

Running Head: SPATIAL WORKING MEMORY AND ATTENTION

Working Memory for Location is Attracted Toward the Focus of Attention

Jeffrey S. Johnson¹ & John P. Spencer^{1,2}

¹Department of Psychology, University of Iowa

²Iowa Center for Developmental and Learning Sciences, University of Iowa

Abstract

Studies examining the relationship between spatial attention and spatial working memory (SWM) have suggested that attention is the rehearsal mechanism for SWM. To explore this possibility, participants completed a color-discrimination task during the delay interval of a spatial recall task. In the critical *shifting attention* condition, the color stimulus could appear either *toward* or *away from* the memorized location relative to a midline reference axis. If spatial attention is responsible for maintenance in SWM, shifting attention during the delay should increase spatial recall errors, regardless of the direction of the shift. Conversely, if attention plays a modulatory role in SWM, there should be systematic changes in the pattern of spatial recall errors depending on the direction of the attention shift. Results showed the latter—spatial memory errors were systematically biased toward the focus of attention. We simulate these findings using a neurally-plausible model of SWM reciprocally coupled to a simple color recognition network. The behavioral results and model simulations highlight the advantage of moving concepts like attention and working memory to a more formal level.

Working Memory for Location is Attracted Toward the Focus of Attention

Attention and working memory are closely related cognitive processes that play an important role in everyday behaviors. For example, attention has been shown to increase the efficiency of perception, cognition, and overt behavior in complex environments by restricting cognitive processes to a task-relevant subset of the information that is available (see review in, Luck & Vecera, 2002). Similarly, working memory plays a role in the comparison of perceptual information obtained from spatially and/or temporally separated objects (Hollingworth, in press), and in maintaining information in an active state for use in complex cognitive tasks such as reading and mental rotation (Hyun & Luck, in press; Just & Carpenter, 1992), among other functions. Although most research has focused on examining these two processes in isolation, exploring and understanding the relationship between attention and working memory has become a central focus of research in visuo-spatial cognition. This can be a challenging endeavor, however, because attention and working memory themselves are not unitary cognitive processes. Consequently, it is important to be precise about the variety of attention and working memory being considered when studying interactions between systems (see discussion in Awh, Vogel, & Oh, 2006).

The present study examines the relationship between attention and working memory within the spatial domain, focusing on interactions between spatial working memory (SWM) and spatial selective attention. A growing body of evidence from studies using a combination of behavioral and brain imaging methods has suggested that there is a substantial amount of overlap between these systems (Awh & Jonides, 1998; Awh, Jonides, & Reuter-Lorenz, 1998; Awh et al., 1999; Cowan et al., 2005; Postle, 2006). Decades of research on spatial selective attention has revealed that observers make faster and more accurate responses to stimuli that fall in attended versus unattended locations (see, Luck & Vecera, 2002; Posner, 1980). Using these findings as a starting point, Awh and colleagues (1998) reasoned that if spatial attention remains focused on memorized locations during memory delays, the processing of items appearing at

those locations should be facilitated compared to items appearing at different, non-memorized locations. This is exactly what they found: performance of a form discrimination task was enhanced at locations that were currently being held in memory, suggesting that spatial attention remained at the memorized location during the delay interval. In an additional experiment, these researchers found that location memory was impaired when attention was shifted away from the memorized location to discriminate the color of an item appearing at a different location. This finding suggests that attention not only remains at the memorized location, but that it directly contributes to the maintenance of location information in SWM.

Functional imaging studies of SWM have likewise suggested a high degree of overlap between systems of attention and SWM (Awh *et al.*, 1999; Postle, Awh, Jonides, Smith, & D'Esposito, 1999). These experiments have revealed that maintaining information in SWM leads to increased activation in many of the same brain areas that are activated in experiments probing spatial attention. Additionally, event-related potential (ERP) studies of SWM (Awh, Anllo-Vento, & Hillyard, 2000) have shown that the presentation of stimuli at locations that are being held in memory elicits ERPs with similar spatio-temporal patterns of amplitude enhancement to those observed in studies of spatial selective attention (Hillyard & Anllo-Vento, 1998; Hillyard, Vogel, & Luck, 1998).

On the basis of these results, Awh and colleagues (Awh & Jonides, 1998; Awh *et al.*, 1998) have suggested that, as with verbal working memory, the maintenance of information in SWM requires an active rehearsal process, and that rehearsal in SWM is mediated by spatial attention. But in what sense is attention the rehearsal mechanism for SWM? Is continued spatial attention necessary to maintain the sustained patterns of neural activation thought to underlie working memory performance? If this were the case, one would expect severe deficits when attention is withdrawn from memorized locations to perform a secondary task. Although Awh *et al.* (1998) showed that performance of a same/different location memory task was impaired when attention was withdrawn to perform a secondary task during the delay interval, performance of the secondary task only produced a 12% drop in performance compared to performance in a

control condition that did not require a shift of attention. One might have expected more severe impairments if focused spatial attention were the sole mechanism responsible for maintenance in SWM.

Another possibility, therefore, is that attention plays a role in modulating and/or enhancing the processes underlying rehearsal (i.e., active maintenance) in SWM, rather than being the rehearsal mechanism itself. In the present study, we explored these issues by combining behavioral work with a neurally-plausible model of SWM. We contend that the conceptual precision afforded by formal models can help clarify concepts like “rehearsal”, as well as the relationship between mechanisms of attention, on one hand, and SWM on the other. The section that follows provides an overview of the theoretical framework we used—the Dynamic Field Theory (DFT). We then use this theory to generate two competing views of how attention might modulate neural mechanisms that maintain location information in SWM.

The Dynamic Field Theory of Spatial Cognition

In the present paper, we use the Dynamic Field Theory (DFT) as a theoretical framework for formalizing how SWM and attention interact. The DFT is a neurally-plausible, processed-based approach to spatial cognition that is in a class of bistable neural networks first developed by Amari and colleagues (Amari, 1977; Amari & Arbib, 1977; see also, Compte, Brunel, Goldman-Rakic, & Wang, 2000; Wilson & Cowan, 1972). This approach provides a basis for characterizing the representational states underlying behavior through the concept of activation, which is familiar from mathematical psychology, connectionism, and theoretical neuroscience (Churchland & Sejnowski, 1992; Williams, 1986). In our dynamic field approach, this concept takes the form of an activation field, defined over the metric dimension represented. For example, in the same/different location memory task used in the experiments of Awh et al. (1998), activation would be distributed across a continuous, metric spatial dimension (e.g., target direction) encompassing the area on the computer monitor where memory targets could appear. A localized peak of activation within this field indicates that a target object (e.g., a small black circle) has been detected at a particular location in the task space. Patterns of activation within

such fields can live in different attractor states, including: a stable *resting state* where activation remains at baseline levels for all field sites; an *input-driven state* where input can form stabilized “peaks”, or “bumps”, of activation in the field, but peaks die away when input is removed; and a *self-sustaining state* where activation peaks remain stable in the absence of continued input—a form of working memory central to our work and the main focus of the present paper.

Figure 1 shows a simulation that implements the core features of the DFT of spatial cognition (for discussion of the full model, see Spencer, Simmering, Schutte, & Schöner, 2007). The model consists of 3-layers: A perceptual, or input, layer (PF; Fig. 1A), a layer of inhibitory interneurons (Inhib; Fig. 1B) and, a spatial working memory layer (SWM; Fig. 1C). In each case, the direction of targets in the task space is shown along the *x*-axis (left-right), the *y*-axis (up-down) shows the activation of each neuron in the field, and the elapsed time from the start to the end of the trial is shown along the *z*-axis, with time moving from front to back. Arrows depict excitatory (solid) and inhibitory (dashed) connections among fields.

One of the most common tasks used in the spatial cognition literature is spatial recall. In this task, participants are briefly shown a single target item and are asked to remember the target’s location. This is followed by a delay interval of variable length (0-20 s), after which the participant indicates the remembered location of the target by, for instance, moving a computer mouse or a joystick. A number of studies of this type have demonstrated that metric memory for location shows delay-dependent biases away from perceived reference axes, such as the midline symmetry axis of a task space (Engelbretson & Huttenlocher, 1996; Huttenlocher, Hedges, & Duncan, 1991; Schiano & Tversky, 1992; Schutte & Spencer, 2003; Spencer & Hund, 2002; Tversky & Schiano, 1989; Werner & Schmidt, 2000). For example, Huttenlocher and colleagues have demonstrated that adults’ recall responses are biased away from the vertical and horizontal symmetry axes of the paper or screen on which a circle is presented (Huttenlocher et al., 1991). Similarly, Spencer and Hund (2002) have shown that adults’ responses are biased away from the midline symmetry axis of a large table.

The simulation shown in Figure 1 illustrates how working memory for location is

implemented in the model, and how biases away from perceived reference axes are explained. At the start of the trial, PF builds a small peak of activation at 180° , reflecting perception of a salient reference frame in the environment (e.g., two dots along the midline of the task space). Next, a target appears at 220° . This creates a peak of activation centered at that location. When the target disappears, a peak re-forms in PF at 180° as the system re-locks onto the reference cues in the task space. This is important because it allows the system to continually calibrate egocentric and object- or table-centered frames of reference during the delay (Spencer et al., 2007).

Panel C shows the effect of coupling PF to SWM. At the start of the trial, SWM receives relatively weak reference input from PF because the reference cue in the task space (i.e., the reference input to PF) is not very salient. Next, the target is turned on, passing strong target-related input into the working memory layer. This event moves the working memory layer into a self-sustaining state where peaks of activation can be maintained in the absence of input.

Maintenance is made possible in the model through reciprocal coupling between the SWM and Inhib layers, which implements a locally-excitatory and laterally-inhibitory form of interaction among neurons (Amari, 1977). Specifically, neurons in SWM have locally excitatory interactions that boost the activation of neurons tuned to similar spatial locations. Additionally, SWM is reciprocally coupled to the layer of inhibitory neurons: neurons in SWM provide excitatory inputs to similarly-tuned neurons in Inhib, which, in turn, project inhibition broadly back to SWM. This pattern of interactions allows localized, self-sustained peaks of activation to be maintained in SWM in the absence of continued input (e.g., when the target object is removed from view).

Importantly, the activation peak in SWM in Figure 1 is stably maintained during the memory delay even though PF has re-acquired the reference frame. This sets the stage for memories to drift *away from the reference frame* because the system is effectively holding onto two stable states at once: it is locked onto the reference frame in PF to stay calibrated with the world, and it is actively maintaining a memory of the target location in SWM. The response bias evident in experiment arises from the place where these two stable states meet—the shared

inhibitory layer. In particular, reference-related input to the inhibitory layer near midline causes the peak of activation in SWM to “drift” away from the midline of the task space because there is stronger inhibition on the midline-side of the SWM peak than the “outer” side. In effect, the reference frame “pushes” the peak outward, exaggerating the “leftness” of the target location.

In light of our model of SWM, what can we say about the origins of the attention effects observed in Awh et al. (1998)? Recall that Awh et al. proposed that the withdrawal of attention from a remembered location disrupts rehearsal, that is, the maintenance of the target location in SWM. This explains the increased spatial memory errors found in their “shifting attention” conditions. How might such effects arise in our model? One possibility closely related to the Awh et al. proposal is that shifting attention away from the memorized location has a global influence on maintenance in SWM. For instance, shifting attention might de-stabilize the maintenance of peaks in SWM through a global change in the resting level of neurons in the SWM field. Less stable WM peaks are more subject to random perturbations, leading to greater response error.

It is also possible, however, that the shifting attention manipulation in these studies caused WM peaks to “drift” either toward or away from the focus of attention. This would occur if, for example, shifting spatial attention served as an additional source of location-specific input to SWM, modulating the patterns of sustained activation during a trial in a manner similar to the reference effects discussed above. If this were the case, then we should find systematic distortions in spatial discrimination or recall responses following the attention manipulation. For instance, increased drift would make it more difficult to determine whether an item’s position had changed or not, particularly on those trials in which the change in the target’s location was in the direction of drift (Simmering, Spencer, & Schöner, 2006). Given the randomization scheme used in Awh et al. (1998), it is not clear whether such “drift” effects occurred.

