

Object-feature binding survives dynamic shifts of spatial attention

Emma Wu Dowd & Julie D. Golomb

Department of Psychology, The Ohio State University

RUNNING HEAD: Object-feature binding under dynamic attention

Correspondence to:

Emma Wu Dowd

Department of Psychology, The Ohio State University

1835 Neil Ave, Columbus, OH 43210

dowd.45@osu.edu

Abstract

Visual object perception requires integration of multiple features; spatial attention is thought to be critical to this binding. But attention is rarely static—how does *dynamic* attention impact object integrity? Here we manipulated covert spatial attention and had participants reproduce multiple properties (color, orientation, location) of a target item. Object-feature binding was assessed by applying probabilistic models to the joint distribution of feature errors: feature reports for the same object could be correlated (and thus bound together) or independent. We found that splitting attention across multiple locations degrades object integrity, while rapid shifts of spatial attention maintain bound objects. Moreover, we document a novel attentional phenomenon, wherein participants exhibit unintentional fluctuations—“lapses” of spatial attention—yet nevertheless preserve object integrity, at the wrong location. These findings emphasize the importance of a single focus of spatial attention for object-feature binding, even when that focus is dynamically moving across the visual field.

Keywords: visual attention; feature binding; attentional shift; lapses of attention

Author Contributions

Both authors were involved in the conceptualization and interpretation of the study. E. W. Dowd designed the experiments, oversaw data collection, analyzed the data, and drafted the manuscript. J. D. Golomb contributed to writing the manuscript and gave feedback throughout the project.

Acknowledgments

We thank Makaela Nartker, Samoni Nag, and Alexandra Haeflein for assistance with data collection and David Osher and Brandon Turner for helpful discussions.

Funding

This study was funded by grants from the National Institutes of Health (R01-EY025648 to J. D. Golomb, F32-028011 to E. W. Dowd) and the Alfred P. Sloan Foundation (to J. D. Golomb). Analyses reported in this publication were supported by an allocation of computing resources from The Ohio Supercomputer Center.

Our visual environment contains a multitude of objects, each composed of multiple visual features (e.g., color, shape, texture) that must be integrated together into a cohesive object-level representation, forming the basis for effective perception, memory, and action (e.g., Kahneman, Treisman, & Gibbs, 1992). Spatial attention is thought to be crucial for this binding process, acting as the “glue” that binds an object’s features together (Nissen, 1985; Reynolds & Desimone, 1999; Treisman & Gelade, 1980). Thus, attending to a given location selects all of the non-spatial features coinciding at that location and binds them into an integrated object representation (e.g., Schoenfeld et al., 2003).

Visual object integrity is critical for interacting with the world. Successful object-feature binding requires integration between different non-spatial features (i.e., feature-feature binding) and between non-spatial features and spatial location (i.e., feature-location binding). Failures of binding can be both detrimental and revealing (Treisman, 1996; Wolfe & Cave, 1999). For example, limited attention can produce “illusory conjunctions”, such as when one views a green square and a red circle, but reports a green circle (Treisman & Schmidt, 1982). Patients with unilateral deficits in spatial attention also report more illusory conjunctions, suggesting that object integrity depends on intact spatial extent (Robertson, 2003). Other studies have demonstrated that spatial crowding can induce greater “swap” errors (see Bays, 2016), i.e., misreporting features from another location altogether (e.g., Emrich & Ferber, 2012). One candidate model of object-feature binding posits that non-spatial features are randomly sampled from a probabilistic window of spatial attention (Vul & Rich, 2010)—when the attentional window is narrower and more precise, features are veridically bound as an integrated object, but when the focus of spatial attention is more diffuse and encompasses multiple object locations, features are encoded independently, resulting in inaccurate groupings of features and thus

failures of object integrity. The strong consensus is that the ability to maintain a precise spatial focus of attention is critical for preserving object integrity.

Visual attention, however, is rarely singular or static—nor do we want it to be. In the real world, multiple objects with multiple features are simultaneously present in the environment, and attention must dynamically shift and split across multiple goals and locations. How do these *dynamic* changes in spatial attention impact visual object integrity?

Here, we examine object-feature binding while manipulating dynamic attention. We applied a recent paradigm measuring distortions in feature perception under different conditions of spatial attention (Golomb, L’Heureux, & Kanwisher, 2014), adapted to probe binding between multiple feature dimensions. Participants were presented with an array of multi-feature objects and were asked to reproduce both the color and the orientation of a target object (i.e., joint continuous-report paradigm; Wilken & Ma, 2004; Bays, Wu, & Husain, 2011). In Experiment 1, we manipulated attention with a spatial pre-cue that either remained stable, shifted from one location to another, or was split across two locations simultaneously. Experiment 2 replicated these conditions with the further addition of a continuous location report. We applied probabilistic mixture models to assess whether errors in recalling multiple features of the same object would be correlated—indicating that features were bound together (Bays et al., 2011). This modeling approach allowed us to evaluate the fate of object integrity under common conditions of dynamic attention (Experiment 1), and moreover, to directly examine the three-way conjunction of feature-feature-location binding (Experiment 2).

We predicted that object-feature binding would be contingent on the spatial extent of attention (Cohen & Ivry, 1989; Vul & Rich, 2010), such that splitting attention simultaneously across two locations would expand the attentional window and disrupt object integrity. But what

happens during shifts of spatial attention? Previous studies have demonstrated that spatial attention does not update immediately, but rather may take 100–200 ms to fully disengage from one location, move, and re-engage attention at a new location (e.g., Müller, Teder-Sälejärvi, & Hillyard, 1998). More recent work suggests that attention can be allocated to a new location *before* it disengages from the previously attended location, resulting in a transient period in which both locations are simultaneously attended (e.g., Eimer & Grubert, 2014; Khayat, Spekreijse, & Roelfsema, 2006). Here, we probe binding during this critical period of spatial updating and test two contrasting theories: If spatial attention shifts as a single “spotlight” from location A to location B (Posner, Snyder, & Davidson, 1980), then the spatial extent of attention should remain narrow, and object-feature binding should survive—for either object A or B. Alternatively, a shift might result in temporarily highlighting both locations—e.g., spreading to encompass both locations (“zooming” out; Eriksen & St. James, 1986) or briefly activating both A and B simultaneously (e.g., Khayat et al., 2006; Golomb, Marino, Chun, & Mazer, 2011)—such that we should expect greater independent feature errors.

We found that dynamic shifts and splits of attention lead to distinct patterns of object-feature binding: rapid shifts of attention preserve object integrity, while splitting attention results in failures of object integrity. We also document a novel attentional phenomenon, wherein participants exhibit unintentional fluctuations—or “lapses”—of spatial attention, akin to momentary fluctuations of sustained attention (Reason, 1984). While lapses of sustained attention are typically defined as slowed responses in a vigilance task (e.g., Rosenberg et al., 2016), lapses of spatial attention are here defined as identifying an incorrect location as the target. By simultaneously modeling color, orientation, *and* location responses in Experiment 2, we show that the vast majority of feature reports are correlated across all three dimensions,

demonstrating that fully-bound objects survive both intentional shifts and inadvertent lapses of spatial attention. Thus, object integrity is preserved when there is a single focus of spatial attention, even if attention is at the wrong location.

METHOD

Participants

Experiment 1 included 23 participants (ages 18 to 30 years; 11 male), and Experiment 2 included 25 new participants (ages 18 to 21 years; 13 male). An a priori power analysis based on effect sizes in previous studies (Golomb et al., 2014) estimated that we would need at least 22 participants to detect feature errors with 80% power, given a .05 criterion of significance. We collected a few extra participants per experiment in anticipation of dropout due to poor eye-tracking or task performance. Additional participants were excluded for not completing the full experimental session (two from Experiment 1; three from Experiment 2) or not successfully performing the task (<50% probability of $pTcTo$ from the Simple Model described below; two from Experiment 2). All participants reported normal or corrected-to-normal visual acuity and color vision, received course credit or a payment of US\$10 per hour, and provided informed consent in accordance with The Ohio State University institutional review board.

Stimuli and Procedure

Stimuli were presented on a 21-inch flatscreen CRT monitor with a refresh rate of 85 Hz and screen resolution of 1280×1024 pixels, using Matlab and the Psychophysics Toolbox (Brainard, 1997). Subjects were positioned with a chinrest approximately 60 cm from the monitor, and eye position was tracked using an Eyelink 1000 eye-tracking system. The monitor was color-calibrated with a Minolta CS-100 colorimeter.

Figure 1 illustrates an example trial sequence for the three critical conditions (Hold, Shift, and Split) for both experiments. In Experiment 1, each trial began with the presentation of a central fixation dot (white, diameter of 0.6°). Once subjects were accurately fixating for 1,000 ms, as determined by real-time eye tracking, the trial continued as follows:

For Hold and Shift trials, a single spatial cue (black 4° × 4° square outline) was presented at one of four possible stimulus locations (the corners of an imaginary square centered on fixation, 7.4° eccentricity) for 250 ms. For Hold trials, the spatial cue was followed by a 1,100-ms delay period in which only the fixation dot was visible on the screen. For Shift trials, the cue was followed by a similar 1,000-ms delay, after which a second cue appeared in a different, adjacent location (never the diagonal location) for 50 ms, followed by a 50-ms delay.

Participants were instructed to covertly attend to the cued location and shift their attention if a second location was cued, such that they were always attending to the most recently cued location. For Split trials, two spatial cues were presented simultaneously at two adjacent locations for 250 ms, followed by a 1,100-ms delay period. Participants were instructed to attend to both cued locations in Split trials.

After the spatial pre-cue(s) and delay period, an array of four colored and oriented bars (0.75° × 4°) appeared for 50 ms. One of these stimuli was the “target” (T), which appeared at the most recently cued location for Hold and Shift trials, or at one of the two cued locations for Split trials (randomly selected and indicated by a post-cue at the time of response). Of the three non-target items, the critical non-target (N1) appeared at a location adjacent to the target, which was either initially cued in Shift trials or simultaneously attended in Split trials; the other adjacent non-target (N2) and the diagonal non-target (N3) were considered control items. In Hold trials,

both adjacent non-targets were never cued, so they were arbitrarily assigned as the “critical” N1 or “control” N2.

