

THE EMERGING SPATIAL MIND

Edited by
JODIE M. PLUMERT
AND
JOHN P. SPENCER

OXFORD
UNIVERSITY PRESS

2007

— -1
— 0
— +1

14

WHAT DOES THEORETICAL NEUROSCIENCE HAVE TO OFFER THE STUDY OF BEHAVIORAL DEVELOPMENT?

Insights from a Dynamic Field Theory
of Spatial Cognition

JOHN P. SPENCER, VANESSA R. SIMMERING,
ANNE R. SCHUTTE, & GREGOR SCHÖNER

The chapters in this edited volume address two questions central to discussions of the emerging spatial mind: *what* changes in spatial cognition occur over development, and *how* these changes come about. By our read of the literature, decades of research on the development of spatial cognitive abilities have yielded many clear answers to the first question but rather limited insights into the second. To illustrate, consider the following three lines of research examining the *what* question. First, evidence shows that the precision of spatial memory changes systematically over development. For instance, there is an increase in the metric precision of children's responses over a broad age range—from 2 to 11 years—in tasks where children either search for objects hidden in a homogeneous space (e.g., Schutte, Spencer, & Schöner, 2003; Spencer & Hund, 2003) or replace objects in

-1 ___
0 ___
+1 ___

previously learned locations (Hund & Plumert, 2002, 2003, 2005; Hund, Plumert, & Benney, 2002; Plumert & Hund, 2001).

Second, there are clear developmental changes in children's ability to coordinate/update spatial reference frames. Three-month-old infants appear to encode locations egocentrically, showing initial evidence of using a world-centered reference frame by 6 months (Gilmore & Johnson, 1997). Beyond infancy, children rely more and more on world-centered reference frames (e.g., Newcombe & Huttenlocher, 2000) but alter how these frames are used. For instance, between 3 and 6 years of age, there is a qualitative shift in children's recall accuracy near symmetry axes such that young children are biased *toward* axes of symmetry while older children are biased *away from* such axes (Schutte, 2004; Schutte & Spencer, 2002).

Third, research demonstrates that children develop the ability to think about space in increasingly complex ways. For instance, according to DeLoache's dual representation hypothesis, there is a dramatic improvement in 2.5- to 3.5-year-olds' ability to think about a scale model both as an object in and of itself and as a symbol for an identical but larger space (see DeLoache, 2004, for review).

In all of these examples, there is a clear understanding—grounded in a rich empirical database and strong theories—of *what* is changing over development. Less clear is *how* these changes take place. For instance, our spatial precision hypothesis (Schutte et al., 2003; Spencer & Hund, 2003) effectively captures the systematic improvement in spatial memory accuracy over development across a range of tasks (see section 14.5). But this hypothesis fails to explain how changes in spatial precision arise over development. Similarly, J. Huttenlocher, Hedges, and Duncan (1991) have proposed a model to explain the shift in children's use of symmetry axes/category boundaries, but this model fails to explain how this shift occurs. Finally, the dual representation hypothesis effectively brings together a diverse set of empirical findings and has generated several counterintuitive predictions (DeLoache, 2004), but this hypothesis says little about how the posited developmental change occurs.

We contend that this state of affairs is not unique to the domain of spatial cognition. Rather, developmental science in general has made relatively modest progress on specifying the mechanisms that underlie behavioral change. In McClelland and Siegler's recent volume *Mechanisms of Cognitive Development*, for instance, Posner (2001) offers the following synthesis of candidate mechanisms: (1) strategic activation of brain regions/circuits, (2) automation, (3) changes in synaptic connections in the brain, (4) changes in cortical maps, and (5) creation of new neural networks. Each of these "mechanisms" was motivated by a rich set of empirical and theoretical work using multiple techniques (e.g., behavioral work, functional MRI [fMRI], computational modeling), many of which explicitly probe brain-behavior relations. And in each case, the proposed mechanisms offer insights about development.

But also in each case, there is a clear tension: as researchers explore/explain the longer time scales of learning and development, it becomes harder and harder to offer a detailed, process-based account of real-time behavior. For instance, many of the advances in our understanding of how synaptic connections contribute to behavioral change come from work with connectionist models (e.g., McClelland & Jenkins, 1991). Although such work has led to fundamental insights about how nonlinear change can emerge from distributed activation in neural networks using relatively simple learning rules (see Munakata & McClelland, 2003, for a review), such models are often weakly linked to the real-time behavior of children in real tasks (see Spencer & Schöner, 2003, for a more detailed discussion). Because such models are often loosely tied to behavior, they can become demonstration proofs of the power of synaptic changes over learning rather than formal theories of particular types of developmental change (see Rogers & McClelland, 2004, for an exception).

A second example further highlights the tension between explanations of real-time and learning/developmental processes. Research by Stiles, Bates, Thal, Trauner, and Reilly (1998) has shown profound reorganization of cortical circuits following brain damage in early development (e.g., during infancy), with impressive recovery of function by 6 years of age. Although there are some constraints on such reorganization (see Stiles et al., 1998, for a discussion), this demonstrates an impressive capacity to recruit new neural networks over the longer time scales of development. Critically, though, we know relatively little about how this is done and how the recruitment/establishment of new networks in the brain emerges from the real-time behavioral experiences of children.

We want to stress that the work we have reviewed thus far has contributed fundamental insights into changes in spatial cognitive development as well as potential mechanisms of development. Our point here is to draw attention to two themes in the literature on the *how* of development: (1) there is a growing trend to search for mechanisms of development in brain-behavior relations, and (2) as researchers explore the longer time scales of learning and development, this often comes at the cost of real-time details. Combining these two themes sets up a key challenge for future work: to develop theories of development grounded in a rich understanding of brain-behavior relations that handles the dual challenges of offering a detailed account of real-time behavior but also incorporates the longer time scales of learning and development.

In this chapter, we describe our pursuit of this goal that integrates empirical work with a theoretical neuroscience perspective. In particular, we describe a dynamic field theory (DFT) of spatial cognition that is grounded in a rich understanding of brain-behavior relations, offers a detailed account of the real-time processes underlying many spatial behaviors, and provides insights into the link between behavior in the moment and learning/development. Although our efforts to integrate short and long time scales are still limited in important respects (see section 14.6), we emphasize two

insights in the present chapter. First, we have uncovered a rich array of behavioral and theoretical constraints because processes that “live” at different time scales in our theory constrain one another. Second, our theory captures a host of developmental changes in spatial cognition with a relatively simple and neurally plausible hypothesis. This has led to the insight that less is required of developmental processes when they are coupled to the complex dynamics at shorter time scales.

The next section provides an overview of the spatial cognitive phenomena we set out to capture with the DFT. Then, we describe how the DFT offers a detailed, constrained view of the real-time and learning-time (e.g., trial-to-trial) processes that underlie spatial performance in multiple tasks. In addition, we discuss how the DFT is constrained not just by behavioral data but also by a rich understanding of brain–behavior relations. These sections highlight that our theory is both task specific and task general. Task specificity is needed for strong ties to the real-time details of behavior. Task generality is needed for generalization across different experiences, which, we contend, plays a role in another type of generalization—generalization over development.

We then tackle the longer time scale of development by discussing our *spatial precision hypothesis*. This relatively simple hypothesis—that neural interactions become stronger and more spatially precise over development—has captured both quantitative and qualitative changes in performance in several tasks, shedding light on both how children use available perceptual structure in spatial tasks and how children construct long-term memories of specific locations in the world. We conclude by discussing the challenges that lie ahead. One critical next step is to “close the loop” on development (see Simon, 1962)—to explain how the spatial precision hypothesis emerges from the reciprocal coupling of processes that live at different time scales. More concretely, we must explain the developmental mechanisms that—in conjunction with real- and learning-time processes—cause developmental change. We speculate on these mechanisms at the end of our chapter and highlight what a theoretical neuroscience perspective can contribute to the more general search for the *how* of development.

14.1 WHAT NEEDS TO BE EXPLAINED IN SPATIAL COGNITION?

The theory that we describe in the present chapter focuses on the dynamics of spatial cognition, that is, the time-dependent processes that underlie coordinated spatial behavior. Such behavior requires that people remember the locations of important objects in the local surrounds with enough fidelity to coordinate a myriad of second-to-second decisions, actions, and attentional shifts. Moreover, the local “map” that is used in one workspace must be coordinated with other maps as people move from context to context—from the office desk, to the office filing cabinet, to the desk at

home. This requires the real-time, contextually specific integration of past and present, of long-term memory with short-term memory.

To illustrate the many challenges involved in achieving coordinated spatial behavior in complex, real-world environments, consider the challenges facing a child during her first day at daycare. When she walks into her “home” room for the first time, she is surrounded by a myriad of new toys, cribs pushed against a wall, a table, some chairs stacked in the corner, and balls hanging in a bag by the door. How is she going to make sense of all of this? She needs to keep in mind and, ultimately, to learn where these objects are in the room, which toys are the really fun ones, and so on. But a static “map” will not do. Rather, her knowledge needs to cope with repeated changes in the local layout of objects in the room as toys are played with, moved, and (sometimes) put back, with changes in the global layout as cribs are moved in and out of place during naptime and chairs are rearranged for snack time, and with changes in the “virtual” layout as some toys became off limits in some contexts (e.g., shakers only available during music). Finally, she needs to learn the layout of separate spaces—the collective play area and the all-important outside play area—each of which share the dynamic properties of the “home” room.

How do people acquire such spatially grounded knowledge? To answer this question, we must first appreciate that people acquire knowledge by interacting with and exploring the world through a richly structured sensorimotor system. This sets up several key challenges. First, people must calibrate/coordinate reference frames in real time. In particular, they must keep egocentric reference frames in register with allocentric or world-centered frames as they move around. This is critical to remembering, for instance, that the books are on the upper shelf of the bookcase currently over to the left. This sets up a related challenge: encoding/remembering the locations of particular items relative to a reference frame. Although it might be useful to know that there is a red ball somewhere in this room, it is much more useful to know that the red ball is in the bag by the door. Finally, remembering the spatial characteristics of many objects can be a daunting task. Thus, people form spatial categories to reduce memory demands, to help them generalize, and to help them retrieve information reliably because categories offer a consistent retrieval cue.

The issue of forming spatial categories highlights that, although the challenges above are all linked to the formation of spatial knowledge in real time, these challenges must be solved in a way that extends beyond initial interactions with objects in the room. In particular, spatial knowledge must be “stored” in a way that it can be successfully retrieved and linked to the current context when the child reenters the home room each morning. This involves activating the “right” memories based on, for instance, the spatial layout of landmarks in the room (i.e., the child must recall object locations from memories of the “home” room and not the play room) and bringing these memories into register with the current scene (e.g., the child might have to mentally rotate retrieved information because she previously viewed objects in the room from a different vantage point).

Finally, we can extend all of the challenges above to the longer time scales of learning and development. These time scales present their own unique challenges. The first is generalization: how does the child generalize knowledge across tasks and contexts to form, for instance, a survey representation of the layout of rooms in a house or town? Similarly, how are developmental changes in spatial cognition linked to the striking array of other interconnected skills that come “on-line” during, for instance, infancy and early childhood: social skills (e.g., which things are off limits in certain situations), language skills (e.g., use of spatial relational terms and object names), and motor skills (e.g., reaching, walking, tool use).

The daycare example does a particularly good job highlighting that the challenges above are real challenges, because it is often painfully evident with children that learning to skillfully interact with the world is hard: they forget where objects are, they go to the wrong room of the house, and so on. But these challenges are not isolated to early development. They emerge when children and adults are taxed in controlled, experimental settings. And they resurface in the behavior of disordered populations across a broad range of development (Barnett et al., 2001; Karatekin & Asarnow, 1998; Landau, 2003; Lawrence, Watkins, Sahakian, Hodges, & Robbins, 2000).

Given the complexity of the spatial challenges people face in real-world situations, it is not surprising that, to date, no theory effectively explains how people handle these challenges. Thus, the goal of our theoretical efforts was to take a first step toward a formal model of spatial cognition capable of integrating some of the diverse processes that underlie spatially grounded behavior. We focused on four “core” aspects of spatial cognition discussed below: (1) how people form metric or “fine-grained” location memories, (2) how people link metric spatial information to particular reference frames and form “geometric” categories, (3) how people integrate spatial information perceived and remembered in the hear-and-now with long-term memories that can be reactivated in context, and (4) how people achieve behavioral flexibility, which, we contend, is related to the generalization of spatial knowledge across tasks.

