

Shifting Expectations: Lapses in Spatial Attention are Driven by Anticipatory Attentional Shifts

Authors: Christopher M. Jones¹, Emma Wu Dowd², Julie D. Golomb¹

Affiliations: The Ohio State University¹, University of Texas at Austin²

Abstract

Attention is dynamic, constantly shifting between different locations – sometimes imperfectly. How do goal-driven expectations impact dynamic spatial attention? A previous study (Dowd & Golomb, 2019) explored object-feature binding when covert attention needed to be either maintained at a single location or shifted from one location to another. In addition to revealing feature-binding errors during dynamic shifts of attention, this study unexpectedly found that participants sometimes made correlated errors on trials when they did not have to shift attention, mistakenly reporting the features and location of an object at a different location. The authors posited that these errors represent “lapses” in spatial attention, which are perhaps driven by the implicit sampling of other locations in anticipation of having to shift attention. To investigate whether these spatial lapses are indeed anticipatory, we conducted a series of 4 experiments. We first replicated Dowd & Golomb (2019)’s original finding of spatial lapses, and then showed that these spatial lapses were not observed in contexts where participants are not expecting to have to shift attention. We then tested contexts where the direction of attentional shifts was spatially predictable, and found that participants lapse preferentially to more likely shift locations. Finally, we found that spatial lapses do not seem to be driven by explicit knowledge of likely shift locations. Combined, these results suggest that spatial lapses of attention are induced by the implicit anticipation of making an attentional shift, providing further insight into the interplay between implicit expectations, dynamic spatial attention, and visual perception.

Acknowledgements: We would like to thank Veronica Olaker & Maurryce Starks for their assistance with data collection.

Funding: NIH R01-EY025648 (JG), F32-EY028011(EWD); NSF 1848939 (JG)

Introduction

At any given time, we utilize a variety of cognitive resources in order to successfully perceive and navigate our world. Because our experience of the world is often visual, resources such as visual attention and memory are crucial to successful perception of the world around us. However, just because these resources are effective does not mean they always function flawlessly.

Attention is not a binary resource that is either on or off; rather, it is flexible and “waxes and wanes” (Esterman, Noonan, Rosenberg, & DeGutis, 2013). Lapses (sometimes referred to as “slips”) of sustained attention are particularly well researched examples of ways attentional resources may fail. Typically, lapses in sustained attention are associated with mind-wandering and attention drifting away from a current task (Cheyne, Carriere, & Smilek, 2006; Reason, 1984; Roca, Lupiáñez, López-Ramón, & Castro, 2013; Smallwood et al., 2004), but they can also have consequences for other cognitive processes, like subsequent memory performance or working memory capacity (deBettencourt, Norman, & Turk-Browne 2018). Other work has found that lapses of sustained attention can be predicted by fMRI activity (Esterman et al., 2013; Rosenberg et al., 2017; Rosenberg et al., 2015), and are correlated with working memory capacity and measures of fluid intelligence (Unsworth, Lakey, & Young, 2010). Importantly, lapses of sustained attention can have strong consequences for tasks requiring sustained attention over long durations, such as driving (Roca et al. 2013).

Distinct from the more commonly studied lapses in *sustained* attention, lapses in *spatial* attention may be thought of as an additional type of attentional failure, where attention briefly highlights a task irrelevant location. If one thinks of attention as a flashlight which shines on a location of interest, a lapse in sustained attention would be like a momentary dimming of the lightbulb, while a lapse in spatial attention would be like the flashlight slipping to illuminate the wrong location. Put another way, if lapses of sustained attention reflect periods of inattention, lapses of spatial attention may reflect periods of mis-attention.

The term “lapse of spatial attention” was first coined by Dowd & Golomb (2019), in describing a somewhat incidental finding. They had participants perform a task in which spatial attention was manipulated in one of three conditions: holding attention at a single spatially cued location, dynamically shifting attention to a second cued location, or splitting attention simultaneously between two locations, and participants were then asked to reproduce the color, orientation, and location of a target object. The main finding of the paper was that object-feature integrity (reporting all three features from the same object) was preserved when shifting goal-directed attention from one location to another, in contrast to a degradation of object integrity when attention was concurrently split between two separate locations. On the shifting attention trials, the participants sometimes made errors where they reported all three features from the object at the initially cued location, as if spatial attention had not yet updated to the new location at the time of the probe. But interestingly, similar “correlated” errors were also observed on a small but reliable portion of single-cue (holding attention) trials, where even though participants had to simply maintain attention at a single location, they sometimes erroneously reported all three features (color, orientation, and location) of a distractor object at a different location. Dowd &

Golomb suggested that these errors on single-cue trials were indicative of lapses in spatial attention, where attention happened to be focused on an incorrect location at the critical point in the trial. Again, these errors were distinct in nature from a more general lapse of sustained attention (inattention), which would have been expected to result in random guessing of features in an unbound (uncorrelated) fashion (Dowd & Golomb, 2019). In other words, if attention was unfocused and participants were not attending anywhere in the object display at the time of stimulus presentation, errors would reflect random guessing of each feature. Instead, participants reported a correctly-bound object (all three features) that was present in the display, indicating that spatial attention was focused somewhere during stimulus presentation, just not at the location they were supposed to be attending to.

What is the nature of these lapses of spatial attention? Some research has shown that attention routinely samples different spatial locations with some rhythmicity, most commonly reported at 7-8Hz theta frequencies for sampling across spatial locations or 4Hz cycles for performance fluctuations at a single location (Fiebelkorn et al., 2013; Re et al., 2019; Landau & Fries 2012; VanRullen, 2007). Similar to how we explore scenes with overt saccadic eye movements 3-4 times per second (Steinman et al., 1973), even with the eyes fixated covert attention might explore space in these rhythmic patterns akin to “attentional saccades” (Gaillard et al., 2020). While rhythmic oscillations may be entrained (Thut et al., 2011) or phase-reset (Landau & Fries 2012; Gaillard et al., 2020) by stimulus events, there is also increasing evidence that oscillations in attention are intrinsic in nature (Fiebelkorn & Kastner 2019). One possibility is that the spatial lapses seen in the Dowd & Golomb (2019) study merely reflect this routine rhythmic sampling of other spatial locations. However, it is also possible that these errors may have occurred

because participants were anticipating having to make an attentional shift on some trials; i.e., that the lapses were not random or automatic, but perhaps more adaptive based on task context.

Previous research has shown that task expectations can have strong implications for how attention is deployed. For instance, incidental learning of spatial probabilities can guide attention, such that participants find search targets faster when they appear in a higher probability (“rich”) location, even when participants are unable to explicitly identify the rich location (Geng & Behrmann, 2005; Jiang, Swallow, & Rosenbaum, 2013). It has been found that spatial expectations (where something may appear) and temporal expectations (when something may appear) can each individually improve performance on trials congruent with these expectations (Rohenkohl, Gould, Pessoa, & Nobre 2014). Rohenkohl et al (2014) also showed that spatial and temporal expectations can synergistically interact to improve target perception. Therefore, if knowing where and when a stimulus might occur is associated with perceptual benefits, perhaps attention preemptively samples likely stimulus locations at the time they would be likely to appear. For example, in the Dowd & Golomb (2019) task, if participants were expecting to have to shift attention to a second cue on some trials, then their spatial attention may have been shifting to sample other locations in anticipation of this second cue.

In the current study we test the speculation that these lapses of spatial attention may be driven by anticipatory sampling of other spatial locations due to the expectation of an upcoming attentional shift. (Note that our primary goal here is to test whether behavioral lapses of spatial attention may be anticipatory in nature; we speculate more on the potential links with rhythmic oscillations in the discussion.) We modified the paradigm used by Dowd & Golomb (2019) and

conducted 4 experiments manipulating the predictability of the second cue to better characterize the nature of these spatial lapses. The first experiment (Experiment 1: “Non-Predictive Second Cue”) was intended as replication of Dowd & Golomb (2019), with an equal mix of single-cue trials on which participants had to maintain attention at a single spatial location and double-cue trials on which participants had to covertly shift their attention from one location to another. On double-cue (shift) trials, the second cue could be located either clockwise or counter-clockwise to the first cue, and thus was not predictable. In Experiment 2 (“Single-Cue Only”) participants only encountered single-cue trials over the course of the experiment. If lapses in spatial attention are driven by the expectation of having to make an attentional shift, then we hypothesized that removing that expectation should reduce or eliminate lapses in spatial attention. In Experiment 3 (“Clockwise Second Cue”), we again included both single- and double-cue trials, but the second cue, when it occurred, was always located in the position clockwise to the initially cued location. This experiment aimed to determine whether participants lapse preferentially to more likely shift locations, if the direction of attentional shifts were spatially predictable. Experiment 4 (“Counter-Clockwise Second Cue”) was designed in a similar way, except the location of the second cue, when present, was always counter-clockwise to the first cue. We focus our analyses primarily on the comparison of lapses to the clockwise versus counterclockwise nontarget positions, which are matched in every way on hold trials, except for the likelihood of shift trials to these locations. Experiment 4 also included an explicit knowledge task to evaluate whether participants were explicitly aware that the shift direction was predictable.

Methods

Open Science Practices

The current study was designed to closely follow methods reported in Dowd & Golomb (2019).

Although Experiments 1-3 were not formally preregistered, the experimental design, participant inclusion criteria, and analyses follow those described in the previous paper as closely as possible, except where noted. These first three experiments were conducted in parallel, with participants randomly assigned among them. Experiment 4 was conducted after analyzing the first three experiments, and was pre-registered at {<https://osf.io/mxq9j> }.

Participants

Dowd & Golomb (2019) included an a priori power analysis estimating a sample size of 22 participants to detect feature errors with 80% power. Based on this, we set 22 participants as our minimum sample size for each experiment. In anticipation of participant exclusions, we collected a few extra participants in each experiment. Thus, in Experiment 1, 23 participants (ages 18 - 23; 13 Male, 10 Female) were included in analysis. Experiment 2 had 26 participants (ages 18 – 26; 17 Male, 9 Female) included for analysis. Experiment 3 had data from 24 participants (ages 18 – 20; 16 Male, 8 Female) included in analysis, and Experiment 4 had 23 participants (ages 18 – 21; 9 Male, 14 Female). In order to participate in this experiment, participants had to report that they had normal or corrected-to-normal vision and visual acuity. Individual trials were considered usable if the participant maintained accurate fixation on the fixation dot (see details below) and made a response for color, orientation, and location. Participants who completed at least 80 usable hold and shift trials were included in analysis. Participants were also excluded for poor task performance, quantified as a probability < 0.5 of reporting the correct color and orientation of the target (pTcTo) on hold trials, consistent with Dowd & Golomb (2019). Of the 128

participants across all experiments who completed the entire session, 10 did not have enough trials for analysis, and 21 were excluded based on the pT_CT_O criteria. An additional participant was excluded because they stated that they always reported the object at the first cue, regardless of trial type. Participants were compensated with either course credit or a payment of \$10 per hour. All participants provided informed consent in accordance with The Ohio State University Institutional Review Board.