In the present paper, we probe these different possibilities by combining our detailed, fine-grained understanding of the mechanisms that underlie SWM for individual locations with the dual-task design used by Awh et al. (1998). Specifically, we used the same task from Awh et

al. with two modifications: (1) we probed memory for a relatively small set of target locations (rather than randomly selecting targets from a large set) and shifted attention either *toward* or *away from* each location relative to a midline reference axis; and (2) we used a spatial recall task after the memory delay (rather than probing memory with a spatial discrimination task). Our results suggest that shifting attention during the delay interval of a spatial recall task *biases WM toward the focus of attention*. We then demonstrate how the DFT can explain these findings.

Experiment 1

Awh and colleagues (1998) showed that accuracy in a same/different location memory task declined when attention was shifted away from the memorized location to complete a secondary task during the delay interval. In the present experiment, we used a similar design to determine whether shifts of spatial attention to non-memorized locations produce systematic distortions in working memory for location. To address this issue, participants completed the location memory task depicted in Figure 2A-B, with the inclusion of a color-discrimination task during the delay interval. In this task, participants are shown a single target item and must remember the target location across a subsequent 10-second delay interval. During the delay interval, two different types of colored stimuli were presented. In one case, color stimuli consisted of small colored circles that were presented near the location of the memory target either toward or away from the midline of the task space (see Figure 2B). Thus, accurate discrimination performance entailed shifts of attention in a particular direction relative to the location being held in memory (i.e., toward versus away from midline). In other conditions, color stimuli were large colored circles that covered a large portion of the task space, including all possible target locations (see Figure 2). In these conditions, color information was available throughout the task space, and no shift of attention was required for accurate discrimination performance.

This design allowed us to test two competing ideas about how shifting attention impacts memory for individual locations. The first idea is consistent with the proposal of Awh and colleagues (1998), and suggests that shifting attention leads to a global reduction in the stability

of WM maintenance. According to this view, attention is fundamentally tied to the rehearsal mechanism in SWM. Therefore, if attention is directed elsewhere, location memory should suffer. In particular, we should see an increase in response error (i.e., absolute error relative to the target), but no systematic shift in the directionality of the error across shifting attention conditions (i.e., whether attention is shifted toward or away from midline). A second possibility derived from the proposal that attention enhances and/or modulates the processes involved in maintenance—rather than being *responsible* for maintenance—is that shifting attention will lead to systematic distortions in WM for location depending on the direction of the attention shift. In this case, we should see a systematic shift in directional error across shifting attention conditions, but little increase in overall error. Note that in both cases, we expect to see drift away from midline given the pervasiveness of reference-related biases in spatial recall tasks (Huttenlocher et al., 1991; Huttenlocher, Hedges, & Vevea, 2000; Schiano & Tversky, 1992; Schutte, Spencer, & Schöner, 2003; Spencer & Hund, 2002; Tversky & Schiano, 1989; Werner & Schmidt, 2000). The critical question, therefore, is how such errors are influenced by the attention manipulation.

Methods

Participants. Twenty-four University of Iowa undergraduate volunteers participated in this experiment (14 females and 10 males; M age = 20.33 yrs, SD = 2.51 yrs). Participants received class credit or monetary compensation (\$7/hr) for their participation. All participants reported normal or corrected-to-normal visual acuity and normal color vision. One participant was replaced due to difficulties with the eye tracker.

Apparatus. Throughout the session, participants were seated on a chair placed within an arc cut out from the side of a large (0.921 m x 1.194 m) table (see Figure 2). The surface of the table was homogeneous, and three edges of the table were covered with a curved border to occlude the corners from participants' view. Sessions were conducted in a dimly lit room with black curtains covering the walls and ceiling. Stimuli were projected onto the surface of the table (a rear-projection surface) from below using a Barco 708 data projector, with a screen resolution of 1024 x 768. Responses were generated using a Microsoft Sidewinder force-feedback joystick

positioned immediately to the right of the participant on the surface of the table. Stimulus presentation and response recording was controlled by a Pentium IV PC-based computer running custom software in Microsoft Visual Studio 6.

Eye-position data was acquired by an ISCAN RK-726 PCI head-mounted eye tracker. Head-position data was acquired by a 3-dimensional optical electronic motion analysis system (OPTOTRAK, Northern Digital, Inc.) that tracked the movements of three small (radius = 3.5 mm) infrared-emitting diodes (IREDs) attached to the brim of the eye tracker. Acquired eye- and head-position data were then used to calculate fixation position using a custom *eye-frame calibration* procedure (Johnson, Liu, Thomas, & Spencer, in press) running in Matlab 6.5 (Mathworks, Inc.). This procedure provides robust eye-position estimates that are tolerant to movements of the head, allowing us to monitor fixation position while allowing relatively unrestricted head movements.

Stimuli. Stimuli were presented against a grey background on the surface of the display table. Throughout each experimental trial, participants maintained fixation on a small (1.66° in diameterⁱ), yellow fixation dot positioned 21.17 cm in front of them on the surface of the display table (x=512 pixels, y=544 pixels). Memory stimuli consisted of small (1.0° at base), light-blue equilateral triangles (“spaceships”) presented at different angular separations on the periphery of an imaginary circle (radius = 10 cm) centered at fixation. Sample arrays consisted of a single memory stimulus presented at one of four possible target locations either $\pm 30^\circ$ or $\pm 60^\circ$ of angular rotation away from the midline of the display table (i.e., to the left (-) or right (+) of midline = 0°; see Figure 2).

As in Awh, et al. (1998), two types of color stimuli were presented in the color-discrimination task that was completed during the retention interval of the memory task: (a) A small colored circle subtending 1.74° of visual angle; and (b) a large colored circle subtending 17.38° of visual angle. In both cases, the colors of these stimuli could be either two different shades of blue, or two different shades of red. For small colored circles, the RGB values of the colors used wereⁱⁱ: blue1 = [0, 0, 255], blue2 = [173, 216, 230]; red1 = [255, 0, 0], red2 = [255, 0,

255]. For large colored circles, the colors were: blue1 = [0, 0, 255], blue2 = [25, 160, 235]; red1 = [255, 34, 0], red2 = [255, 70, 128]. When large colored circles were presented, they were always centered at fixation. Conversely, on half of the trials, the small colored circle was presented at a location 10° of angular rotation closer to midline than the position of the memory stimulus that appeared on that trial, whereas on the other half it appeared at a location 10° in the opposite direction (i.e., away from midline). For example, if the memory stimulus appeared at -30° (i.e., 30° to the left of midline), the small colored circle could appear at either -20° (toward midline) or -40° (away from midline).

Design. We used a within-subjects design with four different experimental conditions (see Figure 2): A dual-task *shifting-attention* condition where the small-colored-circle color discrimination task was completed during the delay interval of the spatial recall task; a dual-task *static-attention* condition where the large-colored-circle discrimination task was completed during the delay interval, and; two single-task control conditions which were identical to the shifting-attention and static-attention conditions, with the exception that the color discrimination task was not completed during the delay interval (i.e., participants were instructed to ignore the colored circles and to focus exclusively on accurate performance of the spatial memory task). There were two different independent variables in the shifting attention conditions: The location of the memory target ($\pm 30^\circ$, $\pm 60^\circ$) and the direction of the color-discrimination stimulus (toward versus away from midline).

Procedures. At the beginning of the experimental session, the eye tracker was calibrated using our standard procedures (Johnson, Liu et al., in press). Next, participants were seated at the table and the experimenter explained the “spaceship” game, playing several different demonstration trials to illustrate each of the different experimental conditions. The sequence of events in a trial is depicted in Figure 2B. Each trial began with the appearance of a yellow “start dot” and a set of green crosshairs that could be moved using the joystick. When the dot appeared, participants were instructed to move the crosshairs to the dot’s position and fixate the dot throughout the subsequent trial. If the eyes moved further than 4.7 cm (approximately half the

distance to the target) off of the start dot at any point during the trial, the trial was halted, the participant was prompted to return to the start location, and the trial was repeated.

Once the crosshairs and the eyes were positioned over the start dot, the computer stated, “Beginning search for enemy spaceships”, after which a single “spaceship” stimulus was presented for 2000 ms at one of the four possible target locations ($\pm 30^\circ$, $\pm 60^\circ$). This was followed by an inter-stimulus interval (ISI) that varied randomly from 1500 to 7000 ms, and a 500-ms presentation of either a small or a large colored circle, depending on the experimental condition. When the colored circle appeared during dual-task trials, participants made a speeded discrimination response, indicating which of two possible colors they saw by pressing one of two buttons on the base of the joystick using the middle and index fingers of the left hand. This was immediately followed by auditory feedback in the form of a laser sound, for correct responses, or a buzzer sound for incorrect responses. Participants also received auditory feedback if they failed to respond to the discrimination target within 1000 ms of stimulus offset. The offset of the colored circle was followed by a variable delay of 1000-7000 ms, at which point the participant heard “ready, set, go”, with the word “go” occurring exactly 10 seconds after the offset of the original memory display. On “go”, participants moved the crosshairs to the remembered spaceship location with their right hand using the joystick, and clicked the trigger on the joystick using the right index finger. Once they made their recall response, the spaceship stimulus reappeared for 500 ms, followed by a 1500 ms feedback screen that gave detailed information about their performance on that trial (see details in, Spencer & Hund, 2002). The next trial began after a 1000 ms inter-trial interval.

The experiment began with a series of single-task practice blocks designed to familiarize participants with the tasks. First, they completed 12 trials of the spatial memory task by itself; 3 trials to each of 4 memory targets. This was followed by one block of 8 trials of the small-colored-circle discrimination task by itself, and another block of 8 trials of the large-colored-circle task by itself. Additionally, participants completed a short practice block at the beginning of each experimental block during the session. For the shifting-attention condition and its

corresponding control condition, practice blocks consisted of a total of 8 trials: 2 trials to each memory target (4), with the color (e.g., blue1 or blue2) and direction (toward versus away from midline) of the small-colored circle occurring with equal likelihood. For the static-attention condition and its corresponding control condition, participants completed 4 practice trials: 1 trial to each memory target (4), with each color occurring with equal likelihood. This resulted in a total of 24 regular practice trials.

For the shifting-attention and control conditions, experimental blocks consisted of 32 trials: 4 trials to each memory target (4) with the color-discrimination target appearing in each direction (2), and with color counterbalanced across trials. For the static-attention and control conditions, participants completed 16 trials: 4 trials to each target (4), with each color appearing equally often. Thus, participants completed a total of 148 trials during the experimental session: 28 single-task practice trials at the beginning of the session, 24 regular practice trials, and 96 total experimental trials. The color sets used in the discrimination task were counterbalanced across participants, with half of participants responding to the red color set and the other half to the blue set. For control conditions, each participant saw the opposite color set. Conditions were blocked and were counterbalanced across participants using a balanced Latin Square design.

Method of analysis. To calculate spatial recall errors for each trial, we recorded the x and y coordinates of the crosshairs when the participant pulled the joystick trigger. Then, we computed the angular separation between the vector from the start location to the response and the vector from the start location to the actual target location (in degrees). These directional errors were computed such that positive errors were in a direction away from the midline symmetry axis of the table, whereas negative errors reflected errors toward midline. Directional errors were checked for outliers by removing trials on which errors were greater than ± 2 standard deviations from the mean directional error for a given participant. A total of 92 outlier trials (3.99% of all trials) were removed. Mean directional errors were computed by averaging across trials to the same target and, for shifting-attention conditions, the same direction of the attentional shift. In addition, before conducting our statistical analyses, we collapsed across

targets positioned the same distance from midline on the left and right side of the screen (i.e., across $\pm 30^\circ$ and $\pm 60^\circ$ targets). We also calculated percent correct accuracy as a measure of performance in the color-discrimination task.

