

From Recall to Discrimination: The Dynamic Neural Field Theory Generalizes Across Tasks and Development

Vanessa R. Simmering (vanessa-simmering@uiowa.edu)

John P. Spencer (john-spencer@uiowa.edu)

Department of Psychology, University of Iowa
Iowa City, IA 52242 USA

Gregor Schöner (gregor.schoener@neuroinformatik.ruhr-uni-bochum.de)

Institut für Neuroinformatik, Ruhr-University
44780 Bochum, Germany

Abstract

Many models developed to account for tasks in cognition and development are well-suited to a single task or phenomenon. In recent years, modelers have begun to generalize their approaches to incorporate other tasks and/or developmental change. Few models, however, have attempted to generalize across both tasks and development. This paper presents one such model, the Dynamic Neural Field Theory (DNFT). The DNFT was developed to account for biases in spatial recall across development. A second line of research, presented here, generalizes the DNFT to a new task, position discrimination, and predicts a complex pattern of performance with both children and adults. We conclude with implications for future theoretical and empirical work.

Specificity vs. Generality in Models of Cognition and Development

Human behavior is rich and complex. As such, a central goal in cognitive science is to develop models that can capture this complexity and predict behavior in a meaningful way. In many cases, models in cognitive science are exclusively linked to particular tasks (for a discussion of this concern in categorization, see, e.g., Murphy, 2002). This approach is sensible, as it is often very difficult to capture the details of performance in a task in a way that sheds light on the processes that underlie behavior. Although theoretical specificity is essential to capturing the details of complex human behavior-in-the-moment, it often comes at the cost of generalization. For instance, there are relatively few detailed, process-based models that generalize across multiple tasks (for one exception, see Love, Medin, & Gureckis, 2004). Similarly, there are few process-based models that generalize across development. And the set of models that lie at the intersection of these two types of generalization is very sparse indeed (for examples, see Huttenlocher, Hedges, & Vevea, 2000; Morton & Munakata, 2002).

The difficulty of balancing specificity on one hand and generality on the other is not simply reflected in the relatively small number of models that generalize; it is also reflected in the nature of models that have effectively moved beyond specific tasks. For instance, the few models that have generalized across tasks and development have relatively weak ties between the model and the details of

behavior-in-the-moment. That is, these models are strong on generalization, but weak on the details so central to process-based models of performance in particular tasks.

The research presented here seeks to achieve both specificity and generality. In particular, we report theoretical and empirical work that generalizes a Dynamic Neural Field Theory (DNFT) of spatial working memory (SWM) beyond the spatial recall task that was central to the development of this theory. We extend the DNFT to capture performance in a position discrimination task, and then use this theory to generate specific behavioral predictions which we test in an experiment with adult participants. Next, we probe whether a core developmental hypothesis implemented in the DNFT—the spatial precision hypothesis—that has captured developmental changes in spatial recall can predict developmental changes in position discrimination. Finally, we test these predictions in an experiment with 3- to 6-year-old children. Results from our theoretical and empirical work demonstrate that the DNFT effectively generalizes across tasks and development. Importantly, such generality does not come at the cost of behavioral specificity due to our consistent emphasis on the processes that create behavior in the moment.

Dynamic Neural Field Theory: Overview

The DNFT is a process-based theory of spatial working memory (SWM) instantiated in a neural network model that captures children's and adults' spatial recall performance (Schutte, Spencer, & Schöner, 2003; Spencer & Schöner, 2003, 2006). In spatial recall tasks, participants are asked to remember a target location within a large, homogeneous space. After a short delay (0 – 20 s), adults show systematic biases away from the midline symmetry axis of the task space (e.g., Spencer & Hund, 2002). In this same task, young children show the opposite pattern of bias, that is, they show systematic bias toward midline (e.g., Schutte & Spencer, 2002). In the following section, we outline the DNFT and how this theory accounts for changes in spatial recall biases over development.

Figure 1 shows a simulation of the DNFT¹ during a single

¹ Note that for simplicity we show only the relevant portion of the full model here. For a complete description, see Spencer, Simmering, Troob, Schutte, & Schöner (2006).

spatial recall trial. Figure 1a shows the excitatory layer of a two-layered perceptual field (PF), and Figure 1c shows the excitatory layer of a two-layered spatial working memory field (SWM). Both of these layers are coupled to a single layer of interneurons (Inhib) in Figure 1b (see reciprocal green/excitatory and red/inhibitory arrows). In addition, PF passes excitatory input to SWM. In each field, the x -axis consists of a collection of spatially-tuned neurons. The y -axis shows each neuron's activation level. Lastly, time is captured along the z -axis, with the beginning of the trial at the front of the figure.