Experimental setup

All stimuli were presented on a 21 in. flat screen CRT monitor (ViewSonic Graphic Series G225f) with a refresh rate of 85Hz and a screen resolution of 1,280 x 1,024 pixels. Each monitor was color calibrated using a Minolta CS-100. Each participant sat with their head in a chin rest approximately 60cm from the monitor. The stimuli were presented in MATLAB using the Psychophysics toolbox (Brainard, 1997; Kleiner et al, 2007). Eye position was tracked using an Eyelink 1000 eye-tracking system. Each participant had their left eye tracked and 500 fixation samples were collected per second. Before beginning the experiment, each participant was calibrated using a 5-point calibration array. Each participant was calibrated so that the average error across calibration points was less than 1° error and no individual point had greater than 2° error. Participants were recalibrated during the experiment as necessary.

Stimuli and Procedure

The general paradigm of all four experiments was nearly identical to Experiment 2 of Dowd & Golomb (2019). Figure 1a illustrates the trial sequences for Hold attention (single cue) and Shift attention (double cue) trials. At the start of each trial participants were required to maintain

fixation on a fixation dot for 1000ms. If the participant moved their eyes before fixating for 1000ms the fixation dot blinked red to encourage participants to return to fixation. The trial did not begin until participants maintained fixation for a consecutive 1000ms. After the trial began, participants received a spatial cue at one of sixteen locations around an invisible circle (7.4° eccentricity) on the screen. The cue was a $4^\circ \times 4^\circ$ black square outline. The cue was presented for 250ms. During single-cue trials there was a blank fixation screen presented for 1100ms before the object array. During double-cue trials there was a blank fixation screen presented for 1000ms followed by a second spatial cue displayed for 50ms, and then a 50ms blank.

In Experiment 1 the second cue could be located at the position either 90° (clockwise) or -90° (counter-clockwise) from the first spatial cue (Fig 1b). In Experiment 2 participants never encountered a second spatial cue. In Experiment 3 the second spatial cue was always at the position 90° (clockwise) from the first spatial cue. In Experiment 4 this second cue was always located at the position -90° (counter-clockwise) from the first spatial cue. The second spatial cue (when present) was always followed by another 50ms fixation delay before the presentation of the stimulus array and mask.

The stimulus array was then presented for 50ms. This array contained four colored and tilted bars ($.75^\circ \times 4^\circ$); one at the cued location (target item), and the others equally spaced 90° along the invisible circle surrounding the fixation dot. The target item was a randomly selected color and orientation, with the other items' colors and orientations independently spaced 90° away in color space and 45° away in orientation space from each other. Participants were instructed to attend to the object appearing at the location of the most recent spatial cue (the target item). The array was

presented for 50ms after which the stimulus locations were masked for 200ms with $4^\circ \times 4^\circ$ masks.

The stimulus masks were squares with a random color assigned to each pixel.

Participants were then asked to reproduce the color and orientation of the target. A probe stimulus with random initial color and orientation values was presented in the center of the screen. This stimulus was the same size as the stimuli presented in the visual array. Participants were instructed to use the keyboard to manipulate the color and orientation of the probe until it matched the target item. Participants used the 'x' and 'z' keys with their right hand to adjust the color of the object and '<' or '>' with their left hand to adjust the orientation of the stimulus. Participants were able to adjust color in steps of 2° and Orientation in steps of 1° . Participants were instructed to press the Space Bar once they thought the probe matched the target item. Participants had 10s to submit their response.

After submitting a response for color and orientation participants were then asked to indicate the target's location (i.e. which location was most recently cued). Participants used the '<' and '>' keys to move a white square outline around a circle and were instructed to press the spacebar to enter their final response. The dimensions of the square were exactly the same as the spatial cue ($4^\circ \times 4^\circ$ visual angle). Participants were able to adjust the location of the black outline in intervals of 2.25° . Participants had 5 seconds to submit their response.

After participants submitted their responses they received visual feedback on their performance for that trial. During feedback the original target object was displayed in its original location and the reported object was displayed in the center of the screen. The white square participants used

to report the target location was also displayed in the location they reported simultaneously during feedback. Participants also received feedback on their eye-tracking performance. Eye tracking feedback was given as the percent of eye-tracking samples within 2° of the fixation dot (out of total eye-tracking samples for that trial). Feedback was given in black text if participants successfully maintained fixation during more than 90% of the trial and red text if less than 90%. Trials were excluded from analysis if participants had more than 15% deviant eye-tracking samples (pre-registered criterion consistent with Dowd & Golomb (2019)). This resulted in the exclusion of an average of 8.8% of trials in Experiment 1, 7.2% Experiment 2, 8.2% Experiment 3, and 9.0% Experiment 4. (Posthoc analyses confirmed that eye-tracking performance was highly accurate on the included trials: on average only 0.53% of eye-tracking samples were deviant when considering the entire 2450ms trial duration, and 0.21% of samples during the critical 50ms stimulus presentation period.) Eyetracking feedback was presented at the same time as color, orientation, and location feedback. All performance feedback was presented for 1500ms. After a 500ms blank inter-trial interval, participants began the next trial.

In all experiments, participants first completed 2 short practice blocks. The first block consisted of only hold trials; participants were given instructions pertaining to the hold trials before this block, so instructions for hold trials were the same across all four experiments. In Experiments 1, 3, and 4 participants, participants were then alerted to the possibility of shift trials in the main task, and received the instructions for the shift trials (identical instructions for each experiment), followed by a practice block consisting of intermixed hold and shift trials. In Experiment 2, participants completed a second practice blocks consisting only of hold trials, with no additional instructions.

Each main experimental block in Experiments 1, 3, and 4 contained 16 single-cue (hold) trials and 16 double-cue (shift) trials, which were randomly intermixed. Each main experimental block in Experiment 2 contained 16 single-cue (hold) trials and no double-cue (shift) trials. For each experiment, participants completed between 5-8 blocks; they were asked to complete up to 8 blocks if time allowed, however their data were included in analysis if they completed at least 5 blocks and stayed for the entire experiment session.

Experiment 4 Exit Questions & Explicit Knowledge Task

Participants in Experiment 4 also completed an additional explicit knowledge task. Participants were first asked a single question: “On trials where there were two cues did you notice a pattern in where the second cue would appear?”, and were given the option to answer “yes” or “no”. Regardless of their answer, they then completed 16 trials of an explicit knowledge task (Fig 1c). Each trial presented an array of four squares in a configuration seen during the main experiment. One square was black; participants were informed that this black square represented the first spatial cue. The remaining 3 squares were white and filled with the digits 1-3. Participants were asked “If there is a second cue, please use the number pad to indicate your guess of where it might occur.” Participants used the number pad to input their response (white square 1, 2, or 3). In order to discourage participants from pressing the same number on all trials, response numbers were randomly assigned to each white square on each trial. The black square representing the first cue appeared at each of the 16 possible stimulus locations once during the 16 trials, and the order in which participants completed the trials was randomized. If participants had explicitly learned the

task structure, the correct response would have been the white square in the counter-clockwise position on every trial.

Error Calculation and Alignment

For the main task, we recorded participants' continuous responses for each of the three feature dimensions (color, orientation, location) on each trial. Error was calculated as the angular deviation between the reported feature value and the actual target value on that trial, such that a perfect report of the target feature would be 0° error. Color and Location had an error range of -180° to 180° and Orientation had an error range of -90° to 90° . For modeling purposes all orientation response errors were multiplied by 2 so that all possible feature errors had a range of -180° to 180° .

So that we could interpret responses jointly across all three dimensions, we coded features in terms of the four items in the display (target: T, and nontargets: N1, N2, N3). On hold trials N1 was always defined as the nontarget object located clockwise on the screen from the target object, N2 was always the nontarget object located counterclockwise in the array from the target, and N3 was always the nontarget located furthest from (opposite) the target. On shift trials, N1 was defined as the object at the initially cued location, N2 was the other object immediately adjacent to the target, and N3 was the object opposite the target. We then directionally aligned the errors for each of the three dimensions such that in each dimension, 0° error represented the correct T feature, errors in the direction of the N1 feature in feature-space were positively signed, and errors in the direction of the N2 were negatively signed. For example: after error alignment, a color response with 90° error would indicate participants reported the color of the N1 object,

while an orientation response with -45° error would indicate participants reported the orientation of the N2 object, and a location response with 0° error would indicate participants reported the location of the T object. Because participants make an independent response for each feature dimension, it was possible to report different features from different objects (i.e. N1 color with N2 orientation and the correct target location).

Statistical Models

We used probabilistic mixture modeling (Bays, Catalao, & Husain, 2009; Dowd & Golomb, 2019, Zhang & Luck, 2008) to determine the probability of reporting a given object's features and spatial location. Similar to Dowd & Golomb (2019), data from each experiment and each trial type (i.e. hold or shift) were analyzed using triple-joint-probabilistic models. We used two triple-feature mixture models (Table 1): a “standard” triple mixture model, which is the same triple mixture model used by Dowd & Golomb (2019), and a simple triple mixture model which focused on a more limited subset of response combinations and allowed us within-participant estimates for the spatial lapse comparisons. The goal was to model the probability that a given feature could be attributed to the target (p_T), a nontarget (p_{N1} , p_{N2} , or p_{N3}), or a random guess (p_U), allowing for different combinations across the three feature dimensions (e.g., reporting all three features of the target object, or reporting the target object's color with N1's location and randomly guessed orientation). The probabilities of reporting the T, N1, N2, or N3 features were modeled with von mises distributions centered around 0° , 90° , -90° , or 180° , respectively, and random guessing was modeled as a uniform distribution across all possible feature responses.

The joint distribution of responses was modeled as follows:

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

where θ_C , θ_O , and θ_L are degrees error between the reported value and the target value for each feature (color, orientation, and location, respectively), m is the number of color-orientation-location response combinations, α_m is the probability of each response combination, and p_m represents the combined probability density, as listed in Table 1. The standard triple model and the simple triple model have different numbers of color-orientation-location response combinations such that $m_{standard} = 1:20$ and $m_{simple} = 1:13$. The standard triple joint model includes 20 response combinations of location (4: T_L , $N1_L$, $N2_L$, $N3_L$) x color-orientation (5: $T_C T_O$, $N1_C N1_O$, $N2_C N2_O$, $N3_C N3_O$, $U_C U_O$); per Dowd & Golomb (2019), unbound color-orientation reports and location-guesses were not included in the model because their probabilities were negligible and of minimal theoretical interest. We further simplified this model into 13 response combinations in the simple triple model, by further combining certain response combinations of negligible theoretical interest and modeling them with single probability estimates (e.g., $N1_C N1_O N3_L$ and $N1_C N1_O N2_L$ were combined into $pN1_C N1_O N23_L$). Statistical tests on parameter estimates from the simple triple model were completed using Jeffreys's Amazing Statistics Program (JASP Team, 2019).

