

CHAPTER 5

Dynamic Field Theory and Embodied Cognitive Dynamics

John P. Spencer, Sammy Perone, and Jeffrey S. Johnson

The goal of this chapter is to explain some of the core concepts of dynamic field theory (DFT) and how this theory provides a formal framework for thinking about embodied cognitive dynamics.

DFT is now 15 years old. In 1993, Gregor Schöner and his colleagues published a paper (in the proceedings of an international conference) that presented a theory addressing how eye movements are planned; the authors also addressed the neural bases for their theory (Kopecz, Engels, & Schöner, 1993). Since that time, DFT has been extended to a range of topics including the planning of reaching movements (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schöner, & Riehle, 2003; Erlhagen & Schöner, 2002), the development of motor planning (Smith & Thelen, 1993; Thelen & Smith, 1994), the perception of motion (Hock, Schöner, & Giese, 2003; Jancke, Erlhagen, Schöner, & Dinse, 2004), the processes that underlie habituation in infancy (Schöner & Thelen, 2006), the control of autonomous robots (Bicho, Mallet, & Schöner, 2000; Steinhage & Schöner, 1998), the processes that underlie visuospatial cognition and spatial language (Lipinski, Spencer, & Samuelson, in press; Spencer, Simmering, Schutte, & Schöner, 2007), the development of visuospatial cognition (Simmering, Schutte, & Spencer, 2008), the processes that underlie visual working memory and change detection (Johnson, Spencer, & Schöner, 2006), the fast learning of object labels and other aspects of word learning (Faubel & Schöner, 2008; Samuelson, Schutte, & Horst, in press), and

the development of executive function (Buss & Spencer, 2008).

Despite the extensive application of DFT, discussion at the 2005 conference whence this edited volume emerged suggested that aspects of DFT are poorly understood. Pockets of researchers are familiar with one model or another, but there is little consensus regarding the broader theoretical concepts that underlie this theory. One barrier to understanding the underlying concepts is that there have been relatively few reviews of DFT pitched squarely at the central concepts. The present chapter takes up this charge.

We begin with some context, describing the milieu in which the theory emerged and the problems it was initially designed to tackle. In particular, we focus on a central problem in cognitive science: How do cognitive systems interface with real-time sensorimotor systems? That is, how can real-time cognitive dynamics be *embodied*? We discuss our answer to this question within DFT where we use neural dynamics and stable peaks of activation as the basic unit of cognition. Peaks can be continuously coupled to sensorimotor systems in real time. But peaks move far beyond a basic feed-forward picture, allowing for a host of cognitive properties that come for free in this theoretical framework. We highlight these properties and also show how peaks provide critical links to two other central challenges in cognitive developmental sciences: the integration of processes over multiple timescales and the origins of behavioral flexibility.

TOWARD EMBODIED COGNITION: THE DYNAMIC APPROACH TO MOTOR CONTROL AND DEVELOPMENT

To understand some of the central innovations of DFT, it is useful to describe the context in which this approach emerged (for reviews, see Johnson, Spencer, & Schöner, 2008; Spencer & Schöner, 2003). In the 1980s and 1990s, several researchers were interested in applying the concepts of *dynamical systems theory*¹ (Braun, 1994) to aspects of human behavior (Port & van Gelder, 1995). Much of this work was initially conducted to investigate motor control (Kelso & Tuller, 1984) and motor development (Thelen, 1995), in part because the continuous, real-time nature of movement mapped nicely onto continuously evolving dynamical systems.

As researchers probed the link between dynamical systems theory and motor behavior more deeply, a set of central concepts emerged (see Spencer & Schöner, 2003). First, action reflects a dynamic balance among stability (stably reaching for a cup of coffee despite a bump on the arm), instability (grabbing the cup a little too vigorously and spilling several drops), and flexibility (managing to drop the cup in order to catch the computer about to fall onto the concrete sidewalk!). Second, behavior is *softly assembled*² from multiple component processes—neural plans to move the hand to a particular location in space, oscillatory processes that regulate the timing of movement, forces generated by the contraction of muscles to move the hand through space, and so on (Schöner, 1995). Third, such processes are often nonlinear in nature, where small changes in one component (movement speed at the end of a reach) can lead to big changes in behavior (knocking the cup over). Finally, the dynamics of action change over multiple timescales, from the millisecond-to-millisecond unfolding of a single reach, to a sequence of actions in context—reach to a coffee cup, then to a

computer—to the myriad reaches to objects that mount up over weeks, months, and years.

These four concepts have been extended to a variety of phenomena probing processes that underlie complex action systems and a host of topics in motor development (e.g., Corbetta & Thelen, 1996; Kelso, 1995; Schöner, 1995; Thelen & Ulrich, 1991). But even early in the development of these concepts a central issue arose: What about cognition?

Early explorations of this question led to a strong, antirepresentationalist stance and to a central challenge—to rethink cognition without using static/rigid symbols (Townsend & Busemeyer, 1995; van Gelder et al., 1998). In this sense, the dynamical systems community had much in common with connectionist ideas taking hold at the time (Rumelhart, McClelland, & PDR Group, 1986). In addition, there was a move to think of cognition as *embodied*, that is, always intimately tied to sensorimotor systems. This line of thinking resonated within the developmental science research community apropos fundamental questions about the origin of cognitive systems (Smith & Thelen, 1993; Thelen & Smith, 1994). Recent work has demonstrated that embodiment plays a central role in adult cognitive activity, as well, including the neural bases of thought (Barsalou & Wiemer-Hastings, 2005; Damasio & Damasio, 1994; Spivey, 2007).

Although some researchers have made strong arguments in defense of embodiment (Glenberg, 1997) and dynamical systems approaches to cognition, there was, at the time, no formal theory for how such interactions could arise. The reason: There are some fundamentally challenging issues that lie at the heart of this appealing idea that perception, action, and cognition are all intimately intertwined.

The initial challenge is that sensorimotor systems evolve continuously in real time, but cognition can *jump* from one state to another, that is, from one thought to another. Consider both sides of this challenge. First, in what sense

¹ Note that *dynamic* systems theory and *dynamical* systems theory are used interchangeably in this volume.

² Behavior is characterized as softly assembled (or soft assembled) if the behavior is assembled to suit a current and specific task/environment. This term is used to distinguish such behavior from that existing prior to the current environment and independent of the specific task.

can't a sensorimotor system jump? When I reach my hand out to pick up a coffee cup, my hand obligatorily must move through the space that intervenes between the computer keyboard and the cup. Although I might move my hand quite quickly, there is no getting around this physical fact. By contrast, what would it mean for a cognitive system to jump? Motor plans in my brain can actually jump from a plan to, say, point to a key on my keyboard at 0° movement direction to a plan to reach 120° to the right to retrieve a cup. Although motor plans can evolve continuously (as is the case in some forms of mental rotation, see Georgopoulos, Lurito, Petrides, Schwartz, & Massey, 1989), they don't have to. Note that claims about *jumpy cognition* do not imply that motor plans and other forms of cognitive activity are somehow removed from continuous, real-time neural processes. The point here is that what these neural processes *represent* in the moment can jump in content.

This creates some serious issues. If cognition is jumping around, but heads, hands, fingers, and legs all move continuously through space and time, how does cognition remain connected? How does cognition remain embodied?

The second challenge is a challenge of *representation*. Clearly there must be something *stored* over time, some form of memory. But what is the nature of that memory? Work within the framework of information processing and artificial intelligence led investigators to hypothesize that cognition happens upon a bed of static, symbolic representations (Newell & Simon, 1976). If we are going to jettison static symbols, as early proponents of dynamical systems theory argued we should (van Gelder et al., 1998), what are we going to replace them with? What types of representational format can effectively interface with real-time sensorimotor systems, yet also do the job of cognition?

Again, there are layers of subtlety here. Consider one example from the domain of spatial cognition. According to the dominant theory of spatial memory, people remember locations over the long-term relative to spatial *prototypes* (Huttenlocher, Hedges, Corrigan, & Crawford, 2004; Huttenlocher, Hedges, & Duncan, 1991). For instance, if I want to remember where my

computer screen is located, I might remember that it is on the right half of my office desk, roughly near the center of the right category, that is, near the right spatial prototype. But what happens when I'm sitting at my desk, and I turn my head to gaze longingly out the window? To keep track of my computer screen, I need to know where the desk is located relative to my current position in the world, and I need to know my orientation in space (so I can reconstruct what *right* means). Things become even more complicated when I get up and leave my office, only to return later. In this case, I need to make sure I enter the correct office, I need to reorient myself in this space, and then I can discover (much to my annoyance) that my computer screen was stolen the night before.

What's the point? The point is that forming an association between the right spatial prototype on my office desk and my computer screen *is only useful if that abstract representation is linkable in context to my real-time sensorimotor activity*. And it's not entirely clear how to make this link robust; particularly because sensorimotor systems rely on continuous, metric representations of space, but prototypes strip away much of this detail (Bridgeman, Gemmer, Forsman, & Huemer, 2000; Jackendoff, 1996). What system decides which details to throw away and which details to keep? And once those details are thrown away, how do we reconnect the abstract representation to the continuous, metric information provided by our senses? Of course, the alternative—storing all the details of daily experience—seems cumbersome, as well. Thus, what is the appropriate middle ground that effectively integrates sensorimotor and cognitive systems yet provides a useful level of abstraction for *higher cognition* (for a discussion of related issues, see Mareschal et al., 2007)?

SETTING THE STAGE: GROUNDING COGNITION IN NEURAL PROCESS

In the early 1990s, Schöner, Thelen, Smith, and colleagues began thinking about how the dynamical systems approach to motor control and development could be extended to capture the dynamics of cognitive systems. The starting

point for this approach was to consider several facts about neural systems. Neural systems are noisy, densely interconnected, and time dependent; they pass continuous, graded, and metric information to one another; and they are continuously coupled via both short-range and long-range connections (Braitenberg & Schüz, 1991; Constantinidis & Steinmetz, 1996; Edelman, 1987; Rao, Rainer, & Miller, 1997). These neural facts raise deep theoretical challenges (for background, see Skarda & Freeman, 1987). How can a collection of neurons *encapsulate* some bit of information (e.g., a static symbol) amidst near-constant bombardment by other neural signals? And how can a collection of neurons integrate information from multiple systems in a stable, reliable way?

A number of neurophysiologists and theoretical neuroscientists have argued that the densely interconnected and dynamic nature of the brain provides poor mapping to the concepts of information processing. Thus, rather than imposing classic views of cognition onto the brain, the challenge is to take inspiration from what we know about neural reality to rethink cognition (Braitenberg & Schüz, 1991; Brooks, 1991). This perspective has taken hold within the field of theoretical neuroscience where theoreticians have embraced the use of complex, dynamic neural networks to capture brain-behavior relations. Although neural networks have architectures that can be depicted as separate systems, they are—at their core—complex, reentrant, densely interconnected, complex systems that violate the separability assumptions of information processing (Palmer & Kimchi, 1986).

