gaussian distribution, from which it was possible to estimate the standard deviation (precision) of responses and whether the mean of this distribution was shifted from

0, as well as a uniform noise distribution, from which the probability of random guessing could be estimated. The simple model with these two components captured the variance well in all conditions, but the full mixture model allowed us to explicitly test for the probability of swapping errors as well, with additional Gaussians centered on the attended and control distractor colors.

In the full mixture model, neither the standard deviation nor the probability of random guessing significantly varied across conditions [$F(3, 57) = 0.86$, $p = .465$, $\eta_p^2 = .04$, and $F(3, 57) = 0.40$, $p = .756$, $\eta_p^2 = .02$, respectively]. For the swapping errors (pMisreport), there was a significant interaction between type of misreport (attended vs. control) and target–distractor similarity [$F(3, 57) = 24.29$, $p < .001$, $\eta_p^2 = 0.56$]; this effect was driven by a large probability of control misreports in the 30° target–distractor condition. Indeed, the difference between attended and control misreports was only significant in this 30° target–distractor condition [$t(19) = -6.07$, $p < .001$, $d = -1.36$]; in all other conditions, the probabilities of misreports were small and not significantly different between attended and control distractors (all ts < 1.20, all ps > .245). As was noted above, in the 30° target–distractor condition it is nearly impossible to dissociate errors driven by an increased probability of misreports versus a shift in the mean, because the small distractor distance falls within the normal standard deviation of the distribution. Figure 3C illustrates the fits of the full mixture model to the data; here, the increase in pMisreport looks nearly identical to a shift for this 30° condition.

The critical parameter for both models is the shift in the mean of the distribution. Both models revealed a consistent effect: The Shift parameter varied significantly as a function of target–distractor similarity [simple model: $F(3, 57) = 36.15$, $p < .001$, $\eta_p^2 = .66$; full model: $F(3, 57) = 4.69$, $p = .005$, $\eta_p^2 = .20$]. For the 30° and 60° target–distractor conditions, responses were shifted *away* from the attended distractor [simple model: $t(19) = -7.91$, $p < .001$, $d = -1.77$, and $t(19) = -2.09$, $p = .051$, $d = -0.47$, respectively; full model: $t(19) = -1.39$, $p = .18$, $d = -0.31$, and $t(19) = -2.45$, $p = .024$, $d = -0.55$, respectively]. For the 90° and 120° target–distractor conditions, on the other hand, responses were shifted *toward* the attended distractor [simple model: $t(19) = 2.76$, $p = .012$, $d = 0.62$, and $t(19) = 2.24$, $p = .037$, $d = 0.50$, respectively; full model: $t(19) = 2.08$, $p = .051$, $d = 0.47$, and $t(19) = 1.49$, $p = .152$, $d = 0.33$, respectively]. The estimates for this parameter are more reliable for the simple model because of the challenge noted above, and indeed, the analyses of model fit revealed better fits (lower AICs; see the Materials and method section) for the simple than for the full model; nonetheless, both models present a similar story, reinforcing the observations noted earlier: that a simultaneously attended distractor can bias the perceived target color, by either attractive or repulsive mixing of features, depending on target–distractor similarity.
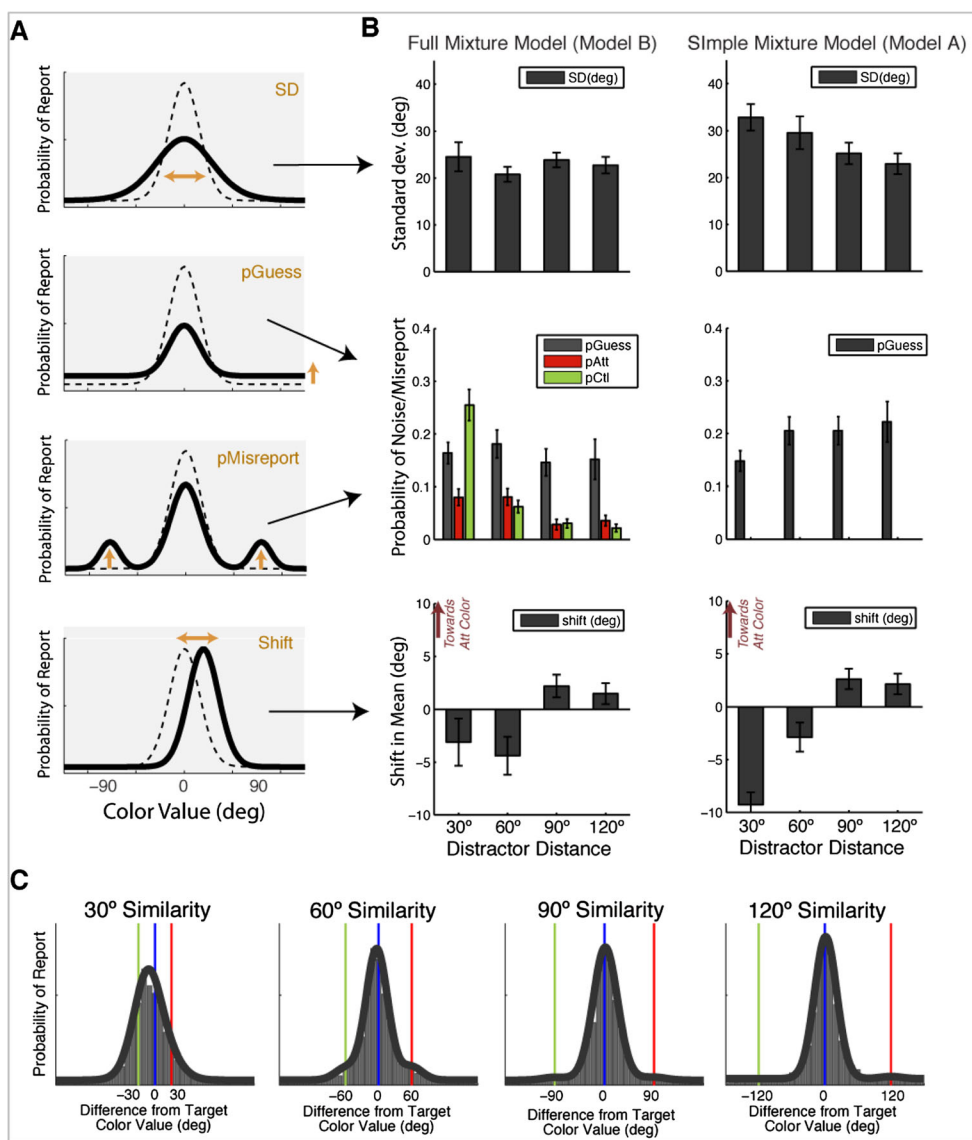
Fig. 3 Model fits. (A) Cartoon models showing different ways that a baseline distribution (dashed lines) could change as a result of increases in different error sources (thick black lines). (B) Maximum-likelihood estimate fits for corresponding parameters of the mixture models: standard deviations, probabilities of noise (guessing) and misreports (swapping of attended or control distractors), and shifts in mean are shown for the full and simple mixture models (see the Materials and