Due to the large number of parameters of the standard triple model, it was fitted across data collapsed across all subjects. The simple triple model was fitted on individual participants. Both models were fitted separately for hold trials and shift trials and separately for each experiment. We used the same model-fitting procedures as described in Dowd & Golomb (2019): Markov

chain Monte Carlo procedure implemented through custom MATLAB scripts (available at osf.io/h2xpu/) using the MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013) through the Ohio Supercomputer Center (1987, <https://www.osc.edu/>). We collected 15,000 postconvergence samples and used the posterior distributions to compute the maximum-likelihood estimates of each parameter as well as its 95% highest-density interval (HDI). Parameter estimates from the standard triple model were considered significantly different if their 95% HDIs did not overlap (Kruschke, 2011). For the simple model, standard within-subjects statistical tests (e.g. paired t-tests, ANOVAs) were used to assess significance.

Results

Figure 2 shows joint-response scatterplots of color, orientation, and location responses for each experiment, separately for hold (single cue) and shift (double cue) trials. Generally speaking, across all experiments, if a participant reported a given location on a trial – whether it was the target location or a non-target location – they also tended to report the features of the object that was presented at that spatial location, consistent with Dowd & Golomb (2019). To quantify differences across our four experiments, especially for lapses of spatial attention, we fit each of these datasets using the triple mixture models described above. Parameter estimates for all model parameters are shown in Tables 2 and 4 for the hold trials and 3 and 5 for the shift trials.

Hold Trials: Overall performance

Across all four experiments participants were generally able to perform the task successfully, reporting the correct color, orientation and location of the target object on the majority of trials.

In Experiment 1, parameter estimates from the standard triple model showed that the probability of correctly reporting the triple bound target (p_{TcToT_L}) on hold trials was .843 (95% HDI = [.820, .853]), which was comparable to the probability reported in Dowd & Golomb (2019). In Experiment 2, p_{TcToT_L} was .840 (95% HDI = [.825, .858]), in Experiment 3 it was .818 (95% HDI = [.804, .834]), and in Experiment 4 it was .870 (95% HDI = [.855, .880]). Values obtained from the simple triple model were similar (Table 4). A 1x4 between-subjects Analysis of Variance (ANOVA) on individual p_{TcToT_L} estimates from the simple triple model revealed no significant difference across experiments in correctly reporting the triple-bound color, orientation, and location of the target object, $F(3, 92) = 0.709$, $p = 0.549$, $\eta^2_p = .023$.

Lapses of spatial attention

Our primary goal in this study was to measure lapses of spatial attention under experiments with differing task contexts and expectations regarding attentional shifts. Following Dowd & Golomb (2019), we defined lapses in spatial attention as triple bound swaps on hold trials (i.e. reporting a non-target object's color, orientation, and location, e.g. $p_{N1cN1oN1L}$). In this set of experiments, our primary comparison is of lapses to the clockwise location ($p_{N1cN1oN1L}$) vs lapses to the counter-clockwise location ($p_{N2cN2oN2L}$), specifically because these are the only possible locations of a second cue on shift trials across Experiments 1, 3, and 4, and they are matched for all other factors (e.g., distance from the target in physical space and in feature space). Stimuli at the N3 location were mainly included to make our feature space more evenly distributed and less predictable; this location was not an a priori focus of our analyses, but we report findings for this N3 location in an exploratory section below.

Experiment 1 was intended as a replication of Dowd & Golomb (2019). Participants completed intermixed hold and shift trials, and the direction of the shift, when it occurred, was not predictable. We predicted that on hold trials participants would sometimes experience lapses in spatial attention and report triple bound non-target objects, and that they would have approximately the same probability of making these triple bound swaps to each of the two neighboring non-target objects (N1 and N2), similar to Dowd & Golomb (2019). Indeed, in Experiment 1 (Figure 3a) the standard triple model results revealed no statistically reliable difference in the probability of making a clockwise lapse of spatial attention ($p_{N1cN1oN1L} = .012$, 95% HDI = [.008, .015]) compared to a counterclockwise lapse of spatial attention ($p_{N2cN2oN2L} = .008$ 95% HDI [.006, .012]), and both probabilities were credibly greater than zero (95% HDIs not overlapping with 0). The simple triple model results confirmed that participants were not significantly more likely to lapse to the N1 vs N2 object, $t(22) = 1.081$, $p = .292$, $d = .225$.

Experiment 2 aimed to determine if lapses in spatial attention are induced by the expectation of having to dynamically shift attention on some trials. We hypothesized that if lapses are induced by the expectation of having to make a shift, then participants would make fewer lapses if they never had to shift attention within a trial. In Experiment 2, 100% of trials were hold (single cue) trials. Parameter estimates from the standard triple model showed that participants had an extremely low probability of making triple bound swaps for all non-target objects in this experiment (Figure 3b). The probability of participants making a triple swap to the N1 object ($p_{N1cN1oN1L} = .000$, 95% HDI = [.000, 0.002]), the N2 object ($p_{N2cN2oN2L} = .001$, 95% HDI = [.000, 0.002]), and the N3 object ($p_{N3cN3oN3L} = .001$, 95% HDI = [.000, .002]) were all

lower than the lowest triple bound parameter across all other experiments, and the HDIs all overlapped with zero and with each other. Additionally, a t-test on parameter estimates from the simple triple model showed that there was no difference in the probability of participants reporting the triple bound N1 object vs the triple bound N2 object, $t(25) = -0.554, p = .585, d = -0.109$.

Moreover, a between-subjects 2x2 repeated measures ANOVA comparing triple-swaps of N1 vs N2 nontargets across Experiments 1 and 2 found a significant main effect of experiment ($F(1,47) = 8.264, p = .006, \eta^2_p = .15$), but no significant main effect of nontarget object ($F(1,47) = 0.757, p = 0.389, \eta^2_p = .016$), nor a significant interaction ($F(1,47) = 1.593, p = .213, \eta^2_p = .033$). Post-hoc between-subjects t-tests showed that participants in Experiment 2 made fewer triple swaps than participants in Experiment 1 for both the N1 ($t(47) = 2.686, p = .01, d = .769$) and N2 ($t(47) = 2.594, p = .013, d = 0.743$) objects.

In Experiments 3 and 4 we next wanted to determine whether participants would lapse preferentially to predictable shift locations. We asked: if the second cue was always predictable in where it would appear, would this affect participants' distribution of lapses? In Experiment 3 the second cue, when it occurred, was always clockwise to the first spatial cue, and in Experiment 4 the second cue was always counterclockwise. On hold trials in both experiments, participants had a higher probability of lapsing to the predictive location than to the other adjacent (control) location (Figure 3c-d). The standard triple model parameter estimates from Experiment 3 show that participants had reliably a higher probability of reporting the features and location of the object located clockwise to the target ($p_{N1cN1oN1L} = .021, 95\% \text{ HDI} =$

[.017, .026] compared to the object counterclockwise to the target ($p_{N2cN2oN2L} = .011$, 95% HDI = [.006, .013]). In contrast, participants in Experiment 4 had a reliably higher probability of reporting the triple bound features of the counterclockwise N2 object ($p_{N2cN2oN2L} = .014$, 95% HDI = [.010, .017]) compared to the clockwise N1 object ($p_{N1cN1oN1L} = .004$, 95% HDI = [.002, .007]).

The simple triple model analyses confirmed that in Experiment 3, participants reported the triple bound N1 (Clockwise) nontarget more often than the N2 (Counter-Clockwise) nontarget ($t(23) = 2.539$, $p = .018$, $d = .518$), while participants in Experiment 4 reported the triple bound N2 more than the triple bound N1 ($t(22) = -2.517$, $p = .02$, $d = -.525$). A between-subjects 2x2 ANOVA comparing non-target object reports (N1 Triple-Swap, N2 Triple-Swap) x Experiment (3, 4) revealed no significant main effect of non-target object report ($F(1,45) = .006$, $p = .937$, $\eta^2_p = 0$), but did show a significant main effect of Experiment ($F(1,45) = 4.738$, $p = .035$, $\eta^2_p = .095$), with the overall rate of lapses higher for the participants in Experiment 3. Importantly, the interaction was significant ($F(1,45) = 12.749$, $p < .001$, $\eta^2_p = .221$), consistent with a different relative pattern of spatial lapses in the two experiments.

Finally, we conducted an exploratory analysis of Experiment 4 (which had the most data per subject) to explore learning effects over the duration of the experiment. We aggregated data across subjects and separately modeled data for each block of trials using the simple triple model. Figure 4a shows the relative rates of lapsing to each of the 3 nontarget locations (N1, N2, N3), as a proportion of the total lapses on that block (Figure 4b). Consistent with an effect driven by learned expectations (here that the N2 location is the predictive shift location), the proportions

of N1 and N2 lapses trended in opposite directions over time. The proportion of N1 lapses significantly decreased over the duration of the experiment ($r(6) = -.803$, $p = .017$), while the proportion of lapses to the predictive N2 location showed a nonsignificant increase ($r(6) = .491$, $p = .216$). Interestingly, the total lapse rate (sum of N1, N2, and N3 lapses) did not significantly change over the course of the experiment ($r(6) = .338$, $p = .412$), nor did the N3 rate ($r(6) = -.149$, $p = .724$; see section below), suggesting that the preferential lapse pattern may have been driven more by a decrease in lapses to the unpredictable N1 location as participants learned the probabilities, though we reiterate that this timecourse analysis was exploratory and likely underpowered.

N3 swaps

Our primary analyses focused on the comparison of attentional lapses to the N1 vs N2 nontargets, since the experiments were designed to equate these two locations across all factors except learned probabilities. As previously discussed, the N3 object was included primarily as a control item of non-interest, intended to make feature values evenly spaced 90° in all directions and thus non-predictable. However, in analyzing the model parameters for the N3 responses, we found an unexpectedly high probability of lapsing to this N3 non-target across multiple experiments (this trend was also present, but not significant, in Dowd & Golomb, 2019), even though it was never a possible shift location in any of the experiments. In Experiment 1 lapses to the N3 location were numerically higher than both N1 and N2 (significantly more so than N2; Table 2), and in Experiments 3 and 4, participants were as likely to lapse to the object at the N3 location as they were the predictive location of the second cue (95% HDIs overlap). Although we don't have a clear explanation for why participants made so many N3 lapses, we note that the

relative likelihood of making N3 lapses was comparable across Experiments 3 and 4 (.326 and .341, respectively; $t(45) = -.248$, $p = .805$, $d = -.072$) – and over time across blocks in the exploratory timecourse analysis (Figure 4a) – unlike the significant cross-over interactions in the relative proportions of N1 and N2 lapses, reported above. We speculate further on these N3 errors in the discussion.