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. The goal of such models is two-fold: (1) to model the behavior of different cell types

within a cortical region, and (2) to capture aspects of behavioral performance via population dynamics.

Although the biophysical approach has led to new insights into brain function and neuronal dynamics, there's a problem: such models do a great job capturing the behavior of neurons, but they do *not* do a great job capturing the behavior of people. This is due, in part, to concrete limitations of such models: they can be very complex, they can have immense computational demands, and they can be difficult to analyze (Haskell, Nykamp, & Tranchina, 2001). Moreover, it is hard to get biophysical networks to behave in realistic ways. Such networks can generate behaviors (e.g., reaction times); however, there have been few attempts to create biophysical networks that are able to generate multiple types of behavior and to flexibly switch behaviors as task demands vary. Thus, while this approach to rethinking cognition is very promising, biophysical network models do well at the level of neurons, but poorly at the level of behavior (Finkel, 2000).

The dynamic neural field framework we have developed shares the same goal of many neural network approaches—to think about cognition as the emergent product of complex, contextually specific neural dynamics—but accomplishes this in a way that is more tightly linked to behavior. To achieve a stronger link to behavior, we focused our analysis at the level of population dynamics (see also Spivey, 2007). That is, rather than building networks that start from a set of, say, integrate-and-fire neurons that retain a great deal of biophysical detail, we have chosen to focus on the emergent product of the dynamics at the neural level—attractors at the level of the neural population.

To situate this work, consider classic research by Georgopoulos and colleagues (Georgopoulos, Kettner, & Schwartz, 1988; Georgopoulos, Taira, & Lukashin, 1993) that established how one could directly link population-based neural activity with behavior. These researchers examined the neural signals in motor cortex that specified the direction of an up-coming arm movement in a simple pointing task. Nonhuman primates

were trained to point toward a lighted target from a central starting location at a *go* signal. Simultaneous neural recordings revealed that each neuron fired maximally when the monkey moved in a *preferred* direction, and this firing profile gradually tapered off as the arm moved in less preferred directions. To construct a population response, Georgopoulos and colleagues transformed the real-time pattern of neural activity from each neuron into a vector. The direction of each neuron's vector pointed in its *preferred* direction, and the magnitude of the vector was scaled by the strength of the neuron's current firing rate. These individual vectors could then be summed to produce a population vector. Results showed that this neural population vector accurately predicted the direction of an up-coming movement before (160 ms before) movement onset. More recent work has verified the predictive nature of the population vector. In a particularly vivid case, population vectors were shown to rotate in real time as nonhuman primates performed a mental rotation task, and the rotation time correlated with the response time (Georgopoulos et al., 1989).

Work by Georgopoulos and others (see, for example, Dayan & Abbott, 2001) suggests that stable patterns of activation at the neural population level directly give rise to behavior. Since this discovery in the 1980s, work has moved forward on two fronts. One line of investigation has tried to understand the link between population dynamics and single cell neurophysiology (Amit & Brunel, 1997; Compte, Brunel, Goldman-Rakic, & Wang, 2000; Durstewitz et al., 2000). The question here is how the activation of individual neurons gives rise to stable activation at the population level. A second line of investigation has tried to understand how population dynamics give rise to behavior. DFT emerged out of this second line of research. In the section that follows, we describe the central concepts that underlie DFT, including how this framework is linked to neurophysiology.

THE REAL-TIME INTEGRATION OF COGNITIVE AND SENSORIMOTOR SYSTEMS: DFT CORE CONCEPTS

Figure 5.1 shows a particular implementation of DFT in the form of a one-dimensional, two-layer

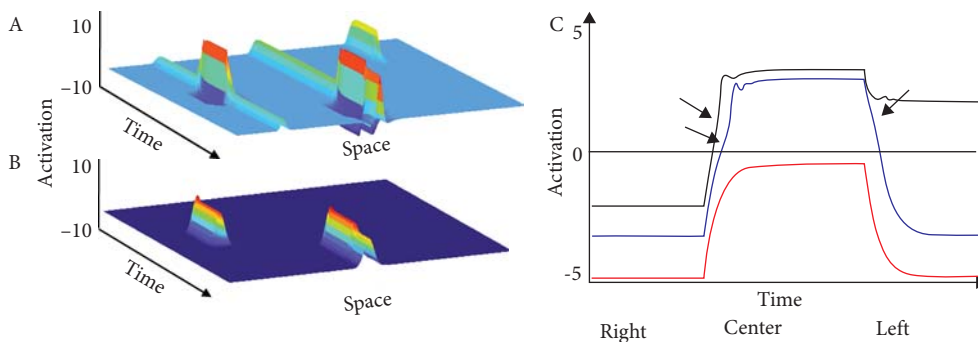


Figure 5.1. Simulation illustrating a one-dimensional, two-layer dynamic neural field

This implementation of the basic model consists of a field of excitatory neurons (panel A) reciprocally coupled to a field of inhibitory neurons (panel B). Neurons in the top layer (panel A) interact via *local excitatory connections* and drive activation in the inhibitory layer (panel B), which, in turn, projects broad inhibition (*longer range inhibition*) back to the excitatory layer. Panel C shows the activation profile through time at the spatial location associated with each of the three peaks in panel A: red trace = right hump (apex is yellow and may be difficult to see); black trace = central peak; and blue trace = left peak. Note that the activation profile for the three peaks are superimposed in time to highlight differences in activation strength. Arrows highlight joints in the activation curve at which the field is moving through a bifurcation from one attractor state to another.

dynamic neural field (DNF). DNFs were first developed by Amari (1977) and are part of a class of bistable attractor neural networks (Camperi & Wang, 1998; Wilson & Cowan, 1972). DNFs use the well-known construct of *activation* from mathematical psychology (Churchland & Sejnowski, 1992), which captures the certainty with which some bit of information is currently known to the nervous system—the greater the activation, the greater the certainty. Activation in DNFs is distributed over continuous dimensions (i.e., *fields*). These dimensions can be spatial locations, movement direction, continuous color dimension, and so on.

A central issue for a dynamical systems approach to neural networks is to understand how patterns of activation change through time and the stability characteristics of these patterns. Activation in DNFs evolves continuously through time as a function of the current inputs impinging on the network, as well as the network's own *intrinsic* dynamic neural interactions. DNFs use a particular form of neural interaction based on the space code principle from neurophysiology. According to this principle, neurons that fire maximally to similar stimuli excite one another. These neurons are considered *close* along the represented dimension (e.g., close along a spatial dimension). By contrast, neurons that fire maximally to very different stimuli inhibit one another. These neurons are considered *far* along the represented dimension. This form of neural interaction gives rise to *local excitatory interactions* and *lateral or longer range inhibitory interactions*.

Note that in Amari's (Amari, 1977) initial work, neural fields reflect activation on the cortical surface because neurons in, for instance, the early visual cortex are lined up in this way (i.e., neurons close together on the cortical surface *code* for similar spatial locations). But other work by Georgopoulos and colleagues (e.g., Georgopoulos & Massey, 1988) revealed that one could construct a type of *functional*

topography mapping. Specifically, neurons in the motor cortex are not arranged topographically: two neurons next to one another in the motor cortex might prefer very different movement directions. Nevertheless, neural activation at the population level actually obeys topographic relations. That is, even though two neurons might be far apart on the cortical surface, they can still prefer the same movement direction and be mutually excitatory. Similarly, these same neurons will inhibit other neurons in the population that prefer very different movement directions. Critically, such interactions give rise to stable *peaks* of activation, even though the peaks are effectively distributed across the cortical surface. One can demonstrate this directly by lining up neurons in a field, not by their positions in the cortex, but rather, by their preferred movement directions (see Bastian et al., 1998; Bastian et al., 2003).

In Figure 5.1, the excitatory interactions all occur among neighboring neurons within the top layer of excitatory neurons (Fig. 5.1A). If these neurons are stimulated such that a local *hump*³ of activation forms and goes above-threshold (i.e., above 0), the excitatory neurons begin to drive up the activation of neurons in the second layer of inhibitory interneurons (Fig. 5.1B). These inhibitory neurons, in turn, project broad inhibitory feedback to the excitatory layer. As excitation and inhibition cycle back and forth in the network, a stable *peak* of activation can form. The basic unit of cognition in DFT is a stable peak of activation.

At this point, it is important to note that only positive activation—*above-threshold activation*—contributes significantly to neural interactions within and between fields. That is, only field sites with activation greater than zero (0) transmit their activation to other neurons and thus contribute to the evolving patterns of activation across the network. This premise is captured by the sigmoidal nonlinearity characteristic of neuronal dynamics (Grossberg, 1973); the same premise holds for all simulations

³ We use the term *hump* here in contrast to the term *peak*. Humps refer to small, sub-threshold patterns of excitation that don't by themselves engage strong interaction. Peaks, by contrast, are supra-threshold activations that do engage strong neural interaction.

described herein. The presence of sub-threshold (i.e., below 0) patterns of activation can still exert a pronounced effect on activation patterns arising within the network, as discussed below. Note also, that all simulations begin with the excitatory and inhibitory neurons in a *resting state*, that is, stably sitting at a negative resting potential (i.e., a negative activation value).

One of the central contributions of the original work by Amari (1977)—and one of the primary reasons we build from this starting point—is that he analyzed the qualitatively different types of behaviors (or attractor states) this class of neural networks can exhibit. Amari identified five attractor states, three of which have been central to DFT: the resting state, the self-stabilized state where peaks of activation are stable but require input for their continued existence, and the self-sustaining state in which peaks of activation can remain stably *on* even in the absence of input.⁴ These three states have a particular relationship to one another because DNFs are bistable. Under one parameter setting, the resting and self-stabilized states coexist; under other settings, the resting and self-sustaining states coexist.

The three attractor states are depicted in Fig. 5.1A. The first hump of activation—the one to the far right—depicts the resting state. The increase in activation around the right location is driven by the onset of a stimulus at the right location. When this input is presented, activation in the excitatory layer builds and comes close to the zero (0) *activation threshold* (see red activation trace in Fig. 5.1C). However, the stimulus strength is not sufficient to create a peak and the field stays in the resting state. This can be observed directly by looking at activation in the inhibitory layer during this period of stimulation; the inhibitory layer is never engaged during the presentation of this first stimulus (see Fig. 5.1B). Consequently, once the input is turned *off*, the excitatory layer near the

right location relaxes back to the negative resting level.

