Biased feedback in spatial recall yields a violation of delta rule learning

JOHN LIPINSKI

Ruhr-Universität Bochum, Bochum, Germany

AND

JOHN P. SPENCER AND LARISSA K. SAMUELSON

University of Iowa, Iowa City, Iowa

This study investigates whether inductive processes influencing spatial memory performance generalize to supervised learning scenarios with differential feedback. After providing a location memory response in a spatial recall task, participants received visual feedback showing the target location. In critical blocks, feedback was systematically biased either 4° toward the vertical axis (*toward* condition) or 4° farther away from the vertical axis (*away* condition). Results showed that the *weaker* teaching signal (i.e., a smaller difference between the remembered location and the feedback location) produced a *stronger* experience-dependent change over blocks in the away condition than in the toward condition. This violates delta rule learning. Subsequent simulations of the dynamic field theory of spatial cognition provide a theoretically unified account of these results.

Spatial cognition research suggests that people divide regions of space into categories to enhance interactions with the local surrounds. Although this can lead to an increase in the precision of spatial memory (Huttenlocher, Hedges, Corrigan, & Crawford, 2004; Huttenlocher, Hedges, & Duncan, 1991), spatial categories also lead to systematic biases in location recall. For example, in one common task, people reproduce the location of a dot in a circle following a short delay. Results show that recall estimates are biased away from the horizontal and vertical axes toward the centers of the quadrants within the circle. Two different spatial memory models have captured these *geometric* biases: the category adjustment model (CAM; Huttenlocher et al., 1991) and the dynamic field theory (DFT; Schutte & Spencer, 2009; Simmering, Schutte, & Spencer, 2008).

Several recent studies have investigated whether spatial cognition is also influenced by a second class of categorical effects: induced category biases. Inductive or experience-dependent processes, in which classification responses come to reflect the distribution of exemplars within a set of items, are an important component of category formation in multiple domains (Huttenlocher, Hedges, & Vevea, 2000). Although Huttenlocher et al. (2004) reported that spatial recall is not influenced by the target distribution, this contrasts with other reports. Spencer and Hund (2002), for example, reported that when a target 40° to the right of a vertical axis at 0° was paired with targets farther from the vertical axis (e.g., 60°), bias away from the vertical increased. When the 40° target was paired with targets closer to the vertical axis (e.g., 20°),

however, bias away from the vertical decreased. In both cases, then, geometric memory biases were adjusted toward the center of the target distribution. Spatial recall is also influenced by the frequency with which participants estimate each location (Spencer & Hund, 2003).

These findings indicate that experience-dependent processes influence spatial memory performance, contrary to the claims of the CAM, which lacks a mechanism to account for such effects. The DFT, by contrast, captures experience-dependent effects through a long-term memory (LTM) mechanism. In this theory, the maintenance of location information in working memory during short-term delays leaves a trace in LTM. Such traces can accumulate and create biases in working memory, including attraction toward an average remembered location (Simmering et al., 2008; Spencer, Perone, & Johnson, 2009).

To date, the LTM mechanism of the DFT has only been formally probed using data from unsupervised learning tasks—that is, tasks in which there is not an explicit teaching signal that specifies the error on each trial. Unsupervised learning is often contrasted with supervised learning, in which an explicit teaching signal is present, and different learning rules have been proposed to account for findings from these two types of tasks (McLeod, Plunkett, & Rolls, 1998). Here, we examine whether experience-dependent effects in spatial recall generalize to supervised learning contexts in which a differential teaching signal is present and, further, whether the LTM mechanism of the DFT can account for participants' performance. Note that the use of precise metric feedback in spatial recall—redisplaying the

target location after the participant has responded—opens the door to a new probe of supervised learning because previous studies have used categorical or forced choice responses (Herzog & Fahle, 1997; Pashler, Cepeda, Wixted, & Rohrer, 2005) and discrete feedback. Precise metric feedback allows for a finer-grained analysis of supervised learning mechanisms, which are not yet fully understood (Pashler et al., 2005; Tsodyks & Gilbert, 2004).