Experiment 4 Explicit Knowledge Task

In Experiment 4 (pre-registered and conducted after the other 3 experiments), we also included an explicit knowledge task at the end of the experiment. In this task participants a) reported if they noticed that the location of the second cue was predictable (yes/no), and then b) performed a series of trials in which they were given the location of the first cue and asked to guess in which location the second cue was most likely to occur if there was a shift. Only 7 participants (out of 23 total) reported yes to the first question, that they noticed that the location of the second cue was predictable. Moreover, participants performed at chance in choosing the location of the second-cue in the explicit report task; across all participants the mean probability of reporting the correct (counterclockwise) N2 location was 29.35%, where chance was 33.33% (one-sample t-test: $t(22) = -.996$, $p = .33$, $d = -.208$). The 7 participants who stated that they noticed a pattern reported the correct N2 location 29.46% of the time, while the 16 participants who stated they did not notice the second cue was predictive reported the correct N2 location on 29.3% of trials. Furthermore, the main pattern of results reported above for Experiment 4 did not differ as a function of response to this question; excluding the 7 participants who reported explicit knowledge, there was still a significantly greater likelihood of lapsing to the N2 location than the N1 location ($t(15) = 2.249$, $p = .04$, $d = 0.562$).

We also explored if there were individual differences between which location participants thought was the likely second-cue location and if these values were correlated with the relative likelihood of spatial lapses to different locations. We computed correlations between the percentage of trials on which each participant reported a given non-target location in the post-test and the relative proportion of triple swaps (spatial lapses) that participant made to the corresponding non-target object in the main task (calculated as the proportion out of total triple swaps for that subject). We found that participants who had a higher proportion of reporting the N2 location in the post-test also had a higher proportion of triple swaps to the N2 object ($r(21) = .533, p = .009$). We found a similar relationship with the likelihood of participants choosing the N3 location and lapsing to the N3 non-target object ($r(21) = .444, p = .034$), but no significant correlation between the likelihood of participants choosing the N1 location and lapsing to the object at that location ($r(21) = -.353, p = .098$).

Shift Trials

Experiments 1, 3, and 4 included a mix of hold (single-cue) trials and shift (double-cue) trials. While the hold trials were the focus of this study to examine spatial lapses of attention, we can also explore whether participants made any differential patterns of errors in the shift trials in the different experimental contexts. In terms of overall performance, in all three experiments with shift trials, participants were able to shift their attention from the initially cued location to the second-cue location to successfully perform the task. In Experiment 1 participants had a .822 probability of reporting the correct color, orientation, and spatial location of the object at the second cue (p_{TcToTL} 95% HDI = [.802, .836]). In Experiment 3, p_{TcToTL} was .896 (95% HDI

$=[.873, .898]$), and in Experiment 4, p_{TC}^{TOTL} was .881 (95% HDI = [.873, .898]). It is possible that the higher probability of reporting the correct color, orientation, and location of the object at the second spatial cue in Experiments 3 & 4 compared to Experiment 1 was driven by the predictability of the second cue in the latter experiments.

In terms of swap errors on shift trials, in Experiment 1, similar to Dowd & Golomb (2019), participants sometimes made triple-bound swap errors reporting the color, orientation, and spatial location of the object at the initially cued location ($p_{N1cN1oN1L} = .055$, 95% HDI = [.045, .062]), while almost never reporting any other triple bound non-target objects ($p_{N2cN2oN2L} = .000$, 95% HDI = [.000, .001]; $p_{N3cN3oN3L} = .000$, 95% HDI = [.000, .001]).

A t-test confirmed that participants in Experiment 1 were more likely to report the features and location of the object at the initially cued location (N1) than the control N2 ($t(22) = 3.232$, $p = .004$, $d = .674$; Figure 5a) location. On shift trials, triple-bound swaps to the initial N1 location likely don't reflect lapses in spatial attention, but instead probably reflect trials where attention had not yet updated to the second-cue location at the time the stimuli were presented (Dowd & Golomb, 2019). Thus, if participants were better able to anticipate the location of the shift in Experiments 3-4, then we might also expect there to be fewer swap errors to the item at the originally cued location in those experiments.

In Experiment 3, the probability of reporting the features and location of the object at the initially cued location ($p_{N1cN1oN1L}$) from the standard triple model was .017 (95% HDI = [.013, .024]), while the probabilities of reporting the triple bound N2 or N3 objects were again very low ($p_{N2cN2oN2L} = .001$, 95% HDI = [.000, .002]; $p_{N3cN3oN3L} = .000$, 95% HDI = [.000, .002];

Figure 5b). In Experiment 4, $p_{N1cN1oN1L}$ was .032 (95% HDI = [.025, .038]), while $p_{N2cN2oN2L}$ was .001 (95% HDI = [.000, .002]) and $p_{N3cN3oN3L}$ was .000 (95% HDI = [.000, .001]; Figure 5c). Despite the non-overlapping HDIs, the simple triple model results revealed a non-significant difference between the probabilities of reporting the triple bound N1 and N2 in Experiment 3 ($t(23) = 1.993$, $p = 0.058$, $d = .407$). Upon further exploration this may have been driven by variability introduced by one participant, whose maximum posterior estimate for $p_{N1cN1oN1L}$ on shift trials was 4.5 standard deviations above the mean. When this subject's data were excluded, there was a significant difference in the probability of swaps to N1 vs N2 ($t(22) = 3.467$, $p = .002$, $d = .723$). In Experiment 4, participants also made significantly more swap errors to the initially cued (N1) location than the control N2 location ($t(22) = 3.384$, $p = .003$, $d = .706$).

To test whether participants made fewer triple swaps to the initially cued location on shift trials in the spatially predictable Experiments 3 and 4 compared to Experiment 1, we compared the swap rates from the simple model across experiments, conducting a 2x3 ANOVA on the probability of reporting the non-target object at the initially cued location ($p_{N1cN1oN1L}$) vs control location ($p_{N2cN2oN2L}$) across Experiments 1, 3, and 4. With all participants included there was a significant main effect of which non-target was reported ($F(1,67) = 23.572$, $p < .001$, $\eta^2_p = .26$), but no significant main effect of experiment ($F(2,67) = 1.506$, $p = .229$, $\eta^2_p = .043$) nor a significant interaction ($F(2,67) = 1.514$, $p = .227$, $\eta^2_p = .043$). When the outlier participant from Experiment 3 was removed, however, there was a significant main effect of experiment ($F(2,66) = 3.437$, $p = .038$, $\eta^2_p = .094$) and a significant interaction ($F(2,66) = 3.498$, $p = .036$, $\eta^2_p = 0.096$), in addition to the significant main effect of which non-target was reported ($F(1,66) =$

25.173, $p < .001$, $\eta^2_p = .276$). Thus, on shift trials in all three experiments, participants were more likely to make triple-bound swap errors reporting the item at the initially cued location than the control location, but when the shift location was spatially predictable and could be anticipated, these errors may have been reduced. Thus, shift predictability may have influenced how effectively participants were able to shift attention from the first cue to the second cue on shift trials, in addition to influencing the likelihood and distribution of lapses of spatial attention on hold trials.

Discussion

In this series of experiments, we set out to explore a recently reported phenomenon: lapses in spatial attention, a type of attentional error that occurs when spatial attention remains in a focused state, but is temporarily focused on a task irrelevant location (Dowd & Golomb, 2019). We defined spatial lapses in our data as trials in which participants reported the color, orientation, and spatial location of a single object in the display (i.e., a triple-bound correlated report indicating focused attention), but the reported features belonged to a non-target object, specifically on hold trials where only a single target location was cued and task-relevant. In the current study, we investigated why these spatial lapses occur by testing: a) if lapses in spatial attention are driven by the anticipation of having to make attentional shifts; b) if participants lapse preferentially to a predictable shift location; and c) if participants have explicit knowledge of likely shift locations that influence where they lapse.

Lapses in spatial attention are related to implicit dynamic spatial expectations

To test whether lapses are driven by the anticipation of having to make attentional shifts, we first replicated Dowd & Golomb (2019)'s findings, confirming that participants made spatial lapses to non-target objects on hold trials, in an experiment where hold and shift trials were intermixed (Experiment 1). We then compared these results to Experiment 2, where participants never encountered any shift trials during the experiment. In this latter experimental context, participants almost never lapsed to a non-target object. These data provide strong evidence that these lapses are related to the expectation of an attentional shift. If participants were reporting a triple bound non-target object for any reason unrelated to the anticipation of the second cue, then participants would have made a similar amount of spatial lapses in Experiment 2 as in Experiment 1. For example, an alternate explanation for these errors discussed in Dowd & Golomb (2019) is that perhaps participants never actually saw the cue on those trials, and instead chose to attend a random spatial location. As Dowd & Golomb (2019) discussed, this explanation seemed unlikely because the spatial cue was presented for 250ms, and in general participants were able to attend to the correct location on the majority of trials across all experiments. The current study definitively rules out this account, since participants did not make spatial lapses in Experiment 2 (and in Experiments 3 & 4, the lapse distributions were not random). The current study also rules out generic influences such as spatial priming (Maljkovic & Nakayama, 1996) or serial dependence (for review see Kiyonaga et al. 2017), where the location of a previous trial's target may bias perception and attention on subsequent trials, since those effects would also be expected to produce similar patterns across all four experiments. Thus, while other attentional phenomena can guide spatial attention and result in different types of errors, the particular pattern of errors participants made here seems reflective specifically of lapses of spatial attention driven by an anticipatory attentional shift.

In Experiments 3 and 4, we further probed the anticipatory nature of these lapses, asking if they were also sensitive to learned spatial probabilities. To test this, we implemented task contexts where the location of the second cue (if it appeared) was 100% predictable. We found that participants indeed lapsed preferentially to the adjacent location where the second cue was more likely to appear. In Experiment 3 the second cue, when it occurred, was always clockwise to the first cue, and participants lapsed to the clockwise non-target more than the counter-clockwise control non-target. Conversely in Experiment 4, where attentional shifts were always counter-clockwise to the first cue, participants lapsed to the counter-clockwise non-target more often than the clockwise control. Our findings demonstrate that spatial lapses are indeed sensitive to dynamic attentional expectations about the likelihood and location of a future event, with our exploratory timecourse analyses further supporting a pattern consistent with learned expectations. Whether these lapses are sensitive to more nuanced task context – for instance, parametric variations of the proportion of shift trials compared to hold trials or more probabilistic manipulations of the second cue’s likely location, or manipulating the timing of the stimulus array relative to the second cue – remains to be seen, but the current study establishes important boundary conditions on this effect.