Results and Discussion

Color discrimination. Performance in the color discrimination task is shown in Figure 3A. Color discrimination performance was uniformly high overall, but was somewhat better for the small-colored-circle task than the large-colored-circle task. This was confirmed by a one-way ANOVA showing a trend toward a main effect of Condition, $F(1,23) = 4.15, p = .053$. This result contrasts with the findings of Awh et al. (1998, Experiment 3), where performance was found to be worse in the small- versus the large-colored-circle task. This discrepancy likely reflects differences in the particular colors used in the color-discrimination task across studies. In particular, for our experiment the colors used in the small-colored-circle task were highly discriminable (e.g., blue and light-blue), whereas in the large-colored-circle task, differences between the colors were more subtle (e.g., two similar shades of blue). Although our results contrast with those of Awh et al., the relative difficulty of the large-colored-circle task ensures that any increase in error found in the shifting- versus static-attention conditions cannot be attributed to differences in the overall difficulty of the color-discrimination task in this condition.

Location memory. Mean directional errors in the static attention and static control conditions are shown in Figure 3B. As in previous studies of spatial recall, participants' estimation responses revealed systematic drift away from the midline of the table (positive errors) in each condition, with substantially larger errors for $\pm 30^\circ$ targets than for $\pm 60^\circ$ targets. Additionally, errors were larger overall in the static-attention condition than the static-control condition; thus, there was a non-specific increase in error when participants had to complete the dual task. In keeping with these findings, a repeated-measures ANOVA with Condition (static attention, static control) and Target ($\pm 30^\circ, \pm 60^\circ$) as factors revealed significant main effects of Condition, $F(1,23) = 5.36, p < .04$, and of Target, $F(1,23) = 48.55, p < .001$. The interaction was not significant, $F < 1$.

Mean directional errors for the shifting attention and shifting control conditions are shown in Figure 3C. As before, participants' estimation responses revealed substantial drift away from the midline of the task space (positive errors), and larger errors for targets near versus far from midline. Additionally, participants made substantially larger errors when the small-colored-circle was presented in a direction away from midline, and smaller errors when the small-colored-circle was presented toward midline. Importantly, this difference was roughly three times larger in the shifting-attention condition (Mean difference = 2.83°), where participants were required to complete the color-discrimination task, than in the shifting-control condition (Mean difference = 1.0°), where they were instructed to ignore the color stimuli. This finding suggests that memory for location was attracted towards the focus of attention when attention was shifted to the location of the color stimulus to complete the discrimination task during the delay.

These data were analyzed in a three-way, repeated-measures ANOVA with Condition (shifting-attention, shifting-control), Color Direction (toward, away), and Target ($\pm 30^\circ$, $\pm 60^\circ$) as factors. This analysis revealed significant main effects of Target, $F(1,23) = 31.49$, $p < .001$, and Color Direction, $F(1,23) = 5.89$, $p < .03$. The Condition \times Target interaction was also significant, $F(1,23) = 4.60$, $p < .05$, as was the Condition \times Color Direction interaction, $F(1,23) = 4.63$, $p < .05$. All additional effects were not significant, all $ps > .20$.

The significant Condition \times Target interaction was analyzed using tests of simple effects comparing mean directional errors in the shifting attention and shifting control conditions at each target location. This analysis revealed no significant effect of attention condition at either the 30° or 60° target locations, all $ps > .11$, although errors were approximately 1° larger overall to targets in the shifting-control versus shifting-attention conditions. Additionally, we compared mean directional errors across targets for the shifting attention and shifting control conditions separately. This analysis revealed significant main effects of Target for both conditions, $F(1,23) > 15.0$, all $ps < .001$, although the slope across targets was somewhat steeper in the shifting attention condition (see Figure 3C).

More central to our experimental hypotheses is the significant Condition \times Color Direction interaction. To explore this interaction, we again conducted tests of simple effects looking at errors when the color target appeared toward versus away from midline in the shifting-attention and shifting-control conditions separately. If increased errors in the shift-away versus shift-toward condition reflect the voluntary shifting of attention to the location of the color stimulus, rather than stimulus-driven orienting to the location of a sudden onset, we would expect a significant effect of color direction in the shifting-attention, but not the shifting-control condition. Consistent with this proposal, these analyses revealed a significant main effect of Color Direction for the shifting-attention condition, $t(23) = 2.92, p < .01$, but not for the shifting-control condition, $p > .20$. This confirms that differences in errors for shifting-away versus shifting-toward trials were significantly larger when attention was voluntarily shifted to the location of the color stimulus.

In summary, results revealed that memory for location was attracted towards the focus of attention when attention was shifted to a nearby, non-memorized location during the delay interval of a spatial recall task. Importantly, shifting attention was found to modulate the patterns of delay-dependent spatial biases seen in spatial recall tasks, rather than producing larger overall recall errors. These findings are consistent with a role for attention in modulating the processes underlying maintenance in SWM, rather than with attention as *the* rehearsal mechanism. In the following section, we describe a dynamic field model that accounts for these findings. In the model, the attention effects evident in our experiment arise as a result of reciprocal coupling between the three-layered neural field model described previously and a topographically organized neural field that represents the color and spatial location of the color target and implements a simple form of color recognition. Additionally, we show that weak, localized excitatory feedback from SWM to this color-space field is sufficient to capture Awh et al's (1998) finding of improved discrimination performance at memorized versus non-memorized locations.

Model Simulations

The pattern of results described above cannot be explained simply as a result of stimulus input at the location of the color target. Although color inputs alone did exert a small attractive influence on memory for spatial location in the shifting-control condition, errors in the shifting-attention condition were nearly three times as large. Can this finding be captured by the DFT?

Recall from our discussion of drift away from reference axes that peaks in WM tend to drift along activation gradients, moving from regions of relatively strong inhibition to areas of weaker inhibition until activation becomes balanced on either side of the peak. In most cases, this property ensures that the peak of activation is centered over the region of maximal stimulation, which most likely reflects the correct position of the stimulus input (a formal analysis of this property of neural fields can be found in Amari, 1977). However, in the presence of other inputs (e.g., from reference frames, LTM, or other stimuli), this can lead to systematic distortions, such as drift away from reference frames (see Figure 1) or, in some cases, drift toward frequently-viewed targets (see, e.g., Hund & Spencer, 2003; Spencer & Hund, 2002; for recent extensions to working memory for color, see Johnson, Spencer, & Schöner, in press). The latter finding can be accounted for if we assume that LTM traces are fed directly into the SWM layer as weak excitatory inputs. In this case, activation peaks in WM move along an excitatory gradient towards the most strongly activated region (i.e., towards the memory trace; see Spencer et al., 2007).

Attraction towards the focus of attention found in the present experiment can be captured in a similar way. Specifically, differences between the shifting-attention and control conditions can be accounted for if we assume that, in addition to requiring stimulus input, performance of the color discrimination task involves the short-term maintenance of information in a spatially organized color system (i.e., in a neural field that represents both the color and spatial location of the target), and if this system also serves as a source of location-specific excitatory input to SWM. This is a reasonable assumption given that many ventral visual cortical areas with cells that respond selectively to color are also spatially organized (see, e.g., Bartels & Zeki, 2000; Xiao, Wang, & Felleman, 2003), and send projections to dorsal stream areas that have been

implicated in WM for spatial information (Glaeys et al., 2004). Thus, in this framework, the small modulation of peak drift found in the shifting-control condition arises as a result of stimulus input alone, whereas the more substantial modulations seen in the shifting-attention condition are produced by stimulus input together with short-term maintenance while making the color-discrimination decision.

Methods

Model architecture. The model used to simulate the results from Experiment 1 consisted of three components: The neural field model of SWM described previously; a single Color-Space field (CS) that detects the presence of particular colors at particular locations in the task space; and, lastly, bi-directional excitatory coupling between SWM and the CS field. The SWM model consists of three layers (two excitatory layers coupled to a single inhibitory layer; see Figure 1), which, for these simulations, each contained 181 bistable neurons that respond preferentially to targets at particular locations in the task space. Interactions between the SWM and Inhib layers makes it possible for sustained peaks of activation reflecting the location of a target stimulus to be maintained in SWM in the absence of continuing input (i.e., once the stimulus input has been turned off). Perceptual inputs to the three-layer model took the form of a Gaussian centered at a particular location in the PF field and having a particular strength and width, whereas projections among the different layers (e.g., from PF to Inhib and/or SWM) were defined by the convolution of a Gaussian kernel with a sigmoidal threshold function. The threshold function determines the degree to which neurons close to threshold (i.e., 0) contribute to the activation dynamics (see Appendix A and Table 1 for detailed discussion of model equations and parameters). For the present simulations, the threshold function was set such that only those neurons very near to or above threshold made a contribution (i.e., transmitted activation to other neurons).

The CS field consisted of a 181×181 topographic array of neurons. Neurons in this field respond preferentially to specific colors appearing at particular locations in the task space. Thus, whereas one neuron might respond to red targets appearing at 40° in the task space, another nearby neuron might respond to red targets appearing at 45° , and so forth. Spatially

organized populations of feature-selective neurons are known to exist in many ventral visual areas (see, e.g., Bartels & Zeki, 2000; Rolls & Deco, 2002; Xiao et al., 2003). However, as one progresses through the ventral visual stream, the receptive fields of neurons become much larger, with an accompanying decrease in spatial resolution (Desimone & Gross, 1979; Gross, Rocha-Miranda, & Bender, 1972). This property, together with the distributed coding of object properties in later visual areas, can make it difficult to determine which features belong together as attributes of a single object when multiple objects are present in the visual field. This is an example of the well known *binding problem* in vision (Damasio, 1989; Luck & Beach, 1998; von der Malsburg, 1995). A DFT approach to the binding problem has been proposed in Johnson, Spencer, and Schöner (in press).

As with perceptual inputs to PF, inputs to the CS field took the form of a Gaussian, but in two dimensions (i.e., color and space), rather than one. Although the width of the Gaussian can be set to different values for each dimension (e.g., narrow in feature and broad in space), for simplicity they have been set to the same value here. To implement decision making in the CS field, neurons responding to the colors used in the color-discrimination task were provided with weak subthreshold “preshape” inputs, a form of feature-based attention. Unlike perceptual inputs, preshape inputs take the form of a ridge of activation that is localized along the color dimension but not the spatial dimension. That is, preshape input provides low levels of excitatory input to neurons coding for particular colors (e.g., the two possible target colors used in the experiment) across all spatial locations. However, such inputs are typically much weaker than direct stimulus inputs and cannot by themselves induce a localized peak of activation. Additionally, the CS field was tuned so that above-threshold peaks of activation were only possible when a color matching one of the preshaped values (i.e., one of the possible color targets) was presented to the model. This prevented the formation of above-threshold peaks in the CS field in response to the memory target (which was a different color than the colors used in the color discrimination task), and to color targets in the control conditions, where the color stimulus was to be ignored.

The final component of the model consisted of reciprocal excitatory coupling between the SWM layer of the three-layer model and the spatial dimension of the CS field. As with projections among the separate layers of the SWM model, projections between CS and SWM were defined by the convolution of a Gaussian kernel with a sigmoidal threshold function. Thus, whenever an above-threshold peak was present in CS, weak excitatory input was transmitted to neurons tuned to the same location in the SWM field. Similarly, an above-threshold peak in SWM provided weak excitation to neurons coding the same location in the CS field. Like preshape input to the color dimension of the color-space field, this input was propagated across all neurons tuned to the stimulated location, regardless of their color-specific tuning. Thus, this input produced a “ridge” of activation centered at the memorized location in the CS field, effectively priming the CS field to respond to targets appearing at those locations, regardless of their color.

Simulation details and example simulation. Simulations were conducted in Matlab 7.4 (Mathworks, Inc.) on a PC with an AMD Athlon 2.61 GHz dual-core processor (the Matlab code is available from the authors on request). The dynamic field equation (model equations and parameters are given in Appendix A and Table 1, respectively) was integrated using the Euler method with one time step = 2 ms. This time step was sufficiently small so that numerical solutions approximated real solutions of the dynamical system given a target presentation of 2 s and a delay interval of 10 s. For the present simulations, our goal was to capture mean results from the shifting attention and shifting control conditions of Experiment 1 to investigate the mechanism that underlies the influence of spatial selective attention on SWM.