method section). Models were fit separately for each subject, and parameter values were then averaged across subjects. Error bars are *SEM*s. (C) Response histograms for each condition, fit with the full mixture model. The thick black lines show the best-fitting Model B, and the vertical lines indicate, from left to right, the control distractor, target, and attended distractor colors. $N = 20$

Hemifield-based effects

A number of studies have demonstrated attentional differences when stimuli are presented within versus across visual hemifields (Awh & Pashler, 2000), suggesting that the two cortical hemispheres may contain independent attentional resources (Alvarez & Cavanagh, 2005). This raises the interesting question of whether feature-mixing errors might be influenced by hemifield effects. The same analyses were conducted as above, but now with the data split into within-hemifield and

across-hemifield attention trials. None of the main effects or interactions with hemifield were significant [for mean reported color: main effect, $F(1, 19) = 2.53$, $p = .128$, $\eta_p^2 = 0.12$; interaction $F(3, 57) = 0.22$, $p = .88$, $\eta_p^2 = 0.01$; for Shift parameter from the simple model: main effect $F(1, 19) = 0.11$, $p = .74$, $\eta_p^2 = 0.01$; interaction $F(3, 57) = 0.25$, $p = .86$, $\eta_p^2 = 0.01$]. The lack of hemifield effects could be due to the fact that only two locations needed to be attended, and they were separated by a large enough spatial distance so as not to interfere or compete with each other (Clevenger & Beck, 2014).