EXPERIMENT

In this experiment, we tested whether experience-dependent location memory processes would generalize to supervised learning scenarios involving differential metric feedback. Although accurate feedback commonly increases response accuracy over time (e.g., Herzog & Fahle, 1997), systematically biased feedback can also decrease response accuracy (Herzog & Fahle, 1999). The present experiment built on this idea to test whether systematically biased response feedback would produce metric distortions in spatial recall.

In our task, participants saw a target, there was a 10-sec delay, and then they indicated the remembered target location by moving a cursor to the remembered location. Visual feedback indicating the “correct” target location was then given while the participants’ response was still visible, enabling the participants to determine the error on each trial. To probe supervised learning processes, we biased the feedback the participants received. In the *toward* condition, feedback was biased 4° toward the vertical axis (e.g., feedback at 26° for a target 30° to the right of the vertical axis at 0°). In the *away* condition, feedback was biased 4° farther away from the vertical axis (e.g., feedback at 34° for a 30° target).

A common mechanism in models of supervised learning is the delta rule (McLeod et al., 1998). According to the delta rule, connection weights are adjusted on the basis of differences between an external teaching signal (i.e., explicit feedback) and the current output (i.e., the response). Small differences lead to small weight changes, and big differences lead to big weight changes and more dramatic learning. Because spatial recall responses are strongly biased *away from* the vertical axis, differences between the memory response (current output) and the feedback location (desired output) should be larger in the toward condition, in which feedback is biased toward the vertical axis. According to the delta rule, large differences in the toward condition should yield greater changes over learning relative to the away condition.

Method

Participants. Thirty-two students (20 of them male, 12 female) participated in exchange for course credit or payment. All were native speakers of English with normal hearing and normal or corrected-to-normal vision. A participant who moved the cursor along L-shaped trajectories was replaced.

Materials and Apparatus. The participants sat at a large table (79 cm [h] × 117 cm [w] × 127 cm [l]) with a homogeneous surface and no explicit landmarks (Figure 1A). Sessions were conducted in a dimly lit room with black curtains covering all external landmarks. The stimuli consisted of three images: a yellow three-digit number

presented along the vertical midline axis of the table 46 cm from the participants, a yellow reference disk (15-mm diameter) located along the midline 30 cm away from the participants, and a blue equilateral triangle “spaceship” (10-mm base). The stimuli were projected onto the table surface from below using a Sony VPL-PX LCD Projector. The participants responded by moving a mouse cursor that was projected onto the surface of the table. The mouse was located on a pad to the right of the participants, out of the field of view.

Procedure. Each trial began with the appearance of the yellow referent disk, which remained visible throughout the trial (Figure 1B). Then a randomly chosen number between 100 and 500 appeared, and the participants began counting aloud backward from that number by ones until the recall response. This prevented verbal encoding of the spaceship position. Next, a blue spaceship target appeared for 2 sec. The participants looked up from the table when the spaceship disappeared and looked back when the computer said “ready-set-go.” Once they heard “go” (10 sec after target presentation), the participants moved the mouse cursor to the remembered location and clicked the button. After responding, a spaceship reappeared, and the computer said “Here’s where the ship was!” followed by a blank screen and the next trial.

The participants completed several demonstration trials (with accurate feedback) at the start of the session, followed by three blocks of test trials. Blocks consisted of three responses to each of the following 19 targets in randomized order at a fixed 15-cm radius from the yellow reference disk: 0° (vertical), ±10°, ±20°, ±30°, ±40°, ±50°, ±60°, ±70°, ±90°, and ±110° (Figure 1A). In Block 1, the participants received accurate feedback for all the targets. In Blocks 2 and 3, feedback was biased either 4° toward the vertical axis or 4° away from the vertical axis for the ±10° to ±70° targets (Figure 1C). Given that memory for locations along cardinal axes is quite accurate (see Spencer & Hund, 2002), targets at 0°, ±90°, and ±110° were presented with accurate feedback in order to reduce the possibility that the participants would detect the feedback manipulation. All the participants were debriefed after the session. There was no indication that the participants had detected the manipulation.