An example simulation is shown in Figure 4, which depicts a single trial in the shifting attention condition. Figure 4A-B show the PF and SWM layers of the three-layer model, with spatial location along the y -axis, and time, from the beginning to the end of the trial, running from left to right. Figure 4C-F show the CS field at different points in time (indicated by brackets), with spatial location plotted along the y -axis and color along the x -axis of each panel. Each simulation trial began with a 200-ms relaxation period to allow the model to reach a stable

resting state. This was followed by the 2-s presentation of a memory stimulus (see Fig. 4A), accompanied by an additional, weaker, reference input to PF that captured perception of the midline of the task space (the midline input to PF stayed on throughout the remainder of the simulated trial). In addition, neurons in the CS field representing the two possible target colors in the color-discrimination task received weak preshape inputs that moved neurons tuned to these colors closer to threshold (see vertical dark grey lines in Fig. 4C).

Presentation of the memory target was followed by a 10-s memory delay, during which a self-sustained peak of activation, reflecting the spatial location of the memory stimulus, was maintained in SWM (see horizontal dark grey line in Fig. 4B). Next, one of two possible color stimuli was presented to both PF and CS for 500 ms (i.e., “RED 2”; see Fig. 4E). Note that for simplicity, the color stimulus was always presented midway through the memory delay, rather than at variable times within a given time interval following memory stimulus offset, as in the experiment. The preshape input to CS was turned off 500 ms after color stimulus offset, which allowed the peak in the CS field to decay to resting level after the color-discrimination response was made. Thus, during the delay interval following color-stimulus offset, the only activation present in the CS field reflects excitatory feedback from SWM (see Fig. 4F). Finally, spatial recall errors were determined by computing the difference between the location of maximum activation in the SWM field at the end of the delay and the location of the original memory stimulus.

Results and Discussion

Simulation results for the shifting attention and control conditions are shown in Figure 3D. Recall that for the control conditions, the three-layer model received weak, localized input reflecting the position of the color stimulus, but no additional input from the CS field (because no above-threshold peak was present in the CS field following color-stimulus onset). This produced a small increase in drift when the color stimulus appeared away versus toward midline, in keeping with our experimental findings. Note, however, that errors to the $\pm 60^\circ$ targets were somewhat reduced compared to our experimental results, whereas errors to the $\pm 30^\circ$ targets were

somewhat elevated overall. This is due in part to the fact that we found higher errors than usual to the 60° targets in the control conditions of the experiment (see Figure 3C and Results above). It is not clear what caused this pattern in the experimental results.

In the shifting attention simulations shown in Figure 4, where the memory stimulus was presented at 30° and the color stimulus was presented at 40° (i.e., away from midline relative to the memory stimulus), SWM received weak localized input from the CS field in addition to the direct stimulus input it received in the shifting control simulations. Additional excitatory input at 40° produced a 3.0° increase in drift compared to when the color stimulus was presented at 20° (i.e., toward midline), in keeping with the empirical results. Similar modulations of drift were seen in simulations probing memory for stimuli at 60° (see Figure 3D). Thus, these simulations effectively capture the key pattern of errors seen in the shifting attention conditions reported here.

Finally, recall that Awh et al. (1998, Experiment 1) also found that form discrimination performance was faster when discrimination targets appeared at memorized versus non-memorized locations. Participants in this experiment completed a same/different location memory task together with a form discrimination task during the delay interval. Specifically, participants remembered the location of a lowercase letter (e.g., the letter 'e') across a 5-s delay interval, and indicated whether a second letter (e.g., the uppercase letter 'E') was at the same (50% of trials) or at a different location. During the delay interval (either 1500, 2500, or 3500 ms after memory stimulus offset), a discrimination target was presented and participants made a speeded response, indicating which of two possible targets was present. The discrimination target could either appear at the location that was being held in memory, or at another non-memorized location. The key finding was a reaction time (RT) advantage when the discrimination target appeared at the memorized versus the non-memorized location. Because similar performance advantages are seen when participants are instructed to attend to a location where a target subsequently appears, the authors took this as evidence that attention remained at the memorized location during the delay.

RT results and example simulations illustrating this finding can be seen in Figure 5. Figure 5A-B show the PF and SWM layers of the three-layer model of SWM, whereas the lower panels show two different snapshots of the CS field (Fig. 5C-D) and associated RT results (Fig. 5E). For these simulations, the model was presented with a single memory stimulus at a given location for 2 s, followed by a 10 s delay interval, as in the simulations described aboveⁱⁱⁱ. During the delay interval, a single discrimination target (e.g., a particular color or form target) was presented at one of the preshaped feature values in the CS field for 500 ms. In the first simulation, the target was presented at the same location that was being maintained in SWM (Fig. 5C), whereas in the second simulation it was presented at a non-memorized location (Fig. 5D). Results revealed that activation in the CS field reached threshold (i.e., 0) faster when the discrimination target appeared at the memorized versus the non-memorized location (bottom right), in keeping with experimental findings.

The simulation results reported here demonstrate that the basic recall effects captured by the DFT (i.e., drift away from midline), together with extra input from a Color-Space field during the delay is sufficient to capture the results reported here. Specifically, we were able to successfully model attraction towards the focus of attention found in the shifting attention and control conditions. Additionally, because SWM and CS are reciprocally coupled in the model, we were also able to capture the finding of an RT advantage at memorized versus non-memorized locations reported by Awh et al. (1998, Experiment 1). Can our model also account for the decline in location memory found in Experiment 3 of Awh et al.?

One of the key differences between the experiment reported here and the studies of Awh et al. (1998) is that they used a sequential same/different location memory task. In contrast to the spatial recall task, this task requires the participant to remember a target location and to determine whether a second target presented a short time later is at the same or a different location. Clearly, systematic distortions in location memory, including the drift effects evident in spatial recall tasks, could also contribute to errors in position discrimination tasks. Such errors would be likely to increase in the presence of an extra source of peak motion, as we have

postulated occurs when attention is shifted to a nearby non-memorized location during the delay. Is there any evidence suggesting that the peak motion postulated to underlie errors in spatial recall also contributes to errors in same/different location memory tasks?

Simmering et al. (2006, Experiments 5-6) conducted a series of same/different position discrimination experiments that addressed this issue. In their experiments, participants viewed a small dot at a given location either to the right or the left of the midline of the task space, and indicated whether a second dot presented a short time later was at the same or at a different location. When the second dot was at a different location, its position could change either in the direction of midline (i.e., the direction opposite the expected direction of drift), or away from midline relative to the original location (i.e., in the same direction as drift). They hypothesized that if participants' location memory begins drifting away from midline following stimulus offset, it should be easier for them to detect position changes at test when they occur in the direction opposite to drift (i.e., when the change in position is *toward* midline) – because the distance between the test location and the location being held in memory will be larger here than when the change occurs in the same direction the memory is drifting. This is precisely what they found: Participants' discrimination thresholds were much higher when the test stimulus was presented away from the midline reference.

These findings confirm that peak drift occurs in position discrimination tasks using relatively short delays (500 ms), and can have a measurable effect on performance. Thus, it seems reasonable to propose that drift could also have an effect on performance in same/different location memory tasks using longer delays. However, to account for the decrease in performance in the shifting attention condition of Awh et al. (1998, Experiment 3), location changes at test would need to occur slightly more often in the direction of drift than in the opposite direction. Unfortunately, the randomization scheme used in their experiments makes it difficult to assess this possibility. Recall, however, that declines in performance in the shifting attention task of Awh et al. were only 5% greater than those seen in the static attention task (12% versus 7%, respectively, compared to control conditions where the color task was not completed during the

delay). As a result, peak motion in the shifting attention condition would need to produce errors on as few as six additional trials (i.e., 5% out of 120 total trials) to account for the observed pattern of results. Thus, increased peak motion represents a reasonable candidate source of increased error in Experiment 3 of Awh et al.

General Discussion

The purpose of the present study was to explore the relationship between spatial selective attention and spatial working memory using a combination of behavioral experimentation and formal modeling. To this end, we conducted a single experiment consisting of a color-discrimination task inserted during the delay interval of a spatial recall task. The critical manipulation was the location of the color stimulus relative to the memorized location. Specifically, in a shifting attention condition, the color stimulus could appear either *toward* or *away from* the memorized location relative to a midline reference axis. We hypothesized that if spatial attention serves as the rehearsal mechanism in SWM, shifting attention away from the location being held in memory should increase errors relative to a condition where attention remains at the memorized location, regardless of the direction of the attention shift. Conversely, if attention serves to modulate the patterns of activation in SWM, we expected to find systematic changes in the pattern of errors in the SWM task depending on the direction of the attention shift. Our results provided support for the latter proposal. Specifically, drift away from midline was found to be nearly 3.0° larger when attention was shifted away from versus toward midline, compared to a 1.0° difference in a control condition where attention did not need to be shifted away from the memorized location. We were able to capture these findings using our neural field model of SWM (see Figure 1) reciprocally coupled to a single Color-Space field, which implements a simple form of color recognition.

Multiple overlapping mechanisms of attention and working memory in the DFT

So what does the model described here have to say about the proposal that attention is the rehearsal mechanism for SWM, and about the relationship between spatial attention and SWM more generally? Clearly, the model is at odds with the proposal that spatial attention is *the*

rehearsal mechanism in SWM. Recall that maintenance (i.e., self-sustained activation) in the model emerges through interactions between excitatory and inhibitory neural populations that respond selectively to spatial location. Although the stability and precision of maintained patterns of activation can be influenced by additional factors (e.g., the presence of reference frames, inputs from LTM, and so forth), they do not strictly depend on them. Nonetheless, the model proposed here is generally consistent with the proposal of Awh and colleagues regarding mechanisms of attention and working memory.

Specifically, we are in agreement with the proposal that attention and working memory share overlapping mechanisms. Indeed, in many cases, it can be difficult to determine where working memory ends and attention begins. For example, the local excitation / lateral inhibition function that underlies maintenance in the model also implements one of the core features of selective attention: the selection of one or more targets or locations out of potentially many possibilities. Specifically, at the limit, when lateral inhibition is broad and relatively strong, the model functions as a winner-take-all network, building stable peaks of activation in response to only the most salient target inputs, and suppressing activation associated with other targets. When the model is functioning in this regime, the relative salience of a given target is determined by a combination of its visible properties (e.g., how similar it is to the other targets in the array), and top-down factors reflecting the relevance of the target or one or more of its features to the task at hand, as in other prominent models of attention (Deco & Rolls, 2004; Mozer & Sitton, 1998; Wolfe, 1994). Thus, prior to peak formation, the interaction function at work in the model can be described as serving an attentional function, whereas once a stable peak of activation is formed, the same interaction function implements sustained activation, which is a core feature of working memory.

Another form of attention in the model, which can also influence maintenance in WM, is preshape input to the CS field in the shifting and static attention conditions. Such inputs capture the requirements of the color-discrimination task that participants completed during the delay interval of the spatial recall task. Specifically, this input produced a spatially non-specific

baseline shift in the firing rates of neurons in the CS field tuned to the colors used in the task, in keeping with studies of feature-based attention (see, e.g., Liu, Slotnick, Serences, & Yantis, 2003). This allowed a weakly self-sustaining peak of activation to be created following stimulus presentation and maintained in the CS field long enough for a response to be generated. Thus, in addition to ensuring that an above threshold peak of activation was present in the CS field only when a task-relevant color was presented to the model, preshape inputs also made it possible for the color and location of the color stimulus to be retained for a brief time in memory. Although the source of such “preshaping” inputs are not specified in the model, a reasonable possibility is that they reflect top-down inputs from another population of neurons responsible for maintaining task-specific information in an active state as proposed by Liu et al. (2003) and captured in several other models of attention (see, e.g., Desimone & Duncan, 1995; Duncan & Humphreys, 1989; Mozer & Sitton, 1998). Interestingly, this simple form of feature-based attention, rather than a specifically spatial form of attention, plays a key role in allowing us to capture the pattern of *spatial biases* reported here.