To facilitate our discussion of these core aspects of spatial cognition and our presentation of the DFT (see section 14.2), we focus primarily on empirical results from studies using our “spaceship” task. Figure 14.1 shows an image of the spaceship table we use in our laboratory (e.g., Spencer & Hund, 2002), with a frame of reference centered at a “start” location directly in front of where participants sit. Within this frame, 180° is along the midline of the task space projecting away from the start location through a second reference dot, with 90° to the right and 270° to the left. Figure 14.1 also shows a sample target (a triangular “spaceship”) at 220°. To study the characteristics of spatial cognition, we have used relatively simple tasks in this task space. For instance, in our spatial recall task, the spaceship appears on the table, then it disappears, there is a short delay (0–20 seconds), and then participants move a mouse cursor from the start location to the remembered location. As we describe below, even this

Figure 14.1 The “spaceship” table used in our laboratory. Two dots mark the midline symmetry axis, which corresponds to 180° in the model, with 90° to the right and 270° to the left. A sample triangular “spaceship” target is shown at 220° .



simple task can reveal key characteristics of the processes that underlie spatial cognition.

14.1.1 Metric Memory for Locations

To successfully act in the world, one must remember relatively precise, spatially continuous, metric information—not just qualitative, categorical information. Evidence suggests that children and adults can remember many types of metric information—direction and distance in a polar coordinate frame (McIntyre, Stratta, & Lacquaniti, 1998), x and y coordinates within a rectangular frame (J. Huttenlocher, Newcombe, & Sandberg, 1994), and so on. Given that spatial memory must be linked to multiple sensory and motor systems as well as to other internal processes (e.g., those dealing with reference frames), and given that the state of these multiple subsystems may vary in time due to the complex behavior of the organism in changing environments, the maintenance of spatial information in working memory requires processes that stabilize metric information in the face of both internal and external perturbations (Spencer & Schönner, 2003).

Although remembering precise metric information seems relatively straightforward for adults, this is not the case for children. Early in development, infants and young children do not always succeed in stabilizing memorized information in our spaceship task. For instance, when 3.5-year-olds were asked to remember a spaceship location to the left of the midline symmetry axis, they “forgot” the location on about 10% of all trials and responded to a target to the right of midline that they had seen on the just previous trial (Schutte, 2004). What is so striking about these errors is that the memory delay was only 10 seconds and the right target was 80° away (for related effects with infants, see Smith, Thelen, Titzer, & McLin, 1999; Thelen, Schönner, Scheier, & Smith, 2001).

Insufficient stability of spatial memory manifests itself in more subtle ways later in development—even into adulthood. For instance, 6- and 11-year-old children and adults showed a systematic, delay-dependent increase in bias relative to midline in the spaceship task during 0–20 second

delays (Spencer & Hund, 2002, 2003). In particular, memory responses to targets to the left of midline showed an increase in leftward errors over delays (e.g., responses to the 220° target were accurate at no delay but were shifted leftward by 4°—to 224°—by the 20-second delay), while targets to the right of midline showed an increase in rightward errors over delays. Responses showed a systematic, delay-dependent increase in variance as well. Thus, in the absence of any perceptual markers indicating the location of a target, children and adults remember roughly where the target is located, but the metric content of this memory is systematically distorted over time.

14.1.2 Metric Memory Relative to Reference Frames and Geometric Categories

A central challenge for any theory of spatial cognition is to explain how metric information remains calibrated and updated within a frame of reference. On one hand, this challenge requires keeping metric information calibrated with on-line changes in sensorimotor reference frames as people move their eyes, head, and body through space over time (e.g., Graziano, Hu, & Gross, 1997; Soechting & Flanders, 1991; see also chapter 4). Similarly, one must keep metric information grounded in external reference frames—the edges of a table, the boundaries of an object, and so on (J. Huttenlocher et al., 1991; Newcombe & Huttenlocher, 2000; Tversky & Schiano, 1989). Several neural mechanisms have been proposed to transform information from one frame into another frame (Deneve & Pouget, 2003; Pouget, Deneve, & Duhamel, 2002). Less understood is how such calibration processes are linked to processes that deal with recognition of reference frames and the real-time maintenance of metric information in memory.

Although the processes that link frames of reference to metric memories are poorly understood, there is strong evidence that metric memory is grounded in perceived reference frames. For instance, studies have shown that adults' memory for precise, metric information is biased away from the visible edges of L- and V-shaped frames, as well as their symmetry axes (Engebretson & Huttenlocher, 1996; Schiano & Tversky, 1992, experiments 1 & 2; Tversky & Schiano, 1989). The “midline” biases discussed above provide a concrete example of this type of error (see also J. Huttenlocher, Hedges, Corrigan, & Crawford, 2004; J. Huttenlocher et al., 1991; Plumert & Hund, 2001; Sandberg, Huttenlocher, & Newcombe, 1996). Such biases have been referred to as *geometric category biases* to highlight that use of visible edges and axes of symmetry carves up space into smaller regions that can help keep metric memory contained within the correct spatial region.

A central question in this domain is how people decide where to put or “impose” category boundaries. Our view is that geometric category boundaries are *always tied to perceived reference frames*. Adults' use of visible edges

in the context of L- and V-shaped frames clearly fits with this picture. But what about the use of symmetry axes such as the midline symmetry axis of the spaceship table? After all, symmetry axes are not visible per se. Although this is certainly the case, evidence from the perceptual literature suggests that symmetry axes are perceived like weak visible lines (Wenderoth & van der Zwan, 1991). Moreover, the accuracy with which people perceive symmetry axes increases with the salience of visible cues (Beh, Wenderoth, & Purcell, 1971; Li & Westheimer, 1997; Simmering, Spencer, & Schöner, in press). Finally, data suggest that symmetry axes have a special developmental status in that the ability to perceive symmetry develops quite early (3- to 4-month-old infants are capable of perceiving vertical symmetry; see Quinn, 2000).

More direct evidence that spatial category boundaries are grounded in perceived reference frames comes from a recent study from our laboratory (Simmering & Spencer, 2006a). We asked adults to reproduce a target location to the left of a virtual axis defined by two dots. This axis was parallel to the midline symmetry axis of the task space but translated to the right by 15 cm (imagine the two dots in figure 14.1 translated to the right). As in previous studies, responses to the left target were biased leftward after a 10-second delay. By contrast, when we removed the dots, responses to the same target location were biased rightward—away from the midline axis. Critically, this happened on a trial-by-trial basis, even when participants were instructed to “act as if the dots are there on every trial” and even when the trial with the dots “off” was preceded by five trials in a row with the dots “on.” Furthermore, none of the participants mentioned using the midline symmetry axis in a postexperiment questionnaire, suggesting that participants were spontaneously locking onto whatever visible structure was present in the task space with little awareness that they were doing so.

14.1.3 Integration of Short-Term and Long-Term Spatial Memories

Thus far, we have focused on a very basic issue: perception of and memory for individual locations over short-term delays within a calibrated reference frame. How are such memories linked to performance over a longer time scale; that is, how are short-term memories integrated with long-term spatial memory?

Work in our laboratory has demonstrated that a simple mechanism can bridge these two time scales: as people remember individual locations over short-term delays, this activity leaves a trace of activation in long-term memory. Such traces can then bias performance toward these previously visited locations on subsequent trials. For instance, older children (6 and 11 years) and adults show biases toward an average remembered location. In particular, when children and adults are asked to repeatedly remember a small set of close locations (e.g., three targets separated by 20°), they show biases toward the central target (Spencer & Hund, 2002, 2003). These

biases can be systematically distorted toward one of the outer targets by increasing the frequency with which people visit an outer target relative to the other two (Hund et al., 2002; Hund & Spencer, 2003). Such distortions in short-term memory toward information in long-term memory are often referred to as *experience-dependent category biases* (for related effects, see Schutte et al., 2003; Smith et al., 1999; Spencer, Smith, & Thelen, 2001; Thelen et al., 2001).¹

14.1.4 Task Generality and Behavioral Flexibility

The final “core” issue we sought to tackle provides a bridge from task-specific behavior toward more task-general behavior. Very few theories have attempted to capture generalization across diverse experiences. This is not surprising because what is carried from one experience to the next is rarely transparent. This is true even if we ignore the complexity of generalizing across experiences in different spaces. For instance, what should generalize across the following three experiences? (1) A child puts down her red, plastic party cup, turns to look at the television, and then turns back a second later to find that Mom has placed an identical red cup next to hers, but she identifies the correct cup. (2) The child leaves the room, but forgets her cup; she then asks her Mom to bring the cup and verbally specifies its location. (3) The child returns to the room; the cup is in a new location, but she quickly recalls where it is and walks over to the television to retrieve her drink. These three experiences require a largely perceptually based “same/different” decision, the generation of a linguistic description, and a recall-based motor response, respectively. What should the child “take” from these experiences? Should they all be “bound” together since they are grounded in the same spatial array? Should they all be isolated because they tap into what have been characterized as different spatial systems (see, e.g., Bridgeman, 1999; Bridgeman, Gemmer, Forsman, & Huemer, 2000; Brungart, Rabinowitz, & Durlach, 2000)?

One way to foster generalization across diverse experiences would be to have a system that is behaviorally flexible—that can generate different types of responses and configure itself to solve multiple tasks as the situation demands. This would naturally lead to some form of generalization because the same spatial system would be involved in multiple behaviors. We have begun to probe this possibility empirically by asking a simple question: how is performance in one spatial task linked to performance in another spatial task? For instance, rather than asking participants to recall a remembered location on our spaceship table, we showed them two dots in quick succession (500 msec apart) and asked them whether the dots were in the “same” or a “different” position (Simmering et al., in press). Importantly, we found that discrimination responses showed evidence of “fast repulsion” from midline: discrimination abilities were impaired when the second dot was presented away from midline relative to the first, that is, in the direction of bias observed in our studies of spatial recall (for related

effects, see Werner & Diedrichsen, 2002). Specifically, our data suggest that the memory of the first dot drifted away from midline during the 500 msec delay. Consequently, when the second dot was presented away from midline relative to the first, the memory of the first dot blended with the presentation of the second dot, and people erroneously said they were in the “same” position. Although this highlights that the use of a midline reference frame has consequences for both recall and discrimination, it has a positive influence, as well: discrimination abilities showed an overall enhancement when the stimuli were presented closer rather than farther from midline (for a second example that probes task generality, see Spencer, Simmering, & Schutte, 2006). These data suggest that a shared spatial system might handle the generation of these different response types, even though, to date, spatial recall and position discrimination have been treated as separate literatures captured by different classes of models (for a discussion, see Simmering & Spencer, 2006b).

14.2 A DYNAMIC FIELD THEORY OF SPATIAL COGNITION

Figure 14.2 shows a simulation that implements our DFT.² This simulation consists of seven layers: a perceptual field in an egocentric frame of reference (PF_{ego} ; figure 14.2A), a system that transforms activation in an egocentric reference frame into an object-centered frame (shift; figure 14.2B), a perceptual field in an object-centered frame (PF_{obj} ; figure 14.2C), a long-term memory field associated with the object-centered perceptual field (LTM_{PFobj} ; figure 14.2D), a layer of (inhibitory) interneurons ($Inhib_{obj}$; figure 14.2E), a spatial working memory field in an object-centered frame (SWM_{obj} ; figure 14.2F), and a long-term memory field associated with the working memory field (LTM_{SWMobj} ; figure 14.2G).

In this section, we step through key characteristics of the theory and how it provides a neurally plausible account of each aspect of spatial cognition discussed above. We first cover some of the basic principles of dynamic fields and how such fields capture short-term “working” memory for locations. Second, we discuss our approach to the dynamics of reference frame calibration, which involves the top four layers of the model (figure 14.2A–D). We then turn our attention to three central layers of the model (figure 14.2C,E,F), showing how our theory accounts for location memory relative to a perceived reference frame and geometric category biases. Third, we discuss how the DFT integrates short-term and long-term time scales through the coupling of two long-term memory fields to PF_{obj} and SWM_{obj} (in figure 14.2C,D,F,G). Fourth, we discuss how this theory captures aspects of task generality and behavioral flexibility.

