Published in final edited form as:

Brain Res. 2008 April 2; 1202: 68–86. doi:10.1016/j.brainres.2007.06.081.

Generalizing the dynamic field theory of spatial cognition across real and developmental time scales

Vanessa R. Simmering^{1,*}, John P. Spencer^{1,2}, and Anne R. Schutte³

¹Department of Psychology, University of Iowa

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

³Department of Psychology, University of Nebraska — Lincoln

Abstract

Within cognitive neuroscience, computational models are designed to provide insights into the organization of behavior while adhering to neural principles. These models should provide sufficient specificity to generate novel predictions while maintaining the generality needed to capture behavior across tasks and/or time scales. This paper presents one such model, the Dynamic Field Theory (DFT) of spatial cognition, showing new simulations that provide a demonstration proof that the theory generalizes across developmental changes in performance in four tasks—the Piagetian A-not-B task, a sandbox version of the A-not-B task, a canonical spatial recall task, and a position discrimination task. Model simulations demonstrate that the DFT can accomplish both specificity—generating novel, testable predictions—and generality—spanning multiple tasks across development with a relatively simple developmental hypothesis. Critically, the DFT achieves generality across tasks and time scales with no modification to its basic structure and with a strong commitment to neural principles. The only change necessary to capture development in the model was an increase in the precision of the tuning of receptive fields as well as an increase in the precision of local excitatory interactions among neurons in the model. These small quantitative changes were sufficient to move the model through a set of quantitative and qualitative behavioral changes that span the age range from 8 months to 6 years and into adulthood. We conclude by considering how the DFT is positioned in the literature, the challenges on the horizon for our framework, and how a dynamic field approach can yield new insights into development from a computational cognitive neuroscience perspective.

Keywords

spatial cognition; neural field model; development; cortical maps

1. Introduction

A central goal of computational cognitive neuroscience is to develop models of cognitive processes that provide insight into the organization of behavior while adhering to neural principles. As such, theorists strive to create models with sufficient specificity—both behavioral and neural—to generate novel, testable predictions. It is also critical, however, that

*Corresponding author Mailing address: E11 Seashore Hall, University of Iowa, Iowa City, IA 52242 Email: vanessa-simmering@uiowa.edu Phone: 319-353-2990 Fax: 319-335-0191

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

theories achieve a sufficient degree of generalizability across tasks. This is not a trivial accomplishment: capturing the details of performance in even a single task can be difficult, let alone generalizing across multiple tasks, multiple behaviors, and, perhaps even more challenging, multiple time scales (e.g., extending into learning and development).

Several theories have confronted the challenges of achieving generality with specificity, with varying levels of success (e.g., Cohen & Servan-Schreiber, 1992; Love, Medin & Gureckis, 2004; McClelland, McNaughton & O'Reilly, 1995; Morton & Munakata, 2002); here, we present one such theory, the Dynamic Field Theory (DFT) of spatial cognition (Spencer, Simmering, Schutte & Schöner, 2007). The dynamic field framework was originally developed to capture the dynamics of neural activation in visual cortex (Amari, 1977). More recently, this framework has been extended to account for the processes that underlie saccadic eye movements (Kopeck & Schöner, 1995; Wilimzig, Schneider & Schöner, 2006), motor planning (Erlhagen & Schöner, 2002; Schutte & Spencer, 2007b), infants' performance in Piaget's A-not-B task (Thelen, Schöner, Scheier & Smith, 2001), the dynamics of neural activation in motor and premotor cortex (Bastian, Riehle, Erlhagen & Schöner, 1998; Bastian, Schöner & Riehle, 2003), and the behavior of autonomous robots (Bicho, Mallet & Schöner, 2000; Iossifidis & Schöner, 2006; Steinhage & Schöner, 1998).

In the present report, we present new simulations of a dynamic field theory of spatial cognition that demonstrate that this theory generalizes across developmental changes in performance in four tasks—the Piagetian A-not-B task, a version of the A-not-B task conducted in a sandbox, a canonical spatial recall task, and a position discrimination task. Moreover, simulations of our theoretical model demonstrate that the DFT can span developmental changes in performance in these tasks across a range of ages from 8 months to 6 years through adulthood with a relatively simple developmental hypothesis. This highlights a novel developmental insight: because we use a richly structured real-time neural system, we can get “more from less” over development, that is, we can produce both quantitative and qualitative changes in performance via a simple developmental mechanism. Critically, the same model can perform these different behaviors with no modification to its basic structure and with a strong commitment to neural principles (see Amari, 1977; Bastian, et al., 1998; Bastian, et al., 2003; Erlhagen, Bastian, Jancke, Riehle & Schöner, 1999; Jancke, et al., 1999). These simulations demonstrate that the DFT can achieve both generality and specificity.

In the section that follows, we provide an overview of the DFT including its foundations in neural principles and our central developmental hypothesis. Next, we describe behavioral signatures in four spatial cognition tasks that have previously been simulated with earlier versions of the framework presented here—but never within a single model with a single parameter setting scaled systematically over development. Simulation results show that the current instantiation of the model can capture development within and across these tasks by scaling a small set of parameters. Thus, the exact same model with slight modification in parameters can show the complex pattern of developmental change observed across these four tasks. We conclude by placing our theoretical framework in the broader literature, considering the challenges on the horizon for our theory, and discussing the novel insights into developmental process offered by the simulations presented here, in particular, how one can get “more from less” over development.

2. Overview of the Dynamic Field Theory

A growing number of researchers have argued that we should take inspiration from the densely interconnected and dynamic nature of the brain to re-think cognition (e.g., Barsalou, 1999; Skarda & Freeman, 1987; Spencer & Schöner, 2003). A centerpiece of this approach is to embrace the use of complex, dynamic neural networks to capture brain-behavior relations.

Although neural networks have architectures that can be depicted as separate systems, they are at their core complex, reentrant, densely interconnected systems that violate core assumptions of encapsulation and separability (for discussion, see Spencer, et al., 2007).

We have contributed to this broader agenda using continuous dynamic neural fields first proposed to capture neural dynamics within the topographic organization of visual cortex (Amari, 1977; Amari & Arbib, 1977). More recently, this approach has been extended to capture the dynamics of neural activity in cortical areas with a non-topographic organization (e.g., motor cortex). For instance, time-dependent changes in neural activation in a dynamic field model of motor planning were compared to single-unit neural activity in motor cortex measured in a precue paradigm using population coding techniques (Bastian, et al., 1998; Bastian, et al., 2003; Erlhagen, et al., 1999). The first step in making this comparison was to map the responses of neurons to basic stimuli and create a continuous field by ordering the neurons based on their “preferred” stimulus. This was followed by a behavioral precuing task that probed predictions of a dynamic field theory of movement preparation (Bastian, et al., 1998; Bastian, et al., 2003; Erlhagen & Schöner, 2002). Note that this same theory has also been tested using ERP techniques (McDowell, Jeka, Schöner & Hatfield, 2002). These studies have reported a robust relationship between predictions of dynamic field models and neural measures, suggesting that this particular marriage between theoretical and behavioral neuroscience is quite promising.

The theory we present here uses additional insights gained from studies of the layered structure of cortex. In particular, we use a multi-layered architecture inspired by the cytoarchitecture of visual cortex (Douglas & Martin, 1998). This has given us an entry point into the dynamics that emerge from interactions among cortical layers. Moreover, our approach to long-term memory is grounded in established neural principles: we use long-term memory fields (described below) that capture a form of Hebbian learning (see Schöner, 2007; Wilimzig & Schöner, 2005; Wilimzig & Schöner, 2007).

2.1 A 5-layer Dynamic Field Model of Spatial Cognition

Our focus in the present report is on a process-based theory of spatial cognition instantiated in a 7-layer dynamic neural field model that captures children’s and adults’ performance in a host of spatial tasks (for a complete description of the theory, see Spencer, et al., 2007). Here, we focus on the dynamics of five layers of the model, shown in Fig. 1 (for equations and parameter details, see Appendix A): a perceptual field (PF; Fig. 1B), a spatial working memory field (SWM; Fig. 1E), a shared inhibitory field (Fig. 1D), and two long-term memory fields coupled to PF (LTM_{PF} ; Fig. 1C) and SWM (LTM_{SWM} ; Fig. 1F). 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, beginning at the front of the figure. These layers pass excitation and inhibition as indicated by solid and dashed arrows, respectively.

Inputs to the full 7-layer model described in Spencer et al. (2007) come from two additional layers not used in the simulations presented here: a perceptual field that receives input in an egocentric frame of reference (e.g., retinal coordinates) and a system that translates spatial information from this egocentric frame into an allocentric frame grounded in perceptual cues in the task space (e.g., the edges of a tabletop). These two layers keep our spatial system calibrated with the world despite, for instance, intervening movements of the head and/or body. To simplify discussion of the model here, we have replaced these two layers with the input layer shown in Fig. 1A. This layer simply presents inputs (i.e., Gaussian activation profiles) to PF and SWM in an allocentric frame with a time structure dictated by events in the task. For instance, the simulation shown in Fig. 1 shows performance in a single spatial recall trial where the target was presented at -40° in the task space (see Target in Fig. 1A) and there were perceptual cues in the task space marking the midline axis of the table (i.e., 0° ; see Midline in

Fig. 1A). Note that input from the midline axis is relatively weak compared to input from the target; this reflects the increased salience of the target item due to task instructions to remember its location. We also note here that inputs to SWM were considerably weaker than inputs to PF (multiplied by 0.2 in all simulations).

A central component of the model are the neural interactions among the 5 layers depicted in Fig. 1B-F. Neurons within PF (Fig. 1B) and SWM (Fig. 1E) have locally excitatory interactions where activated neurons boost the activity of their local neighbors. In addition, neurons in PF and SWM have reciprocal connections to the layer of interneurons (Fig. 1D): neurons in these layers increase the activation of interneurons tuned to “prefer” similar locations, and these interneurons, in turn, project broad inhibition back to PF and SWM. These combined interactions among the PF, Inhib, and SWM layers lead to locally excitatory and laterally inhibitory interactions within PF and SWM that enable these fields to form self-sustaining “peaks” of activation that maintain themselves in the absence of input (Amari & Arbib, 1977). For example, Fig. 1E shows a self-sustaining peak in SWM near -40° that “remembers” the target location after the target input was removed.