Methods of analysis. Directional errors (in degrees) on each trial were computed such that errors *away from the vertical axis* were positive (see Figure 1D for an example). Responses to targets at mirror image locations on either side of the vertical axis (e.g., +30° and −30°) within the same experimental block were grouped together. Outlier trials—directional errors greater than 40° or exceeding the mean by three *SDs* for a specific target—were removed. There were few outliers (0.5%); responses generally clustered around the actual target locations (see Figure 1D for an example). Only responses to the biased feedback locations (±10° to ±70°) were analyzed.

Results and Discussion

We first compared directional errors in Block 1 to determine whether performance was comparable across conditions before the introduction of biased feedback. A two-way ANOVA with condition as a between-subjects factor and target as a within-subjects factor revealed a main effect of target [Wilk’s $\Lambda = .21$; $F(6,25) = 15.6$, $p < .001$, $\eta_p^2 = .79$] but no significant effect of condition or interaction ($ps > .14$). Thus, performance across conditions was comparable in Block 1. Block 1 responses (Figure 2A) were generally biased away from the vertical axis, consistent with previous findings (e.g., Huttenlocher et al., 1991; Spencer & Hund, 2002).

To obtain a within-subjects measure of performance change with the introduction of biased feedback, we computed difference scores by subtracting each participant’s mean error for each target in Block 1 from the mean error for each target in Blocks 2 and 3. A three-way ANOVA on these scores, with condition as a between-subjects fac-

Figure 1. (A) Target distribution (left panel) and apparatus (right panel) used for the spaceship task. The black disk at the origin of the target distribution corresponds to the yellow reference disk used in the experiment. Targets were projected onto the table from beneath. Lights were dimmed during the experiment, and the table appeared black. (B) Time line of the experimental procedure. (C) Sample target location (30°) along with feedback biased either 4° toward the vertical axis (toward condition) or 4° away from the vertical axis (away condition). (D) Mean error calculation method, demonstrated here using the response distribution for the $\pm 30^\circ$ target in the away condition in Block 2 (all participants) as an example.

Figure 2. (A) Location memory performance. Positive errors indicate drift away from the vertical axis. Negative errors indicate drift toward the vertical axis. Block 1 performance is averaged across the toward and away conditions. The gray region reflects the change in memory performance from Block 1 to the average of Blocks 2 and 3 for those in the away condition. The textured region reflects memory performance change from Block 1 to the average of Blocks 2 and 3 for those in the toward condition. (B) Memory performance difference scores (collapsed across targets) for the toward and away conditions. Positive difference scores indicate an increase in drift away from the vertical axis over blocks. Negative difference scores indicate a decrease in drift away from the vertical. Error bars represent standard errors.

tor and target and block comparison (Block 2–Block 1 or Block 3–Block 1) as within-subjects factors, revealed a significant main effect of condition [$F(1,30) = 18.9, p < .001, \eta_p^2 = .39$]: Difference scores for the toward condition (textured region, Figure 2A) were significantly smaller than those for the away condition (gray region, Figure 2A).

Subsequent analysis revealed that the mean toward difference score (collapsed across targets) did not differ significantly from zero [$t(15) = 1.68, p = .11$]. The biased feedback in the toward condition did not, therefore, alter spatial recall performance (Figure 2B). The away condition difference scores, however, showed a significant overall mean *increase* in response bias away from the vertical axis [see Figure 2B; $t(15) = 5.17, p < .001$]. Moreover, follow-up analyses of difference scores (Block 3–Block 1) showed that the *absolute* magnitude of performance change for the away condition was greater than that for the toward condition [$M_{\text{diff}} = 1.3, SE = 0.61; t(30) = 2.14, p = .04$, two-tailed].