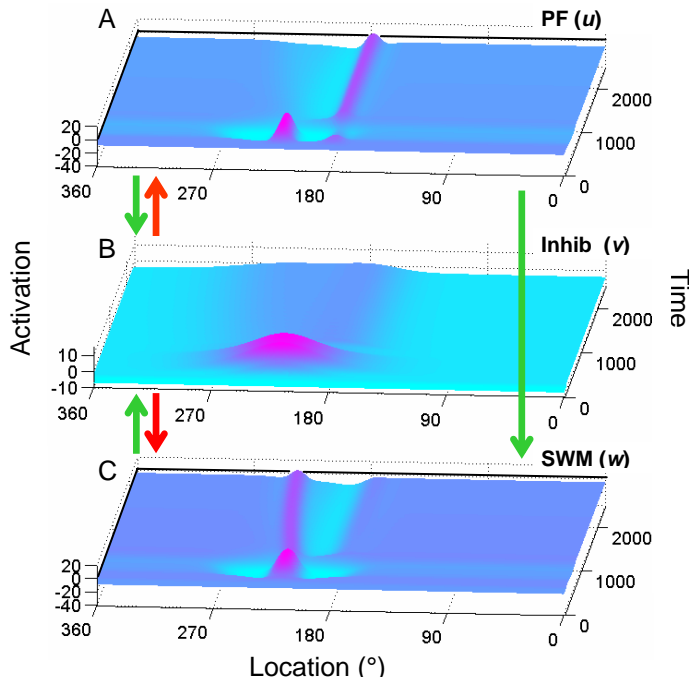


Figure 1. A simulation of a spatial recall trial in the DNFT.

At the start of the trial, the PF brings the midline input into 180° in the object-centered reference frame². Next, the target appears at 220° . This creates a peak of activation centered at this location. When the target disappears, a peak re-forms in PF at 180° as the system locks onto the reference cues in the task space. Figure 1c shows the effect of coupling PF to SWM. At the start of the trial, SWM receives relatively weak reference input from PF. Next, the target is turned on, passing strong target-related input into the working memory field. This event—combined with a boost in the resting level of SWM and Inhib—moves the working memory field into a strongly self-sustaining state. After the target disappears, the SWM field maintains an active memory of the target location during the delay. Importantly, this occurs even though PF has re-acquired the reference frame. This highlights the bi-stability of SWM: because SWM is in a stable attractor state, it can effectively hold on to the target, even while PF deals with the dynamics

² Other components of the model, not described here, address calibrating egocentric to object-centered frames of reference (see Spencer, Simmering, Troob et al., 2006).

of reference frame calibration. Thus, the 3-layer system shown in Figure 1 can achieve the dual goals of remembering the target item and staying calibrated with a perceived reference frame.

This dual ability has both costs and benefits. The benefit, of course, is maintaining a stable memory for the location within the object-centered frame of reference. One “cost”, however, is apparent in Figure 1c: the peak of activation in SWM “drifts” away from 180° during the memory delay. Thus, our model shows an emergent bias away from midline. This occurs because the peak of activation at midline in PF passes activation to the Inhib layer at 180° . This creates greater inhibition on the 180° -side of the target-related peak in SWM, effectively repelling this peak away from the reference frame (see, e.g., Spencer & Hund, 2002).

The same model also produces the bias toward midline early in development. The relatively simple developmental process we have put forth is captured by our *spatial precision hypothesis* (SPH): the spatial precision of neural interactions becomes more precise and more stable over development (Schutte et al., 2003; Spencer & Hund, 2003). This is illustrated in Figure 2. As the interaction functions move from early development (darker lines) to later development (lighter lines), the spatial precision of interaction narrows, and the excitatory / inhibitory gradient becomes steeper. This results in relatively unstable self-sustaining peaks early in development that are sensitive to input across a broad spatial range, as well as stable self-sustaining peaks later in development that are only sensitive to input at narrow separations.

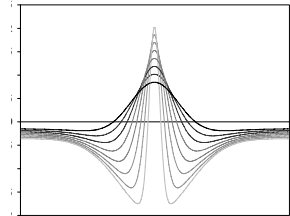


Figure 2. Narrowing interaction over development in the DNFT.