In addition to capturing working memory for a target on a given trial, this 5-layer structure also allows the model to learn from its previous experience through a type of Hebbian learning, incorporated using long-term memory (LTM) fields. Above-threshold activation in PF and SWM passes excitation to associated sites in the respective LTM fields (see Fig. 1C and 1F), leaving corresponding traces of activation. These activation traces gradually accumulate across a slower time scale and decay in the face of competing activation. Importantly, these traces pass excitation back into PF and SWM, which can bias the creation of activation peaks in these layers as well as shift the spatial position of such peaks (an issue we discuss in greater detail in the following sections). Note that each LTM field serves a unique and emergent behavioral function. The LTM associated with PF tracks the use of reference frames in the environment. For instance, the simulation in Fig. 1C has relatively strong activation at 0° , reflecting the continual presence of perceptual cues at midline in the task space—the system comes to remember that midline provides a salient reference axis in this task, which allows it to re-align to the same object-centered frame of reference from trial to trial (see Spencer, et al., 2007, for details). By contrast, the LTM field associated with SWM accumulates traces of previously-remembered locations over trials. In Fig. 1F, the long-term memory shows traces of the targets from the two previous trials at -20° and -50° , respectively, as well as traces built from activation associated with the current target at -40° . Because these targets are relatively close in space and have been presented only once each, the long-term memory traces are relatively weak and blend together. It is this field that allows the model to accurately capture performance in A-not-B-type tasks (see below), as well as effects of experience-dependent spatial categories (see, e.g., Spencer & Hund, 2002).

2.2 Development in the DFT

We have previously captured behavioral changes across tasks and development within the dynamic field framework using a relatively simple developmental hypothesis—the spatial precision hypothesis (SPH, Schutte, Spencer & Schöner, 2003; Simmering & Spencer, 2007; Spencer & Hund, 2003; Spencer, et al., 2007). The SPH posits that neural interactions become stronger and more precise over development. The simulations presented here show how small, quantitative changes in neural interaction are sufficient for the model to reproduce a complex pattern of behavioral performance observed across four tasks, beginning with infancy and spanning childhood into adulthood.

This general hypothesis is consistent with a host of neurophysiological evidence, particularly during the age range that is the focus of the present report—8 month to 6 years (for a related hypothesis about changes in neural processes over development, see Westermann & Mareschal,

2004). For instance, it is likely that the development of the dorso-lateral prefrontal cortex plays a central role in the tasks we discuss because this region of cortex is heavily involved in the on-line maintenance of spatial information (Awh, et al., 1999; di Pellegrino & Wise, 1993; Kessels, Postma, Wijnalda & de Haan, 2000; Nelson, et al., 2000). Moreover, the development of the prefrontal cortex shows a protracted course that continues into the postadolescent years (Gogtay, et al., 2004; Rakic, 1995; Sowell, Thompson, Tessner & Toga, 2001). Synaptic density in the prefrontal cortex reaches its peak at about 2 years of age in humans (Huttenlocher, 1979), and then declines until it finally reaches adult levels at about 16 years of age (Huttenlocher, 1990). Myelination is also still occurring in the frontal lobe between the 2 and 6 years of age (Sampaio & Truwit, 2001). To the extent that these neurophysiological changes lead to stronger and more efficient neural processing in spatial tasks, they are, at least at a qualitative level, consistent with the SPH.

Implementing the spatial precision hypothesis in our model has three components (see Model Parameters in Appendix A and Table A1). First, we have made feed-forward inputs to the model (i.e., activation from the input layer to PF and SWM) stronger and more precise over development (see Scaling of Inputs in Appendix A and Table A1). This reflects both enhanced tuning of feed-forward cortical projections over time due to, for instance, Hebbian processes (Kohonen, 1982), as well as an improved ability to stably align and re-align egocentric and allocentric reference frames (see Spencer, et al., 2007). Note that the values listed in Table A1 were multiplied by the “adult” parameters to arrive at the values used in the developmental simulations. Second, we have made similar modifications to input to the LTM fields (see Scaling of Long Term Memory Projections in Appendix A). For instance, peaks of activation in SWM project activation into LTM_{SWM} via a broader projection early in development relative to the “adult” version of our model. Again, the narrowing of this projection over development reflects improvements in neural precision via learning as well as an improved ability to re-align current memories to past memories.

Third, we have altered the strength and precision of neural interactions among the PF, Inhib, and SWM layers (see Scaling of Local Excitation / Lateral Inhibition in Appendix A and Table A1). In particular, inhibitory projections from the layer of interneurons to both PF and SWM were broad and weak in early development and became more precise to create adult patterns of responding. We also modified local excitatory interactions within PF and SWM, with weaker local excitation early in development and stronger local excitation later. These changes in local excitation and lateral inhibition are depicted in Fig. 2, which shows the locally-excitatory profile within SWM superimposed on the laterally-inhibitory projection from Inhib in early childhood (Fig. 2A) and adulthood (Fig. 2B). In a similar architecture, this type of change in the precision of neural interactions over development predicted developmental changes in BOLD signals (measured with fMRI) during a working memory task (Edin, Macoveanu, Olesen, Tegner & Klingberg, 2007).

In summary, we implemented the SPH by scaling the spatial precision (i.e., widths) and strength of three classes of neural interactions in the model: projections from the input layer into PF and SWM; projections from PF and SWM into the associated LTM fields; and locally excitatory / laterally inhibitory interactions among PF, Inhib, and SWM (see Table A1). Importantly, to systematically move from simulations of infants’ performance in the Piagetian A-not-B task to children’s performance in two types of spatial recall tasks to adults’ performance in spatial recall and position discrimination, we only changed a single scaling parameter: we modulated the strength of local excitation from 0.3 for the young infants in the A-not-B task to 0.4 for older infants in the A-not-B task to 0.5 for all other toddler/child simulations (see Table A1). Critically, these quantitative changes in local excitation combined with the more global shift from the “child” to “adult” parameter sets were sufficient to capture

both quantitative and qualitative changes in performance across tasks and time scales without any other modifications to model parameters.

Note that the four tasks we simulate here have not been explicitly linked in the behavioral literature; rather, they are generally considered to be separate tasks that index separate developing abilities. Concretely, developmental changes in infants' performance in the Piagetian A-not-B task has been used as a measure of the developing object concept (Piaget, 1954), improvements in infants' representation of space (e.g., Acredolo, 1985; Bremner, 1978; Bremner & Bryant, 1977), maturational changes in pre-frontal cortex (e.g., Diamond, 1990a; Diamond, 1990b; Diamond & Goldman-Rakic, 1989), and improvements in infants' memory for objects (e.g., Munakata, 1998; Munakata, McClelland, Johnson & Siegler, 1997), among others. By contrast, performance in the two spatial recall tasks we simulate has been linked to how children use long-term spatial memories and geometric spatial categories to remember locations (e.g., Huttenlocher, Newcombe & Sandberg, 1994; Schutte & Spencer, 2002; Spencer & Hund, 2003). Finally, the position discrimination task we simulate has typically been viewed from a psychophysical perspective (Kinchla, 1971; Palmer, 1986a; Palmer, 1986b) and has only recently been directly linked to phenomena discussed in the spatial recall literature (Simmering, Spencer & Schöner, 2006; see also, Werner & Diedrichsen, 2002).

In the past decade, we have shown that these phenomena can be brought under the same theoretical umbrella—the dynamic field framework—and that this framework can generate novel, testable predictions (Schutte, et al., 2003; Simmering & Spencer, 2007; Simmering, et al., 2006; Spencer, et al., 2007). Here we take this claim one step further by demonstrating—for the first time—that these phenomena can, in fact, be integrated within a single model (the 5-layer version of the DFT shown in Fig. 1) using a well-specified developmental hypothesis.

3. Unifying Behavior across Real and Developmental Time Scales

In the sections that follow, we present simulations of the DFT that capture key behavioral signatures from four tasks that span infancy through adulthood. Note that although we use the same model for all simulations presented here, we focus on different aspects of the model's functioning in the different sections because interactions between SWM and LTM are central to A-not-B-type effects, whereas interactions among PF, Inhib, and SWM are central to effects in recall and discrimination tasks. This shift in focus reflects the changing demands of tasks typically used in infancy compared to those used with older children and adults, but it also reflects the emergence of new abilities in our model that arise due to earlier developmental changes.¹ For instance, more “mature” behaviors such as precise position discrimination depend on prior developing abilities such as the ability to stably sustain a pattern of activation in working memory (though for ties between the mechanism of discrimination in childhood and visual recognition in infancy, see Perone, Spencer & Schöner, in press).

3.1 Piaget's A-not-B Task in Infancy

In Piaget's A-not-B task, an attractive toy is hidden repeatedly at an “A” location, where infants generally search accurately. The toy is then hidden at a nearby (and perceptually similar) “B” location. After a short delay, 8-10-month-old infants tend to search inaccurately at A. Slightly older infants, around 10-12 months, search accurately at the B location. Previous work by Thelen and colleagues (2001) demonstrated that a dynamic field model of reaching behavior could capture infants' A-not-B performance. According to Thelen et al., infants' reaches in this task depend on the interaction between a long-term memory of past reaches to the A

¹We thank an anonymous reviewer for highlighting this aspect of our developmental work.

location and a memory of the cuing event at the B location (see also Diedrich, Highlands, Thelen & Smith, 2001; Diedrich, Thelen, Smith & Corbetta, 2000; Smith, Thelen, Titzer & McLin, 1999). Infants “fail” in this task because the memory of B is not robust and is dominated by the long-term memory of A; infants succeed in this task when they can effectively sustain the memory of B during short-term delays. This developmental shift was originally modeled by changing the resting level of neurons in a dynamic field model. Thelen et al. used a low resting level in early development; consequently, activation patterns in the field were “input-driven”, where activation takes the form of input and decays to a resting level when input is removed. To model later development, a higher resting level made the field dynamics more “cooperative”. Consequently, activation in the field could achieve a self-sustaining state where activation associated with the B location could be sustained in the absence of input.

Although this model effectively captured a host of data and generated novel predictions (e.g., Clearfield, Smith, Diedrich & Thelen, 2006; Diedrich, Clearfield, Smith & Thelen, 2007; Diedrich, et al., 2001; Diedrich, et al., 2000), it is limited in several respects. For instance, it does not address issues of how reference frames are calibrated and re-aligned during the task. Moreover, this account has not been directly integrated with developmental changes in spatial cognition and action that occur beyond the age of 12 months. One way to move beyond these limitations is to integrate this account of A-not-B performance with our more general theory of the development of spatial cognition (Spencer, et al., 2007). We take a step in this direction here by demonstrating that the 5-layer model and our implementation of the SPH can capture developmental changes in the A-not-B task.