The ANOVA also revealed a significant target effect [Wilk's $\Lambda = .40, F(6,25) = 2.79, p = .03, \eta_p^2 = .60$]. Difference scores in the toward and away conditions combined to produce means close to 0 at 10° ($M = -0.4^\circ$), 20° ($M = -0.83^\circ$), 30° ($M = 0.31^\circ$), and 40° ($M = -0.23^\circ$). Targets at 50° ($M = 1.36^\circ$), 60° ($M = 1.07^\circ$), and 70° ($M = 1.98^\circ$), however, had higher average difference scores because the away condition effects generally outstripped those of the toward condition (see Figure 2A). All other effects were nonsignificant ($ps > .11$).

The larger effects for the away condition appear to run counter to delta rule learning, because the size of the teaching signal (i.e., the delta) should, on average, be smaller in the away condition. To explore this, we conducted post hoc analyses on the difference between each location memory response and the subsequent feedback location. This provided an index of the teaching signal magnitude. In Block 1 (where accurate feedback was given), the mean teaching signal for both the away [$M = 2.7; t(15) = 5.5, p < .001$]

and toward [$M = 3.1$; $t(15) = 5.2$, $p < .001$] conditions were significantly greater than zero but did not differ across conditions [$t(30) = 0.53$, $p = .58$]. By contrast, the introduction of biased feedback in Blocks 2 and 3 for the toward condition significantly *increased* [$t(15) = 7.3$, $p < .001$] the mean teaching signal from 3.1 in Block 1 to 6.4 in Blocks 2 and 3. In contrast, the biased feedback for the away condition in Blocks 2 and 3 significantly *decreased* [$t(15) = 6.96$, $p < .001$] the mean teaching signal from 2.7 in Block 1 to 0.43 in Blocks 2 and 3. Critically, the teaching signal was not only significantly smaller in the away condition than in the toward condition [$t(30) = 8.6$, $p < .001$], but also did not differ significantly from zero [$t(15) = 1.2$, $p = .24$]. Thus, the participants showed the largest change over blocks in the condition in which the teaching signal had a magnitude near zero.

We draw two conclusions from these results. First, experience-dependent changes in spatial recall generalize to supervised learning contexts. Second, the weaker teaching signal of the away condition, whose average value was not statistically different from zero, produced

a stronger experience-dependent location memory effect. This violates delta rule learning. In the following sections, we examine whether this learning asymmetry can be captured by the DFT.

THE DYNAMIC FIELD THEORY

The DFT (Schutte & Spencer, 2009; Simmering et al., 2008) is a formal theory of spatial cognition grounded in the principles of neural population dynamics (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999) that captures experience-dependent changes in spatial memory performance (Simmering et al., 2008; Spencer et al., 2009). Here, we probe whether this theory generalizes to a supervised learning scenario and captures the reported learning asymmetry.

We first address how the model accounts for geometric spatial memory biases. Each of the four layers in the DFT (Figure 3) represents location (in degrees) along the x -axis, which is composed of aligned nodes, each preferentially tuned to a given spatial location. The z -axis of each layer

Figure 3. Simulation of the dynamic field theory. (A) The perceptual field. (B) The shared inhibitory field. (C) The spatial working memory (SWM) field. (D) The long-term memory (LTM) field. Green arrows indicate excitatory connections between fields. Red arrows indicate inhibitory connections. In each layer, spatial location (in degrees) is represented along the x -axis, time along the y -axis, and activation along the z -axis. The trial begins at the front of the figure and moves toward the back over the timesteps (ts). See the text for additional details.

represents the activation level of the nodes, and the y -axis represents time (1 timestep = 2 msec). Input enters the perceptual field (Figure 3A) strongly (with weak projections to working memory) and activates sites tuned to the input location. Perceptual field activation then propagates to the shared inhibitory field (Figure 3B) and the spatial working memory (SWM) field (Figure 3C). Importantly, the nodes within the perceptual and SWM fields interact such that above-threshold nodes coding for similar spatial locations excite one another, whereas nodes coding for strongly differing locations inhibit one another (via the shared inhibitory field). With sufficient activation input, local excitation and lateral inhibition can lead to a self-sustaining “peak” in SWM that is maintained over a delay even after stimulus removal.