The second attractor state is depicted to the left (Fig. 5.1A). The same stimulus input was used; now, however, a self-stabilized peak builds. This occurs because the field has a small hump of activation present from the start of the trial. This small hump—which might reflect a contribution from long-term memory (see below)—elevated the resting level of the neurons near the left location (see blue line in Fig. 5.1C). Consequently, the target input moved the system through a dynamic instability and into a new attractor state. (Evidence of the instability can be seen in the small *joints* in the activation function indicated by the arrows in Fig. 5.1C.) As can be seen, the neural population remains in an *on* (i.e., above 0 activation) state while the stimulus input is present. Once the input is turned off, however, activation around the left location goes through another instability (see arrow in Fig. 5.1C) as the system moves from the self-stabilized state back to the resting state (for further discussion, see Spencer & Perone, 2008).

The final attractor state is shown around the central location in Figure 5.1A—the self-sustaining state. As can be seen in the figure, the central location in the field has a stronger hump of activation present at the start of the trial (see Fig. 5.1C). As with the peak around the left location, the peak at the center location grows with the onset of the stimulus. Once the input is turned off, however, the peak remains in the *on* state (note the above 0 level of the black line at the end of the simulation in Fig. 5.1C). In this case, neural interactions are strong enough to keep a memory of the stimulus around, even when the stimulus is no longer present. Note that we did not do anything special to the field to move it into this state. Rather, the stronger input was sufficient to create this shift from the self-stabilized attractor around the left location

⁴ One of the attractor states Amari identified was a state in which the entire field becomes excited. Although this provides a potential mapping to cases of seizure, we avoid this state for obvious reasons. A fifth attractor state is a *spatially periodic* state, where the period varies with the average stimulation to the field. We suspect this state is involved in extracting the axes of symmetry we have probed in our work on the development of spatial recall (see, e.g., Schutte et al., 2003). This state might also have ties to the pattern of *grid cells* observed in the entorhinal and perirhinal cortex (Hafting, Fyhn, Molden, Moser, & Moser, 2005).

to a self-sustaining attractor state around the center location.

In the sections that follow, we provide a more detailed survey of the characteristics of self-stabilized and self-sustaining peaks. This survey highlights the many characteristics of peaks which *come for free* from the neural population dynamics within DNFs. It also provides grist for our claim that peaks provide an appropriate unit of cognition.

We will introduce these characteristics in the context of simple examples from spatial cognition, for instance, actively tracking the location of a moving object on a table or remembering the object's location during short-term delays (e.g., 10 s). We focus on these examples because they are conceptually simple and well-grounded in empirical observations. That said, the properties of fields that we illustrate could operate over any continuous, metric dimension such as object color or orientation (for examples, see Johnson, Spencer, Luck, & Schöner, in press; Johnson et al., 2008). Moreover, these properties can be extended to multidimensional fields, as well as to the coupling or *binding* of multiple cortical fields together (an issue we return to in the conclusion section; see also, chapter 4).

Another reason we focus on examples from spatial cognition is that they provide a link to formal demonstrations that dynamic field concepts are *embodied*. In particular, the characteristics we discuss below have proved useful in autonomous robotics scenarios in which robots make real-time, cognitive decisions (i.e., form peaks within DNFs); robots are able to sense features of the surrounds through sensory channels (such as cameras or infrared sensors that provide input to DNFs) and control their own navigation in real time via motors coupled to DNFs (e.g., Bicho et al., 2000; Faubel & Schöner, 2008; Iossifidis & Schöner, 2006; Steinhage & Schöner, 1997, 1998). We highlight a few specific robotics examples below. More generally, it is important to keep in mind that even though we use simulated fluctuations in input, the field concepts we highlight have proved critical to realizing autonomous and embodied links with perception, cognition, and action.

Finally, we want to emphasize that the simulations below use the same neural architecture *except* in a few cases where noted. Thus, even though we highlight different attractor states of DNFs, the only difference is typically a small change in the parameters of the model (e.g., weaker or stronger excitatory interactions among neurons). We begin with a series of simulations with short durations (1000 ms) to highlight the characteristics and detailed dynamics of each attractor state. Then we move to longer simulations with the types of variations one might experience during an experimental trial (e.g., across a 10-s memory delay). Later, we add long-term memory dynamics and look at the construction of memories on a trial-to-trial timescale (i.e., across thirty 10-s trials). Note that we have fixed these different timescales by mapping the features of DNFs onto empirical observations (e.g., Schutte & Spencer, 2008, submitted for publication).

Self-Stabilized Peaks as a Basic Form of Perceptual Encoding

Self-stabilized peaks are stabilized in the sense that they are robust to noise—they don't go away even though activation fluctuates. Critically, however, they only stay around provided input is present. As Figure 5.1 shows, once the input is removed, the system relaxes back to the resting state.

Given that peaks in the self-stabilized state are dependent on input for their survival, it is tempting to think of them as *input-driven* or as a pattern of activation within a simple feed-forward network. Although these peaks are input-driven, they have some unique features that move them beyond the feed-forward picture.

The first property of self-stabilized peaks is that they are robust to variations in input—they are stable! In Figure 5.2, we use the same two-layer architecture previously shown in Figure 5.1; however, for simplicity, only the excitatory layer is shown (Fig. 5.2B). In this example, we present a single stimulus input at a central location (Fig. 5.2A); we used two types of fluctuations in the input. First, the amplitude

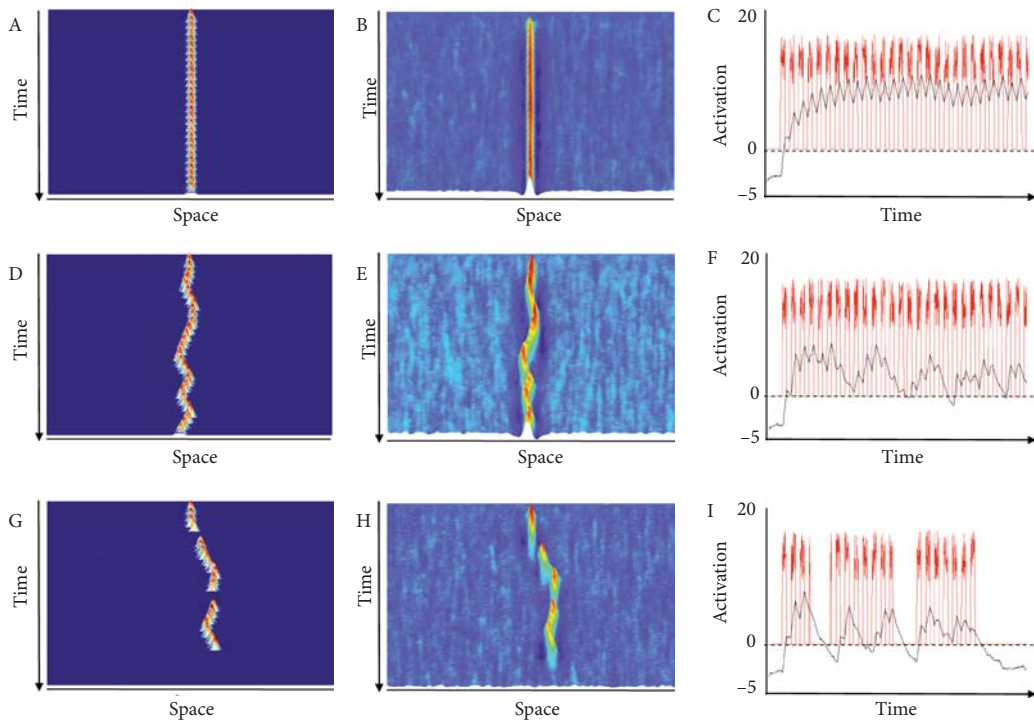


Figure 5.2. Simulations of dynamic neural fields: Self-stabilized peaks

This figure depicts three simulations of activation (panels B, E, and H) in a field with the two-layer architecture described for Figure 5.1. For each simulation, there is a single input (panels A, D, and G) that fluctuates in amplitude and/or spatial location over time. **(Top row of panels)** Despite the fluctuations in input (panel A; colors indicate fluctuating/pulsating input), the field activation (panel B) remains stable in the *on* state throughout the trial; the black activation trace in panel C depicts the field activation through time at the spatial location of the peak in panel B. By contrast, the red line in panel C depicts the fluctuating/pulsating strength of the input through time at the spatial location of the input in panel A. **(Middle row of panels)** Even with strong amplitude fluctuations and spatial variation in the input (panel D and red trace in panel F), the activation peak remains *on*, tracking the input with smooth motion (panel E). **(Bottom row of panels)** When longer interruptions in the input are introduced (panel G and red trace in panel I), the activation peak dies out (panel H and black trace in panel I); but the field quickly reacquires the stimulus when it reappears after each occlusion. The time course of each simulation was 1 s (1000 ms).

of the input varies randomly; note the fluctuations in the red trace in Figure 5.2C. Second, the input is turned off for brief durations as occurs regularly in subjects during eye blinks. As can be seen in the figure, the peak stays in the *on* state (i.e., above 0 activation) throughout the trial (Fig. 5.2B), although the amplitude of the peak fluctuates with the input strength and as a function of the brief occlusions (see black trace in Fig. 5.2C).

This story gets more interesting if we look at what happens when we vary the input in both amplitude and spatial location over time. Recall that one of the goals when we started to use DNFs was to understand the link between perception/action and cognition—to keep cognition close to the sensory surface. Given that DNFs evolve continuously in time and across the continuous dimensions of space, color, and so on, they can provide a real-time

connection to the outside world. The simulation in Figures 5.2D–F, begins with the presentation of a stimulus which builds a self-stabilized peak. Over the course of the simulation, however, the input takes a random walk (panel D) and shows dramatic amplitude fluctuations in real time (red trace in panel F). Despite great variation, the activation peak tracks the input in space and time (panel E).

Interestingly, we can make the situation even more complex by adding another common aspect of real-world neural signals—occlusions of longer durations. This occurs routinely when you move your eyes to scan a visual scene, when you turn your head, when you move your torso, and when you get up to walk around. What happens to peaks when these types of interruptions occur?

Because peaks in the self-stabilized state are *input-driven*, one might think longer interruptions would wreak havoc. The simulation in Figures 5.2G–I, shows that this is not the case—provided the interruption isn't too long. As can be seen in the figure, the self-stabilized peak tracks the input just as before, and it dies out during each long occlusion (see black trace in panel I). But the field is able to quickly reacquire the stimulus when it reappears.

Interestingly, we can make the input pattern even more complex to highlight yet another feature of self-stabilized peaks. Although the input pattern used in Figure 5.2 was noisy in multiple ways, it was still a far cry from the complexity of even simple real-world inputs. What happens to our peak picture when we add complexity in the form of *multiple* inputs appearing and disappearing over a longer timescale (such as a 10-s trial)?