Was this biased distribution of lapses driven by explicit knowledge of the spatial predictability? Experiment 4 included a post-experiment task to determine if participants had learned explicit knowledge of our manipulation. Overall, the majority of participants reported that they did not notice that the location of the second cue was predictable, and when forced to guess, most participants did not correctly guess the location of the second cue. We did find some correlation

where participants who were more likely to guess the correct location tended to have more biased lapse distributions, but interestingly, this was not associated with a greater likelihood to report explicit knowledge of the manipulation in the initial question; if anything, the correlation was weaker in the participants who reported explicitly noticing the manipulation. As reported in the Results, neither group of participants (those who reported noticing the manipulation or those who did not) reliably chose the correct N2 counter-clockwise location as the most likely shift location. However, additional analyses showed that the distribution of responses among the three potential shift locations was not completely random. Among participants who stated they noticed our manipulation, they were, on average, actually most likely to incorrectly report the clockwise N1 location as the likely shift location (46.43%). Among participants who answered “no” to the first question of our explicit knowledge task, they tended to report the diagonal N3 most often (43.75%). It’s unclear whether these results reflect a systematic guessing strategy, a misunderstanding of the post-test task instructions on the part of one or two subjects in the first group, and/or some other tendencies to prefer clockwise or diagonal responses. However, even with this in mind, the post-test results are not consistent with explicit awareness of our manipulation. Thus, our results do not indicate that explicit knowledge of the shift predictability was driving spatial lapses. In other words, spatial lapses do actually seem to be *lapses* where attention erroneously highlights the wrong location, rather than reflecting an explicit strategy to predict the cue.

One puzzling finding was that across our experiments, lapses to the diagonal N3 location were unexpectedly high. The experiments were designed such that the clockwise N1 and counter-clockwise N2 locations were the critical non-target locations and well-matched controls for each

other. When directly comparing these two locations, spatial lapse errors followed a clear and predictable pattern: lapse rates were equivalent for N1 and N2 in Experiment 1 (when both locations were equally probably potential shift targets), significantly greater for N1 than N2 in Experiment 3 (when N1 was the likely shift target), and significantly greater for N2 than N1 in Experiment 4 (when N2 was the likely shift target). But in Experiments 1, 3, and 4, participants also lapsed to the non-adjacent N3 object directly opposite the target, even though the N3 was never a potential shift target. In fact, lapses to the N3 location were consistently higher than would be expected for a baseline location (in Dowd & Golomb, 2019, this pattern was found as well, though was less accentuated). Because the N3 location was never the location of a second cue, and because lapse rates to the N3 location were unexpectedly high regardless of whether the second cue was or was not spatially predictable, it may be that these N3 errors reflect a different attentional mechanism unrelated to a spatial anticipation of where the second cue would appear. For instance, participants may have strategically or implicitly shifted their attention to sample the location opposite the first cue (the furthest location) in order to “cover” the whole display and ensure they don’t miss the second cue. Interestingly, since these errors were not observed in Experiment 2, they may still be related to the *temporal* expectation of having to shift attention.

It is also intriguing that in the post-experiment explicit report test of Experiment 4, participants – especially those who reported being unaware of the manipulation – were more likely to guess that the N3 location was the most likely shift location, and the likelihood of selecting N3 in the post-test was correlated with lapse rates to the N3 location in the main task. Thus, there may be something unique about the diagonal N3 location that captures attention during attentional lapses. One might speculate potential reasons related to inhibition of return (Posner & Cohen,

1984, for review see: Wang & Klein, 2010) or hemispheric division of attentional resources (Alvarez & Cavanagh, 2005), since the N3 location was furthest away and more likely to be in the opposite hemisphere than the N1 and N2 locations, but our set of experiments is not suited to explore this further. Regardless, it seems that this N3 effect is largely independent of the main finding.

How do lapses of spatial attention relate to rhythmic attentional sampling?

A recent focus in the attention literature has been the idea that attention is subject to intrinsic, rhythmic fluctuations (Fiebelkorn et al., 2013; Fiebelkorn & Kastner, 2019; Landau & Fries, 2012; Re et al., 2019; R. VanRullen et al., 2007; VanRullen, 2016). Of particular relevance to the current study is the theory that attention oscillates between two states: (1) a focused “sampling” state where sensory processing is enhanced at the behaviorally relevant location and attentional shifts are suppressed, and (2) an exploratory “shifting” state where perceptual sensitivity at the behaviorally relevant location is diminished and attentional shifts to other locations are more likely to occur (Fiebelkorn & Kastner 2019; VanRullen, 2018).

One interpretation of the spatial lapses of attention we report here is that perhaps they may just be a consequence of the stimulus appearing during a “shifting” oscillatory state. Critically, while the neural oscillations underlying these alternating states may be intrinsic, the rhythmic attention theory poses that the exploratory attentional periods represent merely “windows of opportunity” where shifts of attention to other locations are more likely to, but do not necessarily, occur (Fiebelkorn & Kastner, 2019). A critical finding of the current paper is that spatial lapses are sensitive to implicit expectations about a future task-relevant location. Thus, our data may

provide further evidence that the “windows” for exploratory sampling are regulated by task expectations (e.g. Gaillard et al., 2020). Other work has found that probability manipulations can effectively guide implicit attention (Geng & Behrmann, 2005; Jiang et al., 2013) as well as patterns of eye movements (Jiang et al., 2014), and here we demonstrate that implicit expectations about the likelihood and/or location of anticipated goal-directed shifts of attention can modulate spatial lapses of attention, resulting in errors of object feature perception. If spatial lapses are linked to oscillatory fluctuations, this raises intriguing questions about whether task context and anticipatory expectations interact with intrinsic rhythms, underlying salience maps, or both.

Interestingly, a recent trend in the rhythmic oscillation literature describes alternations between two states: in the case of attention, periods of focused sampling versus unfocused or exploratory periods (e.g. Fiebelkorn & Kastner 2019; VanRullen 2018). However, the distinction raised here and in Dowd & Golomb (2019) between lapses of spatial attention and lapses of sustained attention suggests that attention may be better described as having three states: focused at the behaviorally relevant location, unfocused (lapse of sustained attention), and focused at a different location (lapse of spatial attention). Although there has been some evidence that attention may sample within objects and between objects at different frequencies (Fiebelkorn et al., 2013; Landau & Fries, 2012), or that different cyclic rhythms underlie different perceptual functions (VanRullen 2016), previous studies linking behavioral measures with oscillatory attentional states have primarily used sensitivity measures (d-prime or reaction time), which may be limited to showing that sensitivity rhythmically oscillates between periods of enhanced and diminished sensitivity at the behaviorally relevant location(s) (Fiebelkorn et al., 2013; Landau & Fries,

2012). A benefit of the experimental paradigm used in Dowd & Golomb (2019) and the current paper is that it goes beyond sensitivity measures, such that different types of feature-binding errors can be used to characterize whether one is focusing attention at the target location (triple-bound correlated target reports), experiencing an unfocused lapse of sustained attention (random guess reports), experiencing a focused lapse of spatial attention (triple-bound correlated nontarget reports), or even simultaneously dividing attention between two locations (unbound or illusory conjunction errors). A downside is that we can only probe the attentional state at the time of stimulus presentation (here one presentation per trial, at a relatively fixed point in time), but this paradigm may hold promise for future studies exploring the temporal dynamics of sustained lapses versus spatial lapses of attention and how these relate to rhythmic sampling.

An alternative interpretation is that spatial lapses reflect a separate attentional sampling process independent of ongoing rhythmic sampling. Attention has been shown to be sensitive to temporal expectations (Doherty et al., 2005; Rohenkohl et al., 2014), so perhaps participants' expectation of having to make an attentional shift at a particular point in the trial led to a temporally specific, single-event erroneous allocation of attention. In our data, triple-bound nontarget responses, which we consider to be indicative of spatial lapses, only occurred on an estimated 3-7% of trials. Because our task is only sensitive to spatial lapses happening during that critical point in the trial when the stimulus array was presented, it is unclear whether our spatial lapses are part of a pattern of routine, rhythmic sampling that occurs throughout the trial but is modulated by spatial and temporal expectations, or whether spatial lapses are a more specific type of misallocation of attention. Investigating the links between spatial lapses and intrinsic attentional rhythms – and their perceptual consequences – may be a fruitful direction for future research.

Regardless of the exact mechanism, it seems clear that our results do in fact reflect anticipatory sampling of an incorrect location, rather than an alternative explanation due to anticipatory spreading or dividing of attention to a nearby object. Considering the second cue was always in a location adjacent to the first cue on all shift trials, it is possible that participants could have increased their spatial window of attention in anticipation of a potential shift to encompass both the current and likely future loci of attention; similar to Shioiri, Honjo, Kashwase, Matsumiya, & Kuriki (2016). The likelihood of reporting a non-target object's features increases with closer spatial proximity to the target (Emrich & Ferber 2012), which could be related to attentional spread. However, if attention was spreading to a nearby object, and it was driven by the expectation of having to dynamically shift attention, we would have also expected to see lower precision, measured by standard deviation, on feature reports on hold trials in Experiments 1,3, and 4, versus Experiment 2 (Bays, Wu, & Husain, 2011). It also could be argued that participants may have attended to multiple discrete locations simultaneously in anticipation of a shift cue (i.e. divided attention). However, splitting attention between two locations has been shown to result in increased illusory conjunctions, unbound errors, and other feature mixing errors (Dowd & Golomb 2019; Golomb, L'Heureux, & Kanwisher 2014), which we did not observe here. Because participants reported all three features of a non-target object with high precision, it seems likely that attention was highlighting only one location, even though that location was task irrelevant. To harken back to our original analogy of a slipping flashlight, spatial lapses of attention can be thought of as points in time where the metaphorical flashlight slips from the task-relevant target location to illuminate a non-target location, without dimming or changing its

aperture. That said, while the slipping of the flashlight is an error, it is not simply a random slip, but is more likely to slip under certain circumstances and to certain context-specific locations.

Potential benefits of lapses in spatial attention

The finding that spatial lapses occurred in Experiments 1, 3, and 4 but not in Experiment 2 suggests that lapses of spatial attention may be an adaptive error dependent on task context. Is there a behavioral benefit to these lapses, or are they just a type of error that is more common in certain contexts?

Some insight may come from performance on the shift (double-cue) trials. As observed by Dowd & Golomb (2019), on shift trials participants sometimes apparently failed to shift attention from the first cue to second cue before the stimulus array appeared, and as a result reported the triple bound object at the initially cued N1 location instead of the target object. These errors on shift trials were highly selective for the N1 nontarget, suggesting that they are reflective of attention remaining at the original location, rather than a spatial lapse of attention sampling a location away from the target. Tellingly, participants made these errors more in Experiment 1, when the location of the second cue was spatially unpredictable, compared to Experiments 3 and 4. Moreover, spatial predictability seemed to influence these errors even though the manipulation did not reach explicit awareness. Previous work has shown that implicitly learned spatial probabilities can effectively guide attention (Geng & Behrmann, 2005; Jiang et al., 2013), and bias eye movements towards the “rich” location (Jiang, Won, & Swallow, 2014), so it is reasonable that a predictable shift direction might help participants shift attention to a new location faster as well.