Note that, although we discuss different aspects of the model in turn, we emphasize that the system shown in figure 14.2 is a single, integrated, dynamical system. As an example, consider the two layers shown in figure 14.2, C and D (PF_{obj} and LTM_{PFobj}). These two fields participate in refer-

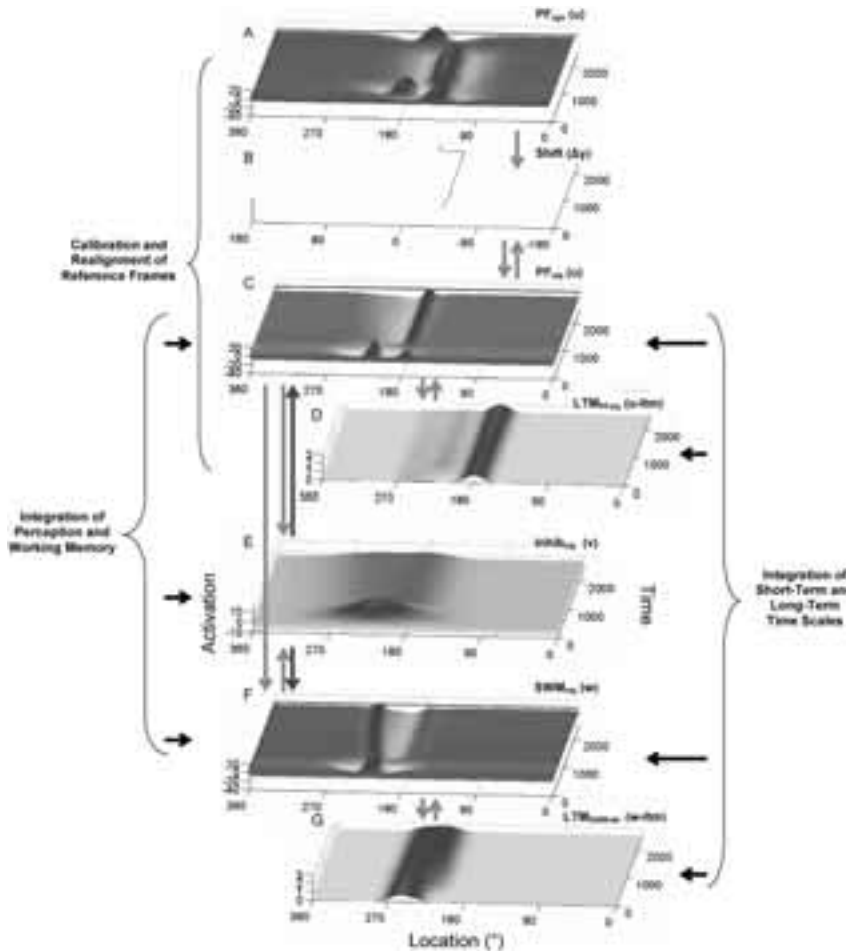


Figure 14.2 A simulation of the DFT performing one spatial recall trial. In each panel, location is across the x-axis, activation on the y-axis, and time on the z-axis. The model consists of seven layers: (A) a perceptual field in an egocentric reference frame, (B) a system that transforms locations from egocentric to an object-centered frame, (C) a perceptual field in the object-centered reference frame, (D) a long-term memory field associated with this perceptual field, (E) a shared layer of (inhibitory) interneurons, (F) a spatial working memory field in the object-centered reference frame, and (G) a long-term memory field associated with the spatial working memory field. Green arrows show excitatory connections between layers, and red arrows show inhibitory connections between layers. Brackets describe the behavioral functions of subsets of the layers. See text for additional details.

ence frame calibration, the real-time coupling of perception and working memory, and the integration of spatial information across short and long time scales. This highlights that our theory does not consist of specialized modules/boxes typical in many models. That said, the different layers do serve conceptually different functions—most of the time! Again, focusing on figure 14.2C, we note up-front that although we call this a perceptual field (PF_{obj}), we can generate behavioral flexibility by having this field serve a decision-making/working memory function during the test phase in specific types of tasks (e.g., position discrimination). This flexibility arises, in part, due to the reciprocal interconnections among the layers. Importantly, such dense interconnection is inspired by the densely connected, reentrant nervous system we are trying to model. Although this presents some theoretical challenges, we contend that thinking of the brain as an integrated dynamical system offers new insights into task generality, behavioral flexibility, and, ultimately, development (for related ideas, see Skarda & Freeman, 1987).

14.2.1 Metric Memory for Locations

There is general agreement that some form of sustained activation is the most plausible neuronal substrate for short-term spatial memory (Constantinidis & Steinmetz, 1996; Fuster, 1995; Miller, Erickson, & Desimone, 1996). This substrate involves graded, metric representations that evolve continuously in time under the influence of current sensory information as well as the current activation state (Constantinidis & Steinmetz, 1996; Rao, Rainer, & Miller, 1997; Smyrnis, Taira, Ashe, & Georgopoulos, 1992). Exactly how sustained activation is neurally realized, however, is not clear. One class of models achieves a stable memory state using bistable networks in which a stable state of sustained activation or “on” state coexists with an “off” state (Amari, 1989; Amari & Arbib, 1977; Compte, Brunel, Goldman-Rakic, & Wang, 2000). Within the “on” state, locally excitatory and laterally inhibitory interactions among neurons create sustained activation patterns.

The DFT is in this class of neural networks (see also Erlhagen & Schöner, 2002; Thelen et al., 2001). To describe the theory, consider an activation field defined over a metric spatial dimension, x (e.g., the angular location of a target). The continuous evolution of the activation field is described by an activation dynamics, that is, a differential equation that generates the temporal evolution of the field by specifying a rate of change, $dw(x,t)/dt$, for every activation level, $w(x,t)$, at every field location, x , and at any moment in time, t . The field achieves stable patterns of activation through time via an inverse relationship between the rate of change and the current level of activation (for a discussion of related dynamic systems concepts, see Braun, 1994). This means that at high levels of activation, negative rates of change drive activation down, while at low levels, positive rates of change drive activation up.

The activation level that emerges from this basic stabilization mechanism is a function of the balance of different inputs to the field (e.g., perceptual inputs) and neural interactions within the field. We use a locally excitatory/laterally inhibitory form of interaction depicted in figure 14.3A. According to this type of interaction, neurons that “code” for similar values along the spatial dimension, x (e.g., similar locations in space), excite one another, while neurons that code for very different values (e.g., different locations in space) inhibit one another. Note that only sufficiently activated sites contribute to interaction. This is achieved by passing activation levels through the sigmoidal function shown in figure 14.3B. According to this function, highly active neurons contribute fully to interaction (they are weighted by a value of 1), while inactive neurons do not participate (they are weighted by a value of 0). Such threshold functions are necessarily nonlinear to achieve the co-existing “on” and “off” states central to bistable networks.

Figure 14.4 shows two simulations (A–C and D–F) that illustrate the consequences of combining these concepts in a dynamic field. In each panel of this figure, the angular location of targets in the task space is along x (recall that 180° is along the midline of the task space), y shows the activation of each site in the field, and z captures time from the start (front of the figure) to the end of a spatial recall trial (1 time step = 5 msec). In figure 14.4, A and D show activation in an input field, whereas B, C, E, and F show activation in a two-layered dynamic field. Because cortical neurons never project both excitatorily and inhibitorily onto target neurons, lateral inhibition must be mediated through an ensemble of interneurons. A generic formulation (Amari & Arbib, 1977) is to introduce a second, inhibitory activation




Figure 14.3 Two mathematical functions that govern neural interactions in the DFT. (A) A local excitation/lateral inhibition function where neurons that code for similar locations excite one another (positive activation) while neurons that code for dissimilar locations inhibit one another. (B) A sigmoidal function that determines which neurons participate in interaction: highly active neurons contribute fully to interaction (weighted by 1), while inactive neurons do not participate (weighted by 0).

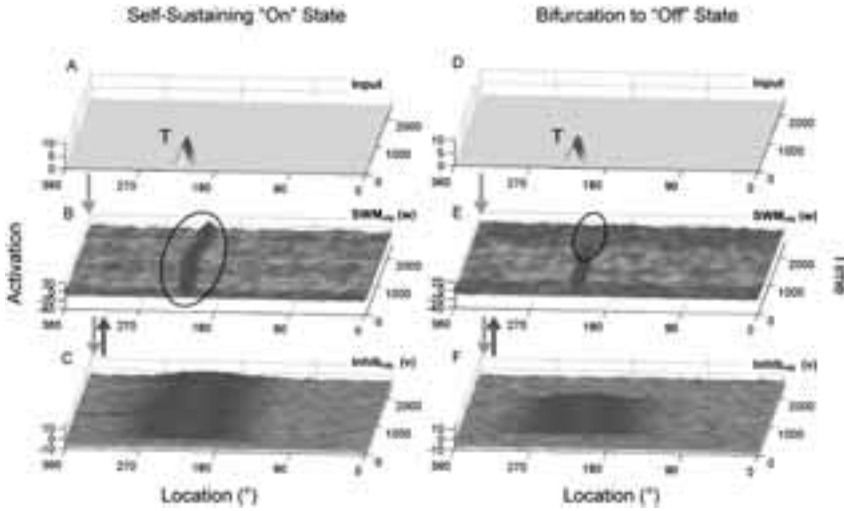


Figure 14.4 The dynamics of the DFT yield qualitatively different attractor states. The simulation in *A–C* shows the self-sustaining “on” state. The target (*T*) is presented at 220° in an input field (*A*), which leads to a peak of activation in the spatial working memory field (*B*), which is sustained during the memory delay (see oval) via locally excitatory and laterally inhibitory interactions with the inhibitory field (*C*). Because the system is close to the transition between the “on” and “off” states with this parameter setting, the simulation in *D–F* shows a spontaneous bifurcation to the “off” state during the memory delay (see oval in *E*) with the same parameters and input (see *D*). In *A–F*, location is across the *x*-axis, activation on the *y*-axis, and time on the *z*-axis. Arrows indicate excitatory (green) and inhibitory (red) connections among layers.

field, $v(x,t)$ (figure 14.4*C,F*), which receives input from the excitatory activation field, $w(x,t)$ (figure 14.4*B,E*), and in turn inhibits that field. These reciprocal interactions are captured by the green (excitatory) and red (inhibitory) arrows connecting paired fields (e.g., between panels *B* and *C*).

The dynamic interaction among the w and v layers combined with a nonlinear threshold function yields bistability—the potential to be in either an “on” or an “off” state. The simulation in figure 14.4*A–C* shows the “on” state. At the start of this simulation, neurons in the w and v layers are at a homogeneous resting level. Next, a target appears in the task space at 220° for 2 seconds, creating the activation profile shown in the input field (*T* in figure 14.4*A*). This event provides localized input to those neurons in the w field spatially tuned to this location, building a “peak” of activation centered at the target location. Next, the target is turned off; nevertheless, neurons in the w field remain in the “on” state—the field actively (and stably!) maintains a “working” memory of the location for the entire 10 second

delay (see circle in figure 14.4B).³ This occurs because of the locally excitatory/laterally inhibitory interactions among neurons in the w and v layers. Specifically, neurons in the w field near the peak of activation excite one another, keeping activation around in the absence of input. At the same time, this local population of neurons excites associated neurons in the v field, which laterally inhibit neurons that code for different locations. This prevents activation from diffusing across the network. Note, however, that in the presence of noise, the location of the peak “drifts” during the delay (see movement of the circled peak in figure 14.4B): at the end of the delay in this simulation, the peak is centered at 223° . Consequently, the model would recall a location 3° to the right of the target. Thus, although the network remains stably in the “on” state, the positioning of the peak is only marginally stable (see also Compte et al., 2000).

The second simulation in figure 14.4D–F shows a bifurcation (i.e., dynamic transition) from the “on” to the “off” state. All the details of this simulation are identical to one described above, but the behavior of the model is quite different. Now, rather than stably remembering the location during the 10 second delay, spontaneous fluctuations cause an initially self-sustaining peak to die out before the end of the trial (see the disappearance of the circled peak in figure 14.4E). This highlights that both “on” and “off” states are simultaneously available to the model; that is, the system is bistable.

14.2.2 Metric Memory Relative to Reference Frames and Geometric Categories

The concepts above describe a neurally plausible, bistable network for spatial working memory; however, what precisely does this model remember? The location the model actively represents must somehow be linked to a concrete location in the task space. We made this trivially easy by assigning a stable reference frame to the field, with the midline of the task space at 180° and the target at 220° . This is, of course, not realistic (although this is the approach taken by the majority of the spatial memory models proposed to date; see, e.g., J. Huttenlocher et al., 1991). What we need is a way to dynamically link spatial memories to a stable reference frame despite changes in the position of the observer (e.g., due to eye movements) or changes in the environment (e.g., due to movement of the reference frame). In this section, we first describe our approach to the dynamics of reference frame calibration and realignment. Then, we describe how we integrate our approach to spatial working memory with this real-time calibration system.

14.2.2.1 THE DYNAMICS OF REFERENCE FRAME CALIBRATION To make the challenges of establishing a stable reference frame concrete, consider the representation of a location in space within two reference frames—an

egocentric, retinotopic frame and an object-centered frame where 180° maps onto a perceived reference axis in the task space (e.g., the two dots shown in figure 14.1; see Simmering et al., in press). How do we stably map the perception of the reference dots—which could land at any spot on the retina—onto 180° in the object frame?