Fig. 3 shows simulations of a younger infant’s performance (e.g., age 8 to 10 months; see Fig. 3B,C) and an older infant’s performance (e.g., age 10 to 12 months; see Fig. 3D,E) in the A-not-B task. Fig. 3A shows the input profiles used in both simulations, and the left and right panels show two layers of the 5-layer model: SWM (Fig. 3B,D) and its associated LTM field (Fig. 3C,E). We focus solely on these two layers because the critical behaviors—reaching to A on the first B trial (i.e., making the A-not-B error) and reaching correctly on the first B trial—arise due to activation in these layers. Note, however, that the full 5-layer model was run in all simulations. Moreover, the model was simulated across the entire sequence of trials in the A-not-B task even though we only show results from the first trial to the B location. Simulation details for the younger and older infant model were identical with one exception: the older model had stronger locally excitatory interactions within PF and SWM (see Table A1).

Each simulated trial began with the presentation of the “task” input to the model that captured perception of the box in the task space with two distinct hiding locations. We implemented this task input with two weak constant inputs at the A and B locations (see Hiding Wells in Fig. 3A). Next, we presented a strong input (i.e., the toy) at the target location (A or B) for 1000 time steps (equivalent to 2 s) and then waited a short delay (3000 time steps, equivalent to 6 s). At the end of the delay in the canonical A-not-B task, the box is pushed forward and the infant is allowed to reach to retrieve the toy. To capture this in the model, we increased the strength of the task input after the delay. The model then “reached” to the location associated with the center of mass of the peak of activation in SWM.

We presented the model with four trials to the A location. Like most infants in this task, the model “reached” successfully to A on these trials, leaving a memory trace of the A location in LTM_{SWM} (see LTM of A at the start of the trial in Fig. 3C,E). Next, we presented the younger and older models with a B trial. The presentation of the target at B (see Target in Fig. 3A) formed a peak of activation at the B location in SWM for both models. However, in the younger model (Fig. 3B), neural interactions in SWM were not strong enough to sustain this peak once the input was removed (i.e., when the toy was hidden). As a result, the peak decayed during the delay and activation began to grow near the A location due to continual input from LTM

(see Fig. 3C). Consequently, at the end of the delay when we increased the strength of the task input to prompt a response, a peak formed at the A location and the younger model reached to this location, making the A-not-B error. The older model, by contrast, effectively sustained the peak at B even when the target input was removed (Fig. 3D). Thus, after the delay, this model correctly reached to the B location. Note that the activation peak remained stably aligned with the B location during the delay despite the LTM at the A location (Fig. 3E) due to the continued presence of the task input. This input effectively provided an anchor for the peak in SWM, an issue we return to in the next section.

To summarize, the simulations in Fig. 3 illustrate that relatively modest changes in neural precision in the 5-layer model can capture the qualitative difference in performance between 8 and 10 months in the Piagetian A-not-B task. Although additional work will be needed to probe whether the 5-layer model can capture the full range of effects described by Thelen et al. (2001), the work presented here takes an important first step in this direction. This has the potential to integrate the account proposed by Thelen and colleagues with our work on the development of spatial cognition, bringing new insights to bear on the A-not-B error in infancy, including how the alignment of egocentric and allocentric information might impact infants' performance (see, for instance, Acredolo, 1985; Bremner, 1978; Bremner & Bryant, 1977).

3.2 A Sandbox Version of the A-not-B Task with Toddlers

According to Thelen and colleagues (2001), the A-not-B error reflects the complex, real-time dynamics that underlie reaching behavior. This theoretical perspective predicts that similar patterns of reaching will be evident in other tasks and at later points in development. Spencer, Smith, and Thelen (2001) examined this possibility by altering one important aspect of the A-not-B task: they removed the visible hiding locations and, instead, hid the toy in a rectangular sandbox. Spencer et al. found that 2-, 4-, and 6-year-old children showed biases toward a previously-remembered A location when searching for a toy hidden at a nearby B location, that is, these children made an A-not-B-type error.

We simulated this result in Fig. 4 using the 5-layer model. To approximate the homogeneous task space provided by the sandbox, we removed the static inputs at the hiding locations (see Fig. 4A). In addition, we increased the local excitation from the previous simulations (see Table A1) to reflect 2-year-olds enhanced WM abilities (Schutte, et al., 2003; Spencer, et al., 2001). As with the infant version of the A-not-B task, we began our simulations with four A trials, which the model performed successfully. Then, on the first B trial (Fig. 4B-C) we presented the target at a nearby B location (see Target in Fig. 4A). Note that the distance between A and B in this simulation was the same as in the infant simulations in Fig. 3. As the B peak sustained in SWM during the delay (Fig. 4B), activation at the A location from LTM_{SWM} (Fig. 4C) began to "pull" the memory peak toward A. Thus, at the end of the trial, the model "reached" to a location about halfway between A and B, comparable to the errors made by children in the sandbox task (Schutte, et al., 2003; Spencer, et al., 2001).

The simulation in the right panel of Fig. 4 shows, however, that we can eliminate this bias by increasing the spatial separation between A and B (Fig. 4D-F). In particular, when the target was presented at a more distant B location (see Target in Fig. 4D), the SWM peak did not overlap with activation from LTM_{SWM} , and the model's memory remained accurate throughout the delay (Fig. 4E). Schutte and colleagues reported the same pattern in children: 6-year-olds showed significant A-not-B-type biases when A and B were 2 inches apart, but not when these targets were 6 and 9 inches apart. Importantly, these metric dependencies change over development as predicted by the SPH. For instance, 2- and 4-year-olds show significant A-not-B-type biases at 2, 6, and 9 inch separations. These data have been quantitatively modeled using a one-layer dynamic field model by varying the separation between targets as

well as the strength and precision of local excitation and lateral inhibition over development (Schutte, et al., 2003).

3.3 Spatial Recall in Children and Adults

These connections between infants' performance in Piaget's A-not-B task and toddlers' spatial memory performance are exciting; however, the developmental changes in these examples are rather intuitive—spatial memory becomes more robust and accurate over time. But the 5-layer DFT and SPH also make the less obvious prediction that changes in memory precision should lead to a qualitative change in a second class of spatial memory effects—biases associated with reference frames. Although other models have been proposed to capture developmental changes in reference frame biases (e.g., Huttenlocher, et al., 1994), such models fail to connect performance in recall tasks to the A-not-B findings described above (a point we return to in the concluding section). Moreover, these accounts have not specified a developmental mechanism that explains what leads to qualitative changes in spatial recall abilities.

In a typical spatial recall task, a target is presented within a homogeneous task space, and participants must remember its location for a short delay (e.g., 5-20 s) before reproducing the remembered location. Studies using these tasks have revealed systematic “geometric” biases that change markedly over development. For young children, around 2-3 years of age, responses aligned with visible lines and symmetry axes are accurate, whereas responses to targets on either side of such reference axes are biased *toward* these axes over delay (e.g., Huttenlocher, et al., 1994; Schutte & Spencer, 2002; Schutte, et al., 2003). Older children (6-11 years) and adults, on the other hand, show biases *away from* reference axis over delay (e.g., Engebretson & Huttenlocher, 1996; Spencer & Hund, 2002; Tversky & Schiano, 1989; Werner & Diedrichsen, 2002).

The 5-layer DFT and the SPH are able to capture this qualitative shift in how children and adults use reference frames in spatial recall. Fig. 5 shows simulations of the DFT performing a single spatial recall trial with broad (A-D) versus narrow (E-H) interactions (see Fig. 2 and Table A1). Three layers of the model (PF, Inhib, SWM) are shown in Fig. 5 to clarify the origin of geometric biases in the model. The trial began with a constant, low level input at 0° , reflecting perception of the midline symmetry axis of the task space (see Midline in Fig. 5A,E; note that, for simplicity, we did not include the edges of the task space in these simulations). Next, a target was presented as a strong input to the model at -40° in the task space for 2 seconds (see Target in 5A,E). Recall that these inputs were broader and weaker in the child model (Fig. 5A) to capture the hypothesized broader feed-forward projections in early development as well as an imprecise ability to align egocentric and allocentric reference frames.

The combination of less-precise midline input into PF and weaker interactions early in development (i.e., the profile shown in Fig. 2A) allow the model to capture young children's biases toward midline in recall tasks. With weak interactions early in development, the midline input was not sufficient to form a peak at 0° in PF (Fig. 5B). Instead, throughout the delay, this input provided weak excitatory input to SWM. Consequently, the memory peak in SWM drifted toward midline to about -32° —an 8° “geometric” error toward midline (Fig. 5D).

With stronger interactions later in development, on the other hand, a self-stabilized peak formed at midline in PF after the target input was removed (see Midline Peak in Fig. 5F). This reference peak in PF produced a strong inhibitory profile around midline in Inhib (Fig. 5G). This reference-related inhibition was then projected into SWM, effectively repelling the memory of the target location away from midline during the delay (Fig. 5H). As a result, at the end of the trial, the model responded at about -46° —a 6° geometric error away from midline.

These simulations demonstrate that the DFT can capture the *qualitative* transition in geometric bias over development via the *quantitative* changes in neural precision specified by the SPH. This highlights that the DFT can move beyond results showing that memory simply becomes better toward findings that are much less intuitive. Indeed, we are currently probing whether a systematic scaling of neural interactions step-by-step over development can capture the complex pattern of changes in geometric bias observed between 3 and 6 years (Schutte, 2004; Schutte & Spencer, 2007a).

3.4 Position Discrimination in Children and Adults

Thus far, we have demonstrated that the DFT and SPH can capture infants' and toddlers' performance in two versions of the A-not-B task, as well as children's and adults' biases in spatial recall tasks. This illustrates how we are able to get more from less with the SPH—a simple developmental hypothesis can capture both quantitative and qualitative shifts in memory performance from infancy through adulthood. Although this level of generalization is impressive, each of these tasks involved relatively simple spatial recall responses. Another type of generalization that is critical to achieving a truly flexible real-time system is generalization across tasks. To explore whether the DFT can handle this challenge, we moved to a new behavior—*same/different* judgments—and a new task—position discrimination (Simmering, et al., 2006).

In a typical position discrimination task, two stimuli are presented in quick succession (i.e., 500 ms apart) and the participant responds whether the locations were in the *same* or *different* locations. This task differs from spatial recall in three key ways that must be captured in our model. First, two stimuli are presented instead of one. Second, the time scale is much shorter than typical recall trials. Third, the task requires a *same/different* decision rather than a pointing response. The DFT can adapt to these three differences across tasks with virtually no modification. Because the system operates in real-time, the first two changes are straightforward to implement in the model: we can simply present two inputs at the appropriate delay. Generating a decision presents a more substantive challenge; as we demonstrate below, however, the dynamic interplay between PF and SWM can lead to emergent *same* and *different* decisions.