So, while dynamic task expectations may have resulted in erroneous deployment of attention when an attentional shift was not required (spatial lapses), it is possible that anticipatory sampling may have actually *facilitated* attentional shifts when they were required, particularly if shift direction was predictable. In some ways this is analogous to how an automated response is usually adaptive, but during a lapse in sustained attention it be detrimental if the response was actually supposed to have been withheld. For instance, (Shalgi, O'Connell, Douell, & Robertson, 2007) found when a go/nogo task is made to be less attentionally demanding (and more prone to lapses of sustained attention compared to when the task requires a high degree of vigilance), accuracy improves, but error awareness is lower. In the experiments in this study we did not find a correlation between increased lapses on hold trials and improved performance on shift trials at a between-subjects level, but future work may be better suited to investigate potential trade offs in behavior under variable expectations.

Further insight into the benefits, consequences, and mechanisms of spatial lapses may come from a better understanding of how lapses of spatial attention are related to other types of cognitive errors, such as lapses in sustained attention and lapses in working memory. Generally, lapses in sustained attention are associated with “zoning-out” or “mind-wandering” and behaviorally can be measured as periods of more variable reaction times (deBettencourt et al., 2018; Esterman et al., 2013). These variable reaction times have been associated with distinct attentional states, which have been described as stable (“in the zone”) or erratic (“out of the zone”), and are associated with different neural signatures (Esterman et al., 2013). Lapses in sustained attention have also been linked to lapses in working memory, which are characterized as instances when

participants perform under capacity during a working memory task, often without awareness of poor performance (Adams & Vogel, 2017), and which may in fact be the result of fluctuations from a shared cognitive resource (deBettencourt, Keene, Awh, & Vogel, 2019). Whether lapses of spatial attention may also be linked to fluctuations of this shared resource, or are driven by an independent mechanism, may help reveal the role of spatial lapses in visual processing. In an exploratory analysis of our data combining Experiments 1,3, and 4, we found that our measure of lapses in spatial attention (sum of the N1N1N1, N2N2N2, and N3N3N3 estimates from the simple triple model for each subject) was not significantly correlated with the random guessing rate (sum of all UU estimates), $r(69) = .105$, $p = .387$, suggesting that at least in this context, lapses in spatial attention seem to occur independently of lapses of sustained attention and other types of errors which can lead to random guessing.

In general, because spatial lapses of attention are a type of error that has thus far received less attention in the literature, the consequences of lapses in spatial attention in everyday life remain to be investigated. One can imagine that during a task which requires frequent attentional shifts, such as driving, attention momentarily highlighting an incorrect location could have potentially dire consequences. It is also worth considering whether certain clinical or developmental populations may be more or less susceptible to lapses in spatial attention (for example, lapses in sustained attention have been linked with Attention-Deficit/Hyperactivity Disorder: Van den Driessche et al., 2017).

Conclusion

In this series of experiments, we have more fully characterized lapses in spatial attention, a novel attentional phenomenon first identified by Dowd & Golomb (2019). Lapses in spatial attention are instances when attention highlights an incorrect spatial location, during trials where spatial attention was clearly cued to a single target location. Lapses of spatial attention result in participants reporting a fully bound (i.e., color, orientation, and location) non-target object, whereas other attentional errors, e.g. lapses of sustained attention, would result in unbound guessing of object features. Here we have shown that these spatial lapses are specifically driven by the expectation of having to make a future attentional shift, and are sensitive to task context including shift likelihood and spatial regularities. Moreover, we observed that these lapses do not seem to be driven by explicit knowledge, supporting the idea that these lapses indeed reflect errors of attentional control.

Open Practices Statement

The current study was designed to closely follow methods reported in Dowd & Golomb (2019).

Although Experiments 1-3 were not formally preregistered, the experimental design, participant inclusion criteria, and analyses follow those described in the previous paper as closely as possible, except where noted. These first three experiments were conducted in parallel, with participants randomly assigned among them. Experiment 4 was conducted after analyzing the first three experiments, and was pre-registered at <https://osf.io/mxq9j>. Materials are available on OSF.

References

Adam, K. C., & Vogel, E. K. (2017). Confident failures: Lapses of working memory reveal a metacognitive blind spot. *Attention, Perception, & Psychophysics*, 79(5), 1506-1523.

Alvarez, G. A., & Cavanagh, P. (2005). Independent Resources for Attentional Tracking in the Left and Right Visual Hemifields. *Psychological Science*, 16(8), 637–643. <https://doi.org/10.1111/j.1467-9280.2005.01587.x>

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), 7–7. <https://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

Bays, Paul M. (2016). Evaluating and excluding swap errors in analogue tests of working memory. *Scientific Reports*, 6(1), 19203. <https://doi.org/10.1038/srep19203>

Brainard, D. H. (1997) The Psychophysics Toolbox, *Spatial Vision* 10:433-436.

Center, O. S. (1987). *Ohio Supercomputer Center*. <http://osc.edu/ark:/19495/f5s1ph73>

Cheyne, J. A., Carriere, J. S., & Smilek, D. (2006). Absent-mindedness: Lapses of conscious awareness and everyday cognitive failures. *Consciousness and cognition*, 15(3), 578-592.

deBettencourt, M. T., Keene, P. A., Awh, E., & Vogel, E. K. (2019). Real-time triggering reveals concurrent lapses of attention and working memory. *Nature Human Behaviour*, 3(8), 808–816. <https://doi.org/10.1038/s41562-019-0606-6>

deBettencourt, M. T., Norman, K. A., & Turk-Browne, N. B. (2018). Forgetting from lapses of sustained attention. *Psychonomic Bulletin & Review*, 25(2), 605–611. <https://doi.org/10.3758/s13423-017-1309-5>

Doherty, J. R., Rao, A., Mesulam, M.M., & Nobre, A.C. (2005). Synergistic Effect of Combined Temporal and Spatial Expectations on Visual Attention. *Journal of Neuroscience*, 25(36), 8259–8266. <https://doi.org/10.1523/JNEUROSCI.1821-05.2005>

Dowd, E. W., & Golomb, J. D. (2019). Object-Feature Binding Survives Dynamic Shifts of Spatial Attention. *Psychological Science*, 30(3), 343–361. <https://doi.org/10.1177/0956797618818481>

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

Esterman, M., Noonan, S. K., Rosenberg, M., & DeGutis, J. (2013). In the Zone or Zoning Out? Tracking Behavioral and Neural Fluctuations During Sustained Attention. *Cerebral Cortex*, 23(11), 2712–2723. <https://doi.org/10.1093/cercor/bhs261>

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>

Fiebelkorn, I. C., & Kastner, S. (2019). A Rhythmic Theory of Attention. *Trends in Cognitive Sciences*, 23(2), 87–101. <https://doi.org/10.1016/j.tics.2018.11.009>

Gaillard, C., Hassen, S. B. H., Di Bello, F., Bihan-Poudec, Y., VanRullen, R., & Hamed, S. B. (2020). Prefrontal attentional saccades explore space rhythmically. *Nature Communications*, 11(1), 1-13.

Geng, J. J., & Behrmann, M. (2005). Spatial probability as an attentional cue in visual search. *Perception & Psychophysics*, 67(7), 1252–1268. <https://doi.org/10.3758/BF03193557>

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. <https://doi.org/10.1177/0956797614522068>

JASP, Team. *JASP (Version 0.11.1)[Computer Software]*. 2019, <https://jasp-stats.org/>.

Jiang, Y. V., Swallow, K. M., & Rosenbaum, G. M. (2013). Guidance of spatial attention by incidental learning and endogenous cuing. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 285–297. <https://doi.org/10.1037/a0028022>

Jiang, Y. V., Won, B.-Y., & Swallow, K. M. (2014). First saccadic eye movement reveals persistent attentional guidance by implicit learning. *Journal of Experimental Psychology: Human Perception and Performance*, 40(3), 1161–1173. <https://doi.org/10.1037/a0035961>

Kiyonaga, A., Scimeca, J. M., Bliss, D. P., & Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences*, 21(7), 493–497. <https://doi.org/10.1016/j.tics.2017.04.011>

Kleiner M, Brainard D, Pelli D, 2007, “What's new in Psychtoolbox-3?” Perception 36 ECPV Abstract Supplement.

Landau, A. N., & Fries, P. (2012). Attention Samples Stimuli Rhythmically. *Current Biology*, 22(11), 1000–1004. <https://doi.org/10.1016/j.cub.2012.03.054>

Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics*, 58(7), 977–991. <https://doi.org/10.3758/BF03206826>

Posner, M., & Cohen, Y. (1984). Components of Visual Orienting. *Attention and Performance X: Control of Language Processes*, 32, 531–556.

Re, D., Inbar, M., Richter, C. G., & Landau, A. N. (2019). Feature-Based Attention Samples Stimuli Rhythmically. *Current Biology*, 29(4), 693-699.e4. <https://doi.org/10.1016/j.cub.2019.01.010>

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.

Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). 'Oops!': Performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, 35(6), 747–758. [https://doi.org/10.1016/S0028-3932\(97\)00015-8](https://doi.org/10.1016/S0028-3932(97)00015-8)

Roca, J., Lupiáñez, J., López-Ramón, M. F., & Castro, C. (2013). Are drivers' attentional lapses associated with the functioning of the neurocognitive attentional networks and with cognitive failure in everyday life?. *Transportation research part F: traffic psychology and behaviour*, 17, 98-113

Rosenberg, M.D., Finn, E. S., Scheinost, D., Constable, R. T., & Chun, M. M. (2017). Characterizing Attention with Predictive Network Models. *Trends in Cognitive Sciences*, 21(4), 290–302. <https://doi.org/10.1016/j.tics.2017.01.011>

Rosenberg, Monica D., Finn, E. S., Constable, R. T., & Chun, M. M. (2015). Predicting moment-to-moment attentional state. *NeuroImage*, 114, 249–256. <https://doi.org/10.1016/j.neuroimage.2015.03.032>

Rohenkohl, G., Gould, I. C., Pessoa, J., & Nobre, A. C. (2014). Combining spatial and temporal expectations to improve visual perception. *Journal of Vision*, 14(4), 8–8. <https://doi.org/10.1167/14.4.8>

Shalgi, S., O'Connell, R. G., Deouell, L. Y., & Robertson, I. H. (2007). Absent minded but accurate: Delaying responses increases accuracy but decreases error awareness. *Experimental Brain Research*, 182(1), 119–124. <https://doi.org/10.1007/s00221-007-1054-5>

Shioiri, S., Honjyo, H., Kashiwase, Y., Matsumiya, K., & Kuriki, I. (2016). Visual attention spreads broadly but selects information locally. *Scientific Reports*, 6(1), 35513. <https://doi.org/10.1038/srep35513>

Smallwood, J., Davies, J. B., Heim, D., Finnigan, F., Sudberry, M., O'Connor, R., & Obonsawin, M. (2004). Subjective experience and the attentional lapse: Task engagement and disengagement during sustained attention. *Consciousness and cognition*, 13(4), 657-690.