In Figures 5.3A and B, we introduced new random, transient inputs at different distances from the central *target* input. (For simplicity, we focus on the dynamics within the fields shown in Figure 5.3 and do not show the input pattern directly.) The self-stabilized peak starts out tracking the input just fine. However, a few milliseconds into the simulation, a second *distractor* input appears. Although this second input is just as strong, the peak actively *suppresses* the distractor. This occurs via the lateral inhibitory interactions generated by the peak. About

a third of the way through the simulation, there is a long occlusion. During the occlusion, a new distractor appears to the left and the field builds a peak at this location. Soon thereafter, the central target reappears along with a second distractor to the right. It takes several hundred milliseconds for the central peak to reach full strength again due to the active competition with the distractor items (see the dip in activation in Fig. 5.3B). Over time, however, the field reacquires the central target and tracks it for the remainder of the trial, suppressing several new distractors toward the end of the simulation. In summary, then, *peaks provide stability against transient local distractions*.

In the simulation depicted in Figure 5.3A, new distracting inputs were generally *on* for only a short period. What happens when distractors are far from the central target peak, sufficiently strong, and remain *on* for a significant duration—that is, when these distractions are discriminably different, sustained, and reliable patterns of input? Figure 5.3C depicts that a new type of behavior emerges: the field enters a multi-peak state and effectively tracks all three input patterns. The possibility for tracking multiple inputs opens up new types of functionality. This makes sense conceptually since the distractions depicted in Figure 5.3C can be considered interesting new inputs to *encode*.

At this point, it should be clear that peaks have some very interesting features that make them more than merely input-driven patterns of activation. They have important stability properties, tolerating a lot of noise in the input including fluctuations in amplitude, fluctuations along the represented dimension, fluctuations in the presence/absence of the signal, and the introduction of new transient inputs. Moreover, DNFs can operate in a mode in which they track multiple peaks simultaneously, provided that the input signals supporting these peaks are sufficiently robust (e.g., are strong enough, *on* for a sufficient duration, and so forth). Critically, the *decision* about which signals are robust and which are transient is emergent from the real-time dynamics of the system (e.g., excitatory and inhibitory interactions), the characteristics of the input, and the

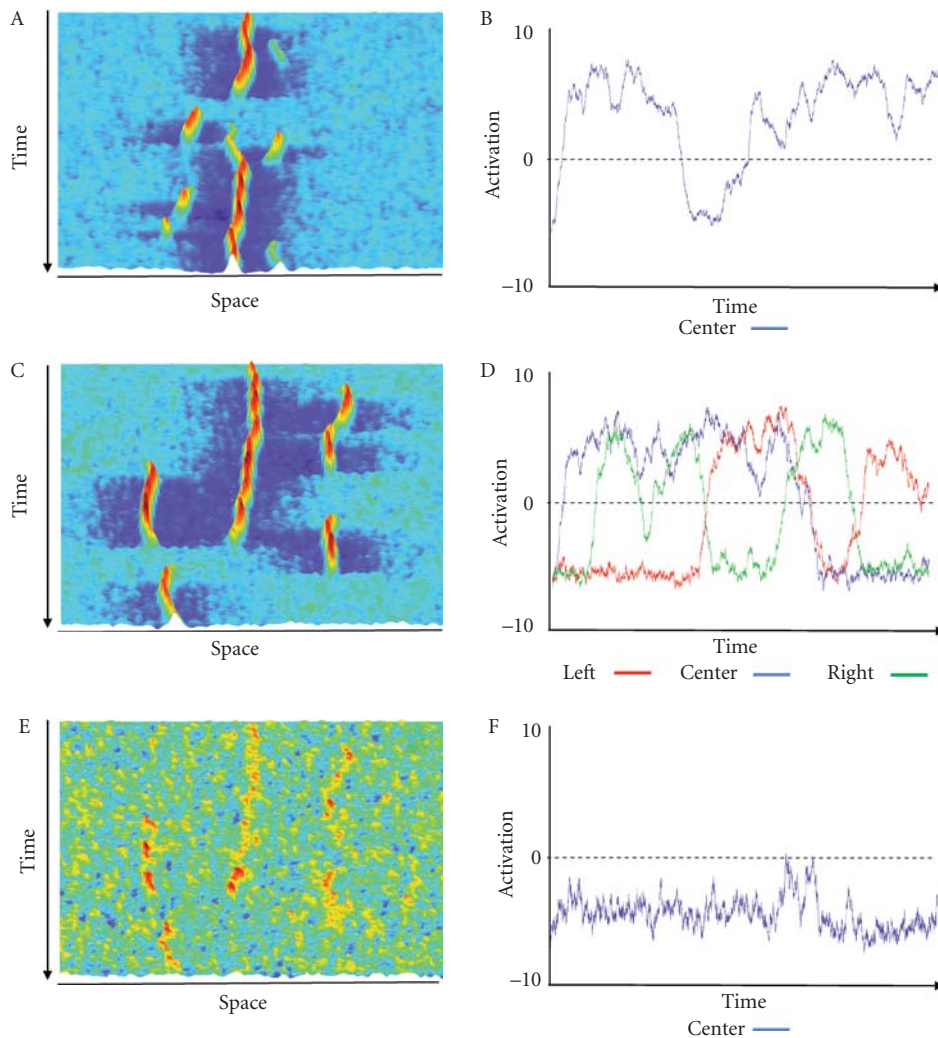


Figure 5.3. Simulations of noisy dynamic neural fields: Self-stabilized peaks

This figure depicts three simulations of activation (panels A, C, and E) in a field with the two-layer architecture as described for Figure 5.1 and depicted in Figure 5.2. The field activations depicted arise in response to target inputs (center peak) in the presence of one or more *distractor* inputs. Despite the presence of nearby distractors, the center activation peak in panel A remains above threshold (see red activation values in panel A) as long as the target input remains *on*. When the target input is *off* and a distractor appears, the center peak briefly goes below threshold (see blue trace in panel B) before reacquiring the target. When distractor inputs are farther away from the target (see left and right activation peaks in panel C), multiple self-stabilized peaks are formed in the field despite substantial variations in amplitude (see traces in panel D) and in spatial location. When neural interactions in the field are too weak to enter the self-stabilized state (panel E), the field functions as a low-pass filter on the same inputs presented in panel C. In this case, activation at the center location never consistently passes above threshold (see panel F). The time course of each simulation was 10 s.

just-previous history of activation in the field. These properties of self-stabilized peaks apparently come for free. All of these properties allow peaks to actively maintain a stable, real-time link to the outside world (or to a complex inside world as in the case of multiple, coupled cortical fields).

We think the characteristics of peaks discussed thus far—the characteristics of stable patterns of activation at the level of a neural population—are critical for the nervous system. Imagine you are a neural field sitting in the posterior parietal cortex. You have just received input from early visual areas that a spot of light has been detected at a particular location on the retina and you have dutifully built an associated pattern of activation to represent this event. At the same time, you begin to receive input from inferior temporal areas that the feature values of a face might be visible. Frontal areas are telling you to keep focused on that computer screen and ignore the people walking around in your periphery at the coffee shop. You are also receiving input from somatosensory areas linked to the stomach saying that the vanilla latte you are drinking tastes really yummy, while your motor cortex wants to tell you about the lovely oscillation of your foot it has started up in response to the influx of caffeine. In short, your little cortical field is being bombarded with noisy inputs—some are relevant, some are irrelevant—and without the stability provided by peaks, you would be in sorry shape.

How sorry? Figure 5.3E depicts a field with the same input pattern used in Figure 5.3C, but with neural interactions too weak to enter the self-stabilized state. As can be seen in this figure, there is a hint of three reliable inputs, but only a hint. Fortunately, the neural dynamics in your parietal cortex operate more like the pattern in Figure 5.3C than in Figure 5.3E... which allows you to stay focused and keep working...

The simulations thus far have highlighted that self-stabilized peaks have some adventurous functionalities: creating stable perceptual decisions that an input is sufficiently strong and reliable to build a peak and subsequently tracking the details of the stimulus in real time. But

there are also more cognitive aspects of the self-stabilized state worth highlighting.

Figure 5.4 depicts the neural field used in all simulations discussed thus far; it can track up to six inputs simultaneously (panel A). Once the number of inputs reaches seven, however, the field *selects* a subset of five inputs to track (panel C). Thus, fields operating in the self-stabilized state have a maximum *capacity*: only so many peaks can be simultaneously built (for discussion, see Johnson et al., 2008). What causes this capacity limit? In panel B, the red line depicts the summed activation in the *excitatory* layer shown in panel A (with 6 inputs); the blue line in panel B depicts the summed activation from panel C (with 7 inputs). Panel D depicts the corresponding summed activation in the *inhibitory* layer from each simulation. As can be seen in these panels, excitation (panel B) and inhibition (panel D) quickly reach their maximal levels when six inputs are presented (red lines). When seven inputs are presented (blue lines), excitation and inhibition increase as before, but the presence of an extra input tries to push excitation too high (the total excitatory strength cannot exceed the total inhibitory strength or activation in the network will grow without bounds). As a consequence, the excitatory and inhibitory strengths settle into a level that is lower than with 6 peaks, but where the difference between these strengths is very small—just large enough to keep inhibition reliably stronger than excitation.

In summary, the simulations in this section illustrate that self-stabilized peaks within DNFs provide a neural basis at the population level for (a) robust perceptual encoding and simple forms of perceptual decision making (e.g., should I build a peak or not?), (b) tracking/updating inputs despite variations in the input signal, (c) simple forms of distractor suppression when multiple inputs are present, and (d) multi-input tracking/updating when a capacity-limited set of reliable input signals are present. These properties of self-stabilized peaks could also be useful in cases in which it is necessary to track multiple moving objects and/or landmarks in the task space.