We illustrate our approach to calibration in figure 14.5, which uses the top four layers in our theory: a single-layer egocentric perceptual field (PF_{ego} ; figure 14.5A),⁴ a system that transforms (e.g., rotates) an angular location in an egocentric frame into a location in an object-centered frame (shift; figure 14.5B), a single-layer object-centered perceptual field (PF_{obj} ; figure 14.5C), and a long-term memory field associated with the object-centered




Figure 14.5 Reference frame calibration and realignment in the DFT. On the first trial (A–D), activation from the reference frame (R) is originally presented at 180° in the egocentric perceptual field (A). Across eye movements (indicated by arrows in A), the reference frame and target inputs (T; presented 40° from the reference) shift; this shift in the egocentric frame is corrected (B) to bring the reference frame input into 180° in the object-centered perceptual field (C). The continued activation of the object-centered reference frame leaves a trace in long-term memory (D). The second trial (E–H) begins with a resting level boost (*b*-boost in G), which creates a peak in the object-centered field (G) due to input from long-term memory (H). After a brief interval when the shift mechanism is off (F), reference input in the egocentric perceptual field (E) is realigned (F) with the reference from the previous trial. Axes and arrows are as in figure 14.4.

perceptual field ($LTM_{PF_{obj}}$; figure 14.5D). This simulation begins with perception of the midline reference frame in the task space at 180° in the egocentric frame (i.e., on the retina; see R at the front of figure 14.5A). Since this is the first trial in this task context, the initial mapping from the egocentric frame to the object-centered frame is arbitrary: the perception of the reference frame on the retina creates a pattern of activation in PF_{ego} that can pass directly into PF_{obj} with no “shift” required.

Once the egocentric and object-centered frames are in register, the system has to maintain this alignment even when the eyes move and the position of the reference frame on the retina shifts. For example, toward the start of the trial in figure 14.5A (at 500 msec), we shifted the position of the reference frame -40° in the retinal frame (reflecting a $+40^\circ$ eye movement; see arrow toward the front of figure 14.5A). Consequently, when activation is passed from PF_{ego} to PF_{obj} , everything must be “shifted” by $+40^\circ$ so that the perception of the reference frame in the retinal field maps onto 180° in the object-centered field (for a neurally plausible approach to such transformations, see Sauser & Billard, 2005). How does the nervous system know the shift amount (Δy) that will keep activation in one frame aligned with activation in the other?

There are many relevant sources of information the nervous system could use to solve this problem provided that an initial alignment has been established (e.g., multisensory integration [see Bard, Fleury, Teasdale, Pailard, & Nougier, 1995] and efference copy signals [see Krommenhoek, van Opstal, Gielen, & van Gisbergen, 1993]). For simplicity, the simulations in figure 14.5 relied on a single, novel calibration mechanism. At each time step, the model computed the spatial correlation between the pattern of activation across the x-axis in PF_{ego} and the pattern of activation across the x-axis in PF_{obj} at every possible shift value (imagine taking the activation profile in PF_{ego} and shifting it to the left and right to find the best “match” with the profile in PF_{obj}). If the maximum correlation was above a criterion threshold, we updated the shift parameter, setting it to the shift amount associated with the maximum correlation. This allowed the model to dynamically adjust the mapping from PF_{ego} to PF_{obj} , compensating for any shift of the reference frame on the retina due to self motion.⁵ For example, the initial eye movement in figure 14.5A caused a change in the shift value (see circle in figure 14.5B), such that the peak of activation in PF_{ego} at 140° was input correctly at 180° in PF_{obj} after the eye movement (see the single reference peak R at 180° in figure 14.5C).

But, of course, this is not quite yet the full picture because the system is not just tracking a reference frame—it is also perceiving targets, distractors, and so on. Thus, the system must be able to distinguish the reference signal from other inputs and update the shift parameter only when the reference is detected. One distinguishing feature of reference signals is that they remain stably visible in the task space (vs., e.g., “target” stimuli that are transiently visible or movable). To capitalize on this feature of reference frames and facilitate reference frame tracking, we gave PF_{obj} weakly self-sustaining

dynamics: this field could hold on to the reference frame for short periods of time.

The simulation in figure 14.5 shows that this produces robust calibration. For instance, we turned a target on at 220° in the retinal frame at 1 second in this simulation (T in figure 14.5A). Because the reference frame is still visible and tracked by PF_{obj} , this event does not cause a realignment. Rather, the system brings the target input into PF_{obj} at the correct location in an object-centered frame— 260° (T in figure 14.5C). Weakly self-sustaining dynamics in PF_{obj} can also help avoid spurious updating that might occur during, for instance, short-term occlusions of the frame. In figure 14.5A, we turned the reference signal off in PF_{ego} at 9 seconds. Despite this occlusion event, PF_{obj} holds on to the current frame (sustained peak in figure 14.5C). When the frame reappears at a new location (second arrow toward the back of figure 14.5A), the system updates the shift parameter (figure 14.5B) and realigns the reference frame.

The final piece of the puzzle is that the system needs some way to re-stantiate the same object-centered frame from trial to trial (or from situation to situation). Recall that we initially passed activation from the retinal frame into the object-centered frame at the start of trial 1. But this will not work on trial 2! Rather, we want to make sure that, wherever the system perceives midline on the retina at the start of trial 2, it correctly brings this input into PF_{obj} at 180° . To achieve this, we can give PF_{obj} a simple form of long-term memory ($LTM_{PF_{obj}}$ in figure 14.5D). Whenever there is above-threshold activation in PF_{obj} , this leaves a trace in $LTM_{PF_{obj}}$. Reversely, activation traces in $LTM_{PF_{obj}}$ feedback as excitatory input into PF_{obj} . Importantly, $LTM_{PF_{obj}}$ has a longer time scale. Thus, activation traces grow slowly in this field and decay slowly in the presence of competing input. As shown in figure 14.5D, the dynamic interaction between PF_{obj} and $LTM_{PF_{obj}}$ on trial 1 results in a robust trace of the midline reference frame at 180° . Importantly, there is virtually no trace of the target location. This reflects the fact that the target input was transient while the reference input was visible in the task space for most of the trial. Thus, $LTM_{PF_{obj}}$ emergently performs an importantly function—it remembers reference frames.

Given this, how can the system use the information in $LTM_{PF_{obj}}$ on trial 2? The simulation in figure 14.5E–H shows our approach to this issue. At the start of trial 2, we turned the shift mechanism off for the first 500 msec (figure 14.5F). In addition, we boosted the resting level (b) of neurons in PF_{obj} (b boost in figure 14.5G). In the presence of input from $LTM_{PF_{obj}}$, this boost causes an activation peak to form in PF_{obj} at the location of maximal input from $LTM_{PF_{obj}}$, that is, at 180° . Next, we relowered the resting level and turned the shift mechanism back on. The system computed the correct shift amount needed to map the midline activation peak at 140° in the retinal field onto 180° in the object-centered field. Thus, the system effectively reacquired the old reference frame on trial 2 with a relatively brief (500 msec) and conceptually simple modulation of the calibration dynamics. From that point on, the simulation functions as before: the

system correctly mapped perception of the target (T in figure 14.5E) to the correct location in the object-centered frame and survived another occlusion/eye movement sequence at the end of the trial (arrow toward the back of figure 14.5E). Note that, by the end of trial 2, the system is starting to construct a relatively strong long-term memory of the object-centered reference frame. This will make the realignment process even more robust on subsequent trials.

In summary, figure 14.5 highlights several key innovations in our theory: (1) the weakly self-sustaining dynamics of PF_{obj} can help maintain calibration despite, for instance, transient occlusions; (2) the stability of reference frames in the task space simplifies the challenges of calibration; (3) given a stable and reliable reference input, a simple form of long-term memory can suffice to build a robust memory of the reference frame across trials; (4) the DFT can use this type of long-term memory to reliably realign reference frames from trial to trial. Although our approach has several innovations, it is, of course, limited. For instance, we did not include processes that would determine if the person or the world moved (e.g., efference copy), nor did we include nonspatial features that will have to play an important role in the identification of reference frames. Nevertheless, the four layers of the DFT described in this section ground the model in a perceived frame of reference, moving beyond several other models of spatial memory (e.g., J. Huttenlocher et al., 1991).

14.2.2.2 SPATIAL WORKING MEMORY WITHIN A STABLE REFERENCE FRAME

Given that the DFT has a mechanism for dealing with the calibration and realignment of reference frames, we turn to the next central question: how does our model form a stable working memory of a location relative to a perceived reference frame? Conceptually, this requires the integration of a spatial working memory field (figure 14.4) with an object-centered perceptual field (figure 14.5). This integration is shown in figure 14.6, which shows three layers from our model: the excitatory layer of a perceptual field (PF_{obj} ; figure 14.6A), the excitatory layer of a spatial working memory field (SWM_{obj} ; figure 14.6C), and a shared layer of interneurons ($Inhib_{obj}$; figure 14.6B). The arrows in figure 14.6 show how these layers are coupled (green arrows indicate excitatory coupling; red arrows indicate inhibitory coupling).

What does this particular type of integration achieve? First, consider the behavior of the perceptual field. As can be seen in figure 14.6A, PF_{obj} operates as before—at the start of the trial, it brings the midline input into 180° in an object-centered reference frame. Next, the target appears at 220° (T in figure 14.6A). This creates a peak of activation centered at this location. When the target disappears, a peak reforms in PF_{obj} at 180° as the system locks onto the stable reference cues in the task space. The effect of coupling PF_{obj} to SWM_{obj} is shown in figure 14.6C. At the start of the trial, SWM_{obj} receives relatively weak reference input from PF_{obj} . Next, the target is turned on, passing strong target-related input into the working memory

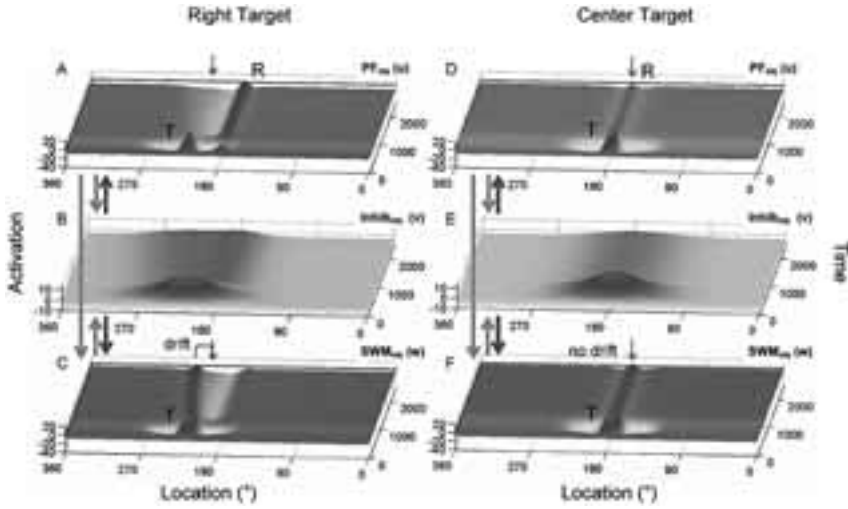


Figure 14.6 Spatial working memory within a stable reference frame. The perceptual field (A) maintains the reference frame (R) throughout the trial, leaving an inhibitory trace in (B). The target (T), presented at 220°, forms a self-sustaining peak in spatial working memory (C); this peak drifts away from 180° (see “drift” in C) over delay due to the inhibition associated with the reference frame. In contrast, when the target is presented at 180° (D–F), activation from the reference frame in the perceptual field (D) aligns with the target peak in spatial working memory (F); this peak does not drift over delay (F). Rather the peak is stabilized by excitatory input (green arrow) from the perceptual field (D). Axes and arrows are as in figure 14.4.

field. This event moves the working memory field into a strongly self-sustaining state. After the target disappears, the SWM_{obj} field maintains an active memory of the target location during the delay. Importantly, this occurs even though PF_{obj} has reacquired the reference frame. This highlights the useful bistability of SWM_{obj} described previously—because SWM_{obj} is in a stable attractor state, it can effectively hold on to the target, even while PF_{obj} deals with the dynamics of reference frame calibration. Thus, the three-layer system shown in figure 14.6 can achieve the dual goals of remembering the target item and staying calibrated with a perceived reference frame.