Discussion

The results from this experiment make two contributions. First, when target and distractor colors are sufficiently different, the experiment replicated the finding that a distractor color can mix, or blend, with the perceived target color when spatial attention is split between the two locations (Golomb et al., 2014). Second, these feature-mixing errors are influenced by target–distractor similarity: Dissimilar colors result in attractive mixing toward the attended distractor color, whereas similar colors result in repulsion away from the attended distractor color. Note that both of these effects exceed any generic influence of other, control (unattended) distractors in the display.

This task was designed to manipulate spatial attention—that is, the extremely brief, masked stimulus presentations required that subjects attend to the precued locations in order to successfully perform the task—and we primarily think of these feature-mixing errors as being driven by attentional mechanisms. In Golomb et al. (2014), the attractive feature-mixing effect was explained as a result of attentional processes: When attention is split across two different spatial locations that are simultaneously occupied by objects containing different features, because both objects share attentional resources, their features are not perfectly distinguished, and may partially blend together. Thus, when the feature of one must be reported, subjects tend to report a color shifted in color space toward the other.

Similar attractive shifts have also been reported for items held serially (Fischer & Whitney, 2014; Huang & Sekuler, 2010) or simultaneously in visual working memory (Brady & Alvarez, 2011). For example, Brady and Alvarez had subjects remember the sizes of a set of circles; when subjects were asked to report the size of a single circle, their reports were biased toward the mean size of the set. Brady and Alvarez argued that this blending makes sense computationally, if you assume that items in the world tend to be similar to other items. It is possible that this optimal-integration account could play a role here as well, with the caveat that it is highly sensitive to selection by spatial attention (items were selectively biased by the attended distractor, relative to the equally differentiable control distractor). But neither this account nor the attentional-resource account described above provides a clear prediction of the finding of repulsion with increased target–distractor similarity.

Three possible explanations are proposed here for the repulsion effect. The first assumes a strategic difference at encoding. Subjects deploy attention to two spatial locations and know they must try to encode both colors (since either could be tested at the time of response). When two sufficiently different colors are presented, subjects may find it easy enough to encode both. But when two very similar colors are presented, subjects may try to make them more discriminable, so that they can correctly assign each to its respective location. For example, if subjects are presented with a greenish-blue color and a bluish-green color, they may—

either consciously or unconsciously—try to encode them as the bluer color at location A and the greener color at location B. Thus, the repulsion errors might be due to shifts at encoding, rather than the attention-splitting process per se disrupting binding. One could question whether the attractive mixing effects could also be driven by the demands of encoding two objects simultaneously, but this explanation could not account for the attractive feature-mixing errors initially reported by Golomb et al. (2014, Exp. 1) in the saccadic remapping context, since in that context subjects were only instructed to encode the color at one single attended location.

A second explanation follows a similar logic, but assumes that these effects manifest neurally during the decision stage. Recent studies have demonstrated that perceptual discriminations between two highly similar stimuli are served better by neurons tuned slightly away from the target (off-channel neurons), because these neurons are actually most informative (Navalpakkam & Itti, 2007; Scolari & Serences, 2009, 2010). Although in the present experiment the task was not to discriminate between the two colors, it is possible that responses could be similarly biased toward these off-target neurons.

The third explanation for the repulsion effect could be that it is driven by dynamic interactions between the items in working memory, as in Johnson et al.'s (2009) dynamic neural field model. Although the task here was designed to tax attention more than working memory, it did involve a small working memory component, and attention and working memory are known to interact (Awh, Vogel, & Oh, 2006). Johnson et al.'s model proposed that items in working memory are represented according to a three-layer model, with activation along the feature dimension for perceptual, inhibitory, and working memory fields. When two similar items are held in working memory, their shared inhibition results in a sharpening of their representations in working memory, which results in enhanced change detection (Johnson et al., 2009). Although the Johnson et al. article focused on change detection performance (as did Lin & Luck, 2009), this dynamic inhibition model makes an additional prediction: that the representations of two similar colors may not only be sharpened, but shifted away from each other. Such a shift could underlie the repulsion effects here, as well as similar repulsion effects reported in a working memory context by Johnson in an unpublished study (Johnson, 2008). Interestingly, the experimental contexts and results of these studies differed in a few key ways: As was noted earlier, in the present study, the emphasis was more on attentional selection than working memory; spatial precues were used to direct subjects to selectively attend to certain items in the display, and the working memory load was lower (only two items). Additionally, here the stimuli were presented for only 50 ms, increasing the likelihood that the reported errors were due to errors in perceptual feature binding rather than to working memory decay. Finally, an intriguing difference is that the Johnson et al. model predicts repulsion

for close target–distractor similarity, but no effect for far items. On the other hand, here robust *attractive* mixing emerged for far items (large target–distractor color differences). An interesting question is whether these findings of joint attraction and repulsion would generalize to more traditional working memory contexts with longer encoding times.