We, the authors of this chapter, think self-stabilized peaks provide a wonderful first step

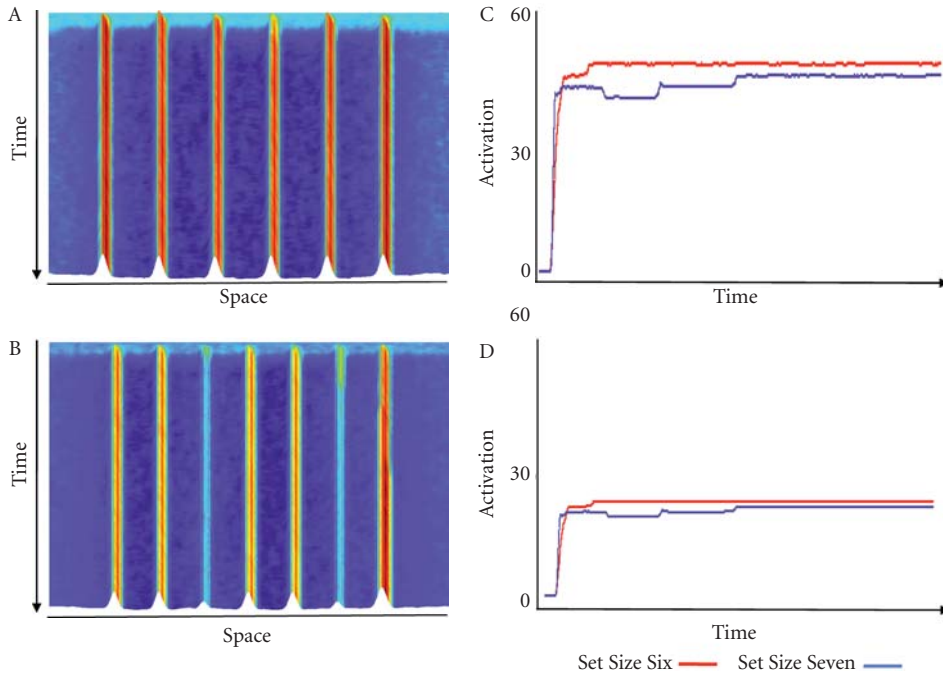


Figure 5.4. Simulations depicting capacity limits of a field in the self-stabilized state

This figure depicts a two-layer field operating in the self-stabilized state. The field successfully tracks up to six simultaneous inputs in panel A. When a seventh input is added in panel C, the field *selects* a subset, forming self-stabilized peaks for five out of the seven inputs. (Blue activation humps are sub-threshold, that is, below 0.) Line plots show summed excitatory (panel B) and inhibitory (panel D) activation, respectively, with six (red line) versus seven (blue line) inputs. The time course of simulations was 10 s.

for linking simple forms of cognition to low-level sensory inputs to obtain states/behaviors that are robust over space and time from inputs that are noisy and transient. But peaks in this self-stabilized state have a fundamental limitation: they must be anchored to input. Once the input is removed for a sustained period, this state is destabilized and the field relaxes back to the resting state. In the next section, we show that this does not always have to occur; DNFs can operate in a different attractor state—the *self-sustaining state*.

Self-Sustaining Peaks and Working Memory

By employing the simulator we've been using in all the simulations thus far and simply making the field a bit more excitable overall,

a new, qualitatively different state emerges. Figure 5.5B depicts a simulation of a *self-sustaining peak*. This simulation uses the same input as in Figure 5.2G (see Fig. 5.5A), *except* that the resting level of the *inhibitory field* has been moved from -5 to -12 . (The inhibitory field is not shown here; refer to Figures 5.1B and 5.1C for depictions of an inhibitory field.) Note that the change in inhibition effectively increases the excitation in the system—less inhibition allows more excitation. This change is evident in the higher activation level in Figure 5.5C (black trace) relative to Figure 5.2I (black trace); compare these traces with the traces in Figure 5.2I. The other critical difference is in how the peak behaves over time when the input is occluded for longer durations: rather than returning to the resting state during the three longer

occlusion events, the peak remains in the *on* state throughout the simulation (Fig. 5.5B).

We have referred to this self-sustaining behavior as a form of *working memory* because it shows how a cortical field can actively maintain information about a stimulus during a memory delay (i.e., when the stimulus is no longer present). There is general agreement that some form of recurrent activation is the neural basis for working memory (Fuster, 1973; Goldman-Rakic, 1995), and DNFs provide one implementation of this concept.

There may also be an important link between self-sustaining peaks and the original proposal by Baddeley (1986) that working memory is a form of active memory that can be used in the service of tasks that require operations on information. In our view, using information in a working memory store in the service of some cognitive operation *requires stability*. That is, the cortical system actively maintaining the information to be used must keep that neural representation stable to support operations on this information conducted by, for instance, some other cortical system. This requires stability in the face of neural noise, stability in the face of potentially *distracting* inputs, and so on. In short, we contend that the working memory construct requires the types of stability we have already shown that peaks have! Of course, the simulations in the previous section showed the stability properties of self-stabilized peaks.

What are the stability properties and unique features of self-sustaining peaks?

As can be seen in Figure 5.5B, self-sustaining peaks can still handle variations in the amplitude, spatial location, and presence/absence of an input. Indeed, in some respects, self-sustaining peaks handle these variations more effectively. For instance, self-sustaining peaks can remain in the *on* state during long occlusions (Fig. 5.5C). Self-sustaining peaks are also very good at tracking multiple inputs with several intervening occlusion events. This is shown in Figure 5.6A. This simulation had the same input pattern as in Figure 5.3C. Now, rather than losing the continuity in the three input patterns, the self-sustaining peak effectively *glues* together the gaps in input. Note that the field does not do this perfectly—the field *expected* the input to the far left to reappear at a location to the left of where it actually reappeared. One can see a quick adjustment, however: the left peak (red line in Fig. 5.6B) dips down, but quickly blends with the new input and returns to a high activation value.

Although self-sustaining peaks do some things quite well, there are also some down sides to this type of cognitive stability. First, self-sustaining peaks are not quite as good at tracking spatial variations in input. Second, when there are large shifts in input, it takes self-sustaining peaks longer to shift from one input source to another. A third consequence is




Figure 5.5. Simulation depicting a self-sustaining peak

This figure depicts a two-layer, dynamic neural field operating in the self-sustaining state (see panel B) when presented with variable input (see panel A; the same input used in Figure 5.2, panel G). A single peak is self-sustained in the field (see panel B and black trace/activation profile in panel C) despite the presence of noise and lengthy interruptions in the input (see red trace in panel C). The time course of simulations was 1 s (1000 ms).




Figure 5.6. Simulation depicting a self-sustaining peak

This figure depicts a two-layer, dynamic neural field operating in the self-sustaining state (see panel A) tracking the same inputs used in Figure 5.3, panel C. Self-sustaining peaks are formed and effectively track the inputs across several intervening occlusion events (see activation profiles through time in panel B associated with each peak location in panel A). The time course of simulations was 10 s.

illustrated in Figure 5.7: fields with self-sustaining peaks have a lower capacity than fields with self-stabilized peaks. Panel A depicts a field with four simultaneously active self-sustaining peaks. In panel C, the input set size increases to six items. Rather than building six self-sustaining peaks, the field is able to sustain only four peaks once the input is removed. This occurs for the same reasons discussed previously: after the input is removed, the field settles into a state that balances the overall level of excitation and inhibition (compare the solid red and blue lines in panel B and panel D). More specifically, in the absence of the strong excitation provided by direct input, the field cannot maintain the level of excitation/inhibition required for six peaks. Thus, two peaks are *killed* and the excitatory/inhibitory balance returns to its maximal level. Note that we have added the overall excitation and inhibition levels for the self-stabilized simulation shown in Figure 5.4 for comparison (see dashed blue lines).

We conclude this discussion with one final note. The simulations in Figure 5.7 leave the impression that self-sustaining peaks will continue to exist forever under the right conditions. This is, of course, not the case in real neural systems; peaks can be de-stabilized by a variety of factors including neural noise, suppression by other inputs, and so on. Indeed, the likelihood that this will occur increases systematically with memory delays, leading to typical

forgetting curves. Although this is the case, the forgetting mechanism in DFT—bifurcation into the resting state—predicts that when peaks drop out, they will drop out in an all-or-none fashion versus gradually degrading in resolution. Recent data from Zhang and Luck (2008) using a color recall task are consistent with this mechanism.

Integrating Perception and Working Memory in Real Time: Multilayered Fields

The simple two-layer networks introduced thus far can serve either a perceptual or a working memory function. For instance, a two-layer network operating in the self-stabilized state, where peaks of activation are formed in response to input but die out when input is removed, could be used to capture the perceptual encoding of an array of items in the world. Additionally, a two-layer network operating in the self-sustaining state, where peaks of activation remain above threshold after input has been removed, can be used to implement a form of working memory. Critically, however, many real-world activities require the *integration of both functions*. For instance, to remember the locations of critical objects in the local surrounds, one needs to actively maintain information about these locations in working memory, but also keep those memories anchored to visible landmarks in the task space (Spencer et al., 2007). Similarly, to

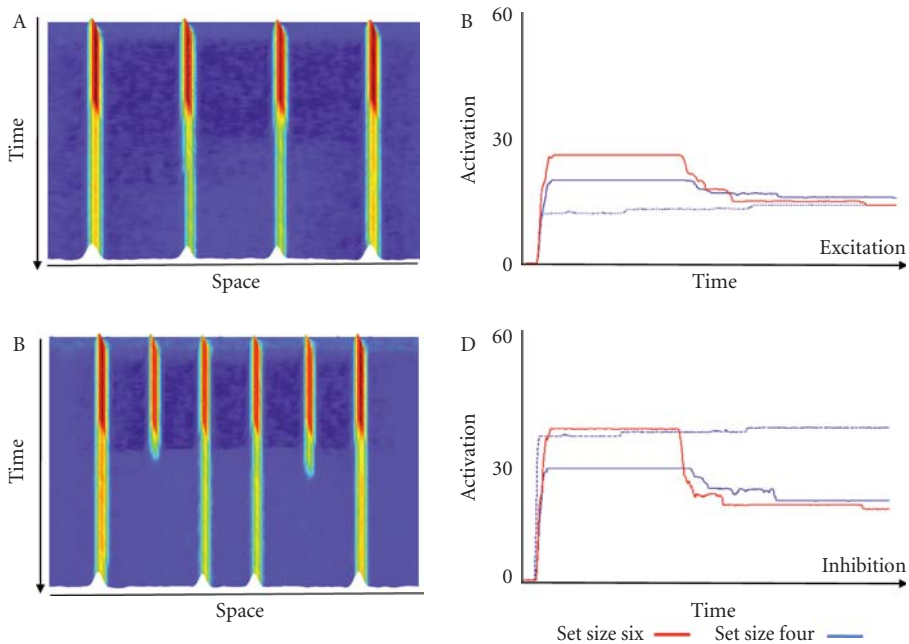


Figure 5.7. Simulations depicting capacity limits of a field in the self-sustaining state

This figure depicts a two-layer field operating in the self-sustaining state. The field successfully forms and maintains four self-sustained peaks of activation in the absence of input (panel A). However, when six inputs are presented, only four are selected and self-sustained throughout the 10-sec simulation (panel C). Line plots show the sum of activation in excitatory (panel B) and inhibitory (panel D) layers for set-size 6 (red lines), set-size 4 (solid blue lines), and, for comparison, for set-size 4 with the field operating in the self-stabilized state, as in Figure 5.4 (dashed lines).

detect changes in the local surrounds, one must compare the currently perceived set of items to the set of items actively maintained in working memory. How are these types of functionality achieved within DFT?