This brings us to another closely related form of attention at work in the model: Modulations of patterns of activation in one population of neurons via inputs from another population to which it is coupled. For example, modulations of peak drift in the model reported here arise as a result of input from a spatially organized population of neurons that maintain color information for use in the color-discrimination task. Similarly, excitatory feedback from SWM to this field allowed us to capture the RT advantages reported in the experiments of Awh et al. (1998). Thus, the key pattern of results here and in the experiments of Awh and colleagues can be viewed as emerging from the reciprocal interaction of ventral and dorsal stream neural populations that specialize in primarily visual and primarily spatial forms of information, respectively (Goodale & Milner, 1992; Ungerleider & Mishkin, 1982). Although semi independent, evidence of substantial crosstalk between these systems has been reported. For example, Glaeys and colleagues (2004) have shown that performance of color-discrimination tasks involves interactions between ventral stream areas that respond selectively to color and

dorsal stream areas that may play a role in selecting an appropriate response. Additionally, Zhong and Rockland (2003) have reported evidence of spatially-specific projections from regions of the parietal cortex to anterior portions of area TE in the macaque, a canonical ventral stream area. Moreover, inputs from oculomotor structures such as the frontal eye fields and portions of the parietal cortex have been proposed as likely sources of spatial reentry signals that could serve as attentional feedback to ventral stream areas such as V4 (Hamker, 2005).

Finally, a variant of attention can also be realized through global modulations of the resting level of neurons within a field. For instance, in the color discrimination task reported here, the resting level in the CS field was set so that self-sustained peaks were only possible when a color matching one of the color targets was present in the task space (i.e., a color appeared at a preshaped location in the field). This prevented the model from building a peak in the CS field in response to the blue “spaceship” memory target, or in response to the color stimuli in each of the control conditions (i.e., to task-irrelevant stimuli). More generally, modulations of the resting level within a given field can make the formation of peaks either more or less likely, and, once a peak has formed, either easier or harder to stabilize (depending on the direction of the modulation). Clearly, therefore, global modulations of resting level can also play a role in the creation and maintenance of self-sustained patterns of activation in working memory—a key function of attention.

The different varieties of attention discussed above lead naturally to a critical question: where precisely is spatial attention in our model? The answer is that spatial attention is not isolated to one mechanism; rather varieties of spatial attention are realized in all of the ways discussed above. This differs from the common conception of spatial attention as a type of mental spotlight that can be focused on different locations more or less independently of the objects occupying those locations (for critiques of this conceptualization of attention, see Driver & Baylis, 1989; Huang & Pashler, 2007; Laberge & Brown, 1989; Laberge, Carlson, Williams, & Bunney, 1997). Given that the attentional effects captured by our model emerged as the model processed localized inputs, this common perspective of attention raises two key issues. First, is

there some type of spotlight in our model? And second, can our model (or people) shift spatial attention to specific locations in the absence of localized input?

With regard to the first question, simulations of Experiment 1 of Awh et al. (1998) illustrate one type of spotlight effect in our model. In this task, attention was initially cued to a particular location by the appearance of a stimulus at a to-be-remembered location. This was followed by a form discrimination decision which was enhanced at the remembered location. Thus, one could say that the spotlight of attention was fixed at the remembered location. Consistent with this, our model showed faster form discrimination when targets were presented at the to-be-remembered location. Analysis of our model, however, shows that such effects do not require attention per se: holding a location in SWM effectively modulates activation in early visual areas, giving a competitive advantage to targets appearing at the memorized location. This is consistent with behavioral and neuroimaging studies of working memory and attention (Awh et al., 2000; Awh et al., 1999; Hillyard & Anllo-Vento, 1998; Hillyard et al., 1998; Postle et al., 1999). Thus, while our model shows spotlight-like effects, such effects do not necessarily reflect explicitly “attentional” mechanisms.

Do such effects always require a localized cue? Several studies demonstrate that people can shift attention to specific locations using *central cues*, such as an arrow or other symbol presented at fixation (see review in, Luck & Vecera, 2002; and Posner, 1980, for a classic example). In most studies of this kind, however, the to-be-attended locations are marked by, for example, outlined squares. Although it seems intuitive that attention can be shifted with some precision to locations in otherwise empty space, we are not aware of any study demonstrating this ability. It is possible in our model to build activation peaks in SWM using long-term memory mechanisms (see Spencer et al., 2007), and such peaks would give rise to the types of spotlight-like enhancement effects shown in Figure 5. This is one possible way that our model could handle results from studies using central cues and unmarked locations. That said, we suspect that people’s ability to stably anchor a spotlight-like focus of attention in empty space is quite limited. If this is indeed the case, it would be consistent with a recent study from our lab

(Simmering & Spencer, 2007) showing that people have very limited ability to mentally impose arbitrary spatial frames of reference in empty space, even for a single trial following several trials with the target frame visible and when the imposition of the frame is consistent with a conceptual context.

Relationship of the DFT to other models of attention and working memory

The simple neural field model reported here shares many features in common with other prominent models of attention and working memory. For example, similar to the neurodynamical models proposed by Deco and Rolls (2004) and Hamker (2005), there is no dedicated “attention field” or stack of fields that uniquely instantiate spatial attention. Rather, attention-like phenomena arise as emergent products of reciprocal coupling among neural populations specialized for different functions, with inputs from one neural population modulating the local intracortical dynamics of the other (see also, Desimone & Duncan, 1995; Duncan, Humphreys, & Ward, 1997; and discussion of attention in Fuster, 2003). For example, in the model of Hamker (2005), spatially specific activation in oculomotor structures such as the frontal eye fields (FEF), which reflects the locations of potential saccade targets, serves as an attentional re-entry signal to spatially organized ventral stream areas such as V4. Top down inputs from the ventrolateral portion of the prefrontal cortex, which is implicated in SWM, can influence saccade plans in the FEF, allowing a spatial focus of attention reflecting the likely location of a subsequent stimulus to affect early visual processing. Similarly, in the model of Deco and Rolls (2004), feedback from prefrontal areas implicated in working memory for stimulus features (e.g., color, shape, etc.) to spatially organized populations of neurons in the ventral stream (e.g., V4) allow neural resources to be focused on locations containing those features. The model proposed here is clearly in line with these neurodynamical approaches to attention and working memory.

Models such as these have been contrasted with the “classical” view of attention, wherein an initial parallel stage of processing is followed by a serial, focal attention stage that results in the serial selection of targets mediated by spatial location (see, e.g., Neisser, 1967; Treisman & Gelade, 1980; Wolfe, 1994). Spatial attention is typically implemented in such models as a

dedicated “saliency map” or “attentional map”, which tags specific locations in the visual field as noteworthy on the basis of top down and bottom up information about the stimuli present (for a formal implementation of the saliency map idea, see Itti & Koch, 2001; Mozer & Sitton, 1998). Although no such dedicated attentional map need be invoked to explain the present findings, something similar could be implemented in our model through the inclusion of an attention or fixation system that gates inputs to the perceptual layer of the 3-layer working memory model that are outside the currently fixated (or attended) region of the visual field, similar to the model of Hamker discussed above. Indeed, Kopecz and Schöner (1995) have developed a simple model of saccade planning using dynamic neural fields that could be easily integrated with the model described here (for a more recent version aimed at accounting for looking behavior in infant habituation tasks, see Golberg & Schöner, 2007, August).

One of the key functions of attention in classical models of attention is to bind the various features present at a given location together into a coherent object representation. In another line of work, we have proposed a similar role for SWM (Johnson, Spencer et al., in press). Specifically, we have suggested that spatially-precise input from SWM to coarsely spatially organized neural populations coding for object properties can help to resolve ambiguous spatial coding in the ventral stream, clarifying which features belong together as parts of the same object. This is similar to the proposed role of attention in the *ambiguity resolution theory* of Luck and colleagues (Luck, Girelli, McDermott, & Ford, 1997), and the *location uncertainty theory* proposed by Ashby and colleagues (Ashby, Prinzmetal, Ivry, & Maddox, 1996).

Finally, with regards to working memory, there is general agreement that some form of sustained activation is the most plausible neuronal substrate for short-term spatial memory. This substrate involves graded, metric representations that evolve continuously in time under the influence of current sensory information as well as the current activation state. Exactly how sustained activation is neurally realized, however, is not clear. The DFT is in a class of neural network models that achieves a stable memory state using bi-stable networks in which a stable state of sustained activation coexists with an “off” state (Amari, 1977; Amari & Arbib, 1977;

Compte et al., 2000). However, most other models in this class have focused on the development of biophysically realistic models of single neurons and populations of neurons that capture key elements of neuronal firing rates, neurotransmitter action, and so on. Although this approach has led to new insights into brain function and neuronal dynamics, such models typically fail to capture many behavioral details. In contrast, with the DFT we have sacrificed some neural realism, but have forged a tight relationship between neurally-inspired modeling on one hand and rigorous behavioral experimentation on the other. A hallmark of this approach is the search for behavioral signatures of proposed underlying neural dynamics. The robust biases seen in spatial recall experiments represent a particularly fruitful example of this approach.

Conclusions

In the present study, we used a combination of behavioral and formal modeling to investigate interactions between spatial attention and spatial working memory. Our investigations suggested the presence of multiple overlapping mechanisms of attention and working memory. Although each of the mechanisms discussed above play a role in generating the phenomena that have been traditionally identified as reflecting the operation of attention or working memory, importantly, none of these map directly onto these psychological constructs in a straightforward manner. This highlights the need to move the investigation of these constructs to a more formal level. In the present report, we took an important first step in this regard, moving in the direction of neurocomputational modeling which allows us to investigate neural mechanisms operating at the population level while retaining strong links to behavioral experimentation.

References

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*, 77-87.
- Amari, S., & Arbib, M. A. (1977). Competition and cooperation in neural nets. In J. Metzler (Ed.), *Systems Neuroscience* (pp. 119-165). New York: Academic Press.
- Ashby, F. G., Prinzmetal, W., Ivry, R., & Maddox, W. T. (1996). A formal theory of feature binding in object perception. *Psychological Review*, *103*, 165-192.
- Awh, E., Anllo-Vento, L., & Hillyard, S. A. (2000). The role of spatial selective attention in working memory for locations: Evidence from event-related potentials. *Journal of Cognitive Neuroscience*, *12*, 840-847.
- Awh, E., & Jonides, J. (1998). Spatial working memory and spatial selective attention. In R. Parasuraman (Ed.), *The Attentive Brain* (pp. 71-94). Cambridge, MA: MIT Press.
- Awh, E., Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 780-790.
- Awh, E., Jonides, J., Smith, E. E., Buxton, R. B., Frank, L. R., Love, T., et al. (1999). Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science*, *10*, 433-437.
- Awh, E., Vogel, E. K., & Oh, S.-H. (2006). Interactions between attention and working memory. *Neuroscience*, *139*, 201-208.
- Bartels, A., & Zeki, S. (2000). The architecture of the colour centre in the human visual brain: new results and a review. *European Journal of Neuroscience*, *12*, 1172-1193.
- Churchland, P. S., & Sejnowski, T. J. (1992). *The Computational Brain*. Cambridge: Bradford Book/MIT Press.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*, 910-923.
- Cowan, N., Elliott, E. M., Saults, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., et al. (2005).