The color of the target item was chosen randomly on each trial from 180 possible colors, which were evenly distributed along a 360° circle in CIE $L^*a^*b^*$ coordinates with constant luminance ($L^* = 70$, center at $a^* = 20$, $b^* = 38$, and radius 60; Zhang & Luck, 2008). The colors of the remaining stimuli were chosen so that the adjacent items (N1 and N2) were equidistant in opposite directions (90° clockwise or counterclockwise deviation along the color wheel, with direction randomly varying from trial to trial), and the item at the diagonal location (N3) was set 180° away in color space. The orientation of the target item was also chosen randomly on each trial from a range of angles 0–180°, and N1 and N2 were likewise equidistant in opposite directions (45° clockwise or counterclockwise deviation), with N3 set 90° away. Feature values for color and orientation were set independently, as was the direction of deviation for each feature. The stimulus array was followed by 200 ms of masks (squares colored with a random color value at each pixel location, covering each of the four stimulus locations).

Participants then made a joint continuous-report response, reporting the color and orientation of the target item. A single probe bar with random initial values for color and orientation was presented at the center of the screen. For Split trials, the probe stimulus was accompanied by a spatial post-cue (white 4° × 4° square outline) indicating which of the two pre-cued locations was the target. Participants were instructed to adjust the color and orientation of the probe item to match the features of the target. The probe’s features were adjusted using two input dials (PowerMate USB Multimedia controllers, Griffin Technology, USA), one operated with each hand (left-color, right-orientation). Turning one dial caused the probe to rotate through the 180° range of possible orientations (steps of 1°); turning the other dial caused the probe’s

color to cycle through the 360° space of possible colors (steps of 2°). Participants could adjust the two dials in any order or simultaneously. To input their response, participants clicked down on either dial. Accuracy was stressed, and there was a time limit of 10 s.

At the end of the trial, participants were shown feedback for 1,500 ms: The reported color-orientation response was shown in the center of the screen, and the actual target item was displayed in its original location. Eye-tracking feedback was presented as the percentage of deviant fixation samples (i.e., when the participant's eye position deviated more than 2° from the central fixation location) between initial cue onset and probe onset. Eye-tracking feedback was intended to motivate participants to maintain central fixation; they were not told about trial exclusion criteria.

Experiment 2 was identical to Experiment 1 except for two key changes (Figure 1B): First, we increased the possible range of stimulus locations, although only 4 locations were ever chosen on each trial. For each trial, the target location was randomly selected from one of 16 locations (along an imaginary circle centered on fixation, 7.4° eccentricity), and the other three locations were arranged at +90°, -90°, and 180° away (four corners of an imaginary rotated square). Second, we added a location report at the end of Hold and Shift trials. (Split trials already included an explicit spatial post-cue, rendering a subsequent location report uninformative.) The location report was presented after the joint color-orientation report, but also as a continuous response: a single location probe (white 4° × 4° square outline) was displayed at a random location along a white circle outline (7.4° eccentricity), and turning the right-hand dial rotated the placeholder through the 360° space of possible locations (steps of 2.25°). Participants were instructed to adjust the location of the probe to match the location of the target (i.e., the most recently cued location). Participants input their location response by clicking down on the

dial, with a time limit of 2 s. For location feedback on Hold and Shift trials, the reported location was displayed as a white outline on the same screen as color-orientation and eye-tracking feedback.

In Experiment 1, each participant completed 7–10 blocks of intermixed Hold and Shift trials (112–160 trials of each condition) and 4–5 blocks of Split trials (128–160 trials), although the order of Hold/Shift or Split blocks was counterbalanced between participants. In Experiment 2, participants completed 5–8 blocks of intermixed Hold and Shift trials (80–128 trials of each condition), always followed by 3–4 blocks of Split trials (96–128 trials). Hold and Shift trials were intermixed to ensure that participants had to attend to the first cue and could not simply wait for the second cue; Split trials were presented in separate blocks for ease of instruction. Both experiments began with fixation training and 8 Hold practice trials, as well as 12 Shift practice trials before the Hold/Shift blocks, and 12 Split trials before the Split blocks. Trials were discarded if they contained over 15% deviant fixation samples between initial cue onset and probe onset (Experiment 1: 3.3%; Experiment 2: 8.8%)¹ or if subjects made no dial adjustments before inputting their response (Experiment 1: 0.1%; Experiment 2: 2.5%).

¹ We checked whether participants maintained fixation immediately before and after array onset for the trials that were *included* by calculating the percentage of deviant eye-tracking samples from 200 ms before array onset to probe onset: Experiment 1: 0.4% and Experiment 2: 0.6%. In contrast, trials that were *discarded* had 52–60% deviant fixations during this period.

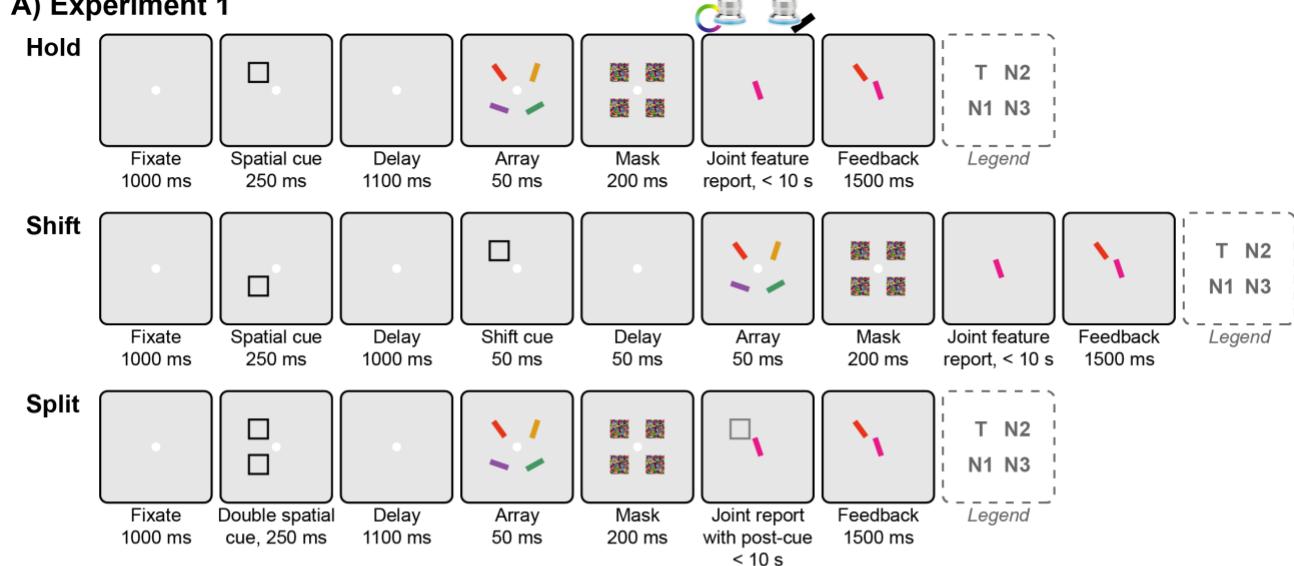
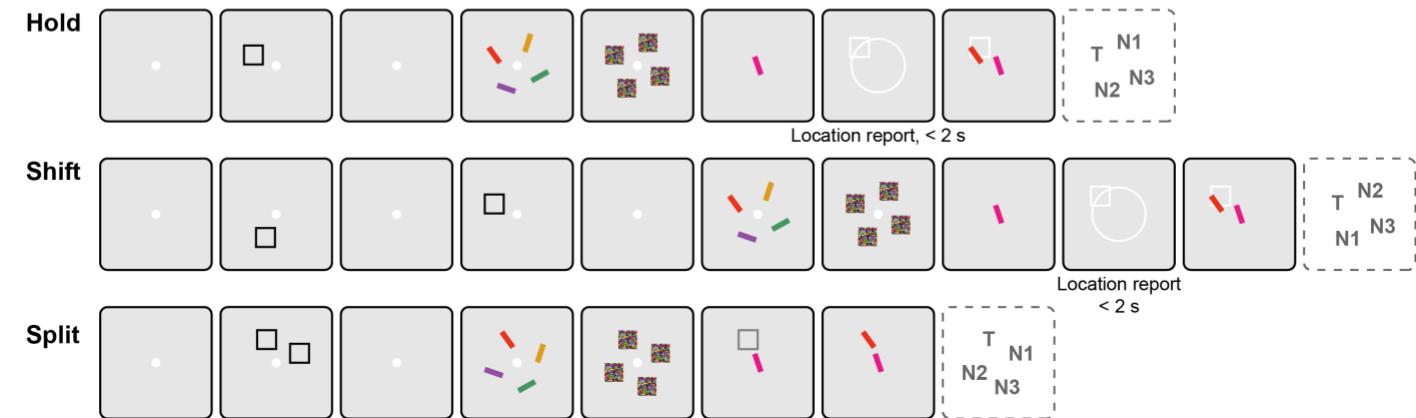
A) Experiment 1**B) Experiment 2**

Figure 1. Example trial sequences for (A) Experiment 1 and (B) Experiment 2. Participants were cued to covertly attend to a spatial pre-cue that remained stable (Hold) or dynamically shifted from one location to another (Shift), before reproducing both the color and orientation (i.e., joint continuous-report) of the cued target item. In Split trials, participants were pre-cued to attend to two locations simultaneously, before reporting the features of the post-cued target item. In Experiment 2, the four array positions could be rotated, and participants additionally reported the location of the target item (on Hold and Shift trials only). For each row, the final panel (dotted outline) denotes which array position corresponds to T, N1, N2, and N3. On Hold trials, N1 and N2 are arbitrarily assigned.

Joint-Feature Analyses

On each trial, response error was calculated as the angular deviation between the continuous probe report and the cued target item, for each feature separately (θ_C = color error, range -180° to 180° ; θ_O = orientation error, range -90° to 90° ; θ_L = location error, range -180° to 180°). In Shift and Split trials, although the direction of N1 varied randomly in relation to T

(clockwise or counterclockwise in terms of spatial location, color, or orientation space), we aligned the responses on each trial so that errors toward the N1 feature were always coded as positive deviations ($+90^\circ$ or $+45^\circ$), and errors toward N2 as negative deviations (-90° or -45°).

To quantify the amount of object-feature binding, we employed a mixture modeling approach (Bays, Catalao, & Husain, 2009; Bays et al., 2011; Golomb et al., 2014; Zhang & Luck, 2008). Within each single feature dimension (Table 1), responses could be attributed to either reporting the target (T , a von Mises distribution centered on the target feature value), misreporting one of the three non-targets ($N1$, $N2$, $N3$, separate von Mises distributions centered on each non-target feature value), or random guessing (U , a uniform distribution across all feature values). Critically, we modeled color and orientation as joint probability distributions, fitting responses from both feature dimensions simultaneously (see Bays et al., 2011).