Fig. 6 shows simulations of *same* (A-D) and *different* (E-H) discrimination trials using the “adult” interaction profile shown in Fig. 2B. *Same/different* responses in the model depend on the spatial overlap between the input to PF as the second stimulus (S2) is presented and the location of the SWM peak associated with the first stimulus (S1). The simulation in the left column shows a case when S1 and S2 overlap, that is, when the two stimuli were presented in the same location. The trial begins with the presentation of S1 at -40° for 1 second (500 time steps; see Fig. 6A). As in the recall simulations in Fig. 5, this event builds a target peak in SWM (Fig. 6D) that is sustained during the brief memory delay (500 ms or 250 time steps). The peak in SWM is supported by inhibition from the Inhib layer (Fig. 6B) which also projects inhibition back to PF. When S2 is presented in the same location as the peak in SWM, this inhibition suppresses the in-coming activation in PF (Fig. 6B)—essentially, the system “decides” that it already has S2 in working memory. Consequently, S2 fails to build a peak in PF (Fig. 6B), the SWM peak is maintained, and the model makes a *same* response (Fig. 6D).

The simulation in the right column shows how the model responds when S1 and S2 are presented in different locations. This simulation is identical to the *same* simulation up to the point where S2 is presented—now S2 is shifted to the right relative to the location of S1 (Fig. 6E). Because S2 does not overlap spatially with the working memory of S1, its activation falls outside of the inhibited range of PF and builds a *different* peak (Fig. 6F). This, in turn, suppresses the SWM peak associated with S1 (see Fig. 6H) due to the shared inhibition

generated by the S2 peak in PF. Consequently, the simulation ends with a peak in PF and no above-threshold peak in SWM—the basis for a *different* response.

The simulations in Fig. 6 highlight that the dynamic interactions among layers in the DFT lead to emergent behaviors that generalize beyond the simple spatial recall behaviors that were central to the development of this theory. Importantly, the mechanism for discrimination illustrated in Fig. 6 has generated a host of novel predictions that have been successfully tested with both children (Simmering & Spencer, 2007) and adults (Simmering, et al., 2006). Moreover, we have recently extended this mechanism to account for performance in a canonical change detection task (Johnson, Spencer & Schöner, in press). Although the details of this work are beyond the scope of the present report, we note that the mechanisms described here that underlie discrimination and self-sustained activation can be effectively generalized to a situation where multiple items (i.e., multiple peaks) must be simultaneously maintained in working memory and compared to a multi-item test array in the task space.

4. Conclusion and Future Directions

The goal of this paper was to demonstrate—for the first time—that a 5-layer dynamic field theory of spatial cognition can capture the real time dynamics of performance in a variety of tasks and connect performance in these tasks across development with a simple developmental hypothesis. We demonstrated that changes in spatial precision can capture developmental changes in A-not-B errors, reference effects in spatial recall, and position discrimination performance, spanning the age of 8 months to 6 years and into adulthood. Model simulations demonstrate that the DFT can account for the details of performance across tasks while maintaining a commitment to neural principles. Considered together, then, the simulations in the present report demonstrate that the DFT can achieve generality with specificity. In the sections that follow, we evaluate our model of spatial cognition by considering three central issues: is the DFT strongly grounded in neural principles, how does this theory compare to other existing models of the phenomena presented here, and what insights does this model offer for understanding development and what are the future challenges regarding developmental change?

4.1 The DFT Is Grounded by Neural Principles

In the overview of the DFT presented earlier, we discussed several ways in which the DFT is grounded by neural principles. We return to that discussion here to link dynamic field models more tightly to issues central to the field of computational cognitive neuroscience.

As discussed above, there is a demonstrated link between a population dynamics approach to cortical activation and patterns of activation in dynamic fields, as well as clear methods that can be used to map single-unit recordings onto dynamic population representations that can be directly compared to dynamic field models (Erlhagen & Schöner, 2002). Importantly, this link has been tested in both motor cortex and visual cortex (Bastian, et al., 1998; Bastian, et al., 2003; Erlhagen, et al., 1999; Jancke, et al., 1999). There is also evidence that space is neurally represented across continuous, metric feature dimensions, though it is important to emphasize that this topography is functional rather than anatomical in most cortical areas (di Pellegrino & Wise, 1993; Georgopoulos, Lurito, Petrides, Schwartz & Massey, 1989; Georgopoulos, Taira & Lukashin, 1993; Wilson, Scalaidhe & Goldman-Rakic, 1993). Similarly, there is strong evidence from studies of cortical neurons for the basic local excitation / lateral inhibition form of neural interaction used in dynamic field models (e.g., Durstewitz, Seamans & Sejnowski, 2000). Moreover, because cortical neurons never project both excitatorily and inhibitorily onto targets, the inhibitory lateral interaction must be mediated through an ensemble of interneurons. We used a generic, two-layer formulation (Amari & Arbib, 1977) to realize this interaction

where an inhibitory activation field receives input from an excitatory activation field and in turn inhibits that field.

Next, as we noted previously, the type of interactivity among layers used in our 5-layer model was inspired by studies of the cytoarchitecture of visual cortex (Douglas & Martin, 1998). This work on visual cortex played a particularly strong role in our implementation of interactions among the three primary layers (PF, Inhib, SWM). That said, interaction among these layers could be implemented in other ways. For instance, the same type of interaction could be achieved using a four layer structure, with separate inhibitory layers associated with PF and SWM. We chose to use a shared layer of interneurons because it is both more parsimonious (i.e., 3 layers vs. four layers) and more strongly constrained (i.e., fewer “free” parameters). Finally, our approach to the integration of dynamics across multiple time scales (i.e., real-time and learning time) is consistent with basic forms of Hebbian learning (see Wilimzig & Schöner, 2005; Wilimzig & Schöner, 2007), though our implementation of Hebbian learning using activation in long-term memory fields (rather than through synaptic weight changes) is somewhat atypical. In each of these cases, then, the dynamic field framework retains strong ties to known neural principles.

Less clear in our approach is a commitment to neural localization. Here we are, frankly, riding the fence. On one hand, several characteristics of dynamic activation patterns in our model reflect processes that have been linked to different cortical areas. For instance, the self-sustaining working memory peaks discussed here can survive intervening presentations of stimuli, a characteristic of spatial representations in dorso-lateral prefrontal cortex (di Pellegrino & Wise, 1993). By contrast, although activation patterns in our perceptual field can be sustained under some conditions (see the generation of a “different” response in Fig. 6), sustained activation in this field is generally disrupted by the presentation of new stimuli, a characteristic of spatial representations in parietal cortex (Constantinidis & Steinmetz, 1996; Steinmetz & Constantinidis, 1995). Given that there is a mapping between the dynamic properties of activation patterns in our model and the localization of brain function, why are we hesitant to assign cortical labels to our layers? This reflects, in part, our grounding in dynamic systems theory with its emphasis on the collective behavior of a system (see Kelso, 1997; Schöner & Kelso, 1988; Thelen & Smith, 1994), as well as our commitment to the massive reentrancy and interconnectedness across cortical circuits in the brain (Fuster, 1995) and the profound potential for reorganization and brain plasticity in early development (e.g., Stiles, Bates, Thal, Trauner & Reilly, 1998). In short, these conceptual biases make us hesitant to embrace a strong form of localization.

Although we remain uncommitted to a strong form of localization, it is still possible to generate predictions from dynamic field models that can be tested using technologies that rely on localized neural signals (Bastian, et al., 1998; Bastian, et al., 2003; Jancke, et al., 1999; see also Edin, et al., 2007, for fMRI predictions using a similar architecture and developmental mechanism). For example, the DFT framework has generated novel predictions regarding links between behavioral and electrophysiological measures of performance in a motor planning task. Specifically, McDowell and colleagues (McDowell, et al., 2002) measured reaction times and event-related potentials (ERPs) during a two-choice pointing task where they varied both the angular separation between the two targets (narrow separation of 20° versus wide separation of 120°) and the probability of the two responses (frequent versus infrequent). According to a dynamic field theory of movement preparation (see also, Erlhagen & Schöner, 2002), reaction times to the wide-infrequent target should be slower compared to the other three targets. Based on the mechanism that underlies this effect in the model, McDowell et al. predicted that there should be a related increase in the P300 ERP component which has been linked to detection of “subjectively rare” events (e.g., Johnson, 1993). Results showed the predicted reaction time-P300 relationship.

Moving beyond the issue of brain localization, our commitment to dynamic systems theory also factors in to our emphasis on sustained activation peaks or “bump” attractors (Amari, 1977; see also, e.g., Vogels, Rajan & Abbott, 2005). One advantage to using dynamic field models is that this class of neural networks have been quasi-formally analyzed (Amari, 1977). Thus, we have a good understanding of the attractor states dynamic field models can realize, and the bifurcations the model moves through as one attractor state becomes unstable (e.g., the input-driven state) and another attractor state forms (e.g., the self-sustaining state). We can (and routinely do!) probe these attractor states across repeated simulations by manipulating, for instance, the presence or absence of inputs (e.g., does activation sustain after input is removed?), as well as evaluating the stability of the attractor (e.g., does the model remain in the attractor state as noise is increased?). Moreover, we know which model parameters are central to the stability of these attractor states (e.g., the relative strengths of excitatory and inhibitory projections among layers), and, thus, understand why the model works the way it does, as well as the critical factors that influence the dynamics of the network. This detailed understanding of the attractor dynamics of the model is not a standard feature of theoretical frameworks within computational cognitive neuroscience. Indeed, this characteristic of dynamic fields is what initially drew us to work by Amari and colleagues (Amari, 1977; Amari, 1989; Amari & Arbib, 1977).

Importantly, the issue of forming attractor states, that is, achieving a stable pattern of activation through time is a central challenge in a massively interconnected nervous system (for discussion see Spencer & Schöner, 2003). In this context, it is important to note that other researchers have probed the stability properties of bump attractors using a more biophysical style of neural modeling (Compte, Brunel, Goldman-Rakic & Wang, 2000; Wang, 2001). In our view, both approaches demonstrate that stable peaks of activation provide a viable neural mechanism for the formation of “working” memories. Indeed, our empirical work shows direct behavioral signatures (e.g., time-dependent “drift”) of this basic stabilization mechanism. In this context, we note that although all of the simulations reported here had a single “bump” present in working memory at any moment in time, the dynamic field framework can be naturally extended to capture multi-item working memory, that is, the simultaneous representation of multiple bump attractors within a single field (Johnson, Spencer & Schöner, 2006; Johnson, et al., in press). More generally, we envision that multiple neural patterns could be simultaneously represented via collections of peaks distributed across many cortical layers. We are currently pursuing this view as an entry point into the representation of multiple objects distributed across both spatial and non-spatial feature dimensions (Faubel & Schöner, 2007; Johnson, et al., in press).