Because the peak in SWM is much broader (due to the broad interactions captured by the SPH), this peak produces broad activation in Inhib. This broad inhibition, combined with the relatively weak self-sustaining dynamics in PF, prevents a peak from building in the perceptual field at midline (i.e., at 180°) during the delay period. Consequently, SWM receives some excitatory input around midline—the midline input is still fed into PF and passed to SWM through the excitatory projection—but no associated inhibitory “push” away from a peak at midline. According to the SPH, this interaction function changes gradually over development (black to gray lines in Figure 2). Spatial recall data with children support this account—between the ages of 3 and 6, recall biases shift gradually from biased toward to biased away from midline (Schutte, 2004; Spencer & Hund, 2003).

Generalizing the DNFT across SWM Tasks: Position Discrimination

As the previous section illustrates, the DNFT is able to capture the moment-by-moment details of spatial recall behavior in both children and adults. The first goal of this

line of research is to generalize the DNFT to a new SWM task: position discrimination. In position discrimination tasks, two stimuli are presented in quick succession (e.g., 500 ms apart) and the participant judges whether the stimuli were in the same or different locations. We chose position discrimination because it is viewed in the literature as a more perceptual task than spatial recall, but is still within the domain of SWM. Additionally, the shorter time scale of discrimination allows for closer examination of reference-related effects.

Position discrimination differs from spatial recall in three important ways. First, two stimuli are presented in discrimination, compared to one stimulus per trial in recall. Second, the delay between stimuli is much shorter in discrimination (500 ms compared to 5 – 20 s). Lastly, a *same/different* decision is required rather than pointing to a remembered location. To capture position discrimination in the DNFT, it must accommodate these three differences. The first two differences are straightforward: a second stimulus can be presented 500 ms after a first stimulus. Critically, the third difference—a *same/different* decision—can be captured by the DNFT by simply modulating a few features of the model (for a related modulation in response to task demands, see Spencer, Simmering, & Schutte, 2006). Specifically, we can make our model a “difference” detector by increasing the resting level of neurons in the PF and Inhib fields when the second stimulus is presented. In some cases, this will result in a self-sustaining peak in the perceptual field (and no peak in SWM) after presentation of the second stimulus—the basis for a *different* response.

Figure 3 illustrates how this process operates in position

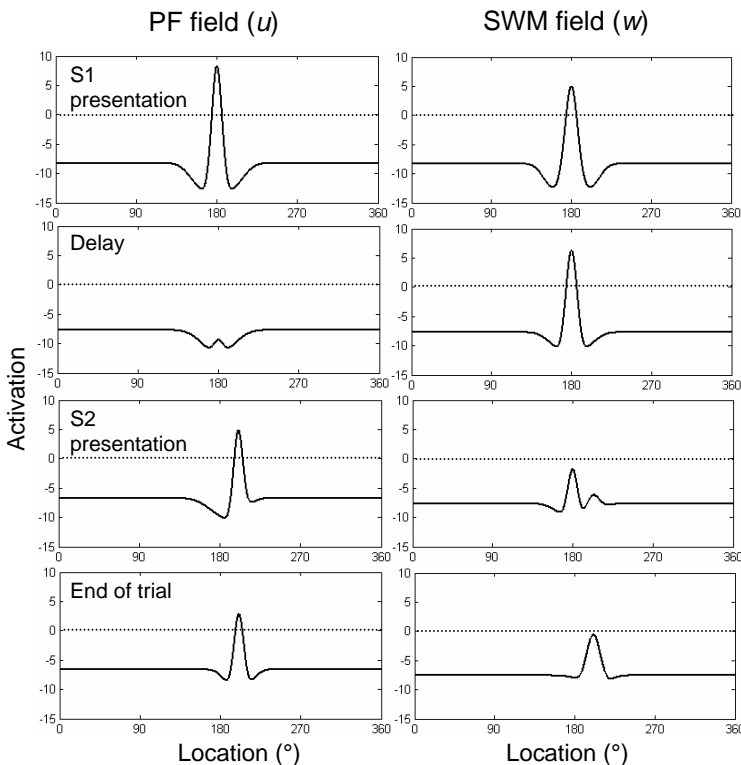


Figure 3. A simulation of a *different* trial in position discrimination.

discrimination. When the first stimulus (S1) is presented to the PF field, it forms a peak of activation, and passes activation to form a corresponding peak in the SWM field. Thus, at the end of presentation of S1, both fields have corresponding peaks of activation. After the stimulus is removed, the SWM peak self-sustains, as the PF peak is essentially deleted by inhibition from the Inhib field (not shown). Next, the second stimulus (S2) is presented in a new location. In the SWM field, S2 falls in the lateral inhibition of the peak from S1, causing activation to grow slowly at the spatial position of S2. In the PF field, on the other hand, S2 is outside of the strong inhibitory region of S1, so activation builds more quickly. After S2 is removed, the peak remains in the PF field, leading to a *different* decision³.