This dual ability has both costs and benefits for the cognitive system. One “cost” is readily apparent in figure 14.6C: the peak of activation in SWM_{obj} “drifts” away from 180° during the memory delay (“drift” toward the back of figure 14.6C). Thus, our model shows an emergent geometric category bias. This occurs because the peak of activation at midline in PF_{obj} passes activation to the $Inhib_{obj}$ layer at 180°. This creates greater inhibition on the 180° side of the target-related peak in SWM_{obj} , effectively repelling

this peak away from the reference frame (for evidence of such effects, see Spencer & Hund, 2002, 2003). Although this produces a systematic bias in memory, J. Huttenlocher and colleagues (J. Huttenlocher et al., 1991; J. Huttenlocher, Hedges, & Vevea, 2000) have pointed out such biases can have adaptive consequences. For instance, geometric bias minimizes the number of cross-category confusions, that is, cases where the target was shown to the left of midline but participants respond to the right. This bias also helps exaggerate the “leftness” or “rightness” of the target item, which could facilitate recall from long-term memory. Indeed, from a Bayesian perspective, geometric category boundaries can help create an optimal spatial memory system, trading off systematic bias for an overall reduction in variance (see J. Huttenlocher et al., 2000).

We agree that this is a viable interpretation, particularly in the context of our model that integrates SWM_{obj} with reference frame calibration: one could have a very accurate memory of an item in memory, but if this memory is not grounded in the world, effectively that “accurate” memory cannot be realized in an accurate response. Thus, although the reference calibration processes in our model lead to systematic bias, they can also lead to enhanced performance. This is captured in the simulation shown in figure 14.6D–F. Here, we presented a target aligned with midline (i.e., at 180°). During the memory delay, SWM_{obj} is stabilized because it is attracted toward/stabilized by the midline input from PF_{obj} . This is consistent with empirical results: children as young as 3 years as well as adults show lower variability for targets aligned with reference frames with no systematic bias (Schutte & Spencer, 2002; Spencer & Hund, 2002, 2003).

14.2.3 Integration of Short-Term and Long-Term Spatial Memories

The theory and model implementation we have sketched thus far offers a neurally plausible account of the real-time processes that construct and maintain a memory of a single target location within a particular, well-calibrated reference frame, as well as how these processes give rise to geometric category biases. The next step is to integrate these real-time processes with processes that live at the trial-to-trial time scale (and beyond). Our use of a long-term memory field (LTM_{PFobj}) to deal with reference frame realignment foreshadowed our approach to this issue.

Figure 14.7 shows a simulation of the bottom five layers in our model (see C–G in figure 14.2), though we show only four of these layers for ease of viewing: PF_{obj} (figure 14.7A), LTM_{PFobj} (figure 14.7B), SWM_{obj} (figure 14.7C), and a long-term memory field coupled to spatial working memory (LTM_{SWMobj} ; figure 14.7D). We have already discussed the first three layers in this figure. The final layer— LTM_{SWMobj} —operates in a manner comparable to LTM_{PFobj} . Specifically, whenever there is above threshold activation present in SWM_{obj} , a trace is left at the associated locations in LTM_{SWMobj} . Conversely, these long-term memory traces feed back as excitatory inputs into SWM_{obj} .




Figure 14.7 Integration of short-term and long-term spatial memories in the DFT. On the first trial (A–D), the perceptual field (A) maintains the reference frame (R) throughout the trial. Activation from both the target (T) and the reference leave traces in the associated long-term memory field (B). The target, presented at 240° , forms a self-sustaining peak in spatial working memory (C), which drifts away from midline over delay. The peak in spatial working memory leaves a trace of activation in the associated long-term memory field (D). By the third trial (E–H), the repeated activation from the reference frame in the perceptual field (E) has left a robust memory for the reference in long-term memory (F). The target, presented at 220° , still drifts away from midline in spatial working memory (G), but now long-term memory (H) shows traces of three target locations—an experience-dependent spatial category. Axes and arrows are as in figure 14.4.

As illustrated in figure 14.7, the reciprocal interplay between SWM_{obj} (figure 14.7C) and LTM_{SWMobj} (figure 14.7D) leads to the emergent formation of experience-dependent spatial categories. The simulation in figure 14.7A–D shows the first trial in a spatial recall task. The target in this simulation was presented at 240° . During the course of a single trial, the model builds up a graded long-term memory of the target location (figure 14.7D) that has an impact on SWM_{obj} . For instance, we ran a simulation to the 240° target with the same model parameters, but without LTM_{SWMobj} . The model showed a stronger bias away from midline in this case (by 2°). Thus, the excitatory input from LTM_{SWMobj} into SWM_{obj} serves to stabilize against drift—even on a single trial.

But this is only the case when the distribution of activation in $LTM_{SWM_{obj}}$ is centered around the target location. If the model visits several close locations in a row and is then shown a target at a different (but not too distant) location, the working memory of the new location will be biased toward the previously responded-to targets. This is shown in figure 14.7E–H. Here, we simulated the model across two trials with close targets (240° and 260° ; see arrows in figure 14.7E,G). Then we probed the model's response when an "inner" target (220°) was presented (figure 14.7E). The model showed a *larger* bias away from midline in this case relative to a simulation where $LTM_{SWM_{obj}}$ was not present. Why?

At the start of this recall trial, the distribution of activation in $LTM_{SWM_{obj}}$ is centered at 255° (circle in figure 14.7H). Consequently, the working memory peak built by the target presentation is both repelled from midline *and* attracted toward 250° . Note, however, that by the end of this trial, the distribution of activation in $LTM_{SWM_{obj}}$ has shifted closer to 220° . This shows the complex dynamics that underlie the formation of experience-dependent spatial categories. Although we have yet to probe these dynamics at this trial-to-trial level, the model effectively captures the buildup of long-term memory effects over blocks of trials, as well as the systematic shift in bias created by having participants move more frequently to an inner or outer target (Hund & Spencer, 2003; Spencer & Hund, 2002, 2003). Finally, we note that while $LTM_{SWM_{obj}}$ is forming a representation of the target distribution, its counterpart— $LTM_{PF_{obj}}$ —is forming a robust memory of the reference frame (see figure 14.7F). This distribution of effort by the model emerges from the different dynamics of PF_{obj} and SWM_{obj} , rather than some a priori knowledge about the difference between targets and reference frames.

14.2.4 Task Generality and Behavioral Flexibility

The preceding sections describe how the DFT accounts for several basic spatial cognitive abilities, but all of our simulation examples were grounded in a particular task—spatial recall. Can the DFT generalize to other spatial cognitive tasks and take a critical step toward task generality and behavioral flexibility?

We illustrate the flexibility of the DFT by showing how this theory generalizes to a new task—position discrimination (for a related example, see Spencer et al., 2006). Recall that in position discrimination tasks, two stimuli (S1 and S2) are presented in quick succession (e.g., 500 msec apart), and the participant judges whether the stimuli were in the "same" or "different" locations (see Simmering et al., in press). To capture performance in this task, the DFT must generate a *same/different* decision. Such decisions can be achieved by increasing the resting level of neurons in PF_{obj} and $Inhib_{obj}$ when the second stimulus is presented. In some cases, this will result in a peak in SWM_{obj} (and no peak in PF_{obj})—the basis for a "same" response. In other cases, this will result in a self-sustaining peak in PF_{obj} (and no peak in SWM_{obj})—the basis for a "different" response.

Figure 14.8 shows a simulation of a “same” (figure 14.8A–C) and “different” (figure 14.8D–F) response (note that we did not include a reference input in these simulations for simplicity). These simulations are identical except for the positioning of the stimuli: we presented two sequential stimuli at the same location in figure 14.8A–C (arrow toward the back of figure 14.8A,C), while we presented two sequential stimuli at slightly different locations in figure 14.8D–F (arrows in figure 14.8D,F). Note that the stimuli were “on” for 1 second each with a 1 second delay between (see Simmering & Spencer, 2006b). As shown in figure 14.8A–C, the presentation of the first stimulus (S1 in figure 14.8A) creates a peak of activation in SWM_{obj} that is sustained during the brief delay. When the second stimulus is presented at the same location (S2 in figure 14.8A), this input blends with the activation peak in SWM_{obj} , which in turn suppresses activation in PF_{obj} at the associated location (via $Inhib_{obj}$). Consequently, when the stimulus turns off, the SWM_{obj} peak remains and the model responds “same” (figure 14.8C).⁶

The same initial events occur in the simulation in figure 14.8D–F. Now, however, the second stimulus is presented to the left of the first (S2 in figure 14.8D). Importantly, this stimulus comes into PF_{obj} at a site that is




Figure 14.8 Position discrimination in the DFT. In A–C, two stimuli (S1 and S2) are presented at 180° in quick succession (A). S1 leads to a trough of inhibition (B) and a peak in spatial working memory (C). When S2 is presented in the same location, activation is maintained in spatial working memory (C), leading to a same peak, and suppressed in the perceptual field (A). In D–F, the stimuli are presented in different locations (D). Because S2 falls outside of the inhibited region in the perceptual field (D) created by the peak in spatial working memory (F), S2 builds a different peak in the perceptual field (D), which suppresses activation in spatial working memory (F). Axes and arrows are as in figure 14.4.

just outside of the inhibitory range of the spatial working memory peak. This allows activation to grow in PF_{obj} . At the same time, the resting level boost in PF_{obj} and $Inhib_{obj}$ makes the perceptual field more strongly self-sustaining. Consequently, a peak forms in PF_{obj} , and this suppresses activation in SWM_{obj} (via $Inhib_{obj}$). At the end of the trial, there is an above-threshold peak in PF_{obj} (figure 14.8D) and a below-threshold peak in SWM_{obj} (figure 14.8F). Note that we can easily “update” spatial working memory by moving the resting levels in PF_{obj} and $Inhib_{obj}$ back to their initial state. When this occurs, a peak builds in SWM_{obj} at the site of S2.

The ability of the DFT to account for recall and discrimination performance within the same framework *and using the same parameters* is unique—we know of no other model of spatial cognition that captures this range of performance (for an example of behavioral flexibility within the domain of object category formation, see Love, Medin, & Gureckis, 2004). Importantly, our theory has also generated several novel behavioral predictions that have been successfully tested with both children and adults (Simmering & Spencer, 2006b; Simmering et al., in press). And we are currently extending our theory to account for performance in other tasks, including multi-item change detection (Johnson, Spencer, & Schöner, 2006).

14.3 IN WHAT SENSE IS THE DFT NEURALLY PLAUSIBLE?

We began this chapter with two observations about the search for developmental mechanisms: (1) there is a growing trend to search for mechanisms of development in brain–behavior relations, and (2) as researchers explore the longer time scales of learning and development, it becomes difficult to offer a detailed, process-based account of real-time behavior. Thus, there is a need to develop theories of development grounded in a rich understanding of brain–behavior relations that offer both a detailed account of real-time behavior and an account of the longer time scales of learning and development. In this section, we position the DFT within the field of theoretical neuroscience and evaluate whether this theory is grounded in an understanding of brain–behavior relations. This sets the stage for our explorations of development in the subsequent sections.

Before discussing the ways in which the DFT is based on neural principles, we first place our theory within the broader field of theoretical neuroscience. A growing number of researchers have argued that we should take inspiration from the densely interconnected and dynamic nature of the brain to rethink 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, recurrent, densely interconnected, complex systems that violate core assumptions of encapsulation and separability. We alluded to this in the overview

of the DFT by noting that PF_{obj} plays a role in virtually every phenomenon we have modeled.

One of the dominant approaches in theoretical neuroscience takes a biophysical approach to the study of brain–behavior relations: theoreticians attempt to build neurally realistic models of single neurons that capture key elements of neuronal firing rates, neurotransmitter action, and so on (Durstewitz, Seamans, & Sejnowski, 2000; Salinas, 2003). These neural models are then coupled together into populations, and populations are coupled together based on evidence of, for instance, cortical–cortical connections. Although the biophysical approach has led to new insights into brain function and neuronal dynamics, such models typically fail to capture many behavioral details. This is partly due to concrete limitations: such models are very complex, have immense computational demands, and are difficult to analyze (Haskell, Nykamp, & Tranchina, 2001). Thus, while this approach to rethinking cognition is very promising, biophysical models do well at the level of neurons but poorly at the level of behavior (Finkel, 2000).

We contend that the DFT overcomes this limitation with its emphasis on the functional organization of behavior while maintaining a strong commitment to neural principles. Dynamic fields were first proposed to account for neural activation in visual cortex (Amari, 1977; Amari & Arbib, 1977). Continuous dynamic fields provided a natural way to capture the neural dynamics of this cortical area given its topographic organization. The advent of population coding ideas extended this application to non-topographically organized cortical areas such as motor cortex (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schöner, & Riehle, 2003). Now, rather than ordering neurons in a field by their position on the cortical surface, one can order them based on the continuous, metric features that each neuron “codes” for—including space (for evidence of spatially tuned neural activity across a range of cortical areas, see di Pellegrino & Wise, 1993; Georgopoulos, Kettner, & Schwartz, 1988; Georgopoulos, Taira, & Lukashin, 1993; Wilson, Scalaidhe, & Goldman-Rakic, 1993). Critically, although neurons in many cortical areas are not topographically ordered, they still interact in a locally excitatory and laterally inhibitory manner (e.g., Georgopoulos et al., 1993).