4.2 Comparing the DFT to Other Models

The fact that the DFT captured performance across tasks and time scales with relatively simple developmental changes suggests continuity across behaviors that have previously been considered in separate literatures. Indeed, in each of these literatures, separate models have been proposed to capture infants’, children’s, and adults’ performance. In this section, we contrast the DFT with several existing models.

One connectionist model that has captured a variety of effects in the canonical A-not-B task is Munakata’s PDP model (Munakata, 1998). This model accounts for infants’ performance in a manner comparable to the account proposed by Thelen et al. (2001) and the simulations reported here. For instance, developmental change in the PDP model is achieved through strengthening connections within the hidden layer, which effectively produces a more robust “active” memory for the current trial. This increased ability to sustain activation is similar to the increase in local excitation in our model.

Although Munakata's model is certainly close in spirit to the model by Thelen et al. and our own, there are some substantive differences (for discussion see Munakata & McClelland, 2003; Spencer & Schöner, 2003). First, activation in the PDP model is "chunked" into discrete spatial nodes rather than being distributed across continuous, metric dimensions. This is a fine simplification in the context of the A-not-B paradigm, but prevents generalization to the other spatial tasks described here. It is possible this limitation could be overcome; however, adding continuous metrics to the PDP framework is a non-trivial endeavor and is likely to be insufficient to capture, for instance, the time-dependent "drift" observed in our simulations of spatial recall (see Fig. 4 and 5). Such drift is a consequence of the local excitation/lateral inhibition form of neural interaction in the DFT (see also Compte, et al., 2000).

Another difference is that the PDP model does not have the same type of interactivity across layers present in our 5-layer model of spatial cognition. Such interactivity plays a central role in the generalization of our model across response types. For instance, mutual interactivity between the perceptual and working memory layers is central to performance in position discrimination (see Fig. 6). Importantly, we have strong evidence for this type of interactivity early in development: as peaks drift in working memory in early development (3 to 5 years), this creates predictable asymmetries in discrimination performance (Simmering & Spencer, 2007).

Two final differences are worth noting because they are linked to the issue of stability discussed above (see Spencer & Schöner, 2003). First, the PDP model provides only a limited view of the response generation process in that, on each trial in the A-not-B paradigm, *group* choice probabilities are estimated from the amount of activation across response nodes after some settling time (see Mareschal, 1998, for a similar critique). By contrast, DFT models of A-not-B performance generate *individual* choice responses on each trial—the model picks A or B. Through simulations of multiple trials in the context of stochastic fluctuations, group response probabilities can be computed and fit to empirical results (Dineva, 2005; Thelen, et al., 2001). It is not clear whether the PDP model could make similar achievements because group and individual choice responses are conflated (for related issues regarding response generation and the need for stability, see Schöner & Dineva, 2007; Spencer & Schöner, 2003).

Geometric biases in spatial recall have been linked to children's and adults' ability to divide space into geometric categories (e.g., Huttenlocher, et al., 1994). The Category Adjustment (CA) model has been proposed by Huttenlocher and colleagues to capture how categorical information influences memory along continuous dimensions, mostly notably space (Huttenlocher, Hedges & Duncan, 1991), but also time (Huttenlocher, Hedges & Prohaska, 1988) and continuous feature dimensions like size and shade of gray (Huttenlocher, Hedges & Vevea, 2000). According to this model, people combine or "weight" fine-grained detail about individual target items with information about the geometric category within which the items were placed. Under conditions of uncertainty (e.g., after long memory delays), categorical information is weighted more heavily leading to a bias in recall responses toward the center of the relevant spatial category. To account for developmental changes in spatial recall, Huttenlocher et al. (1994) have proposed that young children treat large spaces such as a tabletop as one large category with a prototype at the center. Consequently, young children show recall biases toward the center of the space. Older children and adults, by contrast, subdivide space into two categories with prototypes at the centers of the left and right halves of the tabletop. Consequently, they show recall biases away from the center and toward the centers of the left and right halves.

Although the CA model can capture the qualitative pattern of results observed in recall studies with children and adults, this model has several limitations relative to the DFT (see also Spencer, Simmering & Schutte, 2006). First, the CA model is not a process model: it does not

specify the process that gives rise to spatial memories (e.g., perceptual processes involved in the formation of activation peaks in the DFT), how memories for location evolve in real time (e.g., peak drift in the DFT), the process that gives rise to the formation of spatial categories (e.g., the emergence of spatial categories in LTM_{SWM} and the use of perceived reference frames in the DFT), and the process that gives rise to developmental changes in spatial recall abilities (e.g., the spatial precision hypothesis). Moreover, the CA model offers a limited account of the details of behavioral performance. For instance, the model predicts that variance in recall responses near reference axes (i.e., category boundaries) should be high, while variance near spatial prototypes should be low. Empirical data demonstrate that the opposite pattern is observed in adults' recall responses (see Spencer & Hund, 2002). Importantly, the observed empirical pattern of variance is consistent with the DFT: variance is low near reference axes in our model due to the presence of continued reference input during the memory delay which helps stabilize peaks in SWM (much like the task input stabilized SWM peaks in simulations of the A-not-B model presented earlier).

Both Munakata's PDP model and the CA model can capture aspects of the developmental profiles simulated here, and both models have been extended to other tasks not currently addressed within our framework (e.g., Huttenlocher, et al., 2000; Morton & Munakata, 2002). Importantly, however, these cases of generalization are closely tied to tasks that rely on the same or similar response types—selection of a limited number of choices in the case of Munakata's model, and estimation along a continuous dimension for the CA model. The DFT, by contrast, has shown a different type of generalization by also capturing different types of responses across tasks—reaching to a visibly-specified location in the infant A-not-B task, pointing in continuous space for the sandbox and recall tasks, and generating *same/different* decisions in position discrimination (for additional extensions, see Spencer, et al., 2006). This third example is the most dramatic in that the processes that underlie *same/different* discrimination are typically conceptualized using a signal-detection framework. For instance, Kinchla (1971) proposed a signal-detection model to account for some of the discrimination results reported here (e.g., enhanced discrimination near a reference, see Simmering, et al., 2006). Importantly, our process model of discrimination has moved beyond the signal detection framework by generating several novel predictions about how the memory of the first stimulus should “drift” differentially between stimulus presentations depending on the location of the stimulus relative to a reference axis (Simmering, et al., 2006), as well as how such effects should change over development (Simmering & Spencer, 2007).

Thus far, we have contrasted our model with other models of the specific phenomena we presented in this paper. As a final contrast, we consider other approaches to working memory within the computational cognitive neuroscience framework. Specifically, a growing class of models capture working memory using a neural reverberation or neural synchrony mechanism (e.g., Lisman & Idiart, 1995). In these models, sustained oscillations code for the presence or absence of items in working memory. This mechanism has been used to capture elements of working memory such as capacity limitations and feature binding (Raffone & Wolters, 2001). Such models offer a clearly different neural mechanism for working memory than the one used here, and an important future goal is to contrast these mechanisms directly. Toward this end, it is unclear whether models that rely on neural synchrony could account for some of the effects reported in our studies of spatial recall. For instance, it is not clear how a neural synchrony mechanism could produce delay-dependent “drift” in the contents of working memory given that working memory in such models is not defined over a continuous, metric dimension. Clearly, a vigorous comparison between these broad classes of models is needed to reveal the strengths and weaknesses of each.

4.3 Insights into Development and Challenges on the Horizon

The simulations presented here highlight an emerging “more from less” theme in our developmental work: by developing a detailed, real-time account of behavior across multiple tasks, we require less from developmental process. Thus, a coherent set of relatively simple developmental changes effectively captured key aspects of performance across the four tasks discussed here, spanning the ages of 8 months to 6 years and into adulthood with the *same model*, without introducing any new components or processes. Indeed, our simulations required only modest changes in local excitation to capture changes across tasks between 8 months and 6 years. This is, most certainly, an overly simplistic view of the changes in spatial cognition taking place within this age range, and we are confident that more substantive changes would be required to provide a detailed account of all of the phenomena discussed here. That said, the goals of the present report were more modest: to provide the first demonstration proof that the phenomena discussed here could be linked together within a single framework.

Although the simulations reported here to demonstrate a “more from less” view of developmental change, a critical question is whether the “more” we built into our model was reasonable. For instance, is “less” required from development because we have built an overly-complicated model that is simply too powerful or underconstrained? Consider what is known about the structure we built in to the model (as described in Section 4.1): cortical maps are well documented and change over development (e.g., Kohonen, 1982; Miikkulainen, Bednar, Choe & Sirosh, 1999); the layered structure of cortex is evident prenatally (e.g., Rakic, 1995); receptive fields become tuned with experience (e.g., Clark, Allard, Jenkins & Merzenich, 1988). In each of these cases, the structure we built into our model is consistent with studies of both prenatal and postnatal development and is grounded in a neural principles.

Is our model overly complicated? Here, we contend that complexity is in the eye of the beholder. From our vantage point, the DFT is less complicated than many biophysical models that implement bump attractors (e.g., Wang, 2001). That said, the DFT is, in some ways, more complicated than a typical connectionist architecture. To get a better sense of this comparison, consider a few ways in which we could modify the model presented here, yet retain the characteristics of the model shown in the simulations. First, we could limit the number of neurons in each layer to 50 or so. Next, we could get rid of our LTM fields and, instead, implement long-term memory using changes in synaptic weights among neurons in PF and SWM. This would reduce our 5-layer model to a more typical 3-layer structure. The resulting 3-layer model would still have more pre-structure than the standard connectionist model. Note, however, that much of this pre-structure arises due to our commitment to neural principles. For instance, the 2-layered architecture we use to implement lateral inhibition is a direct result of known cortical constraints on how populations of neurons interact. Connectionist models typically do not implement local excitation/lateral inhibition in this neurally-grounded way. Given this, we view much of the apparent complexity of our model as an important feature that moves us closer to understanding brain-behavior relations.