- On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, *51*, 42-100.
- Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, *1*, 123-132.
- Deco, G., & Rolls, E. T. (2004). A neurodynamical model of visual attention and invariant object recognition. *Vision Research*, *44*, 621-642.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193-222.
- Desimone, R., & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. *Brain Research*, *178*, 363-380.
- Driver, J., & Baylis, G. C. (1989). Movement and visual attention: the spotlight metaphor breaks down. *Journal of Experimental Psychology: Human Perception & Performance*, *15*, 448-456.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433-458.
- Duncan, J., Humphreys, G. W., & Ward, R. (1997). Competitive brain activity in visual attention. *Current Opinion in Neurobiology*, *7*, 255-261.
- Engelbreton, P. H., & Huttenlocher, J. (1996). Bias in spatial location due to categorization: Comment on Tversky and Schiano. *Journal of Experimental Psychology: General*, *125*(1), 96-108.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545-572.
- Fuster, J. M. (2003). *Cortex and Mind: Unifying Cognition*. Oxford: Oxford University Press.
- Glaeys, K. G., Dupont, P., Cornette, L., Sunaert, S., Van Hecke, P., De Schutter, E., et al. (2004). Color discrimination involves ventral and dorsal stream areas. *Cerebral Cortex*, *14*(7), 803-822.
- Golberg, J., & Schöner, G. (2007, August). Understanding the distribution of infant attention: A

- dynamical systems approach. In D. S. McNamara & J. G. Trafton (Eds.), *Poster presented at the 29th Annual Cognitive Science Society*. Nashville, TN.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience*, *15*(1), 20-25.
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, *35*, 96-111.
- Hamker, F. H. (2005). The reentry hypothesis: the putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4 and IT for attention and eye movements. *Cerebral Cortex*, *15*, 431-447.
- Hillyard, S. A., & Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences, USA*, *95*, 781-787.
- Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: Electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *353*, 1257-1270.
- Hollingworth, A. (in press). Memory for real-world scenes. In J. R. Brockmole (Ed.), *Memory for the Visual World*. London: Psychology Press.
- Huang, L., & Pashler, H. (2007). A boolean map theory of visual attention. *Psychological Review*, *114*(3), 599-631.
- Hund, A. M., & Spencer, J. P. (2003). Developmental changes in the relative weighting of geometric and experience-dependent location cues. *Journal of Cognition and Development*, *4*, 3-38.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352-376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgement? *Journal of Experimental Psychology: General*, *129*, 220-241.
- Hyun, J.-S., & Luck, S. J. (in press). Visual working memory as the substrate for mental rotation. *Psychonomic Bulletin & Review*.

- Itti, L., & Koch, C. (2001). Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2, 194-203.
- Johnson, J. S., Liu, L., Thomas, G., & Spencer, J. P. (in press). Calibration algorithm for eye tracking with unrestricted head movements. *Behavior Research Methods*.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (in press). Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. In F. Garzón, A. Laakso & T. Gomila (Eds.), *Dynamics and Psychology [special issue]. New Ideas in Psychology*.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 98, 122-149.
- Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural, dynamic fields. *Biological Cybernetics*, 73, 49-60.
- Laberge, D., & Brown, V. (1989). Theory of attentional operations in shape identification. *Psychological Review*, 96, 101-124.
- Laberge, D., Carlson, R. L., Williams, J. K., & Bunney, B. G. (1997). Shifting attention in visual space: tests of moving-spotlight models versus an activity-distribution model. *Journal of Experimental Psychology: Human Perception & Performance*, 23, 1380-1392.
- Liu, T., Slotnick, S. D., Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex*, 13, 1334-1343.
- Luck, S. J., & Beach, N. J. (1998). Visual attention and the binding problem: A neurophysiological perspective. In R. D. Wright (Ed.), *Visual Attention* (pp. 455-478). New York: Oxford University Press.
- Luck, S. J., Girelli, M., McDermott, M. T., & Ford, M. A. (1997). Bridging the gap between monkey neurophysiology and human perception: An ambiguity resolution theory of visual selective attention. *Cognitive Psychology*, 33, 64-87.
- Luck, S. J., & Vecera, S. P. (2002). Attention. In H. P. S. Ed.) & S. Y. V. Ed.) (Eds.), *Steven's Handbook of Experimental Psychology: Vol. 1. Sensation and Perception* (3rd ed. ed., Vol. 1, pp. 235-286). New York: Wiley.

- Mozer, M. C., & Sitton, M. (1998). Computational modeling of spatial attention. In *Attention* (pp. 341-393). London: UCL Press.
- Neisser, U. (1967). *Cognitive Psychology*. New York: Appleton-Century-Crofts.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3-25.
- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *130*, 23-38.
- Postle, B. R., Awh, E., Jonides, J., Smith, E. E., & D'Esposito, M. (1999). The where and how of attention-based rehearsal in spatial working memory. *Cognitive Brain Research*, *20*, 194-205.
- Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.
- Schiano, D. J., & Tversky, B. (1992). Structure and strategy in encoding simplified graphs. *Memory and Cognition*, *20*(1), 12-20.
- Schutte, A. R., & Spencer, J. P. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, *74*(5), 1393-1417.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, *74*(5), 1393-1417.
- Simmering, V. R., & Spencer, J. P. (2007). Carving up space at imaginary joints: Can people mentally impose spatial category boundaries? *Journal of Experimental Psychology: Human Perception & Performance*, *33*, 871-894.
- Simmering, V. R., Spencer, J. P., & Schöner, G. (2006). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception & Psychophysics*, *68*, 1027-1046.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Spatial categories are formed

- using geometric and experience-dependent information *Journal of Experimental Psychology: General*, *131*, 16-37.
- Spencer, J. P., Simmering, V. R., Schutte, A. R., & Schöner, G. (2007). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In J. M. Plumert & J. P. Spencer (Eds.), *Emerging landscapes of mind: Mapping the nature of change in spatial cognitive development*. New York, NY: Oxford University Press.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97-136.
- Tversky, B., & Schiano, D. J. (1989). Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology: General*, *118*(4), 387-398.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale & R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior* (pp. 549-586). Cambridge, MA: MIT Press.
- von der Malsburg, C. (1995). Binding in models of perception and brain function. *Current Opinion in Neurobiology*, *5*, 520-526.
- Werner, S., & Schmidt, T. (2000). Investigating spatial reference systems through memory distortions. In C. Freska, W. Brauer, C. Habel & K. F. Wender (Eds.), *Spatial Cognition II - Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications* (Vol. 1849, pp. 169-183). Berlin, Germany: Springer.
- Williams, R. J. (1986). The logic of activation functions. In D. E. Rumelhart, J. L. McClelland & t. P. R. Group (Eds.), *Parallel distributed processing, Volume 1*. (pp. 423-443). Cambridge: MIT Press.
- Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*, 1-24.
- Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, *1*(2), 202-238.

Xiao, Y.-p., Wang, Y., & Felleman, D. J. (2003). A spatially organized representation of colour in macaque cortical area V2. *Nature*, *421*, 532-539.

Zhong, Y.-M., & Rockland, K. S. (2003). Inferior parietal lobule projections to anterior inferotemporal cortex (areaTE) in macaque monkey. *Cerebral Cortex*, *13*, 527-540.

Author Note

Jeffrey S. Johnson Department of Psychology, University of Iowa. John P. Spencer Department of Psychology, University of Iowa, and Iowa Center for Developmental and Learning Sciences, University of Iowa. This research was made possible by National Institute of Health Grant RO1 MH62480, and National Science Foundation Grant BCS 00-91757 and HSD 0527698 awarded to J.P.S. Portions of this research were presented at the 29th Annual Meeting of the Cognitive Science Society, Nashville, Tennessee, August 2007. Correspondence concerning this article should be addressed to Jeffrey S. Johnson, Department of Psychology, University of Iowa, 11 Seashore Hall E, Iowa City, IA 52242-1407, 319-353-2990 (phone), jeffrey-johnson-2@uiowa.edu (e-mail).

Footnotes

ⁱ Visual angles were estimated using the distance of the stimulus from the front edge of the table and assuming an eye-height of 42 cm above the table surface. These values are approximately correct when head position is fixed using a chin rest attached to the front edge of the table.

However, because head position was not fixed during the experimental session, the exact distance of each participant from the location of the stimuli on the table and the height of the eye above the table likely differed across participants as well as for the same participant at different points during the experimental session.

ⁱⁱ Because the display surface used in the current experiment was a rear projection screen, it was difficult to obtain accurate color coordinates (e.g., CIE x , y , and luminance) using the colorimeter we had available. Therefore, we report RGB values. Additionally, the color values reported for color “blue 2” in the small-colored-circle task are approximate values based on stimuli used during pilot testing. The precise final values were not available due to a hard drive failure.

ⁱⁱⁱ Note that, for simplicity, the midline reference input was left off for these simulations.

Appendix A

Model Equations

In this section, we describe the equations that govern activation in each layer of the 3-layer model and the Color-Space field used in the present report. Although the equations share most components, we describe each separately to highlight the unique contributions to the dynamics of each layer.

Activation in the perceptual field, PF (u), is captured by:

$$\begin{aligned} \tau \dot{u}(x, t) = & -u(x, t) + h_u + \int c_{uu}(x - x') \Lambda_{uu}(u(x', t)) dx' \\ & - \int c_{uv}(x - x') \Lambda_{uv}(v(x', t)) dx' \\ & + s_{ref}(x, t) + s_{tar_{space}}(x, t) + s_{tar_{color}}(x, t) \end{aligned} \quad (1)$$

where $\dot{u}(x, t)$ is the rate of change of the activation level for each neuron across the spatial dimension, x , as a function of time, t . The constant τ determines the time scale of the dynamics (Erlhagen & Schöner, 2002). The first factor that contributes to the rate of change of activation in PF is the current activation in the field, $-u(x, t)$, at each site x . This component is negative so that activation changes in the direction of the resting level h_u .

Next, activation in PF is influenced by the local excitation / lateral inhibition interaction profile, defined by self-excitatory projections, $\int c_{uu}(x - x') \Lambda_{uu}(u(x', t)) dx'$, and inhibitory projections from the Inhibitory layer (Inhib; v), $\int c_{uv}(x - x') \Lambda_{uv}(v(x', t)) dx'$. These projections are defined by the convolution of a Gaussian kernel with a sigmoidal threshold function. In particular, the Gaussian kernel was specified by:

$$c(x - x') = c \exp\left[-\frac{(x - x')^2}{2\sigma^2}\right] - k, \quad (2)$$

with strength, c , width, σ , and resting level, k . The sigmoidal function is given by:

$$\Lambda(u) = \frac{1}{1 + \exp[-\beta u]}, \quad (3)$$

where β is the slope of the sigmoid, that is, the degree to which neurons close to threshold (i.e., 0) contribute to the activation dynamics. Lower slope values permit graded activation near threshold to influence performance, while higher slope values ensure that only above-threshold activation contributes to the activation dynamics. At extreme slope values, the sigmoid function approaches a step function.

Lastly, in the full version of the model, PF receives input from the world via a second perceptual field in a retinal frame of reference (Spencer et al., 2007). For simplicity in the present simulations, we replaced this retinal field with direct input already in an object-centered reference frame. In particular, two inputs passed activation to PF: the reference input (e.g., perception of the midline axis in the recall task), and target input (e.g., the transient presentation of a target object in the task space). All inputs took the form of a Gaussian:

$$S_{tar_{space}}(x, t) = c \exp\left[-\frac{(x - x_{center})^2}{2\sigma^2}\right] \chi(t), \quad S_{tar_{color}}(x, t) = c \exp\left[-\frac{(x - x_{center})^2}{2\sigma^2}\right] \chi(t) \quad (4)$$

centered at x_{center} , with width, σ , and strength, c . These inputs could be turned on and off through time (e.g., the target appears and then disappears). This time interval was specified by the step function $\chi(t)$ (see text for details on the timing of inputs).

The second layer of the model, Inhib (v), is specified by the following equation:

$$\begin{aligned} \tau \dot{v}(x, t) = & -v(x, t) + h_v + \int c_{vu}(x - x') \Lambda_{vu}(u(x', t)) dx' \\ & + \int c_{vw}(x - x') \Lambda_{vw}(w(x', t)) dx'. \end{aligned} \quad (5)$$

As before, $\dot{v}(x, t)$ specifies the rate of change of activation across the population of spatially-tuned neurons, x , as a function of time, t ; the constant τ sets the time scale; $v(x, t)$ captures the current activation of the field; and h_v sets the resting level of neurons in the field. Note that Inhib receives

activation from two projections—one from PF, $\int c_{vu}(x-x')\Lambda_{vu}(u(x',t))dx'$, and one from SWM, $\int c_{vw}(x-x')\Lambda_{vw}(w(x',t))dx'$. As described above, these projections are defined by the convolution of a Gaussian kernel (Equation 2) with a sigmoidal threshold function (Equation 3).