For Experiment 1, we evaluated two types of joint color-orientation mixture models: In the Full Joint Model, we modeled the five types of feature reports described above (T , $N1$, $N2$, $N3$, and U) for each dimension, resulting in 25 response combinations of color (5) and orientation (5) (Table 2). This resulted in a model with 29 parameters (including parameters for the concentrations κ_C and κ_O ($\sigma = \sqrt{1/\kappa}$) and means μ_C and μ_O of the target). Due to this large number of parameters, the Full Model required data to be collapsed across subjects to achieve reliable model fits. Thus, to conduct within-subject statistics for our primary analyses, data from individual subjects were fit with the Simple Joint Model, which only attributes responses to T , $N1$, or U within each dimension, resulting in 9 response combinations of color (3) and orientation (3) (Table 2) plus 4 parameters for concentrations and means. Since our focus here was on the effects of dynamic attention involving the T and critical N1 locations, the Simple Model drops parameters that involve N2 or N3 (which were theoretically less relevant, and had

very low probabilities in the Full Model results; see Table S2). In the Simple Model, the few responses to N2 or N3 items should be absorbed by U . Indeed, the results reported using the Simple Joint Model were confirmed at the population level with the Full Joint Model (Supplement S5).

For both models, the joint distribution of responses was modeled as:

$$p(\theta_C, \theta_O) = \sum_m \alpha_m p_m$$

where θ_C and θ_O are the reported feature errors, and m is the number of joint color-orientation response combinations, with $m_{Full} = 1:25$ or $m_{Simple} = 1:9$. α_m is the probability of each response combination, and p_m represents the joint probability density distribution for that combination. Table 2 lists each of the m combinations and associated probability density functions. For example, the joint probability distribution of reporting the target color and the N1 non-target orientation would be $p_{TcN1o} = \phi_{\mu_C, \kappa_C} \phi_{\frac{\pi}{2}, \kappa_O}$.

For both models, joint-feature response distributions were fit using Markov chain Monte Carlo (MCMC), as implemented through custom Matlab scripts (available on Open Science Framework) using the MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013) on the Ohio Supercomputer Center (Ohio Supercomputer Center, 1987). The MCMC procedure sampled three parallel chains across as many iterations as necessary to achieve convergence, according to the method of Gelman and Rubin (1992). We collected 15,000 post-convergence samples and used the posterior distributions to compute the maximum-likelihood estimates of each parameter, as well as its 95% highest posterior density interval (HDI). For our primary analyses using the Simple Model, we adopted a standard within-subject analytical approach: Parameter estimates were obtained separately for each individual subject and each trial type, then evaluated with frequentist significance testing. Post-hoc tests were evaluated with the appropriate Bonferroni-

correction for multiple comparisons. These standard statistical comparisons were corroborated by analyses using the Full Model: Data were collapsed across all subjects within each experiment, and parameter estimates were obtained separately for each trial type. Parameter estimates from the Full Model were considered significantly different if their 95% HDIs did not overlap (Kruschke, 2011).

For Experiment 2, we first evaluated the joint color-orientation feature reports with both the Simple Model and the Full Model, as described above. With the inclusion of the continuous location report, we also fit single-subject data with a single-dimension Location Model, which attributed location responses to T_L , $N1_L$, $N2_L$, $N3_L$, or U_L , with flexible κ_L and μ_L of the target. We then took the joint modeling approach a step further, modeling the continuous responses from color, orientation, *and* location simultaneously as three-way joint probability distributions. Because a full triple joint model including all possible response types for all three dimensions would include an unwieldy number of parameters (131), we simplified the Triple Joint Model to focus on parameters of theoretical interest. Specifically, we examined whether reporting both the color and orientation of a specific item (feature-feature binding; e.g., $N1cN1o$) was also bound to the location of that specific item (feature-feature-location binding; e.g., location $N1_L$). Thus, the Triple Model included only parameters for correlated feature-feature responses (i.e., $TcTo$, $N1cN1o$, $N2cN2o$, $N3cN3o$) or guessing ($UcUo$), crossed with the different types of location responses. Location reports were attributed to T_L , $N1_L$, $N2_L$, or $N3_L$; we did not include a random guessing component U_L , as the corresponding pU_L parameter from the single-dimension Location Model was found to be negligible (Table S8). Similarly, we did not include flexible parameters for means μ_c , μ_o , and μ_o because the corresponding parameters in the Simple and Location Models were not significantly different from 0 (Tables S1, S8). Consequently, the

Triple Model included 20 response combinations of color-orientation (5) and location (4), and 3 parameters for concentrations κ_C , κ_O , and κ_L (Table 3). The joint distribution of responses was modeled as:

$$p(\theta_C, \theta_O, \theta_L) = \sum_m \alpha_m p_m,$$

where θ_C and θ_O are the reported color and orientation errors, θ_L is the reported location error, α_m is the probability of each response combination, and p_m represents the combined probability density, with $m_{Triple} = 1:20$. For example, the probability distribution of reporting the color and orientation of the target with the location of the critical N1 non-target would be $p_{TcToN1L} = \phi_{0,\kappa_C} \phi_{0,\kappa_O} \phi_{\pi,\kappa_L}$. Given the large number of parameters (23), the Triple Model was fit with data collapsed across all subjects and analyzed in the same way as the Full Model described above. Parameter estimates were obtained separately for Hold and Shift trial types (Split trials were not included in these analyses because there was no location report on these trials).

Table 1

Mixture model response distributions within a single feature dimension

Response type	Color (C)	Orientation (O)	Location (L)
Target (T)	ϕ_{μ_C, κ_C}	ϕ_{μ_O, κ_O}	ϕ_{μ_L, κ_L}
Critical Non-target ($N1$)	ϕ_{π, κ_C}	$\phi_{\frac{\pi}{2}, \kappa_O}$	ϕ_{π, κ_L}
Control Non-target ($N2$)	$\phi_{-\pi, \kappa_C}$	$\phi_{-\frac{\pi}{2}, \kappa_O}$	$\phi_{-\pi, \kappa_L}$
Diagonal Non-target ($N3$)	$\phi_{2\pi, \kappa_C}$	ϕ_{π, κ_O}	$\phi_{2\pi, \kappa_L}$
Random (U)	γ_C	γ_O	γ_L

where ϕ is a von Mises probability density function, with concentration κ_C , κ_O , or κ_L (standard deviation $\sigma = \sqrt{1/\kappa}$), and means of μ_C or μ_L , 90° , -90° , and 180° (color or location) or μ_O , 45° , -45° , and 90° (orientation) for the target, critical non-target, and control and diagonal non-target features, respectively. γ_C , γ_O , and γ_L are uniform distributions that reflects the probability of responding to the respective feature at random (guessing).

Table 2

Simple and Full Joint Mixture Model response distributions combined across both non-spatial feature dimensions

	Response type	m_{Simple}	m_{Full}	Response combination	Joint probability density
Correlated	Correlated target	1	1	$TcTo$	$\phi_{\mu_C, \kappa_C} \phi_{\mu_O, \kappa_O}$
	Correlated swap	2	2	$N1cN1o$	$\phi_{\pi, \kappa_C} \phi_{\frac{\pi}{2}, \kappa_O}$
			3	$N2cN2o$	$\phi_{-\pi, \kappa_C} \phi_{-\frac{\pi}{2}, \kappa_O}$
			4	$N3cN3o$	$\phi_{2\pi, \kappa_C} \phi_{\pi, \kappa_O}$
Independent T*	Illusory conjunction	3	5	$TcN1o$	$\phi_{\mu_C, \kappa_C} \phi_{\frac{\pi}{2}, \kappa_O}$
		4	6	$N1cTo$	$\phi_{\pi, \kappa_C} \phi_{\mu_O, \kappa_O}$
			7	$TcN2o$	$\phi_{\mu_C, \kappa_C} \phi_{-\frac{\pi}{2}, \kappa_O}$
			8	$N2cTo$	$\phi_{-\pi, \kappa_C} \phi_{\mu_O, \kappa_O}$
			9	$TcN3o$	$\phi_{\mu_C, \kappa_C} \phi_{\pi, \kappa_O}$
			10	$N3cTo$	$\phi_{2\pi, \kappa_C} \phi_{\mu_O, \kappa_O}$
	Unbound guess	5	11	$TcUo$	$\phi_{\mu_C, \kappa_C} \gamma_O$
		6	12	$UcTo$	$\gamma_C \phi_{\mu_O, \kappa_O}$
Independent N*	Illusory conjunction	13	$N1cN2o$	$\phi_{\pi, \kappa_C} \phi_{-\frac{\pi}{2}, \kappa_O}$	
		14	$N2cN1o$	$\phi_{-\pi, \kappa_C} \phi_{\frac{\pi}{2}, \kappa_O}$	
		15	$N1cN3o$	$\phi_{\pi, \kappa_C} \phi_{\pi, \kappa_O}$	
		16	$N3cN1o$	$\phi_{2\pi, \kappa_C} \phi_{\frac{\pi}{2}, \kappa_O}$	
		17	$N2cN3o$	$\phi_{-\pi, \kappa_C} \phi_{\pi, \kappa_O}$	
		18	$N3cN2o$	$\phi_{2\pi, \kappa_C} \phi_{-\frac{\pi}{2}, \kappa_O}$	
	Unbound guess	7	$N1cUo$	$\phi_{\pi, \kappa_C} \gamma_O$	
		8	$UcN1o$	$\gamma_C \phi_{\frac{\pi}{2}, \kappa_O}$	
			21	$N2cUo$	$\phi_{-\pi, \kappa_C} \gamma_O$
			22	$UcN2o$	$\gamma_C \phi_{-\frac{\pi}{2}, \kappa_O}$
			23	$N3cUo$	$\phi_{2\pi, \kappa_C} \gamma_O$
			24	$UcN3o$	$\gamma_C \phi_{\pi, \kappa_O}$
	Random guessing	9	25	$UcUo$	$\gamma_C \gamma_O$

where ϕ is a von Mises probability density function, with concentration κ_C or κ_O (standard deviation $\sigma = \sqrt{1/\kappa}$), and means of μ_C , 90° , -90° , and 180° (color) or μ_O , 45° , -45° , and 90° (orientation) for the target T, critical non-target N1, control non-target N2, and diagonal non-target N3 features, respectively. γ_C and γ_O are uniform distributions that reflect the probability of responding at random. Joint response combinations are always written as color response first. Rows 1–4 describe correlated responses (reporting both features of the target object, or misreporting both features of the same non-target [correlated swap errors]); rows 5–12 describe independent T* errors (reporting one target feature and one non-target feature [illusory conjunctions], or reporting one target feature and guessing the other feature [unbound guesses]); and rows 13–25 describe the remaining errors (e.g., misreporting features from different non-targets and guessing). The Simple Joint Model includes only the parameters as numbered by m_{Simple} , whereas the Full Joint Model includes all 25 parameters, as numbered by m_{Full} .