The SWM field is continuously coupled to the perceptual field. Consequently, perceptual structure in the task space such as the vertical symmetry axis (Wenderoth, 1994) aligned with the reference disk can contribute to SWM dynamics. This can help keep spatial information anchored to perceived reference frames under some conditions (Simmering et al., 2008), but it can also systematically bias spatial recall under others. In particular, activation in the perceptual field associated with perceived reference frames (see Figure 3A) and its associated inhibition (see Figure 3B) can create a boost of inhibition in SWM near the vertical axis. Consequently, memory peaks near the vertical axis are “repelled” away from the axis over delay.

In addition to representing remembered locations, SWM peaks also leave activation traces in the LTM layer (see “LTM,” Figure 3D; see also the Appendix). Because the LTM field is reciprocally linked to SWM, experience-dependent traces can, in turn, bias SWM dynamics. Over time, LTM traces build up through this simple form of unsupervised Hebbian learning (Simmering et al., 2008).

How might this same system use feedback to alter spatial recall in a supervised learning task? One possibility is that feedback input, like target input, builds a second, localized peak in SWM that leaves traces at the corresponding sites in LTM. Given that location memories tend to drift away from the vertical axis over time, traces from these drifting memory peaks should overlap more with feedback traces biased away from the vertical axis than with feedback traces biased toward the vertical axis. As a result, feedback in the away condition should lead to a greater change in spatial recall, because two influences conspire to impact LTM. By contrast, feedback in the toward condition should lead to modest learning, because two influences compete. Note that by this scenario, unsupervised and supervised learning occur within the same framework.

Simulation Method and Results

Our model exploration began with parameters from a recent article by Schutte and Spencer (2009), who quantitatively modeled geometric biases in spatial recall from both children and adults. Critically, Schutte and Spencer did not include an LTM mechanism, because their task limited possible LTM contributions. Thus, we took their parameters for the perceptual, inhibitory, and SWM fields and added an LTM mechanism (see Simmering et al.,

2008; Spencer et al., 2009). The question was whether we could qualitatively capture the present findings with this model, including the asymmetry in learning.

Three blocks of trials were presented in each simulation set. Block 1 feedback was presented at the true target location in both conditions. In Blocks 2 and 3, however, the feedback location was biased either 4° toward or 4° away from the vertical axis. Each target was presented three times in a randomized order within each of the three blocks. As in Schutte and Spencer (2009), we simulated only the five targets from 20° to 60°. Each complete simulation run, therefore, consisted of 45 individual trials (5 targets $\times$ 3 trials per block $\times$ 3 blocks). We ran 100 complete simulation runs for each of the two feedback conditions to evaluate model performance in the presence of simulated neural noise. Note that a second set of 100 simulations replicated all the findings reported below.

Block 1 simulations (averaged across conditions) exhibited a memory profile similar to that in the present experiment, with strong biases away from the vertical axis for targets near this axis (20° target = 8.25° bias) followed by weaker biases (30° target = 3.99° bias; 40° target = 2.56° bias; 50° target = 0.17° bias) and an eventual reversal (60° target = -5.25° bias). A two-way ANOVA with condition as a between-subjects factor and target as a within-subjects factor for Block 1 revealed a main effect of target [Wilk's $\Lambda = .02$; $F(4,195) = 2,137.3$, $p < .001$, $\eta_p^2 = .98$] but no significant effect of condition or interaction ($ps > .47$).

Difference scores were calculated to compare within-simulation performance change from Block 1 to Blocks 2 and 3 (see Figure 4). Difference scores collapsed across targets for the toward condition were significantly smaller than those for the away condition. The absolute magnitude of the bias effect was also stronger in the away condition than in the toward condition [$t(198) = 4.8$, $p < .001$]. These simulations thus capture the asymmetric feedback effect observed in the experiment.