The third layer, SWM (w), is governed by the following equation:

$$\begin{aligned} \tau\dot{w}(x,t) = & -w(x,t) + h_w + \int c_{ww}(x-x')\Lambda_{ww}(w(x',t))dx' \\ & - \int c_{wv}(x-x')\Lambda_{wv}(v(x',t))dx' + \int c_{wu}(x-x')\Lambda_{wu}(u(x',t))dx' \\ & + \int c_{ww_{c,s}}(x-x')\Lambda_{ww_{c,s}}(w_{c,s}(x',t))dx' + c_s s_{ref}(x,t) + c_s s_{tar_{space}}(x,t) + c_s s_{tar_{color}}(x,t). \end{aligned} \quad (6)$$

Again, $\dot{w}(x,t)$ is the rate of change of activation across the population of spatially-tuned neurons, x , as a function of time, t ; the constant τ sets the time scale; $w(x,t)$ captures the current activation of the field; and h_w sets the resting level. SWM receives self excitation, $\int c_{ww}(x-x')\Lambda_{ww}(w(x',t))dx'$, lateral inhibition from Inhib (v), $\int c_{wv}(x-x')\Lambda_{wv}(v(x',t))dx'$, and input from PF(u), $\int c_{wu}(x-x')\Lambda_{wu}(u(x',t))dx'$. SWM also receives input from the spatial dimension of a single Color-Space field ($w_{c,s}$), $\int c_{ww_{c,s}}(x-x')\Lambda_{ww_{c,s}}(w_{c,s}(x',t))dx'$, and direct reference, $S_{ref}(x,t)$, and target, $S_{tar}(x,t)$, inputs scaled by c_s .

Finally, for the Color-Space field ($w_{color,space}$), the basic concepts of the field dynamics and the mechanisms that support the maintenance of localized peaks of activation were generalized to two dimensions (i.e., color and space; see Erlhagen & Schöner, 2002). The dynamical equations governing activation in the Color-Space field have the same general form as Equations 1, 5, and 6 above:

$$\begin{aligned} \tau\dot{w}_{color,space}(x,y,t) = & -w_{c,s}(x,y,t) + h_{w_{c,s}} \\ & + \iint c_{w_{c,s}w_{c,s}}(x-x',y-y')\Lambda_{w_{c,s}w_{c,s}}(w_{c,s}(x',y',t))dx'dy' \\ & + \int c_{w_s w}(x-x')\Lambda_{w_s w}(w(x',t))dx' \end{aligned}$$

$$+s_{pre}(y,t) + s_{tar_{space}}(x,y,t) + s_{tar_{color}}(x,y,t). \quad (7)$$

As before, $\dot{w}_{color,space}(x,y,t)$ is the rate of change of activation across the population of spatially-tuned neurons, x , and the population of color-tuned neurons, y , as a function of time, t ; the constant τ sets the time scale; $w_{c,s}(x,y,t)$ captures the current activation of the field; and $h_{w_{c,s}}$ sets the resting level. For this field, the local excitation / lateral inhibition form of interaction is captured by a single equation,

$\iint c_{w_{c,s}w_{c,s}}(x-x',y-y')\Lambda_{w_{c,s}w_{c,s}}(w_{c,s}(x',y',t))dx'dy'$. In contrast to the three-layer model, where inhibition is mediated by a separate field of inhibitory neurons, both the excitatory and inhibitory components of the interaction are defined by the convolution of a Gaussian kernel with a sigmoidal threshold function.

Specifically, the 2D Gaussian kernel was specified by:

$$c(x-x',y-y') = c_{excite,x} \exp\left[-\frac{(x-x')^2}{2\sigma^2}\right] + c_{excite,y} \exp\left[-\frac{(y-y')^2}{2\sigma^2}\right] - c_{inhib,x} \exp\left[-\frac{(x-x')^2}{2\sigma^2}\right] - c_{inhib,y} \exp\left[-\frac{(y-y')^2}{2\sigma^2}\right] - w_{inhib,xy}, \quad (8)$$

with strength, c , width, σ , and an additional inhibitory term, $w_{inhib,xy}$, which provides global inhibition to the field. Finally, the C-S field receives excitatory feedback from SWM (w), $\int c_{w,w}(x-x')\Lambda_{w,w}(w(x',t))dx'$, direct preshape input along the color dimension, $s_{pre}(y,t)$, and direct target input, $s_{tar_{space}}(x,y,t)$ and $s_{tar_{color}}(x,y,t)$, reflecting the color and spatial location of the memory target and the color target, respectively.

Table A1

Parameter values for simulations

Layer	τ	h	self-excitation	excitatory projection(s)	inhibitory projection(s)	reference input	preshape input	target input
u (PF)	80	-7	$c_{uu} = 1.75$ $\sigma_{uu} = 5$		$c_{uv} = 1.1$ $\sigma_{uv} = 5$	$c_{ref} = 8$ $\sigma_{ref} = 3$		$c_{tar}^1 = 40$ $\sigma_{tar} = 3$
v (Inhib)	10	-12		$c_{vu} = 4.88$ $\sigma_{vu} = 5$ $c_{vw} = 2.2$ $\sigma_{vw} = 6$				
w (SWM)	80	-4	$c_{ww} = 2.0$ $\sigma_{ww} = 6$	$c_{wu} = 1.5$ $\sigma_{wu} = 3$ $c_{wwe,s} = 0.04$ $\sigma_{wwe,s} = 6$	$c_{wv} = 0.615$ $\sigma_{wv} = 40$ $k_{wv} = 0.05$		[all inputs scaled by $c_s = 0.2$]	
$w_{c,s}$ (CS)	80	-7	$c_{w_{c,s}w_{c,s}} = 1.69$ $\sigma_{w_{c,s}w_{c,s}} = 3$	$c_{w_{c,s}sw} = 0.025$ $\sigma_{w_{c,s}sw} = 6$	$c_{inhib,xy} = 1.25$ $\sigma_{inhib,xy} = 9$ $w_{inhib,xy} = 0.10$		$c_{pre}^2 = 12$ $\sigma_{pre} = 8$	$c_{tar} = 5.5$ $\sigma_{tar} = 5$

¹ These values reflect target input strength to u for the spatial recall task. For the color-discrimination task, input to u was scaled by 0.125.

² Note that preshape input was only provided to the CS field in the shifting and static attention conditions, where participants were required to complete the color-discrimination task during the delay interval.

Figure Captions

Figure 1. Simulation of the dynamic field theory (DFT) performing a single trial in a spatial recall task.

Figure 2. Stimuli and procedures used in Experiment 1.

Figure 3. Experimental data and simulation results from Experiment 1 and model simulations. Percent correct accuracy in the color-discrimination task completed during the delay interval of the static and shifting attention conditions (A), results from the static-attention and static-control conditions (B), the shifting-attention and shifting-control conditions (C), and model simulations of performance in the shifting attention and control conditions (D).

Figure 4. Simulation of the model performing a single trial in the shifting attention condition of Experiment 1. In this simulation, the memory target is presented at 30° in the task space, and the color-discrimination target is presented 10° further away from midline (i.e., at 40° in the task space – see Figure 2). Weak excitatory input from the Color-Space field (CS) to SWM at the location of the color target increases drift of the WM peak away from the midline reference input. See text for additional details.

Figure 5. Simulation showing faster rise times for activation at memorized versus non-memorized locations in response to the color target in a color-discrimination task. See text for additional details.

Figure 1.

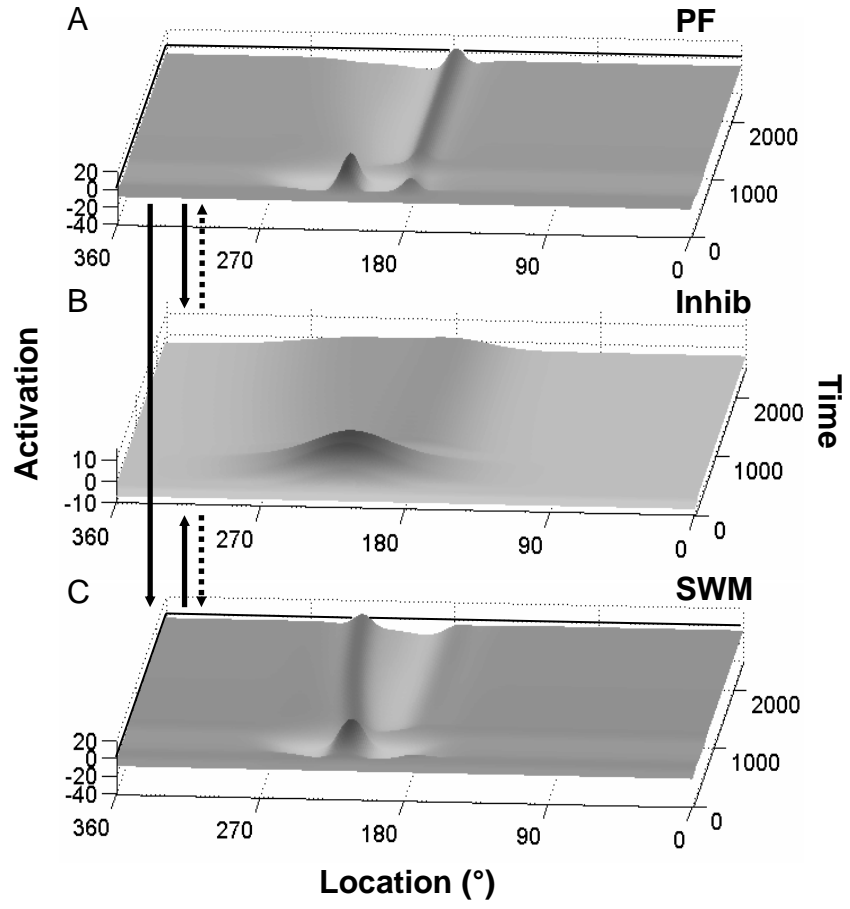


Figure 2.