Steinman, R. M., Haddad, G. M., Skavenski, A. A., & Wyman, D. (1973). Miniature eye movement. *Science*, 181(4102), 810–819.

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, 1–8. doi:10.1167/13.10.9.

Thut, G., Schyns, P. G., & Gross, J. (2011). Entrainment of Perceptually Relevant Brain Oscillations by Non-Invasive Rhythmic Stimulation of the Human Brain. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00170>

Unsworth, N., Redick, T. S., Lakey, C. E., & Young, D. L. (2010). Lapses in sustained attention and their relation to executive control and fluid abilities: An individual differences investigation. *Intelligence*, 38(1), 111–122. <https://doi.org/10.1016/j.intell.2009.08.002>

Van den Driessche, C., Bastian, M., Peyre, H., Stordeur, C., Acquaviva, É., Bahadori, S., Delorme, R., & Sackur, J. (2017). Attentional Lapses in Attention-Deficit/Hyperactivity Disorder: Blank Rather Than Wandering Thoughts. *Psychological Science*, 28(10), 1375–1386. <https://doi.org/10.1177/0956797617708234>

VanRullen, R., Carlson, T., & Cavanagh, P. (2007). The blinking spotlight of attention. *Proceedings of the National Academy of Sciences*, 104(49), 19204–19209. <https://doi.org/10.1073/pnas.0707316104>

VanRullen, R. (2018). Attention Cycles. *Neuron*, 99(4), 632–634. <https://doi.org/10.1016/j.neuron.2018.08.006>

Wang, Z., & Klein, R. M. (2010). Searching for inhibition of return in visual search: A review. *Vision Research*, 50(2), 220–228. <https://doi.org/10.1016/j.visres.2009.11.013>

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




Fig 1. a) Example hold and shift trial sequences for Experiments 1-4. On hold trials participants were instructed to covertly attend a single location, and on shift trials participants had to shift their attention from the initially cued location to a newly cued spatial location. Participants were instructed to report the color, orientation, and spatial location of the object at the most recently cued location. b) We conducted 4 Experiments. In Experiment 1 the second cue, when it occurred, could be clockwise or counter-clockwise to the first cue. In Experiment 2, participants only completed hold trials. In Experiment 3 the second was always clockwise to the first cue, and in Experiment 4 it was always counter-clockwise. c) Experiment 4 also included a post-experiment task. Participants were asked to guess the location of the second cue (one of the white squares), given the location of the first cue (the black square).




Fig 2: Visualizations of color-orientation-location reports in joint-feature space, plotted as error relative to actual target feature values: color responses are shown along the x-axis, orientation responses are shown along the y-axis, and location responses are indicated by dot color. Each dot represents the color-orientation-location response for a single trial, aggregating across subjects. For visualization purposes, we have discretized the continuous location responses into four bins (defined below), and have flattened joint-feature space; all feature dimensions were in fact circular and continuous, such that $+180^\circ$ is identical to -180° in feature space. **a-d:** Scatterplots plot trial-by-trial error distributions separately for Hold and Shift trials, for each experiment. For Hold Trials, location errors in the range $[-45^\circ, 45^\circ]$ were coded as target (T; yellow) location reports, location errors in the range $[45^\circ, 135^\circ]$ were coded as clockwise non-target (N1; red) location reports, location errors in the range $[-135^\circ, -45^\circ]$ were coded as counter-clockwise non-target (N2; blue) location reports, and location errors in the range $[-180^\circ, -135^\circ]$ or $[135^\circ, 180^\circ]$ were coded as diagonal non-target (N3; gray) location reports. For Shift Trials, we used a similar convention, but as described in the text, errors were aligned differently, such that N1 reports reflect the initially cued nontarget location (green) and N2 the adjacent control nontarget location (purple). **e:** Color-coded cartoons showing hypothetical response distributions for triple-bound (correlated) object reports. Central yellow clusters reflect $T_C T_O T_L$ (correlated target) responses, and diagonal clusters reflect correlated swap errors indicating lapses of spatial attention (Hold trials) and failures to shift attention (Shift trials).




Fig. 3. Violin plots on the left illustrate the posterior sample distributions from hold trials of each parameter from the standard triple model representing triple bound swaps over 15,000 post-convergence samples. Error bars on violin plots represent 95% highest density intervals. The standard triple model was fitted on data collapsed across subjects. The bar plots on the right show the average maximum likelihood estimates on parameters from the simple triple model, which was fitted on individual participants. Error bars on bar plots represent the standard error of the mean (SEM). a) Experiment 1, b) Experiment 2, c) Experiment 3, d) Experiment 4. T = target, N1 = clockwise non-target, N2 = counter-clockwise non-target and N3 = diagonal non-target.




Fig. 4. Exploratory timecourse analysis for Experiment 4 Hold trials, illustrating the proportion of spatial lapses to different locations plotted by trial block. The simple triple model was fit on aggregated data across subjects, separately for each block of trials (8 blocks total). a) Each dot indicates the relative proportion of spatial lapses to each non-target location on a given block. Proportions were calculated by dividing the maximum likelihood estimate of a spatial lapse to a given location (e.g., $pN1_cN1_oN1_L$) by the sum of spatial lapses to all locations ($pN1_cN1_oN1_L + pN2_cN2_oN2_L + pN3_cN3_oN3_L$). b) Sum probability of spatial lapses (regardless of location) across blocks. Sum probability calculated as ($pN1_cN1_oN1_L + pN2_cN2_oN2_L + pN3_cN3_oN3_L$) on that block.




Fig. 5. Violin plots on the left illustrate the posterior sample distributions from shift trials of each parameter from the standard triple model representing triple bound swaps over 15,000 post-convergence samples. Error bars on violin plots represent 95% highest density intervals. The standard triple model was fitted on data collapsed across subjects. The bar plots on the right show the average maximum likelihood estimates on parameters from the simple triple model, which was fitted on individual participants. Error bars on bar plots represent the standard error of the mean (SEM). a) Experiment 1, b) Experiment 3, c) Experiment 4. T = target, N1 = initially cued location, N2 = adjacent control location and N3 = diagonal non-target.

Table 1. Response Combinations of Color, Orientation, and Location from the Triple Joint Mixture Models

Response types grouped by location response	$m_{standard}$	m_{simple}	Response combination	Joint probability density
Target Location				
Correlated target (triple bound)	1	1	T _C T _O T _L	$\Phi_{0,\kappa_C} \Phi_{0,\kappa_O} \Phi_{0,\kappa_L}$
Correlated N1	2	2	N _{1C} N _{1O} T _L	$\Phi_{\frac{\pi}{2},\kappa_C} \Phi_{\frac{\pi}{4},\kappa_O} \Phi_{0,\kappa_L}$
Correlated N2	3	3	N _{2C} N _{2O} T _L	$\Phi_{-\frac{\pi}{2},\kappa_C} \Phi_{-\frac{\pi}{4},\kappa_O} \Phi_{0,\kappa_L}$
Correlated N3	4	4	N _{3C} N _{3O} T _L	$\Phi_{\pi,\kappa_C} \Phi_{\frac{\pi}{2},\kappa_O} \Phi_{0,\kappa_L}$
Other	5	5	U _C U _O T _L	$\gamma_C \gamma_O \Phi_{0,\kappa_L}$
Nontarget N1 location				
Correlated target	6	6*	T _C T _O N _{1L}	$\Phi_{0,\kappa_C} \Phi_{0,\kappa_O} \Phi_{\frac{\pi}{2},\kappa_L}$
Correlated N1 (triple bound)	7	7	N _{1C} N _{1O} N _{1L}	$\Phi_{\frac{\pi}{2},\kappa_C} \Phi_{\frac{\pi}{4},\kappa_O} \Phi_{\frac{\pi}{2},\kappa_L}$
Correlated N2	8	8*	N _{2C} N _{2O} N _{1L}	$\Phi_{-\frac{\pi}{2},\kappa_C} \Phi_{-\frac{\pi}{4},\kappa_O} \Phi_{\frac{\pi}{2},\kappa_L}$
Correlated N3	9	8*	N _{3C} N _{3O} N _{1L}	$\Phi_{\pi,\kappa_C} \Phi_{\frac{\pi}{2},\kappa_O} \Phi_{\frac{\pi}{2},\kappa_L}$
Other	10	9*	U _C U _O N _{1L}	$\gamma_C \gamma_O \Phi_{\frac{\pi}{2},\kappa_L}$
Nontarget N2 location				
Correlated target	11	6*	T _C T _O N _{2L}	$\Phi_{0,\kappa_C} \Phi_{0,\kappa_O} \Phi_{-\frac{\pi}{2},\kappa_L}$
Correlated N1	12	10*	N _{1C} N _{1O} N _{2L}	$\Phi_{\frac{\pi}{2},\kappa_C} \Phi_{\frac{\pi}{4},\kappa_O} \Phi_{-\frac{\pi}{2},\kappa_L}$
Correlated N2 (triple bound)	13	11	N _{2C} N _{2O} N _{2L}	$\Phi_{-\frac{\pi}{2},\kappa_C} \Phi_{-\frac{\pi}{4},\kappa_O} \Phi_{-\frac{\pi}{2},\kappa_L}$
Correlated N3	14	10*	N _{3C} N _{3O} N _{2L}	$\Phi_{\pi,\kappa_C} \Phi_{\frac{\pi}{2},\kappa_O} \Phi_{-\frac{\pi}{2},\kappa_L}$
Other	15	9*	U _C U _O N _{2L}	$\gamma_C \gamma_O \Phi_{-\frac{\pi}{2},\kappa_L}$
Nontarget N3 location				
Correlated target	16	6*	T _C T _O N _{3L}	$\Phi_{0,\kappa_C} \Phi_{0,\kappa_O} \Phi_{\pi,\kappa_L}$
Correlated N1	17	12*	N _{1C} N _{1O} N _{3L}	$\Phi_{\frac{\pi}{2},\kappa_C} \Phi_{\frac{\pi}{4},\kappa_O} \Phi_{\pi,\kappa_L}$
Correlated N2	18	12*	N _{2C} N _{2O} N _{3L}	$\Phi_{-\frac{\pi}{2},\kappa_C} \Phi_{-\frac{\pi}{4},\kappa_O} \Phi_{\pi,\kappa_L}$
Correlated N3 (triple bound)	19	13	N _{3C} N _{3O} N _{3L}	$\Phi_{\pi,\kappa_C} \Phi_{\frac{\pi}{2},\kappa_O} \Phi_{\pi,\kappa_L}$
Other	20	9*	U _C U _O N _{3L}	$\gamma_C \gamma_O \Phi_{\pi,\kappa_L}$

The standard triple joint model includes 20 response combinations of location (4) x color-orientation (5), as numbered by m_{full} . The simple triple model also allows for 20 response combinations, however theoretically negligible responses (such as N_{1C}N_{1O}N_{3L} and N_{1C}N_{1O}N_{2L}) are averaged into a single probability estimate (pN_{1C}N_{1O}N_{23L}) to ease fitting of the model, thus $m_{simple} = 13$. In the rightmost column, φ is a von Mises probability density function, with concentration κ_C , κ_O , or κ_L (standard deviation = $\sqrt{1/\kappa}$) and means of 0°, 90°, -90°, and 180° (color or location) or 0°, 45°, -45°, and 90° (orientation) for the target (T), critical nontarget (N1), adjacent nontarget (N2), and diagonal nontarget (N3) distributions, respectively; γ_C and γ_O are uniform distributions that reflect the probability of responding at random.