To address the real-time integration of perception and working memory, we have developed a three-layer architecture consisting of two layers of excitatory neurons coupled to a single layer of inhibitory neurons (Johnson et al., 2006; Simmering, Schutte, et al., 2008; Spencer et al., 2007). Figure 5.8 shows the architecture of this system. The model consists of an excitatory perceptual field (PF; panel A), an excitatory working memory field (WM; panel C), and a shared inhibitory layer (Inhib; panel B). Excitatory and inhibitory connections among the layers are indicated by green and red arrows, respectively. Note that this multilayered architecture was

inspired, in part, by the cytoarchitecture of the visual cortex (Douglas & Martin, 1998). That said, it is possible to implement the same functionality we highlight below using two coupled two-layer fields where activation of the single, shared inhibitory layer in these networks inhibits *both* excitatory layers.

As its name suggests, PF is the primary target of perceptual input to the model; however, neurons in the WM field also receive weak, direct input. In addition, PF provides excitatory input to WM. Finally, both PF and WM pass input to the inhibitory layer and receive inhibition projected broadly back. These pairs of two-layer interactions are identical to the local excitatory and lateral inhibitory neural interactions evident in previous simulations.

What does this particular type of integration achieve? The three-layer architecture has been

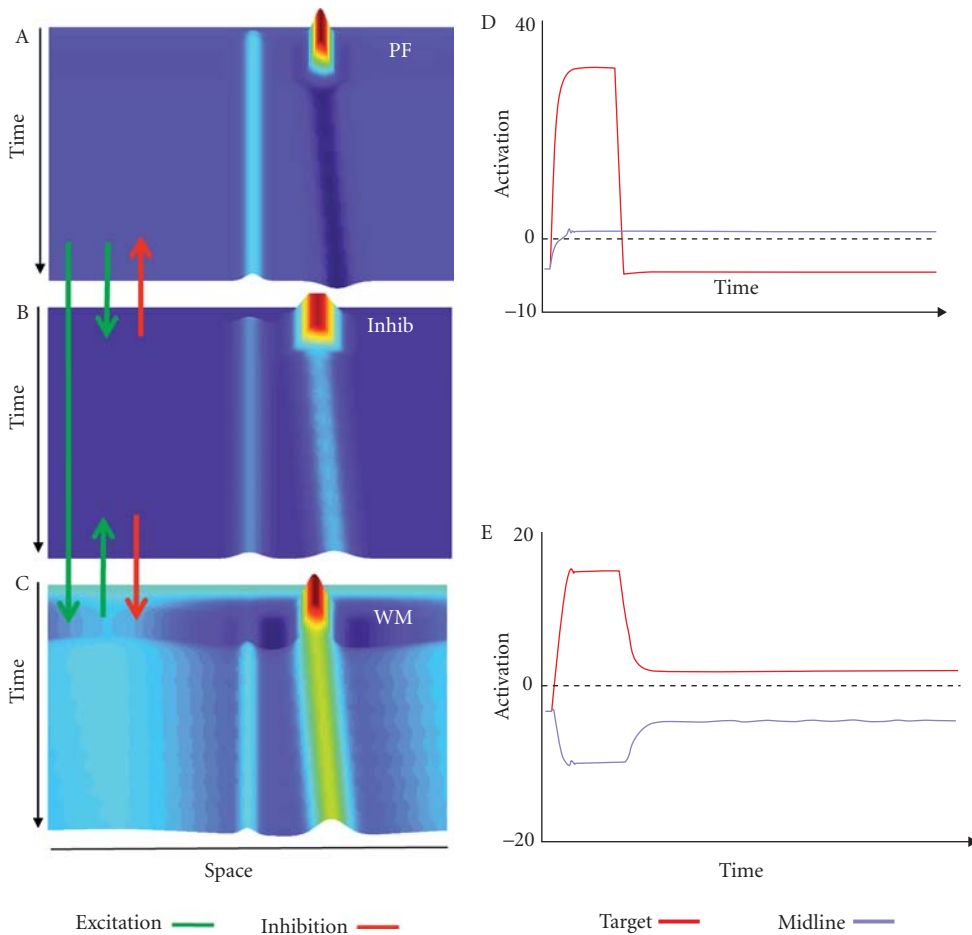


Figure 5.8. Simulation illustrating a one-dimensional, three-layer dynamic neural field

This figure depicts a simulation of a one-dimensional, three-layer dynamic neural field model depicting the anchoring of spatial working memory to the real-time perception of landmarks in a task space. The model consists of an excitatory perceptual field (PF; panel A), an excitatory working memory field (WM; panel C), and a shared inhibitory layer (Inhib; panel B). Excitatory and inhibitory connections among the layers are indicated by green and red arrows, respectively. Line plots (panels D and E) depict activation levels in PF and WM, respectively, in response to the target (red line) and midline (blue line) inputs. The time course of each simulation was 10 sec. See text for further details.

used to capture the anchoring of spatial working memory (SWM) to the real-time perception of landmarks in the task space (Schutte, Spencer, & Schöner, 2003; Spencer et al., 2007). An example can be seen in Figure 5.8, which shows a simulation of a three-layer model remembering a target location relative to a visible landmark during a delay. At the start of this simulation, two inputs are present: a landmark along the midline axis

of the task space and a salient target input to the right of midline. In response to these inputs, PF (panel A) creates two self-stabilized peaks (see panel D), while WM (panel C) creates a single, self-sustaining peak associated with the stronger target input (see panel E).

Next, the target disappears from view. Consequently, the target peak returns to the resting level in PF, but remains in a self-

sustaining state in WM. Thus, during the memory delay, PF is actively locked onto the visible landmark, while WM is actively maintaining a representation of the target. The online maintenance of a peak in PF and a peak in WM has both costs and benefits. One *cost* is readily apparent in panel C: the peak of activation in WM *drifts* rightward—away from midline—during the memory delay, consistent with empirical observations (see, for example, Spencer & Hund, 2002, 2003). This occurs because the peak of activation at midline in PF passes activation to the Inhib layer at midline (see panel B). This creates greater inhibition on the midline side of the self-sustaining peak in WM, effectively repelling this peak away from the landmark (for additional evidence of such effects, see Simmering & Spencer, 2008; Simmering, Spencer, & Schöner, 2006).

Although these interactions produce systematic bias in memory, Huttenlocher and colleagues (Huttenlocher et al., 1991; Huttenlocher, Hedges, & Vevea, 2000) pointed out that such biases can have adaptive consequences. For instance, this type of bias minimizes the likelihood of cases/errors that cross over to the wrong side of the landmark (that is, cases in which the target is shown to the right of midline but participants respond to the left). This bias also exaggerates the *left-ness* or *right-ness* of the target item, which could facilitate recall from long-term memory. Indeed, from a Bayesian perspective, such bias can help create an optimal spatial memory system, trading off systematic bias for an overall reduction in variance (see, Huttenlocher et al., 1991; Huttenlocher et al., 2000).

The simulation in Figure 5.8 shows how perceptual and working memory processes can be integrated in a real-time neural system using the concepts of DFT. Recent empirical and theoretical work has shown that the particular type of integration shown in Figure 5.8 has several emergent properties that can be linked to specific behavioral phenomena. For instance, we have recently demonstrated that the architecture shown in Figure 5.8 can capture children's (Simmering, Johnson, & Spencer, 2008, in preparation) and adults' (Johnson et al.,

2006) performance in a standard *change detection task* (see Luck & Vogel, 1997). In this task, subjects are shown an array of objects on a computer screen that vary in, say, location. The task is to remember these objects and report whether they are in the same or different locations relative to a second array of items presented 1 second later.

How does the three-layer model (as in Fig. 5.8) capture performance in this task? The initial presentation of the array creates multiple self-sustaining peaks of activation in working memory. During the delay, these peaks create *troughs* of inhibition in PF. An example is evident in panel A; during the memory delay, there is an inhibited region in PF associated with the self-sustaining peak in WM (panel C). What is the behavioral consequence of these troughs? Is there any behavioral evidence supporting them? If an item is presented in the same location as an item currently held in WM, this trough can suppress the build up of a self-stabilized peak in PF. By contrast, if an item is presented outside this trough, a self-stabilized peak can quickly form in PF. We use this difference—the absence of a peak in PF on *same* trials versus the presence of a peak in PF on *different* trials—to drive decision making in our model of change detection performance (Johnson et al., 2006; Johnson, et al., in press). And this mechanism for generating same/different decisions has led to a host of novel behavioral predictions that have been successfully tested experimentally (Johnson et al., in press; Johnson et al., 2006; Simmering & Spencer, 2008).

In summary, these examples highlight the adaptive consequences of integrating perceptual and working memory processes in a real-time neural system. But this provides only a beginning for thinking about embodied cognitive dynamics. In the next section, we add another novel ingredient—the integration of dynamics across multiple timescales.

Integrating Multiple Timescales: Long-Term Memory Dynamics

As demonstrated in Figure 5.1, patterns of activation within DNFs can be influenced by the presence of sub-threshold inputs, that is, inputs

that are too weak to drive neurons in a field from their negative resting level to above-zero activation values. Recall that in the absence of sub-threshold activation, the field shown in Figure 5.1 remains in the resting state. With a weak sub-threshold input, the field forms a self-stabilized peak. And with a stronger sub-threshold input, the field forms a self-sustaining peak.

One source of sub-threshold activation within DFT is input activation from a long-term memory (LTM) process. The formation of LTM traces in DFT relies on a very simple type of dynamics, which implements a form of Hebbian learning (see chapter 18). Specifically, whenever activation at a given field site climbs above threshold (i.e., above 0), activation traces begin to accumulate in an LTM field at that site. At subsequent points, such traces are fed back to the field of origin as an additional source of excitatory input. When there is no above-threshold activation in a field, no LTM traces are laid down, and previously accumulated traces gradually decay.

The accumulation and decay of LTM traces occur more slowly than activation evolves in the excitatory and inhibitory fields discussed thus far. This allows patterns of activation that reflect events unfolding over longer timescales to accumulate in LTM. For instance, over the course of an experimental session in the spatial paradigm we have been considering, LTM traces can be formed that reflect the probability of targets appearing at particular locations. Importantly, because LTM traces are only laid down when above-threshold peaks are formed, the memory here is not about the statistics of stimulation, but instead reflects the statistics of peaks, that is, the *decisions* made second-to-second by the dynamics within a neural field. We discuss this in greater detail below. For now, note that sometimes the history of a field's decisions in a given context diverges from the statistics of the input (see chapter 18 for discussion).