Table 3

Triple Joint Mixture Model response distributions combined across color, orientation, and location dimensions

Location response	Non-spatial feature responses	m_{Triple}	Response combination	Joint probability density
Target	Correlated Target	1	$TcToTL$	$\phi_{0,\kappa_C} \phi_{\mu_O,\kappa_O} \phi_{\mu_L,\kappa_L}$
	Correlated N1	2	$N1cN1oTL$	$\phi_{\pi,\kappa_C} \phi_{\frac{\pi}{2},\kappa_O} \phi_{\mu_L,\kappa_L}$
	Correlated N2	3	$N2cN2oTL$	$\phi_{-\pi,\kappa_C} \phi_{-\frac{\pi}{2},\kappa_O} \phi_{\mu_L,\kappa_L}$
	Correlated N3	4	$N3cN3oTL$	$\phi_{2\pi,\kappa_C} \phi_{\pi,\kappa_O} \phi_{\mu_L,\kappa_L}$
	Other	5	$UcUoTL$	$\gamma_C \gamma_O \phi_{\mu_L,\kappa_L}$
Non-target N1	Correlated Target	6	$TcToN1L$	$\phi_{\mu_C,\kappa_C} \phi_{\mu_O,\kappa_O} \phi_{\pi,\kappa_L}$
	Correlated N1	7	$N1cN1oN1L$	$\phi_{\pi,\kappa_C} \phi_{\frac{\pi}{2},\kappa_O} \phi_{\pi,\kappa_L}$
	Correlated N2	8	$N2cN2oN1L$	$\phi_{-\pi,\kappa_C} \phi_{-\frac{\pi}{2},\kappa_O} \phi_{\pi,\kappa_L}$
	Correlated N3	9	$N3cN3oN1L$	$\phi_{2\pi,\kappa_C} \phi_{\pi,\kappa_O} \phi_{\pi,\kappa_L}$
	Other	10	$UcUoN1L$	$\gamma_C \gamma_O \phi_{\pi,\kappa_L}$
Non-target N2	Correlated Target	11	$TcToN2L$	$\phi_{\mu_C,\kappa_C} \phi_{\mu_O,\kappa_O} \phi_{-\pi,\kappa_L}$
	Correlated N1	12	$N1cN1oN2L$	$\phi_{\pi,\kappa_C} \phi_{\frac{\pi}{2},\kappa_O} \phi_{-\pi,\kappa_L}$
	Correlated N2	13	$N2cN2oN2L$	$\phi_{-\pi,\kappa_C} \phi_{-\frac{\pi}{2},\kappa_O} \phi_{-\pi,\kappa_L}$
	Correlated N3	14	$N3cN3oN2L$	$\phi_{2\pi,\kappa_C} \phi_{\pi,\kappa_O} \phi_{-\pi,\kappa_L}$
	Other	15	$UcUoN2L$	$\gamma_C \gamma_O \phi_{-\pi,\kappa_L}$
Non-target N3	Correlated Target	16	$TcToN3L$	$\phi_{\mu_C,\kappa_C} \phi_{\mu_O,\kappa_O} \phi_{2\pi,\kappa_L}$
	Correlated N1	17	$N1cN1oN3L$	$\phi_{\pi,\kappa_C} \phi_{\frac{\pi}{2},\kappa_O} \phi_{2\pi,\kappa_L}$
	Correlated N2	18	$N2cN2oN3L$	$\phi_{-\pi,\kappa_C} \phi_{-\frac{\pi}{2},\kappa_O} \phi_{2\pi,\kappa_L}$
	Correlated N3	19	$N3cN3oN3L$	$\phi_{2\pi,\kappa_C} \phi_{\pi,\kappa_O} \phi_{2\pi,\kappa_L}$
	Other	20	$UcUoN3L$	$\gamma_C \gamma_O \phi_{2\pi,\kappa_L}$

where ϕ is a von Mises probability density function, with concentration κ_C , κ_O , or κ_L (standard deviation $\sigma = \sqrt{1/\kappa}$), and means of 0° , 90° , -90° , and 180° (color or location) or 0° , 45° , -45° , and 90° (orientation) for the target T, critical non-target N1, control non-target N2, and diagonal non-target N3 features, respectively. γ_C and γ_O are uniform distributions that reflect the probability of responding at random.

RESULTS

By probing both color and orientation on each trial, we examined whether errors in recalling multiple features of the same object were correlated (and thus bound together) or independent (and unbound) under different conditions of covert spatial attention. Figure 2 visualizes the joint distribution of responses by plotting individual trials in joint-feature space, where the vertical and horizontal axes correspond to the color and orientation errors, respectively. Figure 2A illustrates the predicted distributions for different types of hypothetical responses, and Figures 2B–C show the actual response distributions for each attention condition.

“Object integrity” was inferred from contrasting *correlated* responses (i.e., reporting both the color and orientation of the same item) to *independent* responses (i.e., reporting only one feature of the target item) (Figure 2A). Correlated responses could stem from: 1) reporting both features of the correct target item (correlated target, TcTo), which would be represented as a two-dimensional Gaussian density centered on the origin (0° error), or 2) misreporting both features of the same non-target item (correlated swap, denoted N1cN1o, N2cN2o, and N3cN3o), which would be represented as two-dimensional Gaussian densities along the positive-slope diagonal of joint-feature space. Failures in object-feature binding, on the other hand, would result in *independent* target errors, represented as a distribution of responses along the horizontal and vertical axes of joint-feature space (i.e., centered on zero error in one dimension but not the other). Independent target errors could be due to misbinding the features of a target and a non-target item (illusory conjunction; e.g., TcN1o) or reporting only one feature of the target and guessing the other (unbound guess; e.g., TcUo). Finally, pure guessing (UcUo) would be represented as a two-dimensional uniform distribution of responses across the entire joint-feature space.

We quantified each of the error types above with joint-feature probabilistic models (Table 2). The probabilistic models also allow for independent non-target errors, such as reporting the color of one non-target and the orientation of a different non-target (e.g., N1cN2o), or reporting only one feature of a non-target and guessing the other (e.g., N1cUo); however, such errors were relatively rare and not discussed further, but see Tables S1 and S2 for all parameter estimates.

Scatterplots of the empirical data for each experiment, collapsed across all participants, for Hold, Shift, and Split trials, are presented in Figures 2B–C. The corresponding parameter estimates from the Simple Model are shown in Figure 3; below, we report within-subject statistics for the Simple Model, but comparison of 95% HDIs from the Full Model confirm these results and are reported in Table S2 (see also Supplement S5).

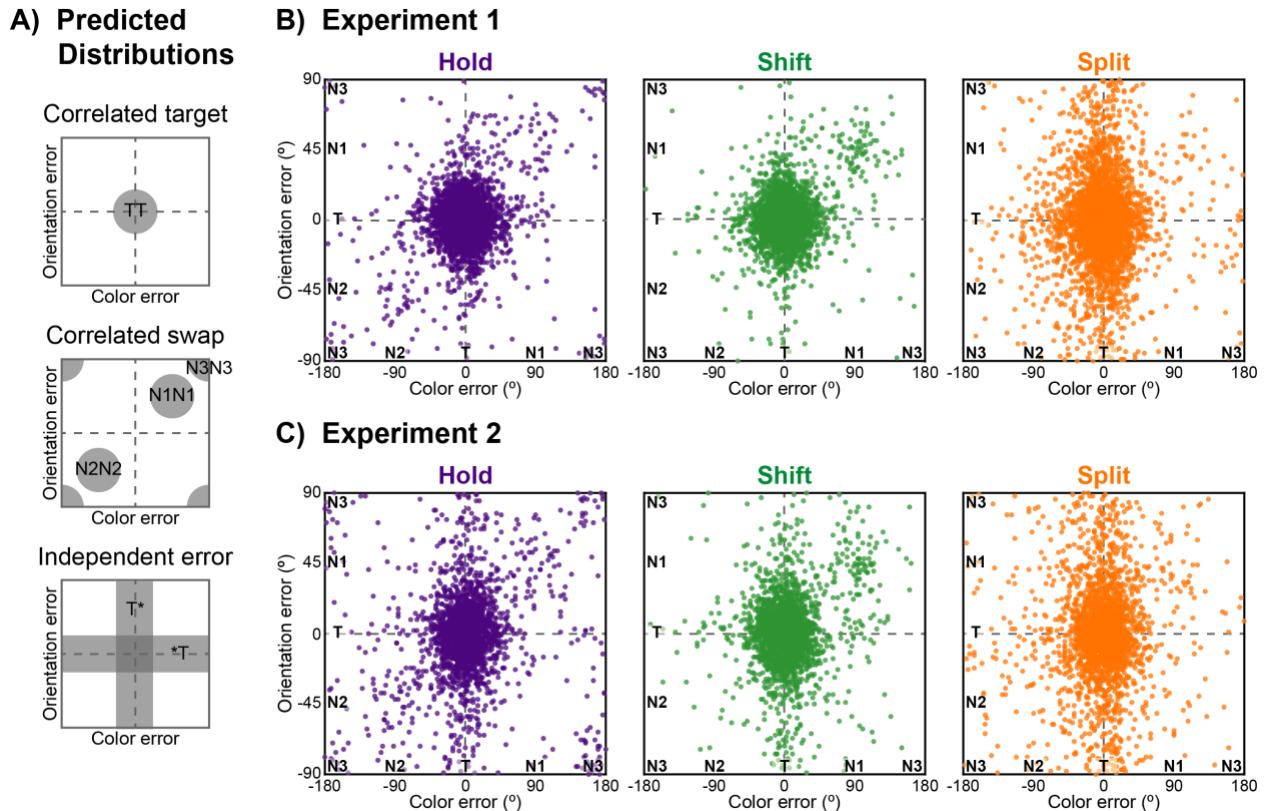


Figure 2. Visualizations of color-orientation reports in joint-feature space, plotted as error relative to actual target feature values: color responses are shown along the x-axis, and orientation responses are shown along the y-axis. For visualization purposes, we have flattened joint-feature space; both feature dimensions are in fact circular, such that $+180^\circ$ is identical to -180° in color space. The schematics in (A) show predicted distributions for possible response

types: Correlated target responses reflect reporting both features of the target item. Correlated swaps reflect reporting both features of the same non-target; see Figure 1 for legend of N1/N2/N3 labels across different conditions. Independent errors reflect reporting only one target feature. The scatterplots in (B–C) plot trial-by-trial error distributions separately for Hold, Shift, and Split trials, for Experiment 1 (B) and Experiment 2 (C). Each dot represents the corresponding color and orientation response for a single trial, aggregating across subjects.