* indicates response combinations that were modeled together with a single parameter estimate.

Table 2. Standard Triple Joint Model Parameter Estimates on Hold Trials

	Experiment 1 n = 23	Experiment 2 n = 26	Experiment 3 n = 24	Experiment 4 n = 23
TcToTL	.843 [.820 .853]	.840 [.825 .858]	.818 [.804 .834]	.870 [.855 .880]
N1cN1oTL	.002 [.000 .003]	.002 [.000 .007]	.001 [.000 .003]	.001 [.000 .002]
N2cN2oTL	.000 [.000 .004]	.003 [.000 .005]	.002 [.000 .006]	.000 [.000 .003]
N3cN3oTL	.000 [.000 .001]	.000 [.000 .002]	.000 [.000 .001]	.000 [.000 .003]
UcUoTL	.092 [.080 .109]	.143 [.119 .151]	.103 [.085 .112]	.085 [.072 .097]
TcToN1L	.001 [.000 .003]	.003 [.001 .005]	.001 [.000 .003]	.001 [.000 .003]
N1cN1oN1L	.012 [.008 .015]	.000 [.000 .002]	.021 [.017 .026]	.004 [.002 .007]
N2cN2oN1L	.000 [.000 .001]	.000 [.000 .001]	.000 [.000 .002]	.000 [.000 .002]
N3cN3oN1L	.000 [.000 .002]	.001 [.000 .002]	.000 [.000 .001]	.000 [.000 .001]
UcUoN1L	.009 [.006 .014]	.002 [.000 .004]	.005 [.002 .008]	.004 [.002 .008]
TcToN2L	.002 [.000 .002]	.000 [.000 .001]	.000 [.000 .001]	.000 [.000 .001]
N1cN1oN2L	.001 [.000 .001]	.001 [.000 .002]	.000 [.000 .001]	.000 [.000 .001]
N2cN2oN2L	.008 [.006 .012]	.001 [.000 .002]	.011 [.006 .013]	.014 [.010 .017]
N3cN3oN2L	.000 [.000 .001]	.001 [.000 .002]	.000 [.000 .001]	.001 [.000 .001]
UcUoN2L	.003 [.001 .007]	.002 [.000 .003]	.005 [.002 .009]	.003 [.002 .008]
TcToN3L	.000 [.000 .001]	.000 [.000 .002]	.000 [.000 .001]	.001 [.000 .002]
N1cN1oN3L	.001 [.000 .003]	.000 [.000 .001]	.001 [.000 .001]	.000 [.000 .001]
N2cN2oN3L	.001 [.000 .003]	.000 [.000 .001]	.000 [.000 .002]	.000 [.000 .001]
N3cN3oN3L	.018 [.013 .024]	.001 [.000 .002]	.021 [.014 .027]	.010 [.007 .013]
UcUoN3L	.005 [.004 .010]	.001 [.000 .003]	.009 [.006 .014]	.005 [.002 .006]
σ_c	26.497 [25.132 27.036]	27.623 [26.588 28.303]	28.624 [27.653 30.111]	21.400 [20.306 21.755]
σ_o	35.463 [34.651 37.649]	30.629 [30.102 32.822]	27.688 [25.928 28.055]	32.762 [31.588 34.865]
σ_L	10.382 [10.183 10.751]	10.524 [10.207 10.694]	8.632 [8.446 8.834]	9.465 [9.218 9.767]

Maximum-likelihood estimates, with 95% highest density intervals presented in brackets. σ_c and σ_L range from -180° to $+180^\circ$, while σ_o ranges from -90° to $+90^\circ$ ($\sigma = \sqrt{1/\kappa}$)

Table 3. Standard Triple Joint Model Parameter Estimates on Shift Trials

	Experiment 1 n = 23	Experiment 3 n = 24	Experiment 4 n = 23
T _c T _o T _L	.822 [.802 .836]	.896 [.873 .898]	.881 [.866 .896]
N _{1c} N _{1o} T _L	.000 [.000 .002]	.002 [.000 .006]	.000 [.000 .002]
N _{2c} N _{2o} T _L	.001 [.000 .001]	.000 [.000 .002]	.000 [.000 .001]
N _{3c} N _{3o} T _L	.000 [.000 .002]	.000 [.000 .001]	.000 [.000 .001]
U _c U _o T _L	.084 [.073 .101]	.069 [.062 .083]	.061 [.047 .070]
T _c T _o N _{1L}	.000 [.000 .002]	.001 [.000 .002]	.000 [.000 .002]
N _{1c} N _{1o} N _{1L}	.055 [.045 .062]	.017 [.013 .024]	.032 [.025 .038]
N _{2c} N _{2o} N _{1L}	.001 [.000 .001]	.000 [.000 .001]	.000 [.000 .001]
N _{3c} N _{3o} N _{1L}	.000 [.000 .004]	.000 [.000 .001]	.000 [.000 .001]
U _c U _o N _{1L}	.032 [.020 .038]	.009 [.006 .014]	.018 [.009 .020]
T _c T _o N _{2L}	.001 [.000 .001]	.001 [.000 .002]	.001 [.000 .002]
N _{1c} N _{1o} N _{2L}	.001 [.000 .003]	.000 [.000 .001]	.000 [.000 .002]
N _{2c} N _{2o} N _{2L}	.000 [.000 .001]	.001 [.000 .002]	.001 [.000 .002]
N _{3c} N _{3o} N _{2L}	.000 [.000 .001]	.000 [.000 .001]	.001 [.000 .001]
U _c U _o N _{2L}	.001 [.000 .002]	.001 [.000 .001]	.003 [.000 .003]
T _c T _o N _{3L}	.001 [.000 .002]	.001 [.000 .002]	.000 [.000 .002]
N _{1c} N _{1o} N _{3L}	.001 [.000 .001]	.000 [.000 .001]	.000 [.000 .001]
N _{2c} N _{2o} N _{3L}	.000 [.000 .001]	.000 [.000 .001]	.001 [.000 .004]
N _{3c} N _{3o} N _{3L}	.000 [.000 .001]	.000 [.000 .002]	.000 [.000 .001]
U _c U _o N _{3L}	.001 [.000 .002]	.001 [.000 .003]	.002 [.000 .003]
σ_c	25.006 [24.041 25.570]	27.937 [27.136 28.665]	20.787 [19.680 21.192]
σ_o	30.971 [29.875 32.350]	24.087 [23.409 24.823]	30.691 [29.552 31.555]
σ_L	10.413 [10.110 10.707]	9.151 [8.968 9.342]	9.289 [9.161 9.624]

Maximum-likelihood estimates, with 95% highest density intervals presented in brackets. σ_c and σ_L range from -180° to $+180^\circ$, while σ_o ranges from -90° to $+90^\circ$ ($\sigma = \sqrt{1/\kappa}$)

Table 4. Simple Triple Joint Model Mean Parameter Estimates on Hold Trials

	Experiment 1 n = 23	Experiment 2 n = 26	Experiment 3 n = 24	Experiment 4 n = 23
T _c T _o T _L	0.851	0.844	0.819	0.853
T _c T _o N123 _L	0.006	0.006	0.007	0.007
U _c U _o T _L	0.047	0.076	0.05	0.053
U _c U _o N123 _L	0.015	0.006	0.012	0.011
N1 _c N1 _o T _L	0.008	0.021	0.008	0.007
N1 _c N1 _o N1 _L	0.014	0.005	0.031	0.007
N1 _c N1 _o N23 _L	0.004	0.005	0.002	0.004
N2 _c N2 _o T _L	0.011	0.014	0.009	0.009
N2 _c N2 _o N2 _L	0.011	0.005	0.018	0.019
N2 _c N2 _o N13 _L	0.003	0.004	0.006	0.004
N3 _c N3 _o T _L	0.006	0.006	0.006	0.004
N3 _c N3 _o N3 _L	0.022	0.004	0.028	0.018
N3 _c N3 _o N12 _L	0.003	0.004	0.004	0.005
σ_c	27.821	29.741	32.454	21.569
σ_o	39.468	36.972	30.901	36.678
σ_L	10.328	10.334	8.586	9.321

Group means, σ_c and σ_L range from -180° to $+180^\circ$, while σ_o ranges from -90° to $+90^\circ$ ($\sigma = \sqrt{1/\kappa}$)

Table 5. Simple Triple Joint Model Mean Parameter Estimates on Shift Trials

Parameter	Experiment 1	Experiment 3	Experiment 4
	n = 23	n = 24	n = 23
T _c T _o T _L	0.831	0.877	0.863
T _c T _o N123 _L	0.006	0.007	0.006
U _c U _o T _L	0.037	0.041	0.046
U _c U _o N123 _L	0.011	0.008	0.009
N1 _c N1 _o T _L	0.008	0.009	0.005
N1 _c N1 _o N1 _L	0.067	0.031	0.04
N1 _c N1 _o N23 _L	0.005	0.004	0.003
N2 _c N2 _o T _L	0.01	0.007	0.007
N2 _c N2 _o N2 _L	0.003	0.003	0.003
N2 _c N2 _o N13 _L	0.005	0.003	0.006
N3 _c N3 _o T _L	0.008	0.004	0.004
N3 _c N3 _o N3 _L	0.003	0.006	0.004
N3 _c N3 _o N12 _L	0.006	0.002	0.005
σ_c	26.816	30.483	21.425
σ_o	38.402	26.686	33.186
σ_L	10.228	9.104	9.284

Group means, σ_c and σ_L range from -180° to $+180^\circ$, while σ_o ranges from -90° to $+90^\circ$ ($\sigma = \sqrt{1/\kappa}$)