This basic form of interaction leads to different types of self-sustaining dynamics. For instance, evidence suggests that patterns of activation in posterior parietal cortex are weakly self-sustaining, while patterns in prefrontal and premotor cortex are more strongly self-sustaining. As an example, patterns of activation in both parietal and prefrontal cortical areas are sustained during delays in spatial versions of delayed match-to-sample tasks; however, activation in posterior parietal cortex does not survive the presentation of new stimuli (Constantinidis & Steinmetz, 1996; Steinmetz & Constantinidis, 1995), while activation in prefrontal cortex does (di Pellegrino & Wise, 1993). We were informed by these data when we made PF_{obj} weakly self-sustaining and SWM_{obj} strongly self-sustaining.

But in our theory, we link these two fields together via a shared layer of interneurons. This seems to violate the distant cortical–cortical connections that link parietal and prefrontal cortex. Although this is the case in terms of the architecture of the model, we prefer the three-layered perceptual field–inhibitory field–spatial working memory structure because it more strongly mirrors the cytoarchitecture of visual cortex (Douglas & Martin, 1998), giving us an entry point into the dynamics that emerge from the layered structure of cortex.

In addition to these ties to real-time cortical dynamics, our approach to the modulation of resting levels in, for instance, $\text{Inhib}_{\text{obj}}$ and SWM_{obj} during a recall trial is neurally plausible. Such modulation of resting levels can be achieved, for instance, by global excitatory input from one cortical area to another. Finally, our approach to long-term memory is based on established neural principles: the long-term memory fields in the DFT capture a form of Hebbian learning (Schöner, 2006; Wilimzig & Schöner, 2006). Because this raises the issue of how we approach the integration of time scales, including the time scales of development, we save further discussion of these issues for section 14.6.

Connections between the DFT and neurophysiology provide important constraints for our theory. But they also raise the exciting possibility of testing our theory using neurophysiological techniques. Indeed, several recent studies have demonstrated that dynamic fields can be directly estimated using single-unit recording (Bastian et al., 1998, 2003; Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999; Jancke et al., 1999). The first step in these studies was to map the responses of neurons to basic stimulus and create a field by ordering the neurons based on their “preferred” stimulus. This was followed by a behavioral task that probed, for instance, several reaction-time predictions by a DFT of movement preparation (Bastian et al., 1998, 2003; Erlhagen & Schöner, 2002). Note that this same theory has also been tested using event-related potential (ERP) techniques (McDowell, Jeka, Schöner, & Hatfield, 2002). These studies all reported a robust relationship between predictions of dynamic field models and neural measures, suggesting that this particular marriage between theoretical and behavioral neuroscience could be extremely generative in the future.

14.4 THE DEVELOPMENT OF SPATIAL COGNITION

One issue central to *any* theory of spatial cognition is to explain the mechanisms that underlie developmental changes in the diverse phenomena described previously (see section 14.1). We have already alluded to some of these changes. First, there is an improvement in children’s ability to actively and stably maintain spatial information in working memory beginning around 10 months (Thelen et al., 2001) and continuing throughout childhood. This results in a systematic reduction in errors over development in the spaceship task. Specifically, 3-year-olds typically made 10–20° errors to targets to the left and right of midline following delays of 5–10 seconds

(Schutte & Spencer, 2002). By contrast, 6- and 11-year-old children typically made 6–8° errors over delays (Spencer & Hund, 2003), whereas adults' errors generally ranged between 3° and 5° (Spencer & Hund, 2002; for related effects, see Hund & Plumert, 2002, 2003, 2005; Hund et al., 2002; Plumert & Hund, 2001).

Second, there is a qualitative shift in how younger versus older children remember locations near axes of symmetry in spatial memory tasks. Two- to 3-year-olds show biases *toward* symmetry axes, while 6- to 11-year-olds show biases *away from* the same axes (Hund & Spencer, 2003; Schutte & Spencer, 2002; Schutte et al., 2003; Spencer & Hund, 2003). Note that these age ranges differ slightly across tasks (see, e.g., J. Huttenlocher et al., 1994; Sandberg et al., 1996). According to the category adjustment model proposed by J. Huttenlocher et al. (1991), the shift in geometric biases reflects a change in how children categorize space: younger children form large geometric categories and are biased toward a spatial prototype at the center of the space (e.g., at midline), while older children subdivide space into smaller regions, showing biases away from a category boundary at midline (see J. Huttenlocher et al., 1991, 1994). As we alluded to above, this model does not offer an explanation for *how* this qualitative developmental shift occurs.

In addition to changes in geometric category biases, young children show developmental changes in experience-dependent category biases. Specifically, as spatial working memory becomes more precise, the long-term traces created by short-term memories also become more precise. Evidence for this comes from a sandbox version of the canonical Piagetian A-not-B task (Schutte et al., 2003; Spencer & Schutte, 2004; Spencer et al., 2001). In this task, a toy is hidden in a large rectangular sandbox and covered up so that there are no visible cues marking the hiding location. After a short delay, children are allowed to dig up the toy. If the toy is repeatedly hidden at an "A" location, and then hidden at a nearby "B" location, 2- to 6-year-old children will search for the toy at locations shifted systematically toward "A," that is, they search near A and not B. Importantly, there is a systematic reduction in the spatial range across which such errors occur over development: 2- and 4-year-olds show biases toward an A location that is 6–9 inches away from B, while 6-year-olds only show biases toward an A location 2 inches away, and 11-year-olds show only a marginal bias toward A at this separation (Schutte et al., 2003).

Finally, there are improvements in behavioral flexibility over development (see, e.g., Hund & Plumert, 2005). In a recent study, for example, we examined children's and adults' ability to modulate performance across recall and recognition versions of a simple paper and pencil task (Spencer et al., 2006). The recall task was similar to the recall tasks described previously: participants were shown a dot inside a small rectangle, the dot was covered up, there was a short delay, and then people placed an "X" at the remembered location in a blank rectangular frame (see also J. Huttenlocher et al., 1994). The recognition task was identical through the delay interval; however, after

the delay, participants were shown a set of colored dots—the “choice” set—inside a rectangular frame. Their task was to say the color of the dot that was in the same location as the original “target” dot. Children and adults showed comparable geometric biases across tasks as long as the choice set presented in the recognition task had a large number of items that spanned a relatively large spatial range. With a smaller number of items to choose from, children were strongly influenced by the local details of the choice set. Although this was maladaptive on some trials, it led to more accurate performance on others. In some cases, children selected a dot that was more accurate than their recall response because their memory had “drifted” beyond the range of possible choices. In other cases, children showed a bias against selecting the outermost choices in favor of a more accurate selection. This points toward developmental differences in how effectively children can “rise above” the local details of the situation, as well as the potentially adaptive and maladaptive consequences of this lack of behavioral flexibility.

14.5 THE DEVELOPMENT OF SPATIAL COGNITION, THE DFT, AND THE SPATIAL PRECISION HYPOTHESIS

We have made promising advances in understanding the development of spatial cognition using a relatively simple hypothesis—the *spatial precision hypothesis* (SPH). This hypothesis states that the spatial precision of neural interactions becomes more precise and more stable over development (Schutte et al., 2003; Spencer & Hund, 2003). This hypothesis is illustrated in figure 14.9 using the same interaction function shown in figure 14.3A. Each curve in this figure shows an example of an interaction function at some point in development, with early development shown with darker lines and later development with lighter lines. As the interaction functions move from early to later development, the spatial precision of locally excitatory interactions narrows and excitation increases. At the same time, laterally inhibitory interactions expand their spatial range and grow stronger. These changes result in relatively unstable self-sustaining interactions early in development that are sensitive to input across a broad spatial range, as well as stable self-sustaining peaks later in development that are sensitive only to input at narrow separations. This can explain the reduction in A-not-B-type effects in the sandbox task over development: with enhanced precision, A and B must be spatially close for the working memory associated with B to “feel” the effects of the long-term memory of A (for details, see Schutte et al., 2003). The enhanced stability of spatial working memory captured by the SPH also explains the reduction in constant and variable errors evident in a number of spatial recall tasks across a broad range of ages from 3 years to adulthood (Hund & Spencer, 2003; Plumert & Hund, 2001; Schutte & Spencer, 2002; Spencer & Hund, 2002, 2003).

Although the SPH captures a number of changes in performance across a range of ages, many of these developmental effects are quite

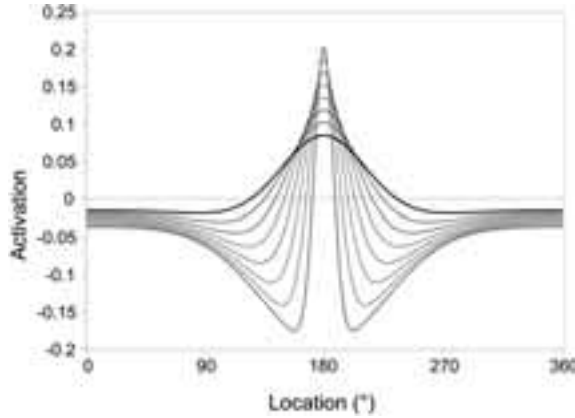


Figure 14.9 Illustration of the SPH. Early in development (black line) interaction is broad and primarily excitatory. Later in development (light gray line) excitation is narrower and stronger with stronger lateral inhibition. Intermediate lines show the hypothesized gradual nature of this change over development.

intuitive—spatial memory gets better over development! Does the SPH capture something less obvious? For instance, one of the more dramatic developmental changes in spatial cognition is the shift in geometric category biases: young children show biases toward symmetry axes, while older children show biases away from such axes (J. Huttenlocher et al., 1994; Schutte & Spencer, 2002; Spencer & Hund, 2003). Can we capture this *qualitative* developmental transition using the same hypothesis?

Figure 14.10 shows that the answer is yes. This figure shows the performance of a “child” (figure 14.10A–C) and “adult” (figure 14.10D–F) version of the DFT for a single trial in a spatial recall task, focusing on three layers of the model— PF_{obj} , $Inhib_{obj}$, and SWM_{obj} . Note that the simulation in figure 14.10D–F is identical to the simulation in figure 14.6A–C. We reproduced this simulation here for comparison. As shown in figure 14.10A–C, the “child” model forms a stable spatial working memory of the target at the start of the trial (T in figure 14.10C). Interestingly, the formation of this peak takes more time than with the adult model. During the delay, this peak drifts toward midline, rather than away from midline (see “drift” toward the back of figure 14.10C). Why does this occur? Because the peak in SWM_{obj} is much broader (reflecting the broad interactions specified by the SPH), this peak produces broad activation in $Inhib_{obj}$, which, combined with the relatively weak self-sustaining dynamics in the perceptual field, prevents a reference peak from building in PF_{obj} (at 180°; see circle in figure 14.10A). Consequently, SWM_{obj} receives some excitatory input around midline (due to the excitatory projection from PF_{obj} to SWM_{obj}), but no associated inhibitory “push” away from a reference peak at midline.

It is important to emphasize that the simulations in figure 14.10 capture a *qualitative* shift in geometric category biases in our task via a *quantitative* shift in model parameters. In particular, to shift from the “child” to “adult” parameters, we simply scaled the inputs and projections among the layers as specified by the SPH. This is consistent with other examples of nonlinear, dynamical systems that capture both quantitative and qualitative developmental change without positing new control processes. Rather, qualitative developmental changes can emerge via a reorganization of the system during the transition (e.g., Thelen & Smith, 1994; van der Maas & Molenaar, 1992; Van Geert, 1998). This stands in contrast to the dominant explanation of the developmental shift in geometric biases. According to the category adjustment model, this shift reflects a change in children’s ability to subdivide space into smaller categories (J. Huttenlocher et al., 1994). Thus, an explicit new ability—spatial subdivision—is used to explain children’s performance. Importantly, however, this model offers no account of how this new ability comes on-line.

Our model takes a critical step closer to explaining the developmental process at work in this case, offering a view of developmental continuity




Figure 14.10 Developmental change in geometric bias in the DFT. When the target (T) is presented at 220° (A), the “child” model (A–C) forms a broad, self-sustaining peak in spatial working memory (C) due to broader and weaker neural interactions early in development. During the delay, the reference input does not create a sustained reference peak in the perceptual field (A). Consequently, the peak in spatial working memory is attracted toward midline (C) due to excitatory input (green arrow) from the perceptual field (A). The right column shows a simulation of the “adult” model drifting away from 180° (F) for comparison (see figure 14.6A–C). Axes and arrows are as in figure 14.4

even as qualitatively new behaviors emerge over time. This continuity has led to a set of novel predictions regarding how the transition in geometric effects unfolds over development. In particular, as excitation narrows and becomes stronger, the “child” model in figure 14.10 shows a narrowing of the attraction effect; that is, responses are biased toward midline only when targets are presented relatively close to this axis. Eventually, neural interactions in the perceptual field become strong enough to create a reference-related peak and repulsion emerges, but only at a few select locations to the left and right of midline. As interactions strengthen further, the repulsion effect increases and extends across a broad spatial range. All of these predictions were recently confirmed with 3- to 5-year-old children (see Schutte, 2004).