Experiment 1

Across all trial types, the vast majority of responses were attributed to reporting both features of the correct target item (correlated target responses; Figure 3A), as reflected in the scatterplots as a central density of responses at the origin (Figure 2B). However, the patterns of errors outside that central density differed as a function of attention condition. As predicted, Split attention degraded performance. Splitting attention across two locations resulted in greater feature errors, with significantly lower correlated target responses ($pTcTo$, Simple Model) in Split trials compared to Hold trials, $t(22) = 4.20, p < .001, d = 0.88$, and Shift trials, $t(22) = 6.78, p < .001, d = 1.42$. The standard deviations of both color (σ_c) and orientation (σ_o) responses were also greater for Split trials, $ps < .001$ (see Table S3 for all comparisons), indicating less precise feature reports when splitting spatial attention (Figure 3A). In contrast, shifts of attention from one location to another did not impact $pTcTo$ or feature precision; pairwise t -tests revealed that these parameters were not significantly different between Hold and Shift trials, $ps > .5$ (Table S3).

Next we compared the different types of errors produced by shifts or splits of attention, comparing the probability of correlated N1cN1o swap errors (Simple Model) versus independent T* errors (e.g., TcN1o, UcTo, Simple Model). Because Split trials produced more non-TcTo responses overall, to compare across conditions we calculated these as proportions of “all errors” (non-TcTo responses) for each attention condition (Figure 3B). A repeated-measures ANOVA across condition (Hold, Shift, Split) and error type (correlated N1cN1o swap, independent T*)

revealed a significant interaction, $F(2, 44) = 23.3, p < .001, \eta^2 = .515$. Split trials had the greatest proportion of independent T* errors ($M = 72.8\%$ of non-TcTo responses, $SD = 17.4\%$) compared to Hold trials ($M = 44.8\%$, $SD = 22.9\%$) and Shift trials ($M = 44.2\%$, $SD = 19.6\%$), $ps < .001$ (see Table S4 for all comparisons). In contrast, Shift trials had the greatest proportion of correlated N1cN1o swaps ($M = 25.1\%$ of non-TcTo responses, $SD = 18.9\%$) compared to Hold trials ($M = 6.5\%$, $SD = 7.4\%$) and Split trials ($M = 2.4\%$, $SD = 5.4\%$) trials, $ps < .001$ (Table S4). This interaction is illustrated in the Figure 2B scatterplots where a sizeable cluster of errors is visible around N1cN1o for Shift trials, whereas for Split trials, errors are reflected as noise around the vertical and horizontal axes. In other words, dynamic splits of spatial attention resulted in failures of object-feature binding (more independent T* errors), while dynamic shifts of attention were more likely to maintain object integrity (more correlated TcTo responses and correlated N1cN1o swaps).

When splits of spatial attention break down object-feature binding, what kinds of independent errors do participants make? In Split trials, unbound guesses ($M = 60.2\%$ of non-TcTo responses, $SD = 19.8\%$, Simple Model) occurred significantly more often than illusory conjunctions between T and N1 ($M = 12.7\%$, $SD = 13.7\%$, Simple Model), $t(22) = 7.78, p < .001, d = 1.62$. This suggests that participants were *not* simply encoding and reporting both features of both cued objects independently, as might be expected if the constituent features were processed as loose bundles of features and illusorily misbound (Vul & Rich, 2010). Instead, splitting attention across multiple objects seemed to induce a feature-load strategy in which one feature dimension was prioritized (e.g., Fougner, Asplund, & Marois, 2010; Woodman & Vogel, 2008), such that participants reported only one feature of the target and guessed the other (Supplement S6).

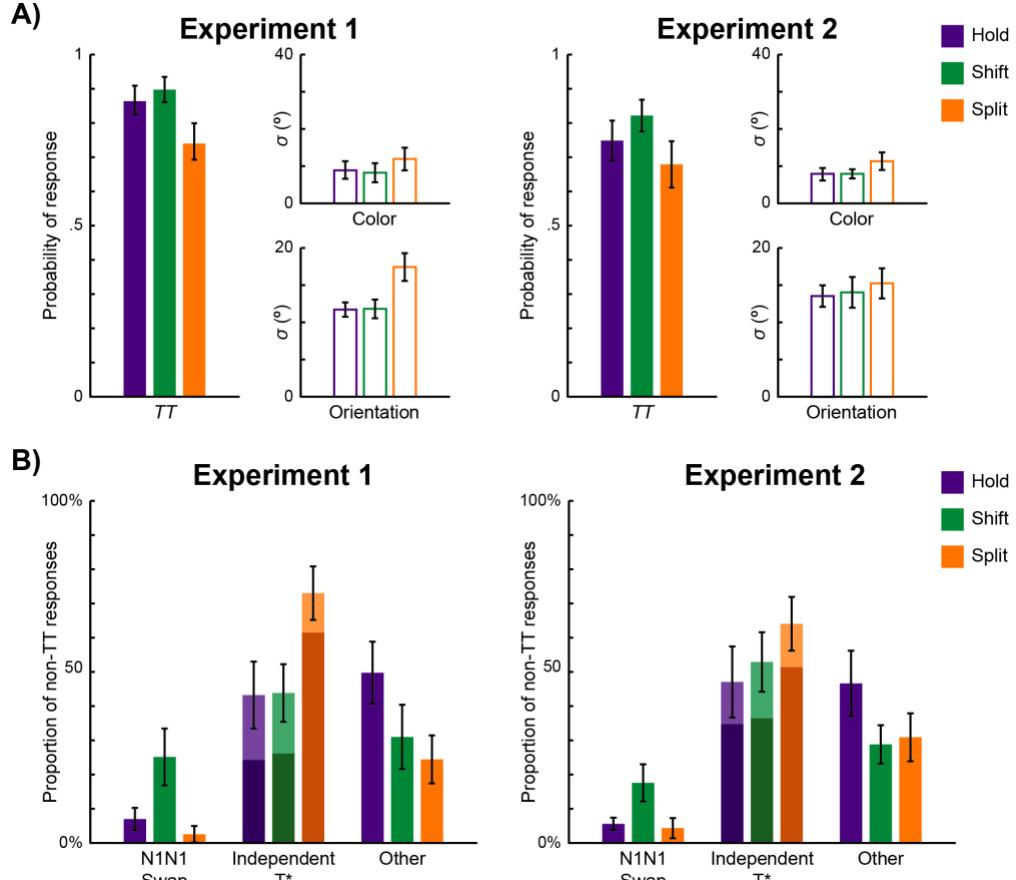


Figure 3. Simple Joint Model maximum-likelihood parameter estimates for both experiments. The graphs in (A) present best-fit estimates for the probability of correlated target responses ($T_c T_o$) and the standard deviations of the $T_c T_o$ distribution for color (σ_c) and orientation (σ_o) for each condition. The graphs in (B) present the proportion of erroneous responses (i.e., non- $T_c T_o$) that can be attributed to correlated $N1cN1o$ swaps, independent target errors (illusory conjunctions in light shading or unbound guesses in darker shading), and other errors (e.g., random guessing) for each condition. Shift trials had more $N1cN1o$ swaps, while Split trials had more independent target errors. The Simple Model was fit separately for each subject and each condition, and then parameter estimates were averaged across subjects. Error bars represent 95% confidence intervals. Results are shown separately for Experiments 1 ($N=23$) and 2 ($N=25$).

In Shift trials, we found a greater proportion of correlated $N1cN1o$ swap errors, consistent with the hypothesis that spatial attention had not yet shifted from the initially-cued N1 location to the correct target location on some trials. To confirm that correlated swap errors in Shift trials were specific to $N1cN1o$, we used the Full Model parameter fits to compare correlated swaps across all possible non-target locations (i.e., $N1cN1o$, $N2cN2o$, $N3cN3o$). In Shift trials,

participants were indeed more likely to misreport the entire feature-bound object at the initially-cued location ($pN1cN1o = .028$, 95% HDI [.021 .032]) compared to the other non-target locations ($pN2cN2o = .001$, 95% HDI [0 .004], $pN3cN3o < .001$, 95% HDI [0 .002]); parameter estimates were considered significantly different if their 95% HDIs did not overlap (Kruschke, 2011). In contrast, in Hold trials, correlated swaps were distributed equally across the three non-targets (Figure 4; 95% HDIs overlap).

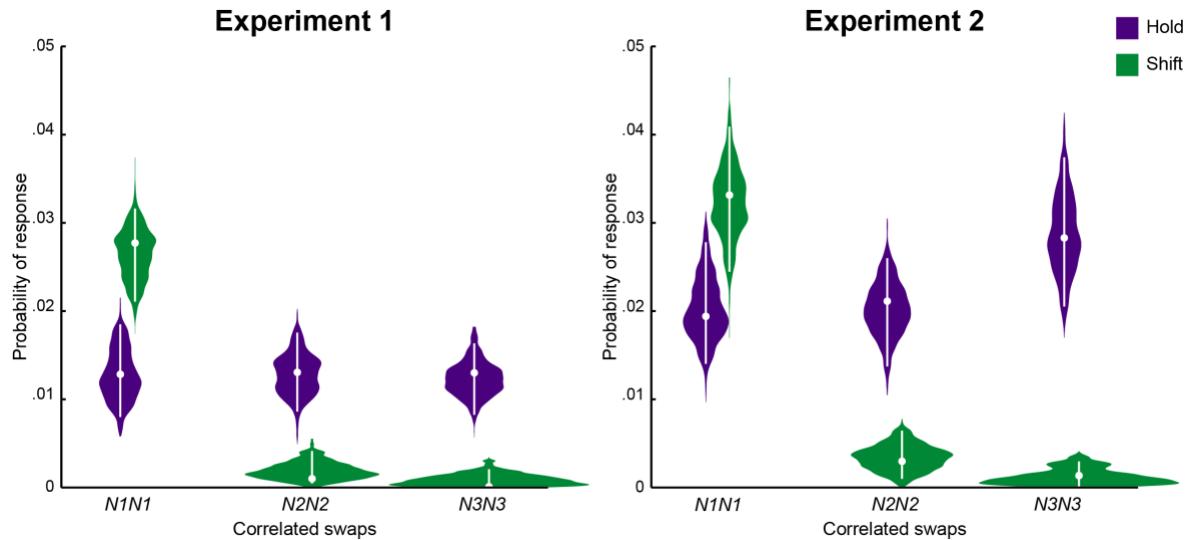


Figure 4. Full Joint Model maximum-likelihood estimates confirm that correlated swaps for Shift trials were highest for the initially-cued non-target (N1), while correlated swaps for Hold trials were equally distributed across all non-target items. For each correlated swap ($N1cN1o$, $N2cN2o$, $N3cN3o$), a violin plot illustrates the posterior distribution of each parameter over 15,000 post-convergence samples. The white dots mark each parameter's best-fit estimate, and the whiskers represent the 95% highest density interval. The Full Model was fit for each condition separately, collapsed across all subjects. Results are shown separately for Experiments 1 and 2.