In summary, then, we contend that the model presented here does shed new light on how a richly-structured real-time system requires less from developmental process. Although this is an exciting insight into development, the DFT provides only a limited window on this issue—our model identifies what is changing, but not precisely how those changes are taking place. As we look to the future, therefore, a clear challenge will be to “close the loop” on development (Simon, 1962) within our framework to explore the origins of the spatial precision hypothesis. Put concretely, the parameter changes reported in Table A1 were all implemented “by hand” for the present simulations. It will be important in future efforts to probe whether such changes can emerge from, for instance, Hebbian processes involved in the tuning of cortical maps (Bednar & Miikkulainen, 2007; Kohonen, 1982; Miikkulainen, et al., 1999). We have taken steps in this direction by implementing a reference frame alignment and calibration process in

the “adult” model (Spencer, et al., 2007). This process will play a central role in exploring the consequences of changes in children’s ability to align and re-align egocentric and allocentric reference frames.

Acknowledgements

We would like to thank our collaborator on this work, Gregor Schöner; earlier collaborators Esther Thelen and Linda Smith; as well as Evelina Dineva, Alycia Hund, Jeff Johnson, John Lipinski, Sammy Perone, Larissa Samuelson, and Wendy Troob for their valuable contributions to the development of the dynamic field theory and related empirical studies. We would also like to thank the children, parents, and members of the community who participated in the research described here, as well as the numerous undergraduate research assistants who helped with data collection.

This research was funded by NIMH R01 MH62480, NSF BCS 00-91757, and NSF HSD 0527698, all awarded to John P. Spencer; NIMH 1 F31 MN066595-01 awarded to Anne R. Schutte; and NICHD R01 HD22830-20 awarded to Esther Thelen. Presentation of this work at the 2nd Computational Cognitive Neuroscience Conference was made possible by a conference student travel fellowship awarded to Vanessa R. Simmering.

Appendix A

A.1 Model Equations

In this section, we describe the equations that govern activation in each layer of the 5-layer model 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:

$$\tau \dot{u}(x, t) = -u(x, t) + h_u + \int c_{uu}(x - x') \Lambda_{uu}(u(x', t)) dx' - \int c_{uv}(x - x') \Lambda_{uv}(v(x', t)) dx' + c_{uim} u_{lim}(x, t) + S_{ref}(x, t) + S_{task}(x, t) + S_{tar}(x, t) \quad (1)$$

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 determine 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.

PF also receives input from an associated long-term memory field, u_{lrm} , with strength c_{uultm} . 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 (see 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, three inputs passed activation to PF: the reference input (e.g., perception of the midline axis in the recall and discrimination tasks), the task input (e.g., the hiding wells in the infant A-not-B 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:

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

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 next layer of the model, LTM_{PF} (u_{lrm}), is governed by the following equation:

$$\tau_{ulrm} \dot{u}_{lrm}(x, t) = [-u_{lrm}(x, t) + \int c_{ulrm} u(x - x') \Lambda_{ulrm}(u(x', t)) dx'] \Theta \left(\int \Theta(u(x'', t)) dx'' \right) \quad (5)$$

This equation specifies the rate of change of activation, $\dot{u}_{lrm}(x, t)$, for each neuron in the long-term memory layer across the spatial dimension, x , as a function of time, t . A separate constant τ_{lrm} determines the time scale of the dynamics and is set such that the LTM field builds activation much more slowly (i.e., $\tau_{lrm} \gg \tau$). As above, the first component that contributes to the rate of change is the current activation in the field, $-u_{lrm}(x, t)$, at each site x . Projections to this field come from PF (u), as defined by $\int c_{ulrm} u(x - x') \Lambda_{ulrm}(u(x', t)) dx'$. The contribution of LTM_{PF} to the activation profile in PF is one-to-one where c_{uultm} specifies the strength of the input from $u_{lrm}(x, t)$ to PF (see Equation 1). The final term, $\Theta \left(\int \Theta(u(x'', t)) dx'' \right)$, is multiplied by the equation to essentially turn the memory trace mechanism on and off. Specifically, if activation in PF is above zero at any location x , then the long-term memory field is engaged and $\Theta(u) = 1$; otherwise, this field is held constant and $\Theta(u) = 0$. This implements a form of Hebbian learning (Schöner, 2007; Wilimzig & Schöner, 2007; Wilimzig & Schöner, 2005).

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

$$\tau \dot{v}(x, t) = -v(x, t) + h_v + \int c_{vu} v(x - x') \Lambda_{vu}(v(x', t)) dx' + \int c_{vw} v(x - x') \Lambda_{vw}(w(x', t)) dx'. \quad (6)$$

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,

$\int c_{vu} v(x - x') \Lambda_{vu}(v(x', t)) dx'$, and one from SWM,

$\int c_{vw} v(x - x') \Lambda_{vw}(w(x', t)) dx'$. As described above, these projections are defined by the convolution of a Gaussian kernel (Equation 2) with a sigmoidal threshold function (Equation 3).

Next is the SWM layer (w) governed by the following equation:

$$\begin{aligned} \tau \dot{w}(x, t) = & -w(x, t) + h_w + \int c_{ww} w(x - x') \Lambda_{ww}(w(x', t)) dx' \\ & - \int c_{wv} w(x - x') \Lambda_{wv}(v(x', t)) dx' + \int c_{wu} w(x - x') \Lambda_{wu}(u(x', t)) dx' \\ & + c_{wlm} w_{lrm}(x, t) + c_{sref} S_{ref}(x, t) + c_{stask} S_{task}(x, t) + c_{star} S_{tar}(x, t). \end{aligned} \quad (7)$$

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,

$\int c_{wv} (x - x') \Lambda_{wv} (v(x',t)) dx'$, and input from PF,

$\int c_{wu} (x - x') \Lambda_{wu} (u(x',t)) dx'$. This field also receives input from an associated long-term memory field, w_{ltm} , with strength c_{wwlwm} , and direct reference, $S_{ref}(x, t)$, task, $S_{task}(x, t)$, and target, $S_{tar}(x, t)$, inputs scaled by c_s .

Lastly, $LTM_{SWM}(w_{lwm})$ is governed by the following equation:

$$\tau_{wlm} \dot{w}_{lwm}(x,t) = [-w_{lwm}(x,t) + \int c_{wlmw} (x - x') \Lambda_{wlmw} (w(x',t)) dx'] \Theta \left(\int \Theta(w(x'',t)) dx'' \right) \quad (8)$$

This equation is identical to Equation 5. As with LTM_{PF} , the contribution of $LTM_{SWM}(w_{lwm})$ to the activation profile in SWM (w) is one-to-one where c_{wwlwm} specifies the strength of the input from $w_{lwm}(x,t)$ to w (see Equation 7). The final term, $\Theta \left(\int \Theta(w(x'',t)) dx'' \right)$, is multiplied by the equation to turn the memory trace mechanism on and off as described above.

A.2 Model Parameters

This section provides the parameter values used in all simulations. Table A1 shows the parameter values for the primary layers of the model for the “adult” simulations: PF (Equation 1), Inhib (Equation 6), SWM (Equation 7). The β parameter, used to threshold activation in all projections, was always set to 5. In addition, the resting level for the Gaussian kernel, k , was generally set to 0 with the few exceptions noted in Table A1. Lastly, the projections to and from each long-term memory layer were as follows: $c_{ultmu} = c_{uultm} = c_{wltmw} = c_{wwlwm} = 0.05$; $\sigma_{ultmu} = \sigma_{uultm} = \sigma_{wltmw} = \sigma_{wwlwm} = 10$.

To capture development in our model, we implemented the spatial precision hypothesis by scaling the spatial precision (i.e., widths) and strengths of three classes of interactions in the model: projections from the input layer to PF and SWM; projections from PF and SWM into the associated LTM fields; and locally excitatory / laterally inhibitory interactions among PF, Inhib, and SWM. The details of these changes are described separately below.

A.2.1 Scaling of Local Excitation / Lateral Inhibition

To capture developmental changes in neural precision, we modified the local excitation / lateral inhibition profile as shown in Fig. 2A. For this change, we added the scaling parameter, dev_c_{self} , to scale self-excitatory projections. Local excitation was modified for both PF (c_{uu}) and SWM (c_{ww}). We also added to parameters, dev_c_v and dev_c_v , to scale the strength and width of inhibitory projections, respectively. The lateral inhibition scaling parameters were multiplied by the corresponding projection parameters from Inhib to PF (c_{vu} , σ_{vu}) and from Inhib to SWM (c_{vw} , σ_{vw}). The scaling parameter values are specified in the bottom portion of Table A1. Note that lateral inhibition did not change across developmental simulations; only local excitation increased with age.

A.2.2 Scaling of Inputs

We added two developmental scaling parameters to each set of inputs to modulate the strengths (c) and widths (σ) of inputs: dev_c_{ref} and dev_c_{ref} for reference inputs; dev_c_{task} and dev_c_{task} for task inputs; dev_c_{tar} and dev_c_{tar} for target inputs. Values of these parameters are shown in the bottom portion of Table A1. Note that the reference and task inputs were used only in specific simulations given the experimental details of the tasks we modeled and known

empirical effects. Moreover, note that the strength and width of the reference input for the recall task were weak and broad relative to the task input. This reflects the fact that the reference in this task is defined by the midline symmetry axis of the task space. Research has shown that symmetry axes are perceived more weakly than visible lines (e.g., Palmer & Hemenway, 1978; Wenderoth & van der Zwan, 1991), which is reflected by the weaker dev_c_{ref} . In addition, we propose that children have relatively imprecise (compared to adults) alignment of egocentric and allocentric reference frames (Spencer, et al., 2007); as a result, the reference input is spread out in space (relative to the more precisely specified task input). We approximate this by using a broader dev_c_{ref} .

A.2.3 Scaling of Long-Term Memory Projections

One final developmental parameter was added to approximate the decreased precision of long-term memory in early development: $dev_c_{lrm} = 2$ was multiplied by the widths of the Gaussian projections into the two long-term memory fields, $\sigma_{u_{lrm}u}$ and $\sigma_{w_{lrm}w}$. This parameter did not vary across developmental simulations.

Table A1
Parameter values for simulations

Layer	τ	h	self-excitation	excitatory projection(s)	inhibitory projection(s)
u (PF)	80	-7	$c_{uu} = 1.25$ $\sigma_{uu} = 3$		$c_{uv} = 1.1$ $\sigma_{uv} = 5$ $k_{uv} = 0.05$
v (Inhib)	10	-12		$c_{vu} = 4.38$ $\sigma_{vu} = 5$ $c_{vw} = 2.2$ $\sigma_{vw} = 6$ $c_{wu} = 2.05$ $\sigma_{wu} = 5$	
w (SWM)	80	-4	$c_{ww} = 1.64$ $\sigma_{ww} = 5$		$c_{wv} = 38$ $\sigma_{wv} = 0.665$ $k_{wv} = 0.05$
developmental scaling parameters			$dev_c_{self} = 0.3$ (8-10 mos.) $dev_c_{self} = 0.4$ (10-12 mos.) $dev_c_{self} = 0.5$ (>12 mos.)		$dev_c_v = 0.1$ $dev_c_w = 0.1$

Note. The top portion of the table shows parameter values used for all simulations described as “adult” parameters (see Fig. 1, Fig. 2B, Fig. 3E-H, and Fig. 4). The bottom portion shows developmental scaling parameters for the remaining simulations, which were multiplied by the “adult” parameters listed above.