By modulating a few features, the DNFT is capable of performing the position discrimination task, showing the behavioral flexibility captured in this single model of SWM. Critically, if the same SWM system operates in both spatial recall and position discrimination tasks, then we should be able to use the principles of the DNFT established in spatial recall to make predictions for position discrimination performance. As illustrated in Figure 3, position discrimination performance depends on overlap between the peaks associated with S1 and S2. In the DNFT, overlap between peaks will be affected by two factors, both of which stem from the strong reference-related inhibition.

The first factor that should influence peak overlap is *peak width*: narrower peaks are less likely to overlap, leading to enhanced discrimination performance (i.e., *different* responses at smaller separations between S1 and S2). In the DNFT, peak width is determined through local excitation and lateral inhibition—stronger inhibition leads to narrower peaks. Recall that biases away from midline in spatial recall result from strong inhibition near the midline symmetry axis. This same reference-related inhibition should lead to narrower peaks. Because this inhibition is stronger near midline, the DNFT predicts that position discrimination will be enhanced near versus far from this axis.

Although narrower peaks near midline are a consequence of reference-related inhibition, the dominant inhibitory effect seen in spatial recall is *directional drift*—this is the second important factor in position discrimination. Figure 4 shows how directional drift should influence position discrimination in the DNFT. Recall from Figure 3 that the presentation of S1 forms a peak in SWM, which self-sustains during the delay before S2 is presented. During this delay, reference-related inhibition repels the peak away from midline (as seen in spatial recall). If S2 is then presented in the direction that S1 has drifted (Figure 4a), the peaks are more likely to overlap, leading to more *same*

³ We have implemented an explicit *same/different* response system using two bi-stable neurons that are dedicated to these labels (i.e., when the *same* node receives strong input, it goes into a self-sustaining state that represents the generation of this response). This allows us to generate an explicit response and to generate realistic reaction time curves as well.

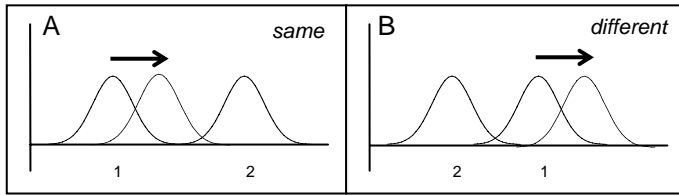


Figure 4. The effect of directional drift of S1 on discrimination when S2 is presented in the same (A) or opposite (B) direction of drift.

responses at larger separations (i.e., impaired performance). On the other hand, if S2 is presented in the opposite direction (Figure 4b), then the peaks are less likely to overlap and will lead to more *different* responses at smaller separations (i.e., enhanced performance). Thus, the DNFT predicts that discrimination will be impaired when S2 is presented in the direction of drift, or away from midline for adults. Moreover, directional drift should affect discrimination performance most at spatial locations where recall biases are greatest, around 20-30° (Spencer & Hund, 2002).

In summary, then, the DNFT predicts two novel effects of reference-related inhibition on position discrimination: enhanced performance near midline, and impaired performance when S2 is presented away from midline. Because these two effects arise through the same mechanism in the DNFT, they should combine to form a specific pattern of performance, shown in Figure 5. Based on peak width, performance will be enhanced (i.e., lower discrimination thresholds) near versus far from midline, regardless of the direction of S2 (Figure 5a). At these same locations, directional drift should not influence performance near midline, where drift is minimal (Figure 5b). By contrast, around 20-30° from midline, drift should impair performance when S2 is presented away from midline, but improve performance if S2 is toward midline (as shown in Figure 4). These effects should combine to form the pattern of performance shown in Figure 5c: comparable performance across locations when S2 is presented toward midline, but impaired performance far from midline when S2 is away from midline. Experiment 1 tested these predictions with adults.