The series of simulations depicted in Figure 5.9 illustrates the build up of LTM traces every five trials across a series of thirty 10-second simulation trials, as well as the consequences of LTM input on patterns of activation

in DNFs. For each simulated trial, a single input was presented at one of two possible mean locations (a left location and a right location). The inputs had the same types of variation used in previous simulations (i.e., amplitude, spatial location, presence/absence). As can be seen in Figure 5.9A, the field operated in a self-stabilized state in the first simulation, tracking the input on the left side, but also dipping into the resting state during two occlusions (see Fig. 5.9B).

Figures 5.9E–J depict how patterns of activation in LTM evolve over the course of 30 simulations, and how this can give rise to category formation and the emergence of *exemplars*. As can be seen in Figure 5.9E, after a few trials to each target (i.e., on trial 5), the activation in LTM is clustered around a central location. This occurs because, although inputs were presented at two different locations *across* trials, variations in the spatial positioning of the inputs *during* a given trial produce significant overlap and a consequent blending of LTM traces. Thus, early in learning the distribution of activation in LTM is unimodal, reflecting an average or *prototypical location*. Over many simulations, this centrally biased activation becomes stronger. By Figure 5.9H, however, LTM begins to differentiate, and by trial 30 (Fig. 5.9J), two distinct targets have emerged with humps centered at the actual mean target locations. As can be seen in Figure 5.9C (which depicts the field on trial 30) the LTM extracted these *targets* despite the continued presence of noise in the input. Figure 5.9C also illustrates an emergent effect that arises from the increased strength of activation from LTM into the field; with this added input, the field entered the self-sustaining state, surviving an occlusion event late in the trial (see Fig. 5.9D).

In summary, the simulations in Figure 5.9 illustrate how the LTM mechanism in DFT captures the emergence of a prototype-like representation early in learning; late in learning, this shifts to an exemplar-like representation, which correctly reflects the true mean input locations. This occurs in an unsupervised manner. Moreover, this form of learning has emergent consequences on the types of active decisions formed in real time during a trial. Thus, the two timescales of cognitive dynamics depicted

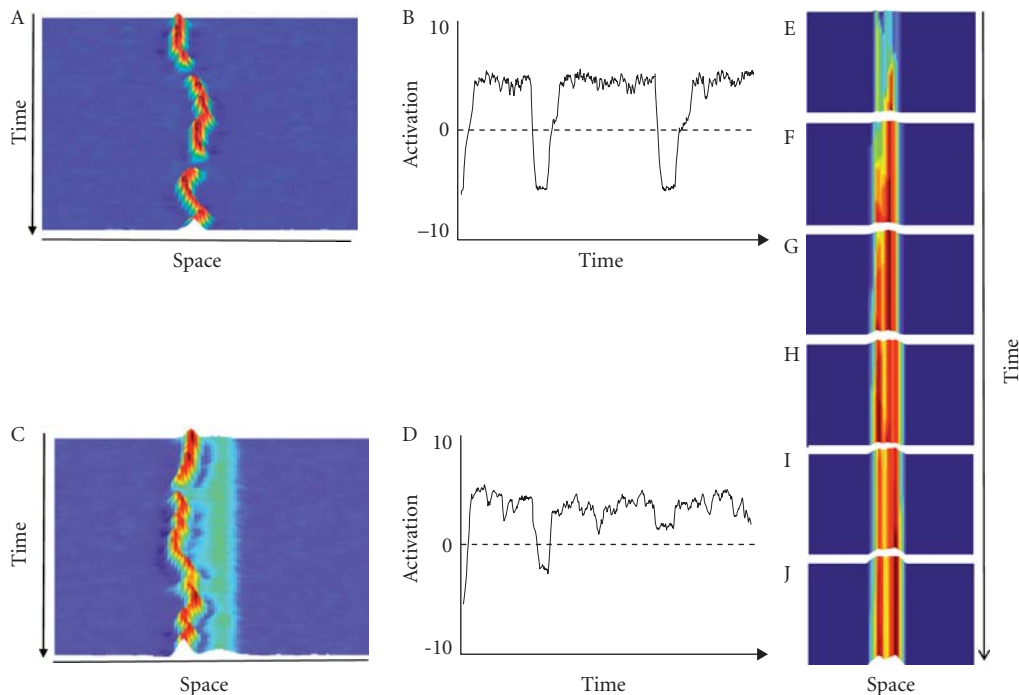


Figure 5.9. Simulated buildup of LTM traces during a series of trials

In this figure, each of the thirty 10-s trials consisted of the presentation of a single input at a mean left or right location. Inputs (not shown) were interrupted periodically and varied in amplitude and spatial location within a trial. At the beginning of the series, a self-stabilized peak tracks the input as long as it is present (panel A), but dies out when the input is interrupted (see activation profile in panel B). By the end of the series (panels C and D), the presence of robust LTM input to the field allows the peak to self-sustain despite numerous brief interruptions. Panels E–J depict activation in the LTM field during every fifth trial. Across the simulation runs, the distribution of activation in LTM shifts from a unimodal (prototypical) distribution (panels E, F, and G) to a bimodal (exemplar) distribution (panel J), which reflects the two target locations presented across the series of simulations.

here—the fast timescale where peaks form and the slower timescale of LTM formation—are truly interactive and lead to new phenomena over the trials.

In addition to contributing to a shift in the attractor states a field can enter, LTM can have other adaptive influences on the properties of peaks within an attractor state. For instance, Figure 5.10A depicts a simulation of a self-stabilized peak tracking an input amid distracting items. This simulation is identical to the simulation shown in Figure 5.3A with one exception—the field shown in Figure 5.10A is coupled to an LTM field. As can be seen by comparing the two figures, the simulations

are quite similar initially. As the trial unfolds, however, the increased activation provided by input from LTM leads to a more robust self-stabilized peak. As a consequence, distractor suppression late in the trial is stronger and extends to distractors farther from the central peak (see arrows).

The simulations described thus far show that LTM input can be quite beneficial, making it more likely for peaks of activation to form in response to stimulation, increasing the stability of peaks in the face of variable and intermittent input, and making peaks more robust to distractors. In addition, the accumulation of traces in LTM can give rise to category formation and




Figure 5.10. Simulation demonstrating distractor suppression

Replication of the simulation depicted in Figure 3, panel A, with the addition of input from an LTM field. LTM traces build up over the course of the 10-s simulation to produce a more robust self-stabilized peak (see panel A) that effectively suppresses distractors. (Compare suppressed distractors in Fig. 5.10, panel A, to activation values in Fig. 5.3, panel A.) Note that yellow and red activation values are above zero in this case.

can pull out exemplars when several inputs are tracked across multiple trials. Nevertheless, the stabilizing influence of LTM is sometimes achieved at the cost of flexibility in the face of changing inputs and task demands.

To appreciate this, consider the simulation shown in Figure 5.11, which captures a type of perseverative responding. At the beginning of the simulation, a single input is briefly presented and a stable, self-sustained peak forms at the location of the input (see Fig. 5.11A). Shortly thereafter, an activation trace begins to accumulate in the LTM field (Fig. 5.11B). The feedback from LTM serves to stabilize and strengthen the peak depicted in Figure 5.11A. The next event in the simulation is the presentation of a second, task-relevant input at a right location. Although there is now only a single stimulus present in the task space, the self-sustaining peak effectively treats this new input as a distractor and prevents a new peak from forming at the right location (see sub-threshold blue trace in Fig. 5.11C). As a consequence, the field perseverates—the ongoing state is determined by its past history rather than by the current input.

The simulations described in this section highlight a fundamental tension between stability and flexibility that arises in the nervous system. On one hand, it is highly advantageous for ongoing behavior to be open to the influence of past experiences. Such influences can have a

stabilizing effect by, for instance, allowing the system to operate in a self-sustaining versus self-stabilized state. However, adaptive behavior is about more than stabilizing a particular behavior or mental state; adaptive behavior involves flexibly switching between stable states as the situation demands. This is the focus of the following section.

Toward Behavioral Flexibility: Going into and out of (vs. near) Stable Attractors

Tensions among stability, instability, and flexibility have been dramatically illustrated in research on the Piagetian A-not-B error (Thelen, Schöner, Scheier, & Smith, 2001). In the A-not-B task, infants are shown a toy, the toy is hidden at location A, and the infants are allowed to search for the toy after a brief delay. This sequence is repeated several times. Next, the toy is hidden at nearby location B. After a 3–5-second delay, 8- to 10-month-old infants reliably search for the toy at the original location A. By contrast, 10- to 12-month-old infants search correctly at location B (Smith, Thelen, Titzer, & McLin, 1999).

The original DFT account of infants' performance in this task used a DNF coupled to an LTM or *pre-shape* field (Thelen et al., 2001). When the target was presented on each A-trial, the model formed a self-stabilized peak. Due to the short delay, the model was able to correctly *search* at this location. On the B-trials, however,

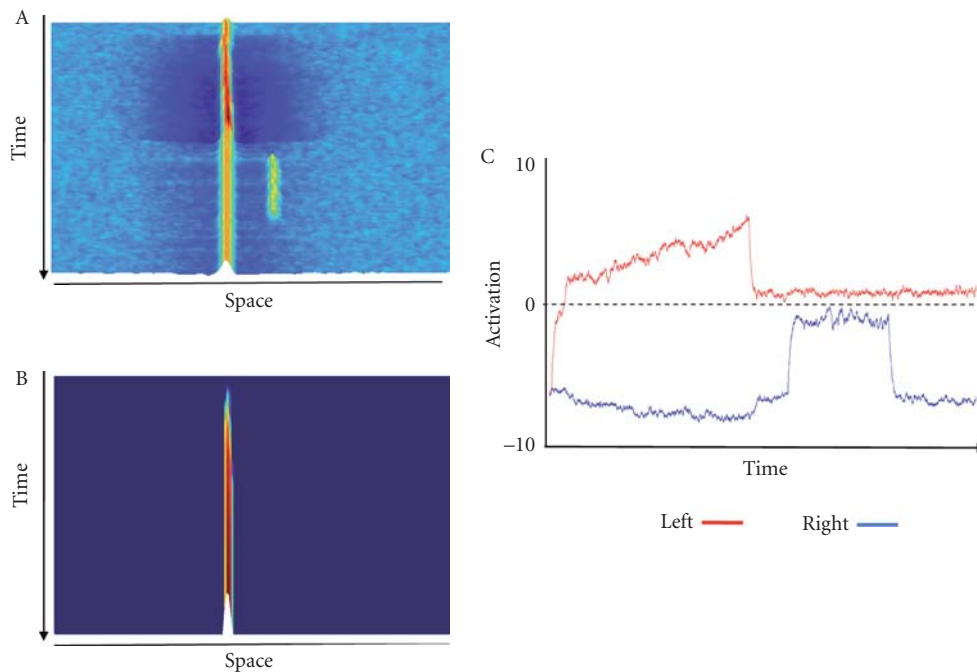


Figure 5.11. Simulation highlighting tension between stability and flexibility in neural systems

At the beginning of the simulation, a self-sustained peak is formed (see panel A and red trace in panel C) in response to a transient task-relevant input in the task space. As the 10-s simulation progresses, the peak is further stabilized by LTM traces building up at the left location (panel B). Thus, when a second task-relevant input appears to the right (see right activation hump in panel A), it is treated as a distractor and is suppressed by the robust peak at the left location—a form of perseveration. Line plots in panel C, show the activation profile through time at the left (red) and right (blue) locations.

the peak returned to the resting state during the longer memory delay. As a consequence, the model perseveratively searched at A, that is, the model's performance was dominated by the LTM of previous searches/reaches to A rather than a robust memory of the target at B. Thelen and colleagues proposed that the shift from perseverative responding to correct responding at 10–12 months of age reflected an increase in the strength of neural interactions. In terms of the dynamic field model, this represents a shift from the self-stabilized state to the self-sustaining state.