We had expected Hold trials to have very few errors overall, since those trials involved only a single relevant spatial location. However, surprisingly, the Full Model parameter fits revealed that in Hold trials, participants made just as many *total* correlated swap errors ($pN1cN1o + pN2cN2o + pN3cN3o = .039$, 95% HDI [.029 .046]) as in Shift trials (.029, 95% HDI [.023 .035]). Why would Hold trials produce substantial correlated swap errors? One possible explanation is that participants occasionally made inadvertent attentional shifts on Hold trials;

because all non-target locations were equally irrelevant, these random “lapses” of spatial attention should lead to correlated swap errors for all non-target locations. To directly probe possible lapses of spatial attention and their impact on object-feature binding, we added a continuous location report to Experiment 2.

Experiment 2

In Experiment 2, we directly evaluated the role of spatial attention in the binding of non-spatial visual features. Specifically, is an integrated object (i.e., intact feature-feature binding) anchored to its spatial location (i.e., intact feature-feature-location binding)? In addition to reporting the target’s color and orientation, participants also reported the target’s location on Hold and Shift trials. When participants correctly reported both features of the target, we expected them to correctly report its location. But when participants made a correlated swap error (i.e., reporting the color and orientation of a non-target), would they also report the *incorrect* location of that same non-target?

For color and orientation feature reports, Experiment 2 replicated the same pattern of results as reported for Experiment 1 (Figures 2C, 3, 4; Tables S1, S2, S5). Figure 5A–B depicts the response distributions for the location reports that followed Hold and Shift trials; Split trials did not include a location report because location was already post-cued. Based on a single-dimension Location Model (Table S8), participants primarily reported the target location (Hold: $p_{TL} = .887$, Shift: $p_{TL} = .920$) with high precision (Hold: $\sigma = 9.0^\circ$, Shift: $\sigma = 9.3^\circ$). Even on Shift trials, when the second cue was only flashed for 50 ms, participants were highly accurate at reporting that second location, indicating that participants were able to perceive the shift cue. Neither the probability of reporting the target location nor standard deviation was significantly

different between Hold and Shift trials, $t(24) = 1.20, p = .240, d = 0.24$ and $t(24) = 1.07, p = .293, d = 0.22$. However, the difference in the *pattern* of location errors is evident: a repeated-measures ANOVA across condition (Hold, Shift) and error type (N1, N2, N3, U) revealed a significant interaction, $F(3, 72) = 8.40, p < .001, \eta^2 = .259$. Post-hoc paired *t*-tests confirmed that in Shift trials, participants were more likely to misreport the initially-cued N1 location compared to N2, $t(24) = 3.66, p = .001, d = 0.73$, or N3, $t(24) = 3.74, p = .001, d = 0.75$, while in Hold trials, participants were equally likely to misreport any of the three non-target locations, $ps > .14$ (see Table S9 for all comparisons). Thus, the pattern of location report errors mimicked the pattern of correlated swap errors for non-spatial features.

Next, we directly investigated how these lapses of spatial attention (using reports of non-target locations as a proxy) impact object-feature binding. We jointly modeled the three-way conjunction of color, orientation, and location with the Triple Model (see Table S10 for all parameter estimates) to test whether reporting both the color and orientation of a specific item (e.g., $N1cN1o^*$) was bound to the location of that specific item (e.g., $N1cN1oN1L$). Figure 5C illustrates the joint color and orientation reports for Hold trials in the same scatterplot form as in Figure 2, but here plotted separately according to location response (coarsely coded as T, N1, N2, or N3) on that trial. Figure 6 shows the parameter estimates from the Triple Model, also grouped by their shared location component (e.g., $*N1L$), such that each feature-feature response type is plotted as a proportion of all the responses containing that location component.

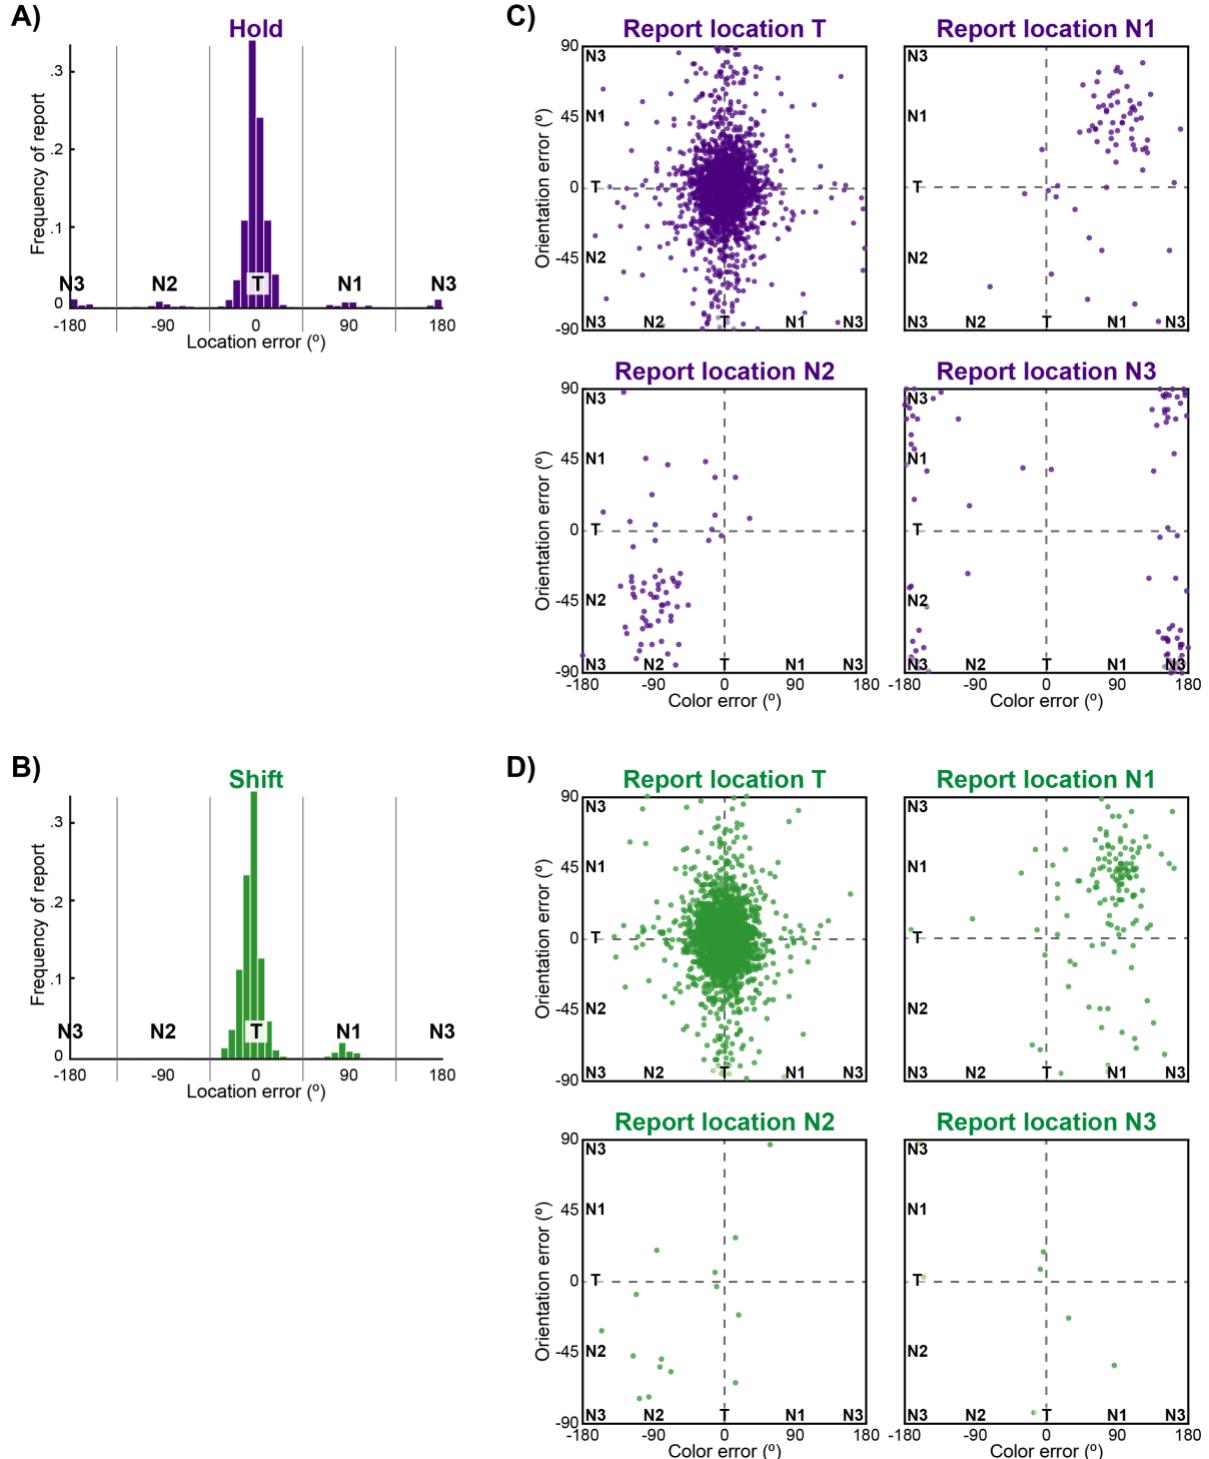


Figure 5. Visualization of color, orientation, and location report errors. The graphs on the left show the location report data from Experiment 2, plotted as histograms showing frequency of report as a function of location errors relative to target location values, for Hold trials (A) and Shift trials (B). For both conditions, each trial was then coded into separate response bins (vertical gray lines) to produce the scatterplots on the right, which plot the trial-by-trial color and orientation error distributions of Hold (C) and Shift (D) trials in joint-feature space (x-axis: color; y-axis: orientation), split by each trial's coded location. Data are combined across subjects.