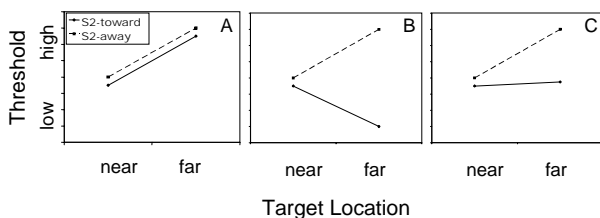


Figure 5. DNFT predictions based on peak width (A), directional drift (B), and both (C).

Experiment 1⁴

Method Fifteen University of Iowa undergraduates

⁴ This experiment is presented in Simmering, Spencer, & Schöner (in press) as Experiment 5.

participated in this study for research exposure credit. Participants were seated at a large, homogeneous table. Stimuli were projected onto the surface of the table from below. Each trial began with a 100 ms warning tone, followed by a 1 s delay and then a 500 ms presentation of S1 (2 pixels in diameter). Next came a 500 ms delay, then a 500 ms presentation of S2. After S2 disappeared, participants entered a *same* or *different* response by pressing corresponding buttons on a keypad. Across trials, S1 was presented a two target locations: 5° and 25° to the right of the midline symmetry axis of the table. S2 could then appear 0-7 pixels to the right (away from midline) or left (toward midline) of S1. Trials were presented in random order.

Results Figure 6 shows mean discrimination thresholds across targets separately for each S2-direction. As this figure shows, discrimination performance did not differ across S2-directions for the location near midline. Far from midline, however, thresholds were impaired when S2 was presented away from midline, that is, in the direction of drift. These data confirm the predicted pattern of results from Figure 5c.

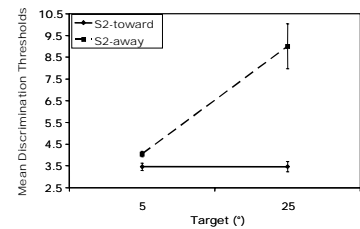


Figure 6. Results of Experiment 1.

Generalizing across Tasks and Development: Development of Position Discrimination

In the previous section, we used our model of spatial recall to predict a specific pattern of performance in a position discrimination task. Here, we probe the generality of the DNFT on a second front. In particular, we use our account of the development of SWM to predict developmental changes in position discrimination.

According to our spatial precision hypothesis (SPH), neural interactions in SWM become stronger and more precise over development. In spatial recall, broad interactions lead to biases toward midline in young children, whereas narrow interactions lead to biases away from midline in older children and adults. Because directional drift led to important predictions for position discrimination in adults, this developmental shift in the direction of drift has implications for the development of position discrimination as well. Specifically, the advantage shown in adults' performance based on the direction of S2 should flip with children who show drift toward midline (i.e., around 3 years of age): performance should now be better when S2 moves *away* from midline. Moreover, because recall biases shift gradually between the ages of 3 and 6, this S2-away advantage should gradually shift to an S2-toward advantage by 6 years of age (the age when children show the adult pattern in spatial recall). Lastly, the narrowing interactions posited by the SPH should also affect peak width over development, with broad peaks early in development. This

global change in peak width should produce overall decreases in discrimination thresholds over development.

Experiment 2

Method Participants in this experiment were 41 3-year-olds, 42 4-year-olds, 39 5-year-olds, and 39 6-year-olds from the Iowa City community, and 19 undergraduates. Stimuli were projected onto the same large, homogeneous table top. The procedure was similar to Experiment 1, with some changes to make the task more appropriate for children. Stimuli were increased to 7 pixels in diameter, and presented for 1 s each (rather than 500 ms). Across trials, S1 was presented at 10° or 30° to the right of midline. For children, S2 could then appear 0-20 pixels from S1, in steps of 4 pixels. For adults, S2 could appear 0-10 pixels from S1, in 1-pixel steps. Two other adaptations decreased the total number of trials required of children. First, children were randomly assigned to S2-toward or S2-away conditions; adults still completed both within a single session, although trials were now blocked by S2-direction. Second, we used a stair-casing procedure to reduce the total number of trials necessary to estimate discrimination thresholds (see Simmering & Spencer, 2006, for details).