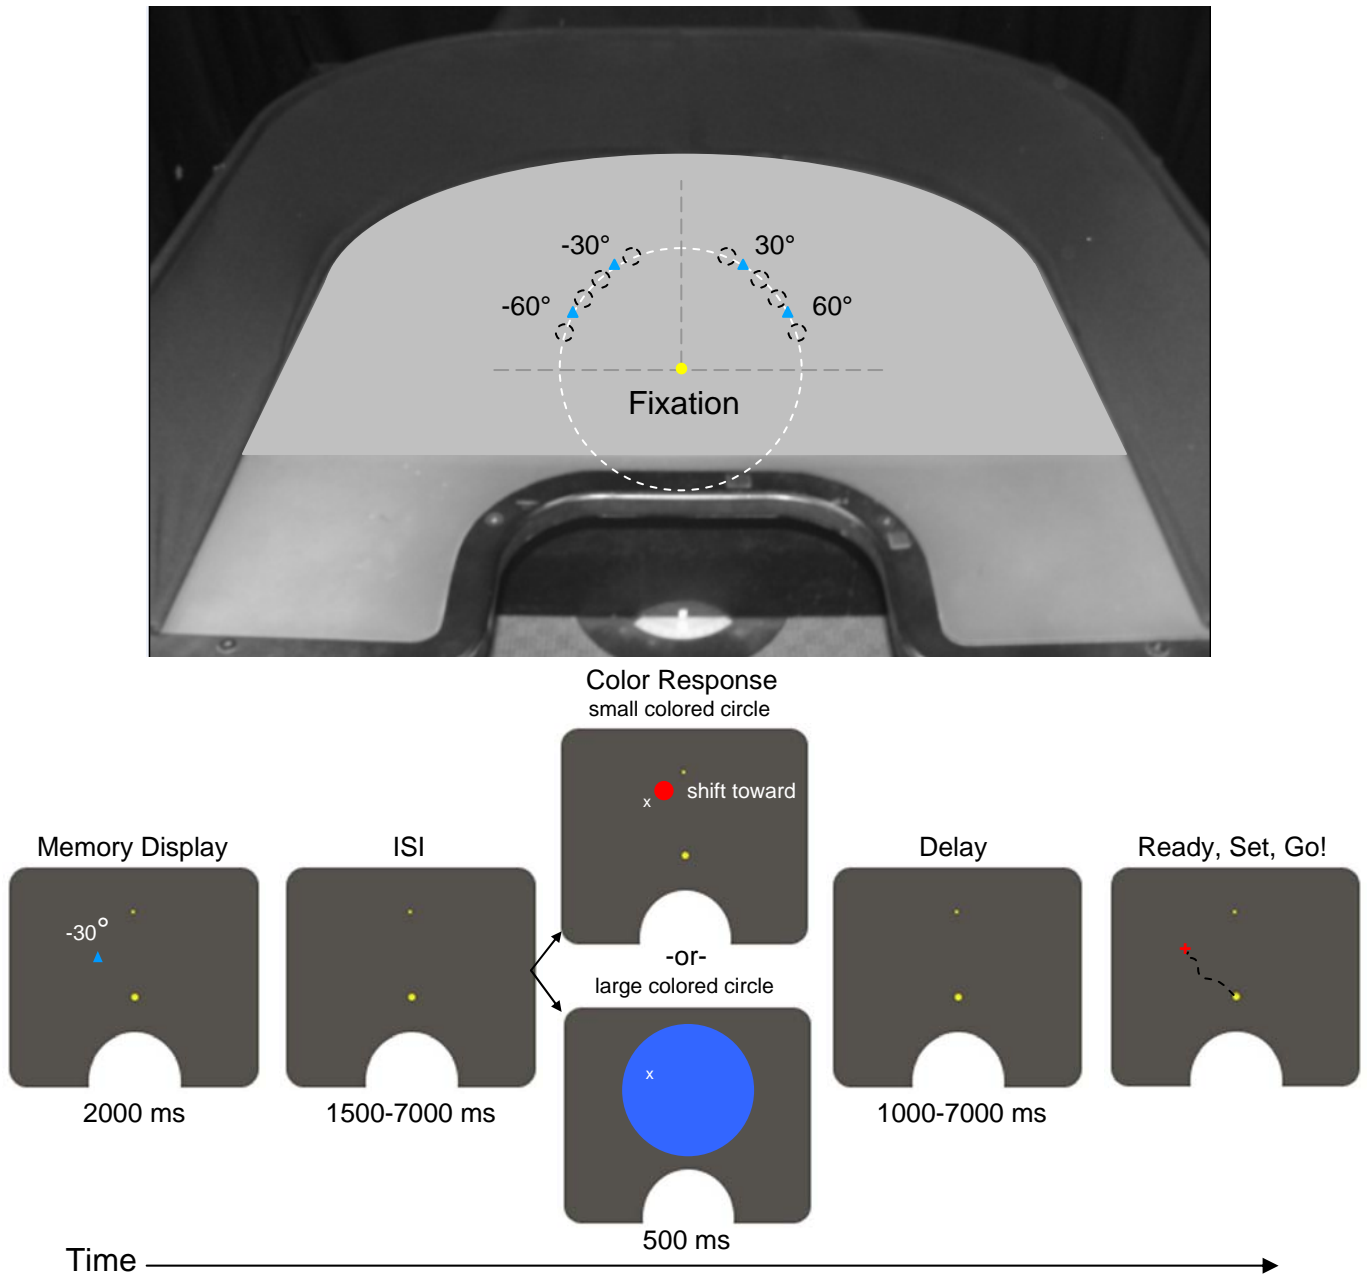


Figure 3.

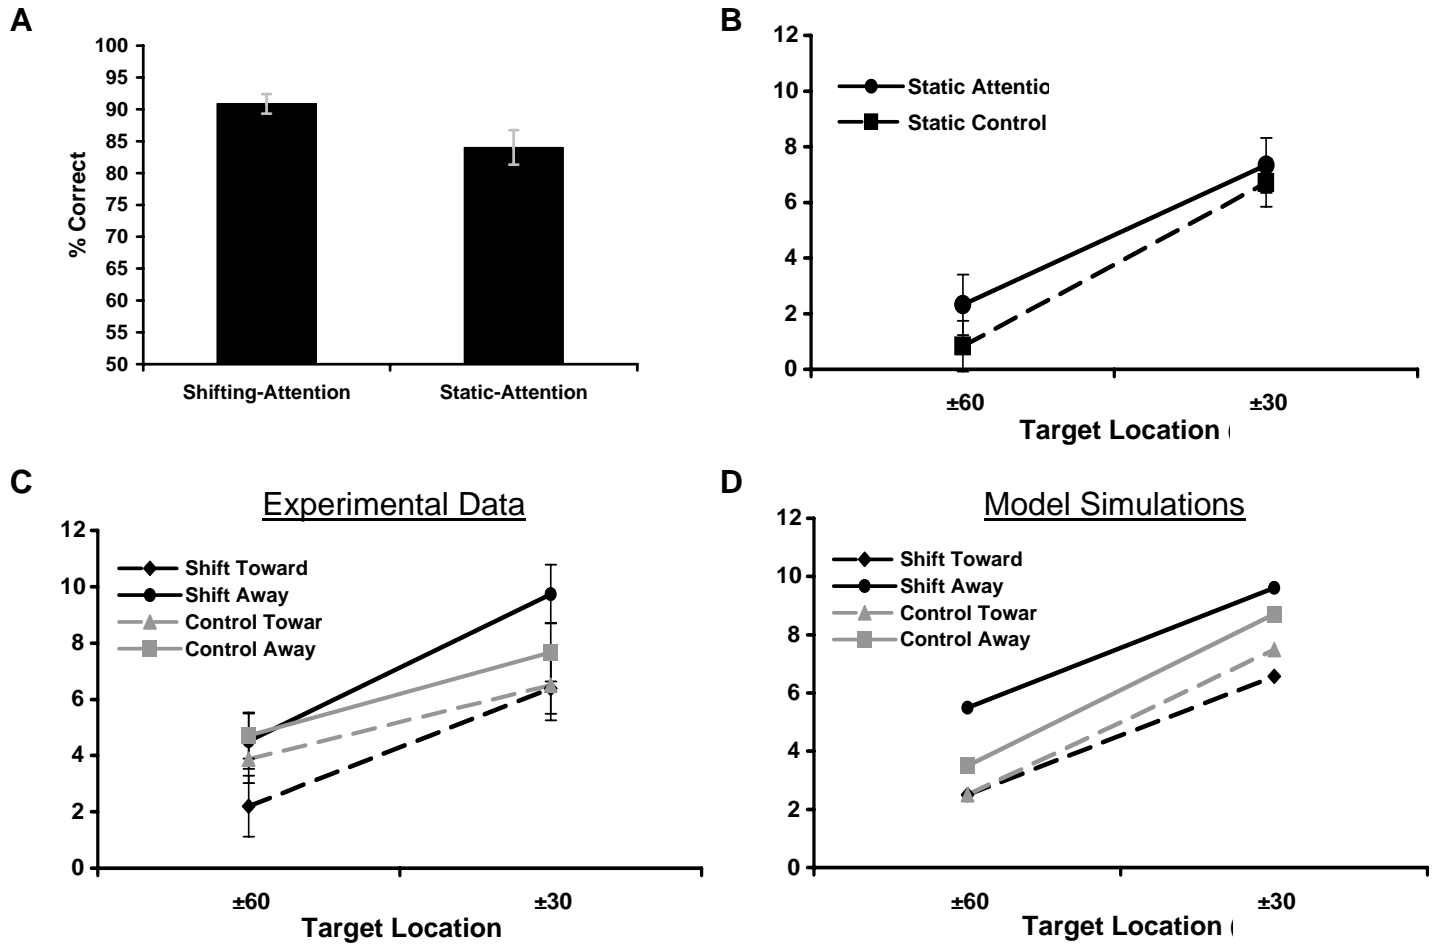


Figure 4.

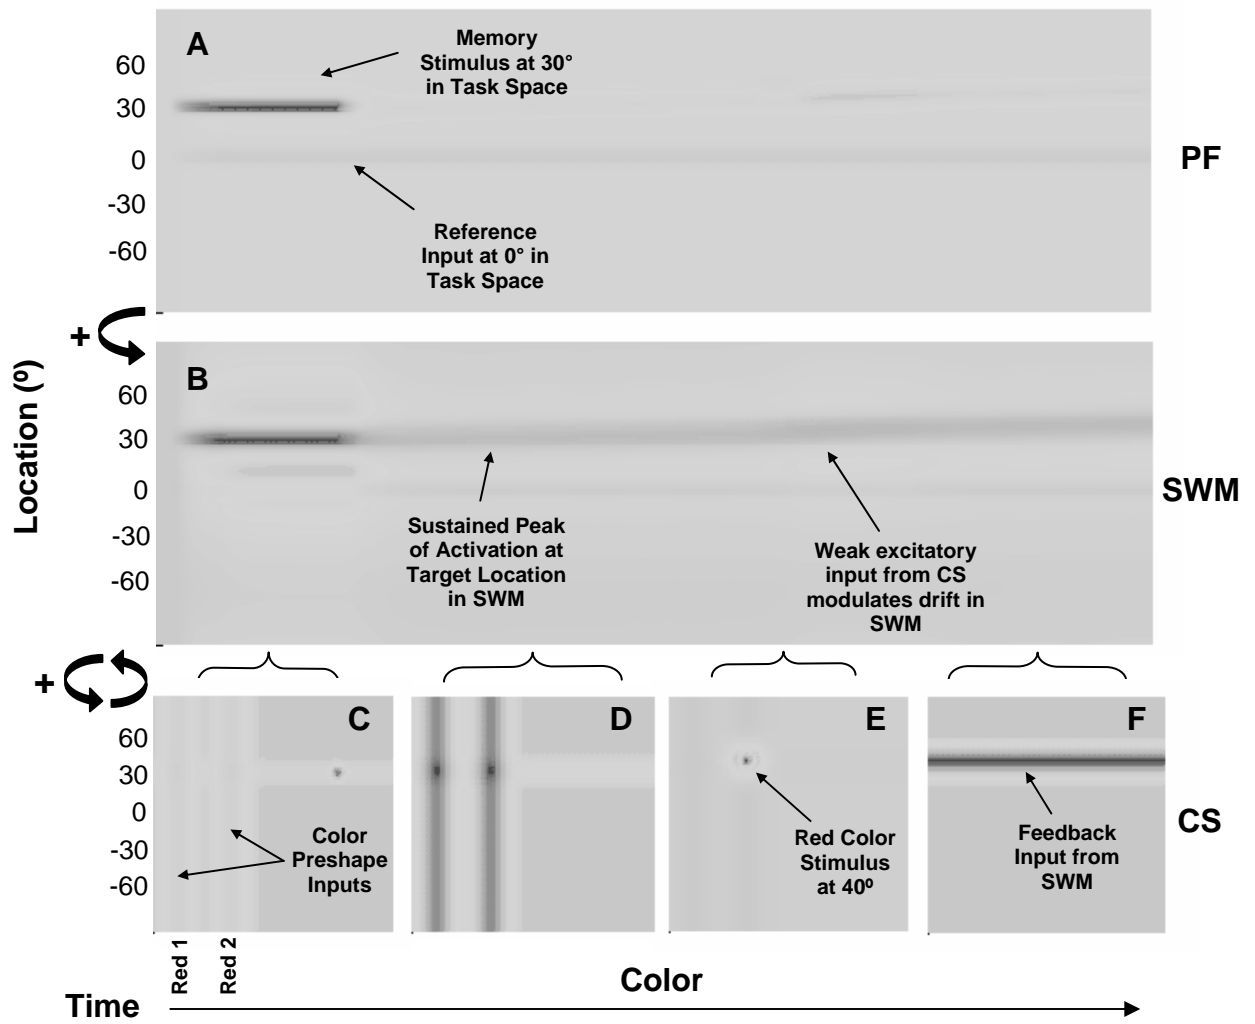


Figure 5.