An important question for future research will be whether a single mechanism can explain both the attraction and repulsion effects, or whether multiple mechanisms contribute to them. Furthermore, what determines where exactly this transition from repulsion to attraction takes place?

Both the attractive and repulsive feature-mixing errors reported here raise important implications for feature binding in the context of multiple simultaneously attended items. Spatial attention is typically thought of as a way to *aid* feature binding, by allowing features within the spatial locus of attention to be bound together (Treisman, 1988, 1998; Treisman & Gelade, 1980). But this idea assumes that features within an attentional locus should belong to the same object. Is this an assumption that our visual system makes as well? If so, it would explain feature-binding errors stemming from ambiguities in the allocation of spatial attention: that when features belonging to two different objects in distinct spatial locations are both spatially attended, they can blend together or repulse each other in feature space. Of course, in the real world, multiple objects and spatial locations are often of interest, so it is unclear whether we regularly experience subtle versions of these errors in everyday functioning, or whether our visual systems adopt additional, compensatory mechanisms to avoid or minimize these errors.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, AC-19,* 716–723. doi:10.1109/TAC.1974.1100705

Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science, 16,* 637–643. doi:10.1111/j.1467-9280.2005.01587.x

Awh, E., & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 834–846. doi:10.1037/0096-1523.26.2.834

Awh, E., Vogel, E. K., & Oh, S. H. (2006). Interactions between attention and working memory. *Neuroscience, 139,* 201–208. doi:10.1016/j.neuroscience.2005.08.023

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*(10):7, 1–11. doi:10.1167/9.10.7

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. doi:10.1177/0956797610397956

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

Clevenger, J., & Beck, D. M. (2014). Refining the resource model: Cortical competition could explain hemifield independence. *Visual Cognition, 22,* 1022–1026. doi:10.1080/13506285.2014.960668

Donk, M. (1999). Illusory conjunctions are an illusion: The effects of target–nontarget similarity on conjunction and feature errors. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 1207–1233. doi:10.1037/0096-1523.25.5.1207

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458. doi:10.1037/0033-295X.96.3.433

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience, 17,* 738–743. doi:10.1038/nn.3689

Golomb, J. D., L'Heureux, Z. E., & Kanwisher, N. (2014). Feature-binding errors after eye movements and shifts of attention. *Psychological Science, 25,* 1067–1078. doi:10.1177/0956797614522068

Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision, 10*(2):24, 1–27. doi:10.1167/10.2.24

Johnson, J. S. (2008). *A dynamic neural field model of visual working memory and change detection. Theses and Dissertations*. Retrieved from http://ir.uiowa.edu/etd/12

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. doi:10.1111/j.1467-9280.2009.02329.x

Lin, P.-H., & Luck, S. J. (2009). The influence of similarity on visual working memory representations. *Visual Cognition, 17,* 356–372. doi:10.1080/13506280701766313

Makovski, T., & Jiang, Y. V. (2009). Feature binding in attentive tracking of distinct objects. *Visual Cognition, 17,* 180–194. doi:10.1080/13506280802211334

Navalpakkam, V., & Itti, L. (2007). Search goal tunes visual features optimally. *Neuron, 53,* 605–617.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal, 7,* 308–313. doi:10.1093/comjnl/7.4.308

Scolari, M., & Serences, J. T. (2009). Adaptive allocation of attentional gain. *Journal of Neuroscience, 29,* 11933–11942. doi:10.1523/JNEUROSCI.5642-08.2009

Scolari, M., & Serences, J. T. (2010). Basing perceptual decisions on the most informative sensory neurons. *Journal of Neurophysiology, 104,* 2266–2273. doi:10.1152/jn.00273.2010

Shapiro, K. L., & Miller, C. E. (2011). The role of biased competition in visual short-term memory. *Neuropsychologia, 49,* 1506–1517.

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision, 13*(10), 9. doi:10.1167/13.10.9

Treisman, A. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology, 40A,* 201–237. doi:10.1080/02724988843000104

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society B, 353,* 1295–1306. doi:10.1098/rstb.1998.0284

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136. doi:10.1016/0010-0285(80)90005-5

Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14,* 107–141.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453,* 233–235. doi:10.1038/nature06860