As suggested by the scatterplots and confirmed with the Triple Model, the vast majority of all responses were attributed to fully-bound objects: correlated color-orientation-location reports (yellow outlined bars, Figure 6). When the correct target location was reported, participants overwhelmingly reported both non-spatial features of the target object ($pTcToT_L = 95.1\%$ of T_L responses, 95% HDI [.934 .960]; Figure 5C, top left). While there was a small proportion of random feature guessing ($pUcUoT_L$), participants did not make correlated swap errors when spatial attention was properly maintained (e.g., $pN1cN1oT_L$).

But what about during lapses of spatial attention? When participants incorrectly reported the location of a non-target, the color and orientation errors appear to systematically cluster around *that specific non-target's* features (Figure 5C). Thus, when spatial attention lapsed to a non-target location, the non-spatial feature reports reflected these location errors, with participants mostly reporting both the color and orientation of the item at the lapsed spatial location ($pN1cN1oN1_L = 73.5\%$ of all $*N1_L$ responses, 95% HDI [.662 .831]; $pN2cN2oN2_L = 73.4\%$ of all $*N2_L$ responses, 95% HDI [.614 .858]; $pN3cN3oN3_L = 86.9\%$ of all $*N3_L$ responses, 95% HDI [.747 .898]; Figure 6). On these lapse trials, while there were small proportions of correct target reports ($pTcTo^*$) and random guesses ($pUcUo^*$), there were almost never correlated swap errors associated with another location. In other words, after lapses of spatial attention, participants still primarily reported fully-bound properties (color, orientation, and location), but for the lapsed non-target item instead of the correct target item.

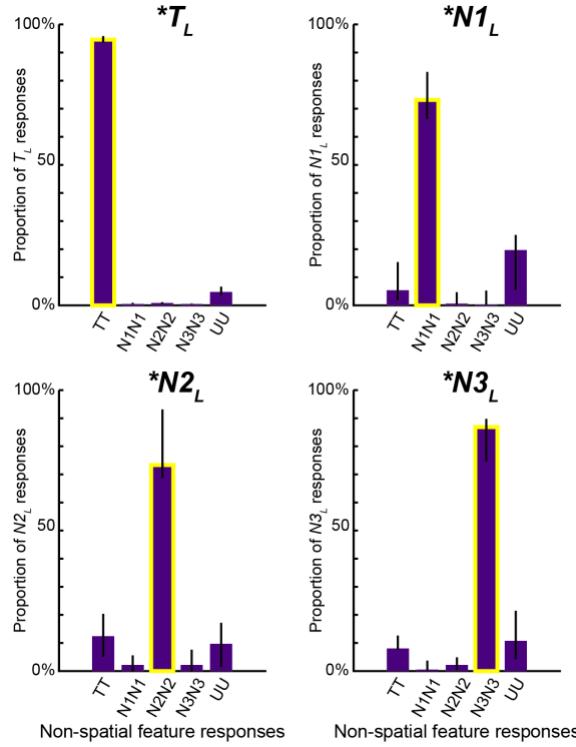


Figure 6. Triple Joint Model maximum-likelihood estimates reveal that non-spatial feature-feature binding (i.e., color-orientation) is linked to the reported spatial location, regardless of whether it is the correct target location. Color, orientation, and location responses from Hold trials in Experiment 2 were fit simultaneously as joint distributions. Parameter estimates are plotted separately, grouped according to each parameter's location component: T_L (top left), $N1_L$ (top right), $N2_L$ (bottom left), or $N3_L$ (bottom right). For each graph, each non-spatial feature response type ($N1cN1o$, $N2cN2o$, $N3cN3o$, $UcUo$) is calculated as a proportion of all responses specific to that location component. For each graph, the bar with the yellow outline signifies reporting the fully-bound color, orientation, and location of the same object. The Triple Model was fit for each condition separately, collapsed across all subjects. Error bars represent 95% highest density intervals.

By simultaneously modeling all feature dimensions (i.e., color, orientation, location), we demonstrate that the locus of spatial attention seems to drive object-feature binding, even when spatial attention mistakenly shifts to or lingers at an incorrect, non-target location. How important, then, is the *precision* of spatial attention for successful object-feature binding? A strong interpretation of a spatially-driven binding mechanism would predict that the degree of successful feature-feature binding should increase with the precision of spatial attention. To test this, we performed a supplementary analysis that took location response error as a proxy for the

precision of spatial attention on a given trial, and then fit data from an expanding window of location error with a basic joint-feature mixture model (Supplement S11). As the magnitude of location error increased, correlated target responses decreased and independent target errors increased—suggesting that the degree of non-spatial object-feature binding was indeed related to spatial precision, which could provide further support for the idea that the spatial extent of visual attention is critical to the successful integration of non-spatial features.

DISCUSSION

Our fundamental question was whether visual object integrity survives conditions of *dynamic* attention, by cueing covert spatial attention to shift and split across multiple objects with multiple features. We observed distinct patterns of object-feature binding: splitting attention across multiple locations degrades object integrity, while rapid shifts of spatial attention maintain bound objects, even when misreporting the wrong features. These reliable effects were consistent across both Simple and Full Models and replicated across two independent experiments, with large effect sizes. Moreover, we document a novel attentional phenomenon: inadvertent shifts, or “lapses” of spatial attention, result in erroneous feature reports, but object integrity is still preserved *at the wrong location*. Together, these findings emphasize the importance of a single focus of spatial attention for object-feature binding, even when that focus is dynamically moving across the visual field.

When spatial attention was cued to shift to a new task-relevant location, object-feature binding for the newly-cued target item was generally successful. We probed binding amidst the dynamic process of shifting attention, since shifts were cued almost immediately (50 ms) before array presentation; thus, intact object integrity for the target suggests that attention had rapidly

shifted and fully updated to the new task-relevant location by the time the array appeared on those trials. Critically, on trials with feature errors, there was an increase in reporting *both* features of the non-target that appeared at the initially-cued location (i.e., correlated N1cN1o swap errors). While these swap errors could stem from never seeing the shift cue on those trials, overall performance for location reports was highly accurate, with participants reporting the newly-cued location on 92% of Shift trials. In other words, participants were very good at seeing the shift cue. Thus, a more likely explanation is that these errors reflect the dynamic process of attentional updating, such that attention had not yet fully disengaged from the initially-cued location and re-engaged at the newly-cued location. For instance, previous work has demonstrated that incomplete shifts of attention can result in misreporting a single visual feature of a non-target (e.g., Golomb et al., 2014), but single-feature swaps could present as either correlated or independent errors when multiple features are probed. Here, we show that an incomplete shift of spatial attention preserves object integrity at the initially-attended location. There was no increase in independent errors, such as illusory conjunctions (i.e., swapping just one feature), contrary to what might be expected if attention spread across the two locations (Cohen & Ivry, 1989), or briefly activated both locations simultaneously (e.g., Khayat et al., 2006). Instead, spatial attention seems to shift rapidly from one discrete location to another, binding together the visual features at each single, attended location (Nissen, 1985; Treisman & Gelade, 1980).

In contrast, splitting attention across two locations resulted in more independent feature errors (i.e., reporting only one target feature) rather than correlated feature errors. The lack of correlated feature errors suggests that participants were not simply attending to one of the two cued locations, in hopes that it would be post-cued as the target, nor were they rapidly shifting

attention back and forth between the two cued locations, which would have also been unlikely given that the stimulus array was only presented for 50 ms (Jans, Peters, & de Weerd, 2010). Eye movements were also restricted throughout presentations of the cue, array, and mask. Instead, participants likely attended to both locations simultaneously (Cave, Bush, & Taylor, 2010), perhaps as separate and parallel foci (e.g., Eimer & Grubert, 2014) or within a larger, diffuse attentional window (e.g., Eriksen & St. James, 1986). Contrary to probabilistic accounts of independent feature sampling (Vul & Rich, 2010), binding errors here consisted primarily of unbound guesses rather than illusory conjunctions, as if only one feature dimension was prioritized (e.g., Woodman & Vogel, 2008). Even when participants did report both features of the target item, the cost of splitting attention was decreased precision for both color and orientation. These results are consistent with previous findings in visual working memory (which employ longer encoding and delay periods), in which remembering multiple multi-feature objects also results in feature independence (e.g., Bays et al., 2011; Fougner & Alvarez, 2011) and decreased feature precision (e.g., Bays et al., 2011; Fougner et al., 2010; Park, Sy, Hong, & Tong, 2017).

Surprisingly, even when only one location was cued (Hold trials), participants sometimes failed to sustain spatial attention at that target location. By including a continuous location report in Experiment 2, we demonstrated that participants had occasional lapses of spatial attention and mistakenly identified a random non-target location as the true target location. We emphasize that these lapses of spatial attention are not simply lapses of sustained attention (i.e., momentarily disengaging from the task at hand; Reason, 1984); indeed, the systematic pattern of correlated swap errors during lapses of spatial attention demonstrates that participants remained on task and were not simply mind-wandering. In other words, participants were attending to a location, just

to the wrong location. Thus, simply seeing a spatial cue does not mean that spatial attention is perfectly maintained at that cued location. For example, when monitoring a red traffic light to turn green, a lapse of *sustained* attention might cause you to miss the light change and react more slowly, but it does not mean that you forgot the task or never understood the task. Likewise, a lapse of *spatial* attention does not necessarily mean participants forgot or never saw the cue, but rather that there was a momentary spatial fluctuation away from the actual cue's location. While it is possible that they never saw the cue and randomly picked a location to attend to, this seems unlikely because the initial cue was presented for 250 ms; it seems more likely in the current task that lapses of spatial attention may have been anticipatory in nature—because Hold and Shift trials were randomly intermixed within blocks, participants may have inadvertently but proactively shifted covert attention to a non-cued location in anticipation of a second spatial cue. The lapses of spatial attention in the current task may also reflect the rhythmic nature of visual attention (see VanRullen, 2016). For instance, Fiebelkorn and colleagues (2013) demonstrated that when participants were cued to attend to a single location, they periodically monitored an uncued location at a frequency of 4 Hz. Importantly, rhythmic spatial attention in that study may have also been anticipatory in nature, as cue validity was 75% (i.e., the uncued location was still relevant).