These data highlight that the SPH offers a detailed account of how real-time behavior changes over the time scale of development—one of the central goals we highlighted in the introduction to this chapter. Is the SPH also grounded in a rich understanding of brain–behavior relations? Put differently, is there neurophysiological evidence of this type of development in cortex around the ages probed in our studies (e.g., 3–6 years)? It is likely that the development of the prefrontal cortex (PFC) is related to the spatial recall tasks we discussed. Data from several studies demonstrate that the PFC plays a central role in spatial working memory (Awh et al., 1999; di Pellegrino & Wise, 1993; Kessels, Postma, Wijnalda, & de Haan, 2000; Nelson et al., 2000), and maturation of PFC is still occurring between 3 and 6 years of age. In fact, there is evidence that the development of the PFC continues into the postadolescent years (Gogtay et al., 2004; Rakic, 1995; Sowell, Thompson, Tessner, & Toga, 2001). Synaptic density in the PFC reaches its peak at about 2 years of age in humans (P. R. Huttenlocher, 1979) and then declines until it finally reaches adult levels at about 16 years of age (P. R. Huttenlocher, 1990). Myelination is also still occurring in the frontal lobe between the 2 and 6 years of age (Sampaio & Truwit, 2001). Thus, there is dramatic brain development happening during the time period when we see dramatic changes in spatial cognition. At this point, however, there is no strong neurophysiological support for the type of changes captured by the SPH. That said, to the extent that the changes noted above (e.g., myelination and pruning) lead to more efficient neural processing, they are, at least at a qualitative level, consistent with the SPH.

To summarize, the SPH effectively captures a broad array of developmental changes in spatial cognition, including changes in A-not-B-type effects (Schutte et al., 2003), changes in the stability of spatial working memory (Schutte, 2004; Spencer & Schöner, 2003), developmental differences in spatial recall versus spatial recognition performance (Spencer et al., 2006), and the qualitative developmental shift in geometric category biases (Schutte, 2004). Thus, this relatively simple, intuitive hypothesis about how neural interactions change systematically over the longer time scale of development has been quite generative. This is no accident: in our view, the

richly structured real-time and learning-time dynamics of the DFT require less from development because they have much to offer. Importantly, we contend that the relatively modest view of development captured by our theory brings an understanding of developmental mechanism within reach.

14.6 THE DFT, THEORETICAL NEUROSCIENCE, AND THE *HOW* OF DEVELOPMENT

The DFT offers both a rich view of brain–behavior relations and a detailed account of real-time behavior. In addition, this theory specifies a mechanism that leads to the formation of long-term memory across a trial-to-trial time scale. Finally, the SPH offers a relatively simple view of change at an even longer time scale that effectively captures a range of phenomena. This brings us back to the issue of developmental mechanism: does the SPH offer a developmental mechanism that explains key changes in spatial cognitive development? Have we successfully arrived at the *how* of development?

From one perspective, the answer is “no” because we have yet to “close the loop” on development (Simon, 1962). That is, we have yet to formalize how our model can *change itself* over the longer time scale of development to yield the changes specified by the SPH. One candidate process is Hebbian learning. Hebbian learning has been shown to produce the type of continuous, graded, metric changes captured by the SPH (e.g., Kohonen, 1982). Concretely, we could implement a Hebbian learning rule that would gradually modify the strength with which neurons in our model interact. We contend that this type of account is within reach.

Although we have not yet arrived at this developmental goal, we want to stress where we have come to date. The SPH specifies precisely what we need from a developmental mechanism—metric changes in the precision of neural interactions. And we have shown the developmental outcomes of this type of change—both quantitative and qualitative changes in performance in spatial cognitive tasks. Any account of developmental mechanism must achieve this level of precision. As we stated at the outset, this has been a challenge for theories of developmental change: as researchers explore the longer time scales of learning and development, it becomes difficult to offer a detailed, process-based account of real-time behavior.

We contend that our approach that builds from the shortest time scales to the longer time scale of development will ultimately arrive at both a detailed account of behavior-in-the-moment as well as specific developmental mechanisms. In this context, however, we acknowledge the utility of taking different approaches to the mechanism question. For instance, if a Hebbian account can effectively “close the loop” on development in our model, this would be an ideal marriage between our dynamic systems approach and previous work on Hebbian learning using connectionist models.

With this diversity in mind, we conclude by addressing the issue emphasized in our title: what does theoretical neuroscience have to offer the study of behavioral development? Our answer: constraints and clarity.

With regard to the former, the neural principles that underlie the DFT provide strong constraints on the real-time behaviors of our model, and these real-time dynamics constrain what is learned. Reversely, learning dynamics constrain the link between past and present, and the dynamics that live at both time scales provide strong constraints for what might develop. What we have discovered, however, is that these constraints—far from limiting—provide clarity on development: because of the rich dynamics at the shorter time scales, developmental change might be relatively simple.

This level of clarity is certainly tied to the formalization we have achieved within the dynamic field framework. Although formal models are often useful for clarify and constraining concepts, we have a sense that this is particularly important when studying development given the tough challenge of thinking about dynamics over multiple time scales. As an example, when we started exploring the development of geometric biases, we had little sense that a gradual, quantitative change might underlie what has been described in the literature as a dramatic shift in how children “conceptualize” space. But because we started by asking questions about behavior-in-the-moment, we were drawn to the issue of what might be special about symmetry axes. This, in turn, led us to probe the coupling between perception and working memory and the utility of reference frames for spatial cognition. Once we had developed a framework that could integrate spatial memory with perceived reference frames, we explored whether the SPH might account for the developmental shift in geometric effects. To our excitement, the answer was positive: we arrived at a complex pattern of behavioral development using a relatively simple, neurally grounded hypothesis.

Looking to the future, this neural grounding might ultimately provide a bridge between our theoretical work and the emerging field of developmental cognitive neuroscience (Nelson & Luciana, 2001). To date, theoretical neuroscience is an underutilized approach in this field which has been dominated by the use of particular methods—ERPs, fMRI, and so on. Ultimately, one exciting direction would be to marry theoretical work with these new techniques. For instance, there is the potential to test DFT-derived hypotheses using ERPs (McDowell et al., 2002), and we noted ties between the SPH and results from studies using structural MRI to map developmental changes in PFC in the previous section (see, e.g., Gogtay et al., 2004). There are also ways to test predictions of the SPH using fMRI. For instance, we suspect there are key changes in posterior parietal cortex during the transition in geometric bias. Recall that evidence suggests parietal cortex has weakly self-sustaining dynamics (Constantinidis & Steinmetz, 1996; Steinmetz & Constantinidis, 1995) and is involved in reference frame calibration (Pouget et al., 2002). Thus, this cortical area is a good candidate for the functions subserved by our object-centered perceptual field. According to our account of the transition in geometric bias, this field undergoes dramatic changes over development with the emergence of a reference-related peak during the memory delay in recall tasks around

4 years, 4 months. It should be possible to detect such a transition in parietal areas using fMRI (for related ideas, see chapter 13).

In conclusion, the dynamic field theory (DFT) we outline in the present chapter takes an important step toward achieving an elusive goal in developmental science: providing an account of developmental mechanism that retains strong ties to behavior in context. Although we are not yet at the *how* of development, our work on the DFT might specify a path for how to get there.

NOTES

1. This type of spatial category bias is related to work on the formation of object categories. Several researchers have shown that recall of metric object features can be biased based on the distribution of exemplars to which one is exposed (e.g., J. Huttenlocher et al., 2000). Thus, contrary to a recent report (J. Huttenlocher et al., 2004), spatial memory is sensitive to the distribution of exemplar locations in the task space.

2. Earlier versions of this theory are described in Schutte et al. (2003) and Spencer and Schöner (2003).

3. There is some controversy about when it is appropriate to refer to “working” memory versus, for instance, short-term memory (see, e.g., Miyake & Shah, 1999). According to Baddeley (1986), working memory is a specific form of memory that can be used in the service of other, typically complex tasks where memory is just one component. An overlooked aspect of memory necessitated by this definition is stability: the nervous system cannot “use” a memory in the service of some task if the memory is unstable. Since stability is at the core of the “on” state we describe in our theory, we have adopted the “working” memory terminology. Note that although our use of this term has important ties to Baddeley’s original formulation, the DFT is not isomorphic with Baddeley’s theory.

4. We used single-layer perceptual fields in these simulations for simplicity. Such fields show the same dynamic properties as the two-layered dynamic field shown in figure 14.3, but they are computationally easier to simulate. The cost is that our approach to reference frame calibration is less grounded in neural reality. This was warranted by the complexity of the calibration problem and our desire to focus on specific aspects of establishing a stable reference frame.

5. This calibration process can be integrated with other approaches (e.g., efference copy) within the dynamic field framework (for a demonstration of this in the domain of autonomous robotics, see Steinhage & Schöner, 1997, 1998).

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

ACKNOWLEDGMENTS This research was funded by NIMH RO1 MH62480, NSF BCS 00-91757, and NSF HSD 0527698 awarded to J.P.S., as well as support from the Obermann Center for Advanced Studies at the University of Iowa awarded to J.P.S. and G.S. We thank the parents, children, and undergraduates who participated in this research, as well as the many undergraduate research assistants

who helped with data collection. We also thank Jeff Johnson, John Lipinski, Sammy Perone, Larissa Samuelson, and Wendy Troob for their valuable contributions to the manuscript and the development of the dynamic field theory.

REFERENCES

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*, 77–87.
- Amari, S. (1989). Dynamical stability of formation of cortical maps. In M. A. Arbib & S. Amari (Eds.), *Dynamic interactions in neural networks: Models and data* (pp. 15–34). New York: Springer.
- Amari, S., & Arbib, M. A. (1977). Competition and cooperation in neural nets. In J. Metzler (Ed.), *Systems neuroscience* (pp. 119–165). New York: Academic Press.
- Awh, E., Jonides, J., Smith, E. E., Buxton, R. B., Frank, L. R., Love, T., et al. (1999). Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science*, *10*, 433–437.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Oxford University Press.
- Bard, C., Fleury, M., Teasdale, N., Paillard, J., & Nougier, V. (1995). Contribution of proprioception for calibrating and updating the motor space. *Canadian Journal of Physiology & Pharmacology*, *73*, 246–254.
- Barnett, R., Maruff, P., Vance, A., Luk, E. S. L., Costin, J., Wood, C., et al. (2001). Abnormal executive function in attention deficit hyperactivity disorder: The effect of stimulant medication and age on spatial working memory. *Psychological Medicine*, *31*, 1107–1115.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral & Brain Sciences*, *22*, 577–660.
- Bastian, A., Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information pre-shapes the population representation of movement direction in motor cortex. *NeuroReport*, *9*, 315–319.
- Bastian, A., Schöner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, *18*, 2047–2058.
- Beh, H., Wenderoth, P., & Purcell, A. (1971). The angular function of a rod-and-frame illusion. *Perception & Psychophysics*, *9*, 353–355.
- Braun, M. (1994). *Differential equations and their applications* (4th ed.). New York: Springer Verlag.
- Bridgeman, B. (1999). Separate visual representations for perception and for visually guided behavior. In G. Aschersleben, T. Bachmann, & J. Müsseler (Eds.), *Cognitive contributions to the perception of spatial and temporal events. Advances in psychology* (pp. 316–327). Amsterdam: Elsevier Science.
- Bridgeman, B., Gemmer, A., Forsman, T., & Huemer, V. (2000). Processing spatial information in the sensorimotor branch of the visual system. *Vision Research*, *40*, 3539–3552.
- Brungrart, D. S., Rabinowitz, W. M., & Durlach, N. I. (2000). Evaluation of response methods for the localization of nearby objects. *Perception & Psychophysics*, *62*(1), 48–65.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*, 910–923.