^aBecause we only simulated the spatial recall task with the adult parameters, task input was not used. We provide the values here that were used in conjunction with the developmental scaling parameters for the infant A-not-B task.

^bTarget inputs for the discrimination task were weaker, reflecting the difference in salience of the dots used in this task (compared to attractive toys in A-not-B). For the discrimination simulations, $c_{tar} = 20$. All other parameters remained unchanged.

References

- Acordolo, L.P. Coordinating perspectives on infant spatial orientation. In: Cohen, R., editor. *The Development of Spatial Cognition*. Erlbaum; Hillsdale, NJ: 1985.
- Amari S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics* 1977;27:77–87. [PubMed: 911931]
- Amari, S. Dynamical stability of formation of cortical maps. In: Arbib, M.A.; Amari, S., editors. *Dynamic Interactions in Neural Networks: Models and Data*. Springer; New York, NY: 1989.

- Amari, S.; Arbib, MA. Competition and cooperation in neural nets. In: Metzler, J., editor. *Systems Neuroscience*. Academic Press; New York: 1977.
- Awh E, Jonides J, Smith EE, Buxton RB, Frank LR, Love T, Wong EC, Gmeindl L. Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science* 1999;10:433–437.
- Barsalou LW. Perceptual symbol systems. *Behavioral & Brain Sciences* 1999;22:577–660. [PubMed: 11301525]
- Bastian A, Riehle A, Erlhagen W, Schöner G. Prior information preshapes the population representation of movement direction in motor cortex. *NeuroReport* 1998;9:315–319. [PubMed: 9507975]
- Bastian A, Schöner G, Riehle A. Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience* 2003;18:2047–2058. [PubMed: 14622238]
- Bednar, JA.; Miikkulainen, R. Constructing visual function through prenatal and postnatal learning. In: Mareschal, D.; Johnson, M.; Sirois, S.; Spratling, M.; Thomas, MSC.; Westermann, G., editors. *Neuroconstructivism: How the Brain Constructs Cognition*. Oxford University Press; 2007.
- Bicho E, Mallet P, Schöner G. Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research* 2000;19:424–447.
- Bremner JG. Spatial errors made by infants: Inadequate spatial cues or evidence of egocentrism? *British Journal of Psychology* 1978;69:77–84. [PubMed: 626806]
- Bremner JG, Bryant PE. Place versus response as the basis of spatial errors made by young infants. *Journal of Experimental Child Psychology* 1977;23:162–177. [PubMed: 839157]
- Clark SA, Allard T, Jenkins WM, Merzenich MM. Receptive fields in the body-surface map in adult cortex defined by temporally correlated inputs. *Nature* 1988;332:444–445. [PubMed: 3352741]
- Clearfield MW, Smith LB, Diedrich FJ, Thelen E. Young infants reach correctly on the A-not-B task: On the development of stability and perseveration. *Infant Behavior & Development* 2006;29:435–444. [PubMed: 17138296]
- Cohen JD, Servan-Schreiber D. Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review* 1992;99:45–77. [PubMed: 1546118]
- Compte A, Brunel N, Goldman-Rakic PS, Wang X-J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex* 2000;10:910–923. [PubMed: 10982751]
- Constantinidis C, Steinmetz MA. Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology* 1996;76:1352–1355. [PubMed: 8871242]
- di Pellegrino G, Wise SP. Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *Journal of Neuroscience* 1993;13:1227–1243. [PubMed: 8441009]
- Diamond, A. Development and neural bases of AB and DR. In: Diamond, A., editor. *The Development and Neural Bases of Higher Cognitive Functions*. National Academy of Sciences; New York: 1990a.
- Diamond, A. Developmental time course in human infants and infant monkeys, and the neural bases of inhibitory control in reaching. In: Diamond, A., editor. *The Development and Neural Bases of Higher Cognitive Functions*. National Academy of Sciences; New York: 1990b.
- Diamond A, Goldman-Rakic PS. Comparison of human infants and rhesus monkeys on Piaget's AB task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research* 1989;74:24–40.
- Diedrich FJ, Clearfield MW, Smith LB, Thelen E. Motor memory influences performance in the A-not-B task: Effects of arm-loading and posture. Manuscript in preparation. 2007
- Diedrich FJ, Highlands T, Thelen E, Smith LB. The role of target distinctiveness in infant perseverative reaching errors. *Journal of Experimental Child Psychology* 2001;78:263–290. [PubMed: 11222002]
- Diedrich FJ, Thelen E, Smith LB, Corbetta D. Motor memory is a factor in infant perseverative errors. *Developmental Science* 2000;3:479–494.
- Dineva, E. Institut für Neuroinformatik. Ruhr-Universität Bochum; Bochum, Germany: 2005. Dynamical field theory of infant reaching and its dependence on behavioral history and context.
- Douglas, R.; Martin, K. Neocortex. In: Shepherd, GM., editor. *The Synaptic Organization of the Brain*. Oxford University Press; New York: 1998.

- Durstewitz D, Seamans JK, Sejnowski TJ. Neurocomputational models of working memory. *Nature Neuroscience Supplement* 2000;3:1184–1191.
- Edin F, Macoveanu J, Olesen P, Tegnér J, Klingberg T. Stronger synaptic connectivity as a mechanism behind development of working memory-related brain activity during childhood. *Journal of Cognitive Neuroscience* 2007;19:750–760. [PubMed: 17488202]
- Engelbretson PH, Huttenlocher J. Bias in spatial location due to categorization: Comment on Tversky and Schiano. *Journal of Experimental Psychology: General* 1996;125:96–108. [PubMed: 8851739]
- Erlhagen W, Bastian A, Jancke D, Riehle A, Schöner G. The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods* 1999;94:53–66. [PubMed: 10638815]
- Erlhagen W, Schöner G. Dynamic field theory of movement preparation. *Psychological Review* 2002;109:545–572. [PubMed: 12088245]
- FaubelCSchönerGLearning to recognize objects on the fly: a neurally based dynamic field approach. 2007, Manuscript submitted for publication.
- Fuster, JM. *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*. MIT Press; Cambridge: 1995.
- Georgopoulos AP, Lurito JT, Petrides M, Schwartz AB, Massey JT. Mental rotation of the neuronal population vector. *Science* 1989;243:234–236. [PubMed: 2911737]
- Georgopoulos AP, Taira M, Lukashin AV. Cognitive neurophysiology of the motor cortex. *Science* 1993;260:47–52. [PubMed: 8465199]
- Gogtay N, Giedd JN, Lusk L, Hayashi KM, Greenstein D, Vaituzis AC, Nugent TF III, Herman DH, Clasen LS, Toga AW, Rapoport JL, Thompson PM. Dynamic mapping of human cortical development during childhood through early adulthood. *Proceedings of the National Academy of Sciences* 2004;101:8174–8179.
- Huttenlocher J, Hedges L, Prohaska V. Hierarchical organization in ordered domains: Estimating the dates of events. *Psychological Review* 1988;95:471–484.
- Huttenlocher J, Hedges LV, Duncan S. Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review* 1991;98:352–376. [PubMed: 1891523]
- Huttenlocher J, Hedges LV, Vevea JL. Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General* 2000;129:220–241. [PubMed: 10868335]
- Huttenlocher J, Newcombe NS, Sandberg EH. The coding of spatial location in young children. *Cognitive Psychology* 1994;27:115–147. [PubMed: 7956105]
- Huttenlocher PR. Synaptic density in human frontal cortex—developmental changes and effects of aging. *Brain Research* 1979;163:195–205. [PubMed: 427544]
- Huttenlocher PR. Morphometric study of human cerebral cortex development. *Neuropsychologia* 1990;28:517–527. [PubMed: 2203993]
- Iossifidis, I.; Schöner, G. Dynamical systems approach for the autonomous avoidance of obstacles and joint-limits for an redundant robot arm. *IEEE International Conference On Intelligent Robots and Systems (IROS)*; Beijing, China. 2006;
- Jancke D, Erlhagen W, Dinse HR, Akhavan AC, Giese M, Steinhage A, Schöner G. Parametric population representation of retinal location: neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience* 1999;19:9016–9028. [PubMed: 10516319]
- Johnson, JS.; Spencer, JP.; Schöner, G. A dynamic neural field theory of multi-item visual working memory and change detection. *The 28th Annual Conference of the Cognitive Science Society*; Vancouver, BC. 2006; Lawrence Erlbaum Associates; p. 399-404.
- JohnsonJSSpencerJPSchönerGGarzónFLaaksoAGomilaT, in press. Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. *Dynamics and Psychology [special issue]*
- Johnson R Jr. On the neural generators of the P300 component of the event-related potential. *Psychophysiology* 1993;30:90–97. [PubMed: 8416066]
- Kelso, JAS. *Dynamic patterns: The self-organization of brain and behavior*. The MIT Press; Cambridge, MA: 1997.