Importantly, the joint-feature errors that arose from these lapses of spatial attention were predictable: Participants systematically reported both the color and orientation of the non-target (i.e., correlated swap error) corresponding to that misidentified location. Just as correlated swap errors occurred after *intentional* shifts of spatial attention, correlated swap errors also occurred after *inadvertent* shifts—or lapses—of spatial attention. Lapses of spatial attention may, however, differ mechanistically from cued shifts of attention in this paradigm, given that

inadvertent shifts in Hold trials were self-directed, while intentional shifts in Shift trials were exogenously cued. The attentional effects of exogenous cueing are thought to be more transient and efficient, although the sustained effects of endogenous cueing may have more time to become more precise and focused (see Jans et al., 2010). Future work could investigate how exogenously- and endogenously-cued shifts of spatial attention might differentially affect object-feature binding.

Importantly, these fully-bound (feature-feature-location) object errors do not reflect traditional failures of feature binding (cf. Bays, 2016), but instead reflect the stubborn persistence of object-feature binding whenever *any* single location is attended, regardless of spatial relevance. Previous studies have argued that spatial location serves as the anchor for object-feature binding (e.g., Nissen, 1985; Reynolds & Desimone, 1999; Pertzov & Husain, 2013; Schneegans & Bays, 2017; Treisman & Gelade, 1980; see also Wolfe & Cave, 1999). For example, Nissen (1985) presented color-shape-location combinations and cued either location (i.e., report color and shape) or color (i.e., report shape and location), while more recently, Schneegans and Bays (2017) presented colored and oriented bars in a continuous-report visual working memory task, and cued the target item by either orientation (i.e., report color and location) or color (i.e., report orientation and location). In both of these studies, when participants reported the incorrect location (i.e., a swap error), their reports of the non-spatial feature were strongly linked to that incorrect location. Although the current study was not designed to examine the specific mechanism by which binding occurs, the results of Experiment 2 reinforce the critical role of spatial location in feature-binding, in addition to revealing the consequences for binding when spatial attention is *dynamic*. We also found little evidence of correlated feature-feature reports without also reporting their corresponding location, consistent

with a spatial-binding mechanism in which the non-spatial features of an object are each bound to its location and only transitively bound to each other via that shared location (e.g., Schneegans & Bays, 2017). Interestingly, some studies have reported feature-feature representations unbound from spatial location in visual working memory (e.g., Logie, Brockmole, & Jaswal, 2011; Saiki, 2016; but see Pertzov & Husain, 2013). Combined with our results, this suggests that when multi-feature objects are perceived and encoded, object integrity is initially anchored to spatial location—but as object representations are uploaded to memory, direct feature-feature binding may survive without the location information.

These findings may have broad implications for object perception, memory, and attentional mechanisms, shedding light on how we manage coherent representations of objects in the world. As spatial attention shifts from location to location, the visual features at each location can be rapidly bound together as cohesive objects. However, object integrity depends on a single and precise focus of spatial attention; splitting attention across multiple locations or increasing the spatial extent of attention comes at a cost. Thus, to avoid potential failures of object integrity, we may default to processing the visual world rapidly and serially (Jans et al., 2010). An intriguing question is whether such processing strategies might be adaptively learned. Our results raise potential implications for other types of dynamic attention (e.g., remapping across eye movements; Golomb et al., 2014), as well as individual differences in attentional abilities (Rosenberg et al., 2016) and attentional deficits (e.g., ADHD, cognitive aging).

References

Bays, P. M. (2016). Evaluating and excluding swap errors in analogue tests of working memory. *Scientific Reports*, 6, 19203. <http://doi.org/10.1038/srep19203>

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), 1–11. <http://doi.org/10.1167/9.10.7>

Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49(6), 1622–1631. <http://doi.org/10.1016/j.neuropsychologia.2010.12.023>

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. <http://doi.org/10.1163/156856897x00357>

Cave, K. R., Bush, W. S., & Taylor, T. G. G. (2010). Split attention as part of a flexible attentional system for complex scenes: Comment on Jans, Peters, and De Weerd (2010). *Psychological Review*, 117(2), 685–695. <http://doi.org/10.1037/a0019083>

Cohen, A., & Ivry, R. (1989). Illusory conjunctions inside and outside the focus of attention. *Journal of Experimental Psychology: Human Perception and Performance*, 15(4), 650–663. <http://doi.org/10.1037/0096-1523.15.4.650>

Eimer, M., & Grubert, A. (2014). Spatial attention can be allocated rapidly and in parallel to new visual objects. *Current Biology*, 24(2), 193–198. <http://doi.org/10.1016/j.cub.2013.12.001>

Emrich, S. M., & Ferber, S. (2012). Competition increases binding errors in visual working memory. *Journal of Vision*, 12(4), 12. <http://doi.org/10.1167/12.4.12>

Eriksen, C. W., & St James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, 40(4), 225–240.
<http://doi.org/10.3758/BF03211502>

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013). Rhythmic sampling within and between objects despite sustained attention at a cued location. *Current Biology*, 23(24), 2553–2558. <https://doi.org/10.1016/j.cub.2013.10.063>

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision*, 11(12), 3.
<http://doi.org/10.1167/11.12.3>

Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision*, 10(12), 27. <http://doi.org/10.1167/10.12.27>

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457–472. <http://doi.org/10.1214/ss/1177011136>

Golomb, J. D., L'Heureux, Z. E., & Kanwisher, N. (2014). Feature-Binding Errors After Eye Movements and Shifts of Attention. *Psychological Science*, 25(5), 1067–1078.
<http://doi.org/10.1177/0956797614522068>

Golomb, J. D., Marino, A. C., Chun, M. M., & Mazer, J. A. (2011). Attention doesn't slide: spatiotopic updating after eye movements instantiates a new, discrete attentional locus. *Attention, Perception, & Psychophysics*, 73(1), 7–14. <http://doi.org/10.3758/s13414-010-0016-3>

Jans, B., Peters, J. C., & De Weerd, P. (2010). Visual spatial attention to multiple locations at once: The jury is still out. *Psychological Review*, 117(2), 637–682.
<http://doi.org/10.1037/a0019082>

OBJECT-FEATURE BINDING UNDER DYNAMIC ATTENTION

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files. *Cognitive Psychology*, 24(2), 175–219. [http://doi.org/10.1016/0010-0285\(92\)90007-o](http://doi.org/10.1016/0010-0285(92)90007-o)

Khayat, P. S., Spekreijse, H., & Roelfsema, P. R. (2006). Attention lights up new object representations before the old ones fade away. *Journal of Neuroscience*, 26(1), 138–142. <http://doi.org/10.1523/JNEUROSCI.2784-05.2006>

Kruschke, J. K. (2011). Bayesian assessment of null values via parameter estimation and model comparison. *Perspectives on Psychological Science*, 6(3), 299–312. <http://doi.org/10.1177/1745691611406925>

Logie, R. H., Brockmole, J. R., & Jaswal, S. (2011). Feature binding in visual short-term memory is unaffected by task-irrelevant changes of location, shape, and color. *Memory & Cognition*, 39(1), 24–36. <http://doi.org/10.3758/s13421-010-0001-z>

Müller, M. M., Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature Neuroscience*, 1(7), 631–634. <http://doi.org/10.1038/2865>

Nissen, M. J. (1985). Accessing features and objects: Is location special? In M. I. Posner & O. S. Marin (Eds.), *Attention and Performance XI* (pp. 205–219).

Park, Y. E., Sy, J. L., Hong, S. W., & Tong, F. (2017). Reprioritization of features of multidimensional objects stored in visual working memory. *Psychological Science*, 28(12), 1773–1785. <http://doi.org/10.1177/0956797617719949>

Pertzov, Y., & Husain, M. (2013). The privileged role of location in visual working memory. *Attention, Perception, & Psychophysics*, 76(7), 1914–1924. <http://doi.org/10.3758/s13414-013-0541-y>

OBJECT-FEATURE BINDING UNDER DYNAMIC ATTENTION

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109(2), 160–174. <http://doi.org/10.1037/0096-3445.109.2.160>

Reason, J. (1984). Lapses of attention in everyday life. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of Attention* (pp. 515–549). Orlando, FL: Academic Press.

Reynolds, J. H., & Desimone, R. (1999). The role of neural mechanisms of attention in solving the binding problem. *Neuron*, 24(1), 19–29. [http://doi.org/10.1016/s0896-6273\(00\)80819-3](http://doi.org/10.1016/s0896-6273(00)80819-3)

Robertson, L. C. (2003). Binding, spatial attention and perceptual awareness. *Nature Reviews Neuroscience*, 4(2), 93–102. <http://doi.org/10.1038/nrn1030>

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2016). A neuromarker of sustained attention from whole-brain functional connectivity. *Nature Neuroscience*, 19(1), 165–171. <http://doi.org/10.1038/nn.4179>

Saiki, J. (2016). Location-unbound color-shape binding representations in visual working memory. *Psychological Science*, 27(2), 178–190. <http://doi.org/10.1177/0956797615616797>

Schneegans, S., & Bays, P. M. (2017). Neural architecture for feature binding in visual working memory. *Journal of Neuroscience*, 37(14), 3913–3925. <http://doi.org/10.1523/JNEUROSCI.3493-16.2017>

Schoenfeld, M. A., Tempelmann, C., Martinez, A., Hopf, J. M., Sattler, C., Heinze, H. J., & Hillyard, S. A. (2003). Dynamics of feature binding during object-selective attention. *Proceedings of the National Academy of Sciences*, 100(20), 11806–11811. <http://doi.org/10.1073/pnas.1932820100>

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.

<http://doi.org/10.1167/13.10.9>

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2), 171–178.

[https://doi.org/10.1016/S0959-4388\(96\)80070-5](https://doi.org/10.1016/S0959-4388(96)80070-5)

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. [https://doi.org/10.1016/0010-0285\(80\)90005-5](https://doi.org/10.1016/0010-0285(80)90005-5)

Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107–141. <http://doi.org/10.1093/acprof:osobl/9780199734337.003.0019>

VanRullen, R. (2016). Perceptual cycles. *Trends in Cognitive Sciences*, 20(10), 723–735.

<https://doi.org/10.1016/j.tics.2016.07.006>

Vul, E., & Rich, A. N. (2010). Independent sampling of features enables conscious perception of bound objects. *Psychological Science*, 21(8), 1168–1175.

<http://doi.org/10.1177/0956797610377341>

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12), 1120–1135. <http://doi.org/10.1167/4.12.11>

Wolfe, J. M., & Cave, K. R. (1999). The psychophysical evidence for a binding problem in human vision. *Neuron*, 24(1), 11–17. [http://doi.org/10.1016/s0896-6273\(00\)80818-1](http://doi.org/10.1016/s0896-6273(00)80818-1)

Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin & Review*, 15(1), 223–229.

<http://doi.org/10.3758/PBR.15.1.223>

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235. <http://doi.org/10.1038/nature06860>