- Constantinidis, C., & Steinmetz, M. A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology*, *76*, 1352–1355.
- DeLoache, J. S. (2004). Becoming symbol-minded. *Trends in Cognitive Sciences*, *8*, 66–70.
- Deneve, S., & Pouget, A. (2003). Basis functions for object-centered representations. *Neuron*, *37*(2), 347–359.
- di Pellegrino, G., & Wise, S. P. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *Journal of Neuroscience*, *13*(3), 1227–1243.
- Douglas, R., & Martin, K. (1998). Neocortex. In G. M. Shepherd (Ed.), *The synaptic organization of the brain* (pp. 459–509). New York: Oxford University Press.
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience Supplement*, *3*, 1184–1191.
- Engelbreton, P. H., & Huttenlocher, J. (1996). Bias in spatial location due to categorization: Comment on Tversky and Schiano. *Journal of Experimental Psychology: General*, *125*(1), 96–108.
- Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schönner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, *94*, 53–66.
- Erlhagen, W., & Schönner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545–572.
- Finkel, L. H. (2000). Neuroengineering models of brain disease. *Annual Review of Biomedical Engineering*, *2*, 577–606.
- Fuster, J. M. (1995). Memory in the cerebral cortex: An empirical approach to neural networks in the human and nonhuman primate. Cambridge, MA: MIT Press.
- Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, *8*(8), 2928–2937.
- Georgopoulos, A. P., Taira, M., & Lukashin, A. V. (1993). Cognitive neurophysiology of the motor cortex. *Science*, *260*, 47–52.
- Gilmore, R. O., & Johnson, M. H. (1997). Egocentric action in early infancy: Spatial frames of reference for saccades. *Psychological Science*, *83*(3), 224–230.
- Gogtay, N., Giedd, J. N., Lusk, L., Hayashi, K. M., Greenstein, D., Vaituzis, A. C., et al. (2004). Dynamic mapping of human cortical development during childhood through early adulthood. *Proceedings of the National Academy of Sciences of the United States of America*, *101*, 8174–8179.
- Graziano, M., Hu, X. T., & Gross, C. G. (1997). Coding the locations of objects in the dark. *Science*, *277*, 239–241.
- Haskell, E., Nykamp, D. Q., & Tranchina, D. (2001). Population density methods for large-scale modelling of neuronal networks with realistic synaptic kinetics: Cutting the dimension down to size. *Network: Computation in Neural Systems*, *12*, 141–174.
- Hund, A. M., & Plumert, J. M. (2002). Delay-induced bias in children's memory for location. *Child Development*, *73*, 829–840.
- Hund, A. M., & Plumert, J. M. (2003). Does information about what things are influence children's memory for where things are? *Developmental Psychology*, *39*, 939–948.

- Hund, A. M., & Plumert, J. M. (2005). The stability and flexibility of spatial categories. *Cognitive Psychology*, *50*(1), 1–44.
- Hund, A. M., Plumert, J. M., & Benney, C. J. (2002). Experiencing nearby locations together in time: The role of spatiotemporal contiguity in children's memory for location. *Journal of Experimental Child Psychology*, *82*, 200–225.
- Hund, A. M., & Spencer, J. P. (2003). Developmental changes in the relative weighting of geometric and experience-dependent location cues. *Journal of Cognition & Development*, *4*(1), 3–38.
- Huttenlocher, J., Hedges, L. V., Corrigan, B., & Crawford, L. E. (2004). Spatial categories and the estimation of location. *Cognition*, *93*(2), 75–97.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, *129*, 220–241.
- Huttenlocher, J., Newcombe, N. S., & Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, *27*, 115–147.
- Huttenlocher, P. R. (1979). Synaptic density in human frontal cortex—developmental changes and effects of aging. *Brain Research*, *163*, 195–205.
- Huttenlocher, P. R. (1990). Morphometric study of human cerebral cortex development. *Neuropsychologia*, *28*, 517–527.
- Jancke, D., Erlhagen, W., Dinse, H. R., Akhavan, A. C., Giese, M., Steinhage, A., et al. (1999). Parametric population representation of retinal location: Neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience*, *19*, 9016–9028.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (2006). A dynamic neural field theory of multi-item visual working memory and change detection. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, (pp. 399–404), Mahwah: Lawrence Erlbaum Associates, Inc.
- Karatekin, C., & Asarnow, R. F. (1998). Working memory in childhood-onset schizophrenia and attention-deficit/hyperactivity disorder. *Psychiatry Research*, *80*, 165–176.
- Kessels, R. P. C., Postma, A., Wijnalda, E. M., & de Haan, E. H. F. (2000). Frontal-lobe involvement in spatial memory: Evidence from PET, fMRI, and lesion studies. *Neuropsychology Review*, *10*, 101–113.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*, 59–69.
- Krommenhoek, K. P., van Opstal, A. J., Gielen, C. C., & van Gisbergen, J. A. (1993). Remapping of neural activity in the motor colliculus: A neural network study. *Vision Research*, *33*, 1287–1298.
- Landau, B. (2003). Axes and direction in spatial language and spatial cognition. In E. van der Zee & J. Slack (Eds.), *Representing direction in language and space* (pp. 18–38). Oxford: Oxford University Press.
- Lawrence, A. D., Watkins, L. H. A., Sahakian, B. J., Hodges, J. R., & Robbins, T. W. (2000). Visual object and visuospatial cognition in Huntington's disease: Implications for information processing in corticostriatal circuits. *Brain*, *123*, 1349–1364.
- Li, W., & Westheimer, G. (1997). Human discrimination of the implicit orientation of simple symmetrical patterns. *Vision Research*, *37*, 565–572.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*, 309–332.

- McClelland, J. L., & Jenkins, E. (1991). Nature, nurture, and connections: Implications of connectionist models for cognitive development. In K. VanLehn (Ed.), *Architectures for intelligence: The twenty-second Carnegie Mellon Symposium on Cognition* (pp. 41–73). Hillsdale, NJ: Lawrence Erlbaum.
- McDowell, K., Jeka, J. J., Schöner, G., & Hatfield, B. D. (2002). Behavioral and electrocortical evidence of an interaction between probability and task metrics in movement preparation. *Experimental Brain Research*, *144*, 303–313.
- McIntyre, J., Stratta, F., & Lacquaniti, F. (1998). Short-term memory for reaching to visual targets: Psychophysical evidence for body-centered reference frames. *Journal of Neuroscience*, *18*(20), 8423–8435.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, *16*(16), 5154–5167.
- Miyake, A., & Shah, P. (Eds.). (1999). *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge: Cambridge University Press.
- Munakata, Y., & McClelland, J. L. (2003). Connectionist models of development. *Developmental Science*, *6*, 413–429.
- Nelson, C. A., & Luciana, M. (Eds.). (2001). *Handbook of developmental cognitive neuroscience*. Cambridge, MA: MIT Press.
- Nelson, C. A., Monk, C. S., Lin, J., Carver, L. J., Thomas, K. M., & Truwit, C. L. (2000). Functional neuroanatomy of spatial working memory in children. *Developmental Psychology*, *36*, 109–116.
- Newcombe, N. S., & Huttenlocher, J. (2000). *Making space: The development of spatial representation and reasoning*. Cambridge, MA: MIT Press.
- Plumert, J. M., & Hund, A. M. (2001). The development of location memory: What role do spatial prototypes play? *Child Development*, *72*, 370–384.
- Posner, M. I. (2001). Educating the human brain: A commentary. In J. L. McClelland & R. S. Siegler (Eds.), *Mechanisms of cognitive development: Behavioral and neural perspectives* (pp. 387–400). Mahwah, NJ: Lawrence Erlbaum.
- Pouget, A., Deneve, S., & Duhamel, J. R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience*, *3*(9), 741–747.
- Quinn, P. C. (2000). Perceptual reference points for form and orientation in young infants: Anchors or magnets? *Perception & Psychophysics*, *62*, 1625–1633.
- Rakic, P. (1995). The development of the frontal lobe: A view from the rear of the brain. *Advances in Neurology*, *66*, 1–8.
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science*, *276*, 821–824.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Salinas, E. (2003). Background synaptic activity as a switch between dynamical states in a network. *Neural Computation*, *15*(7), 1439–1475.
- Sampaio, R. C., & Truwit, C. L. (2001). Myelination in the developing human brain. In C. A. Nelson & M. Luciana (Eds.), *Handbook of developmental cognitive neuroscience* (pp. 35–44). Cambridge, MA: MIT Press.
- Sandberg, E. H., Huttenlocher, J., & Newcombe, N. S. (1996). The development of hierarchical representation of two-dimensional space. *Child Development*, *67*, 721–739.

- Sausser, E., & Billard, A. (2005). Three dimensional frames of reference transformations using recurrent populations of neurons. *Neurocomputing*, *64*, 5–24.
- Schiano, D. J., & Tversky, B. (1992). Structure and strategy in encoding simplified graphs. *Memory & Cognition*, *20*(1), 12–20.
- Schöner, G. (in press). Development as change of system dynamics: Stability, instability, and emergence. In J. P. Spencer, M. S. Thomas, & J. L. McClelland (Eds.), *Toward a new grand theory of development? Connectionism and dynamic systems theory re-considered*. New York: Oxford University Press.
- Schutte, A. R. (2004). A developmental transition in spatial working memory. Unpublished doctoral dissertation. Iowa City, IA: University of Iowa.
- Schutte, A. R., & Spencer, J. P. (2002). Generalizing the dynamic field theory of the A-not-B error beyond infancy: Three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, *73*, 377–404.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, *74*(5), 1393–1417.
- Simmering, V. R., & Spencer, J. P. (2006a). Carving up space at imaginary joints: Can people mentally impose spatial category boundaries? Unpublished manuscript.
- Simmering, V. R., & Spencer, J. P. (2006b). Generality with specificity: The dynamic field theory generalizes across tasks and time scales. Unpublished manuscript.
- Simmering, V. R., Spencer, J. P., & Schöner, G. (in press). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception & Psychophysics*.
- Simon, H. A. (1962). An information processing theory of intellectual development. In W. Kessen & C. Kuhlman (Eds.), *Thought in the young child* (pp. 150–161). Yellow Springs, OH: Antioch Press.
- Skarda, C. A., & Freeman, W. J. (1987). How brains make chaos in order to make sense of the world. *Behavioral & Brain Sciences*, *10*(2), 161–195.
- Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review*, *106*, 235–260.
- Smyrnis, N., Taira, M., Ashe, J., & Georgopoulos, A. P. (1992). Motor cortical activity in a memorized delay task. *Experimental Brain Research*, *92*, 139–151.
- Soechting, J. F., & Flanders, M. (1991). Arm movements in three dimensional space: Computation, theory, and observation. *Exercise Sports Science Review*, *19*, 389–418.
- Sowell, E. R., Thompson, P. M., Tessner, K. D., & Toga, A. W. (2001). Mapping continued brain growth and gray matter density reduction in dorsal frontal cortex: Inverse relationships during postadolescent brain maturation. *Journal of Neuroscience*, *21*, 8819–8829.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, *131*, 16–37.
- Spencer, J. P., & Hund, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*, *47*(4), 432–480.
- Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamical systems approach to development. *Developmental Science*, *6*, 392–412.
- Spencer, J. P., & Schutte, A. R. (2004). Unifying representations and responses: Perseverative biases arise from a single behavioral system. *Psychological Science*, *15*(3), 187–193.

- Spencer, J. P., Simmering, V. R., & Schutte, A. R. (2006). Toward a formal theory of flexible spatial behavior: Geometric category biases generalize across pointing and verbal response types. *Journal of Experimental Psychology: Human Perception & Performance*, *32*, 473–490.
- Spencer, J. P., Smith, L. B., & Thelen, E. (2001). 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*, *72*, 1327–1346.
- Steinhage, A., & Schöner, G. (1997). Self-calibration based on invariant view recognition: Dynamic approach to navigation. *Robotics & Autonomous Systems*, *20*, 133–156.
- Steinhage, A., & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. *Proceedings of SPIE*, *3523*, 160–180.
- Steinmetz, M. A., & Constantinidis, C. (1995). Neurophysiological evidence for a role of posterior parietal cortex in redirecting visual attention. *Cerebral Cortex*, *5*, 448–456.
- Stiles, J., Bates, E. A., Thal, D., Trauner, D. A., & Reilly, J. (1998). 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*, *12*, 131–164.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral & Brain Sciences*, *24*, 1–86.
- Thelen, E., & Smith, L. B. (1994). A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.
- Tversky, B., & Schiano, D. J. (1989). Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology: General*, *118*(4), 387–398.
- van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stagemwise cognitive development: An application of catastrophe theory. *Psychological Review*, *99*, 395–417.
- Van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, Vygotsky, and beyond. *Psychological Review*, *105*, 634–677.
- Wenderoth, P., & van der Zwan, R. (1991). Local and global mechanisms of one- and two-dimensional orientation illusions. *Perception & Psychophysics*, *50*, 321–332.
- Werner, S., & Diedrichsen, J. (2002). The time course of spatial memory distortions. *Memory & Cognition*, *30*(5), 718–730.
- Wilimzig, C., & Schöner, G. (2006). A dynamic field theory of categorical and continuous response generation. Unpublished manuscript.
- Wilson, F. A. W., Scalaidhe, S. P., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, *260*, 1955–1958.