- Kessels RPC, Postma A, Wijnalda EM, de Haan EHF. Frontal-lobe involvement in spatial memory: Evidence from PET, fMRI, and lesion studies. *Neuropsychology Review* 2000;10:101–113. [PubMed: 10937918]
- Kinchla RA. Visual movement perception: A comparison of absolute and relative movement discrimination. *Perception and Psychophysics* 1971;9:165–171.
- Kohonen T. Self-organized formation of topologically correct feature maps. *Biological Cybernetics* 1982;43:59–69.
- Kopecz K, Schöner G. Saccadic motor planning by integrating visual information and pre-information on neural, dynamic fields. *Biological Cybernetics* 1995;73:49–60. [PubMed: 7654850]
- Lisman JE, Idiart MAP. Short-Term Memories in Oscillatory Subcycles. *Science* 1995;267:1512–1515. [PubMed: 7878473]
- Love BC, Medin DL, Gureckis TM. SUSTAIN: A network model of category learning. *Psychological Review* 2004;111:309–332. [PubMed: 15065912]
- Mareschal D. To reach or not to reach ... that is the question. *Developmental Science* 1998;1:198–199.
- McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 1995;102:419–457. [PubMed: 7624455]
- McDowell K, Jeka JJ, Schöner G, Hatfield BD. Behavioral and electrocortical evidence of an interaction between probability and task metrics in movement preparation. *Experimental Brain Research* 2002;144:303–313.
- Miikkulainen, R.; Bednar, JA.; Choe, Y.; Sirosh, J. Modeling self-organization in the visual cortex. In: Oja, E.; Kaski, S., editors. *Kohonen Maps*. Elsevier; New York: 1999.
- Morton JB, Munakata Y. Active versus latent representations: A neural network model of perseveration, dissociation, and decalage. *Developmental Psychobiology* 2002;40:255–265. [PubMed: 11891637]
- Munakata Y. Infant perseveration and implications for object permanence theories: A PDP model of the AB task. *Developmental Science* 1998;1:161–184.
- Munakata Y, McClelland JL. Connectionist models of development. *Developmental Science* 2003;6:413–429.
- Munakata Y, McClelland JL, Johnson MH, Siegler RS. Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review* 1997;104:686–713. [PubMed: 9337629]
- Nelson CA, Monk CS, Lin J, Carver LJ, Thomas KM, Truwit CL. Functional neuroanatomy of spatial working memory in children. *Developmental Psychology* 2000;36:109–116. [PubMed: 10645748]
- Palmer J. Mechanisms of displacement discrimination with a visual reference. *Vision Research* 1986a; 26:1939–1947. [PubMed: 3617534]
- Palmer J. Mechanisms of displacement discrimination with and without perceived movement. *Journal of Experimental Psychology: Human Perception and Performance* 1986b;12:411–421. [PubMed: 2946798]
- Palmer SE, Hemenway K. Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception & Performance* 1978;4:691–702. [PubMed: 722256]
- PeroneSSpencerJPSchönerGA dynamic field theory of visual recognition in infant looking tasks. To appear in the Proceedings of the Twenty-Ninth Annual Conference of the Cognitive Science Society, in press.
- Piaget, J. *The Construction of Reality in the Child*. Basic Books; New York: 1954.
- Raffone A, Wolters G. A cortical mechanism for binding in visual working memory. *Journal of Cognitive Neuroscience* 2001;13:766–785. [PubMed: 11564321]
- Rakic, P. The development of the frontal lobe: A view from the rear of the brain. In: Jasper, HH.; Riggio, S.; Goldman-Rakic, PS., editors. *Advances in neurology*. Raven Press; New York: 1995.
- Sampaio, RC.; Truwit, CL. Myelination in the developing human brain. In: Nelson, CA.; Luciana, M., editors. *Handbook of Developmental Cognitive Neuroscience*. MIT Press; Cambridge, MA: 2001.
- Schöner, G. Development as change of system dynamics: Stability, instability, and emergence. In: Spencer, JP.; Thomas, MSC.; McClelland, JL., editors. *Toward a New Grand Theory of*

- Development? Connectionism and Dynamic Systems Theory Re-Considered. Oxford University Press; New York, NY: 2007.
- Schöner G, Dineva E. Dynamic instabilities as mechanisms for emergence. *Developmental Science* 2007;10:69–74. [PubMed: 17181702]
- Schöner G, Kelso JAS. Dynamic pattern generation in behavioral and neural systems. *Science* 1988;239:1513–1520. [PubMed: 3281253]
- SchutteARA developmental transition in spatial working memory . Unpublished doctoral dissertation, 2004University of IowaIowa City, IA
- Schutte AR, Spencer JP. 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* 2002;73:377–404. [PubMed: 11949898]
- SchutteARSpencerJPA developmental transition in spatial working memory: Sudden shift or continuous change? 2007a, Manuscript under review.
- Schutte AR, Spencer JP. Planning “discrete” movements using a continuous system: Insights from a dynamic field theory of movement preparation. *Motor Control* 2007b;11:166–208. [PubMed: 17575340]
- Schutte AR, Spencer JP, Schöner G. Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development* 2003;74:1393–1417. [PubMed: 14552405]
- SimmeringVRSpencerJPGenerality with specificity: The dynamic field theory generalizes across tasks and time scales. 2007, Manuscript under review.
- Simmering VR, Spencer JP, Schöner G. Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception & Psychophysics* 2006;68:1027–1046. [PubMed: 17153196]
- Simon, HA. An information processing theory of intellectual development. In: Kessen, W.; Kuhlman, C., editors. *Thought in the young child*. The Antioch Press; Yellow Springs, OH: 1962.
- Skarda CA, Freeman WJ. How brains make chaos in order to make sense of the world. *Behavioral & Brain Sciences* 1987;10:161–195.
- Smith LB, Thelen E, Titzer R, McLin D. Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review* 1999;106:235–260. [PubMed: 10378013]
- Sowell ER, Thompson PM, Tessner KD, Toga AW. Mapping continued brain growth and gray matter density reduction in dorsal frontal cortex: Inverse relationships during postadolescent brain maturation. *Journal of Neuroscience* 2001;21:8819–8829. [PubMed: 11698594]
- Spencer JP, Hund AM. Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General* 2002;131:16–37. [PubMed: 11900101]
- Spencer JP, Hund AM. Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology* 2003;47:432–480. [PubMed: 14642291]
- Spencer JP, Schöner G. Bridging the representational gap in the dynamical systems approach to development. *Developmental Science* 2003;6:392–412.
- Spencer JP, Simmering VR, Schutte AR. 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* 2006;32:473–490. [PubMed: 16634683]
- Spencer, JP.; Simmering, VR.; Schutte, AR.; Schöner, G. What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In: Plumert, JM.; Spencer, JP., editors. *The emerging spatial mind*. Oxford University Press; New York, NY: 2007.
- Spencer JP, Smith LB, Thelen E. Tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of 2-year-olds. *Child Development* 2001;72:1327–1346. [PubMed: 11699674]
- Steinhage, A.; Schöner, G. Dynamical systems for the behavioral organization of autonomous robot navigation. In: Schenker, PS.; McKee, GT., editors. *Sensor Fusion and Decentralized Control in Robot Systems: Proceedings of SPIE*. SPIE Publishing; 1998.
- Steinmetz MA, Constantinidis C. Neurophysiological evidence for a role of posterior parietal cortex in redirecting visual attention. *Cerebral Cortex* 1995;5:448–456. [PubMed: 8547791]

- Stiles J, Bates EA, Thal D, Trauner DA, Reilly J. Linguistic and spatial cognitive development in children with pre- and perinatal focal brain injury: A ten-year overview from the San Diego longitudinal Project. *Advances in Infancy Research* 1998;12:131–164.
- Thelen E, Schöner G, Scheier C, Smith LB. The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral & Brain Sciences* 2001;24:1–86. [PubMed: 11515285]
- Thelen, E.; Smith, LB. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press; Cambridge: 1994.
- Tversky B, Schiano DJ. Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology: General* 1989;118:387–398. [PubMed: 2531198]
- Vogels TP, Rajan K, Abbott LF. Neural network dynamics. *Annual Review of Neuroscience* 2005;28:357–376.
- Wang X-J. Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neuroscience* 2001;24:455–463.
- Wenderoth P, van der Zwan R. Local and global mechanisms of one- and two-dimensional orientation illusions. *Perception & Psychophysics* 1991;50:321–332. [PubMed: 1758764]
- Werner S, Diedrichsen J. The time course of spatial memory distortions. *Memory and Cognition* 2002;30:718–730.
- Westermann G, Mareschal D. From parts to wholes: Mechanisms of development in infant visual object processing. *Infancy* 2004;5:131–151.
- Wilimzig C, Schneider S, Schöner G. The time course of saccadic decision making: dynamic field theory. Special issue: Neurobiology of decision making. *Neural Networks* 2006;19:1059–1074. [PubMed: 16942860]
- Wilimzig, C.; Schöner, G. The emergence of stimulus-response associations from neural activation fields: Dynamic Field Theory. 27th Annual Conference of the Cognitive Science Society; Stresa, Italy. 2005;
- WilimzigCSchönerGHow categories may emerge from continuous distributions of population activation: Dynamic Field Theory of continuous and categorical responding. 2007, Manuscript in preparation.
- Wilson FAW, Scialidhe SP, Goldman-Rakic PS. Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science* 1993;260:1955–1958. [PubMed: 8316836]

Fig. 1.

A simulation of the dynamic field theory performing one spatial recall trial. In each panel, location is across the x-axis, activation on the y-axis, and time on the z-axis. (A) Inputs to the model are presented directly to the model already in an object-centered reference frame (see text for details of the calibration process that transforms spatial information from an egocentric frame of reference in the full model). After this shift, activation is passed to the model consisting of 5 layers: (B) a perceptual field; (C) a long-term memory field associated with this perceptual field; (D) a shared layer of (inhibitory) interneurons; (E) a spatial working memory field; (F) a long-term memory field associated with the spatial working memory field. Solid arrows show excitatory connections between layers, and dashed arrows show inhibitory connections between layers.

Fig. 2. The local excitation (+) / lateral inhibition (-) function used for child (A) and adult (B) simulations. Dashed line indicates the zero threshold.

Fig. 3.

Simulations of the infant A-not-B task. Both simulations began with four trials to the A location (not shown), then used identical input (A) on the first B trial. Note that these simulations include only the relevant layers of the model, SWM (B, D) and LTM_{SWM} (C, E). Local excitatory interactions were lowest for the “young” infant simulations (B-C), which showed an A-not-B error (see peak circled in B). Local excitation was increased for the “older” infant simulation (D-E), which correctly responded at B (see peak circled in D). Axes are as in Fig. 1.

Fig. 4.

Simulations of the sandbox A-not-B task. Both simulations began with four trials to the A location with identical inputs (not shown). On the first B trial, the target was presented either near (A-C) or far (D-F) from A. When B was near A (see A), the peak in SWM (B) drifted toward the activation from LTM_{SWM} (C). When B was far from A (see D), on the other hand, the peak in SWM remained stable and did not drift (E). Axes are as in Fig. 1.

Fig. 5.

Simulations of spatial recall trials with child (A-D) and adult (E-H) parameter settings. To approximate developmental changes in reference frame calibration, the inputs are broader and weaker in the child model (A) relative to the adult model (E). In the child model, midline does not provide enough focused input to form a peak in PF (B), and the target peak in SWM (D) is attracted toward the sub-thresholds activation from midline. In the adult model, on the other hand, midline forms a peak in PF (F), which contributes inhibition via Inhib (G) to SWM (H), resulting in the peak being repelled from midline. Axes are as in Fig. 1.

Fig. 6.

The DFT performing *same* (A-C) and *different* (D-F) responses in position discrimination. In both simulations, S1 was presented at -40° , with S2 presented at -40° for *same* and -30° for *different*. Note that, for simplicity, we did not include reference input in these simulations. Axes are as in Fig. 1.