To determine whether the stronger away condition effect was accompanied by a weaker teaching signal, we computed the magnitude of the teaching signal on each

Figure 4. Memory performance difference scores (collapsed across targets) for the toward and away conditions from simulations of the dynamic field theory. Positive difference scores indicate an increase in drift away from the vertical axis over blocks. Negative difference scores indicate a decrease in drift away from the vertical. Error bars represent standard errors.

simulation trial. Block 1 (accurate feedback) mean teaching signals for the away [$M = 1.88$; $t(99) = 23.92$, $p < .001$] and toward [$M = 1.93$; $t(99) = 18.42$, $p < .001$] conditions were both significantly greater than zero but did not differ across conditions [$t(198) = 0.41$, $p = .67$]. For the toward condition, however, the introduction of biased feedback significantly *increased* [$t(15) = 7.3$, $p < .001$] the mean teaching signal from 1.93 in Block 1 to 5.8 in Blocks 2 and 3. In contrast, the biased feedback for the away condition significantly *decreased* [$t(99) = 37.97$, $p < .001$] the mean teaching signal from 1.88 in Block 1 to -1.62 in Blocks 2 and 3. Although this away teaching signal differed significantly from zero [$t(99) = 38.96$, $p < .001$], the *absolute* size of the teaching signal in the biased feedback trials was nonetheless significantly *smaller* than that in the toward condition [$t(198) = 71.6$, $p < .001$]. Thus, the weaker teaching signal for the away condition produced a stronger learning effect, consistent with the reported empirical results.

Figure 5 shows LTM activation in the model at different points in learning, which explains the origin of the asymmetric learning across conditions. Both simulation sets show a comparable distribution of activation in LTM after the first block, with a center of mass at 43.2° and

a maximum activation near 30° . Critically, at the end of Block 3, the center of mass of LTM in the away condition (Figure 5A) shifted substantially outward to 47.8° , with a maximum activation near 60° . This major shift in LTM activation occurred because traces from memory peaks during the delay *and* traces built from feedback were biased in the same direction—away from the vertical axis. By contrast, the distribution of LTM in the toward condition (Figure 5B) showed only a small shift in the center of mass to 41.2° , with a maximum near 25° . This modest change occurred due to competition between LTM traces from memory peaks that were biased away from the vertical axis and traces built from feedback that were biased toward the vertical axis. In summary, these simulations show how the DFT generalizes to a supervised learning scenario to capture experience-dependent changes in spatial recall that violate delta rule learning.

GENERAL DISCUSSION

The experimental results show that inductive, experience-dependent influences on spatial recall (Spencer & Hund, 2002, 2003) generalize to supervised learning scenarios with differential feedback and can be formally

Figure 5. Comparison of long-term memory activation at the end of Block 1 (dashed line) and the end of Block 3 (solid line) for the away (A) and toward (B) feedback conditions. Asterisks indicate target locations presented during simulation trials.

captured by the DFT. This further delineates between two competing models of geometric spatial memory biases: the CAM and the DFT. Whereas the CAM has no mechanism to account for these effects, the DFT incorporates an LTM mechanism that can capture key qualitative signatures of experience-dependent changes in spatial recall.

In this context, it is important to highlight the novelty of the proposed mechanism for the present results. Connectionist models of supervised learning typically use explicit error correction processes based on the delta rule. Instead, our dynamically integrated system, built using the same model and parameters from Schutte and Spencer (2009), processes feedback in the same manner as target input. This parsimonious account demonstrates how the fine-grained metric details of memory and feedback can account for the counterintuitive finding that a weaker teaching signal produces a larger change in memory performance. Thus, our work presents a novel empirical result in spatial cognition, and it provides a new theoretical mechanism to explore behavioral changes in both supervised and unsupervised learning paradigms.