Results Figure 7 shows mean discrimination thresholds across targets and ages separately for each S2-direction. As this figure shows, the pattern of results changed systematically across ages. Three-year-olds showed the expected reversal relative to adults' performance based on S2-direction with better performance in the S2-away condition at the far target. Across 4-, 5-, and 6-year-olds' data, thresholds in the S2-toward condition decreased, eventually leading to adult-like performance by 6 years. Adults' data replicated the results of Experiment 1, with no difference at the near target, and better performance for S2-toward at the far target. Lastly, across development, we observed an overall decrease in discrimination thresholds.

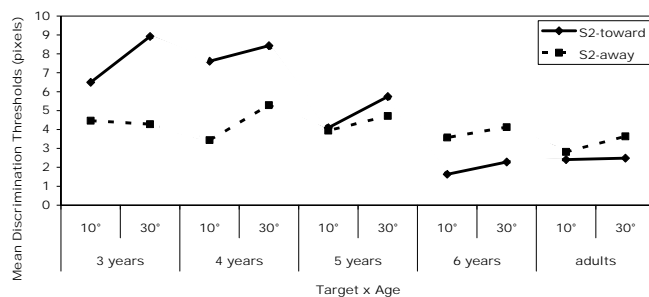


Figure 7. Results across ages for Experiment 2.

Conclusions

A central goal within cognitive science is to develop models that can capture the complexity of human behavior and make meaningful predictions. One of the primary challenges facing theorists who develop these models is the balance between specificity and generality: being specific enough to capture the details of behavior-in-the-moment, yet general enough to still apply across tasks and/or development. The

theoretical and empirical work presented her sought to achieve this balance within the framework of the DNFT.

Previously, this theory was used to capture details of spatial recall performance in both children and adults, and we proposed the SPH to account for developmental changes in reference-related recall biases. The first goal of the current research was to extend the DNFT to a new SWM task. This extension was achieved using the concept of novelty fields, and Experiment 1 confirmed two novel predictions of our model of position discrimination. The second goal was to use the developmental account from spatial recall to predict developmental changes in this new task. Experiment 2 confirmed a specific developmental pattern of performance predicted by the DNFT. Taken together, these experiments provide strong support for the DNFT account of developmental change in SWM.

The next step in this line of research is to use the results from these experiments to further develop the DNFT. In particular, we are currently quantitatively simulating results from these experiments to determine whether the same model can capture the rich set of constraints evident in our data. Although this is a non-trivial undertaking, the work presented here suggests that this theory provides a general framework for thinking about the processes that underlie SWM in multiple contexts and across multiple time scales, without giving up a core strength of process models—tight theory-experiment relations.

References

- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, *129*, 220-241.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*, 309-332.
- Morton, J. B., & Munakata, Y. (2002). Active versus latent representations: A neural network model of perseveration, dissociation, and decalage. *Developmental Psychology*, *40*, 255-265.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: The MIT Press.
- Schutte, A. R. (2004). *A developmental transition in spatial working memory*. Unpublished doctoral dissertation. University of Iowa, Iowa City.
- Schutte, A. R., & Spencer, J. P. (2002). Generalizing the dynamic field theory of the A-not-B error beyond infancy: Three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, *73*, 377-404.
- 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*, 1393-1417.
- Simmering, V. R., & Spencer, J. P. (2006). From recall to discrimination: Generalizing the dynamic field theory across tasks and development. *Manuscript in preparation*.
- Simmering, V. R., Spencer, J. P., & Schöner, G. (in press). Reference-related inhibition produces enhanced position

- discrimination and fast repulsion near axes of symmetry. *Perception & Psychophysics*.
- 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., & Schöner, G. (2003). Bridging the representational gap in the dynamical systems approach to development. *Developmental Science*, *6*, 392-412.
- Spencer, J. P., & Schöner, G. (2006). A dynamic field theory of spatial working memory. *Manuscript in preparation*.
- Spencer, J. P., Simmering, V. R., & Schutte, A. R. (2006). Toward a formal theory of flexible spatial behavior: Geometric category biases generalize across pointing and verbal response types. *Journal of Experimental Psychology: Human Perception & Performance*, *32*, 473-490.
- Spencer, J. P., Simmering, V. R., Troob, W. W., Schutte, A. R., & Schöner, G. (2006). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. To appear in J. M. Plumert & J. P. Spencer (Eds.), *Emerging landscapes of mind: Mapping the nature of change in spatial cognitive development*. New York: Oxford University Press.

This work was made possible by NIMH RO1 MH62480, NSF BCS 00-91757, and NSF HSD 0527698, awarded to John P. Spencer.