AUTHOR NOTE

This research was supported in part by Grant NIMH-NRSA 5F31MH072133 awarded to J.L. and by Grants NSF-HSD 0527698 and NIH-RO1MH62480 awarded to J.P.S. We thank Gregor Schöner for assistance with this study, as well as Sebastian Schneegans and Jeff Johnson for helpful comments on the manuscript. Address all correspondence to J. Lipinski, Institut für Neuroinformatik, Ruhr-Universität Bochum, Universitätsstr. 150, Gebäude ND, Raum NDEF 04/589b, 44780 Bochum, Germany (e-mail: 2johnlipinski@gmail.com).

REFERENCES

- ERLHAGEN, W., BASTIAN, A., JANCKE, D., RIEHLE, A., & SCHÖNER, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, **94**, 53-66.
- HERZOG, M. H., & FAHLE, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Research*, **37**, 2133-2141.
- HERZOG, M. H., & FAHLE, M. (1999). Effects of biased feedback on learning and deciding in a vernier discrimination task. *Vision Research*, **39**, 4232-4243.
- HUTTENLOCHER, J., HEDGES, L., CORRIGAN, B., & CRAWFORD, L. E. (2004). Spatial categories and the estimation of location. *Cognition*, **93**, 75-97.
- 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 judgment? *Journal of Experimental Psychology: General*, **129**, 220-241.
- MCLEOD, P., PLUNKETT, K., & ROLLS, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford: Oxford University Press.
- PASHLER, H., CEPEDA, N. J., WIXTED, J. T., & ROHRER, D. (2005). When does feedback facilitate learning of words? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **31**, 3-8.
- SCHUTTE, A. R., & SPENCER, J. P. (2009). Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. *Journal of Experimental Psychology: Human Perception & Performance*, **35**, 1698-1725.
- SIMMERING, V. S., SCHUTTE, A. R., & SPENCER, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. *Brain Research*, **1202**, 68-86.
- SPENCER, J. P., & HUND, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, **131**, 16-37.
- SPENCER, J. P., & HUND, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*, **47**, 432-480.
- SPENCER, J. P., PERONE, S., & JOHNSON, J. S. (2009). The dynamic field theory and embodied cognitive dynamics. In J. P. Spencer, M. S. Thomas, & J. L. McClelland (Eds.), *Toward a new grand theory of development? Connectionism and dynamic systems theory re-considered* (pp. 86-118). New York: Oxford University Press.
- TSODYKS, M., & GILBERT, C. (2004). Neural networks and perceptual learning. *Nature*, **431**, 775-781.
- WENDEROTH, P. (1994). The salience of vertical symmetry. *Perception*, **23**, 221-236.

APPENDIX

LTM field activation is governed by

$$\dot{l}_w(x, t) = \begin{cases} \frac{1}{\tau_{\text{build}}} [-l_w(x, t) + \Lambda_w \{w(x, t)\}] & \text{if } w(x, t) \geq 0 \\ \frac{1}{\tau_{\text{decay}}} [-l_w(x, t)] & \text{otherwise,} \end{cases}$$

where $\dot{l}_w(x, t)$ is the rate of change of the activation for each node in the LTM field l_w across the spatial dimension, x , at time t . The constants τ_{decay} and τ_{build} (with $\tau_{\text{decay}} > \tau_{\text{build}}$) set the time scale of the dynamics for the decay and building of memory traces in LTM, respectively. Input is determined by the activation in the SWM field w at location x and time t .

The sigmoid

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

where β is the slope, determines the extent to which the activation w of the given node contributes to the network dynamics.

The parameters were set as follows: $\tau_{\text{build}} = 3,000$, $\tau_{\text{decay}} = 100,000$, $\beta = 5$.

The contribution of LTM field activation to location x in the SWM field is determined by the convolution $\int c_w(x - x') l_w(x', t) dx'$, where c is a Gaussian kernel with an amplitude of 0.2 and a width of 6.4.