Although this dynamic field account captured key aspects of infants' performance in the A-not-B task and generated a host of novel predictions (e.g., Diedrich, Highlands, Spahr, Thelen, & Smith, 2001; Diedrich, Thelen, Smith, & Corbetta, 2000), a funny thing happened in the late 1990s. Esther Thelen was a strong

proponent of embodied cognition; she also had a keen suspicion of computational models. Thus, she asked Gregor Schöner to demonstrate the embodiment of the dynamic field account of infant perseverative reaching by implementing the model on a robot with real sensors operating in real time. The result: the original A-not-B model failed this embodiment test! Rather than mimicking infants' performance, the robot spent trial after trial jiggling back and forth, unable to clearly respond to A or to B on any given trial.

What was the problem? The problem was that the original dynamic field model of the A-not-B error failed to generate a stable peak at the end of each trial. Rather, Thelen and colleagues used a basic modeling technique to *read off* the response—the model was said to *respond* to whichever location showed the strongest

activation at the end of the simulation. The problem on the robot was that the field was in a resting state at the end of each trial and, as we showed in Figure 5.3E, fluctuations in amplitude in this state are a far cry from the robust neural signal a nervous system would need to drive a stable response. Thus, instead of reliably moving to A or to B, the robot jiggled back and forth as the location of maximal activation fluctuated from time step to time step.

To solve this problem, Dineva and colleagues (Dineva, Schöner, & Thelen, 2008, in preparation) used a novel feature of DFT; they raised the resting level of the field during the response phase of the A-not-B task (that is, when the box is moved to within an infant's reach). Raising the resting level of a neural field makes the field more excitable overall, allowing small humps of activation to rise above threshold and generate a peak. Simulations of a new model of infant perseverative reaching showed that this could capture not only the detailed distribution of infants' responses in the A-not-B task, but also the pattern of spontaneous errors infants' show (i.e., reaches to B on A-trials) and critical trial-to-trial dependencies in infants' performance (e.g., how a single *spontaneous* reach to B on an A-trial effects the likelihood of reaching to A or B on subsequent trials).

Moreover, Dineva and colleagues (Dineva, Faubel, & Schöner, 2007) have recently demonstrated something possibly more impressive—that a DNF can, in fact, drive the behavior of an autonomous robot placed in an A-not-B task. This robot perceived the locations of colored targets in the task space via a camera. These inputs were fed into a DNF in real time. The field also received inputs from a continuously coupled LTM field which kept track of the robot's decisions from trial to trial. Finally, the DNF was coupled in real time to a motor system that turned the wheels of the robot. In particular, when the robot made a decision (i.e., formed a peak), this created an attractor for the motor system and moved the robot toward the represented location. Dineva and colleagues conducted a series of experiments with the robot and showed that this autonomous system reproduced the behavior of infants in the A-not-B

situation—with real inputs and *live* decision making on the part of the robot. This provides a demonstration proof that DFT offers a coherent framework for thinking about embodied cognitive dynamics.

Dineva and colleagues' (Dineva et al., 2008, in preparation) model of the A-not-B situation demonstrates how a boost in the resting level of a field can move a system into a stable attractor state. This provides an initial view of a critical form of behavioral flexibility that is on firm neurophysiological ground. For instance, recent data suggest that one role of the prefrontal cortex is to selectively boost the gain on some neural populations over others (D'Esposito, 2007). This prefrontal boost could serve the same role as a resting-level boost in a neural field; it makes a neural field more excitable overall, allowing that neural population to make decisions (i.e., form peaks) even in the face of only partially certain information.

Interestingly, however, this is only part of the story. The simulation in Figure 5.12 shows the second part. This simulation is identical to the perseverative simulation in Figure 5.11 *except* that we modulated the resting level of the inhibitory field midway through the simulation (see Fig. 5.12C). By boosting the gain on neurons in the inhibitory layer, we can actively de-stabilize the peak at the left location. This allows the field to form a self-stabilized peak at the right location, effectively responding to this new input and overriding the strong LTM of the left location. Thus, a very simple mechanism—boosting the resting level of a population of neurons—can serve to move a neural field into a peak state or actively destabilize a peak.

We are just beginning to explore the types of behavioral flexibility this simple mechanism enables. Thus far, the picture that is emerging is quite powerful. For instance, Buss and Spencer (2008) used this type of mechanism to explain developmental changes in executive control in the dimensional change card sort task with 3- and 4-year-old children (see Morton & Munakata, 2002; Zelazo, 2006; Zelazo et al., 2003). And, critically, this mechanism makes several novel predictions, which we are currently testing.

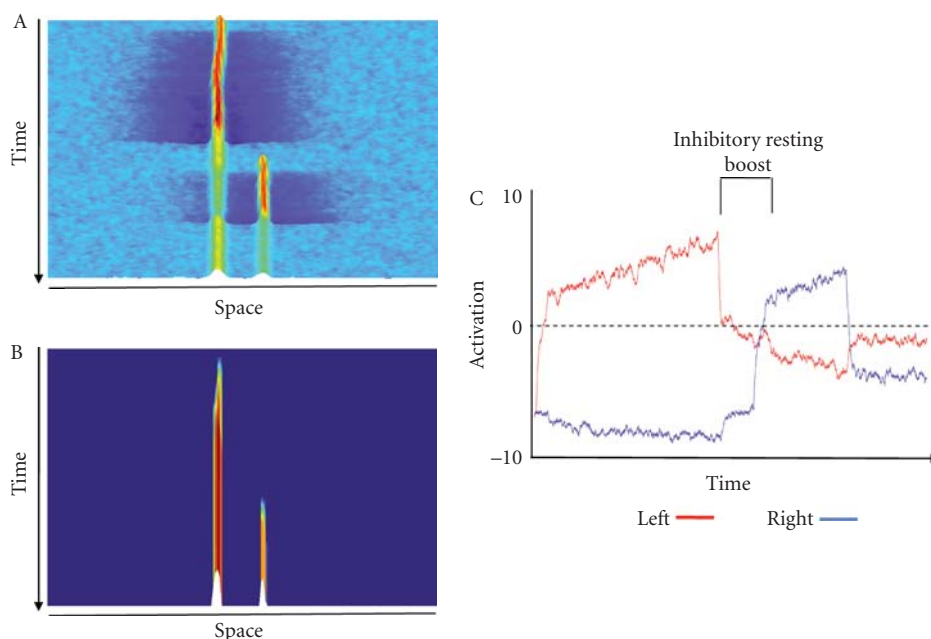


Figure 5.12. A simple way to achieve stability and flexibility in dynamic neural fields

This figure illustrates a simple way in which stability and flexibility can be achieved in dynamic neural fields. The simulation depicted is identical to the simulation depicted in Figure 5.11 *except* that the resting level of the inhibitory field was boosted slightly midway through the trial just before the second input appeared (see panel C). This served to destabilize the left peak (as seen in panel A), allowing a second peak to be built in response to the new task-relevant input to the right. As a result, LTM traces associated with the right target begin to build as the trial ends (see panel B). The time course of the simulation was 10 s.

Beyond specific models, however, we think there are some profound implications of this type of behavioral flexibility for dynamical systems theories more generally. Several recent dynamical systems accounts of cognition have correctly stressed that cognitive dynamics must be flexible to capture the impressive abilities of adult decision making (Kelso, 1995; Spivey, 2007). However, flexibility has been conceptualized as the ability to go near, rather than stably into, attractor states. Why? The reason these approaches have stressed this view of flexibility returns us to the simulation in Figure 5.11—if a system goes too deeply into a stable attractor state, it has difficulty getting *kicked* out of that state. The common image painted here is one of a ball rolling on a complex landscape of wells. If the ball rolls too deeply into a well, it takes a great deal of effort to get out of that well,

and behavioral flexibility is limited. If, however, the ball rolls close enough to a well to support functional behavior, but then quickly moves on to the next behavioral response, the system—in theory—can show flexibility.

From our vantage point, this view of flexibility is missing something critical. Flexibility is not about going *near* attractors; it is about going *into* stable attractors and then having the ability to get back out to move on to the next stable attractor. We think this insight is critical, and our demonstration of embodiment with the robotics example backs up this claim. Simply put, a robot's behavior is poorly served by going near stable attractor states. What an autonomous system requires is stable decision making plus an active mechanism that can destabilize this state when the situation demands. We think peaks provide the stability such a system needs. And we think active changes in the resting level

of a neural population provide the mechanism for getting into and out of such stable states.

This view of behavioral flexibility brings us to a final topic in our review of DFT concepts—development. In our view, a central issue in cognitive development is learning to regulate the dynamic balance among stability, instability, and flexibility (Schöner & Dineva, 2007; Spencer & Schöner, 2003). We discuss what this might mean in the section below.

What about Development?

We have made promising advances in understanding development within the domain of spatial cognition using a relatively simple hypothesis—the *spatial precision hypothesis* (SPH). In the section below, we review this hypothesis and some of the novel, empirical findings it captures. We then address the issue of what might cause this particular type of developmental change.

The spatial precision hypothesis states that the precision of neural interactions becomes greater and more stable over development (Schutte et al., 2003; Spencer & Hund, 2003). This hypothesis is illustrated in Figure 5.13A. Each curve in this figure shows an example

of an interaction function at some point in development with early development shown in darker lines and later development shown in lighter lines. As the interaction functions move from early to late development, the spatial precision of local excitatory interactions narrows and excitation increases. At the same time, lateral inhibitory interactions expand their spatial range and grow stronger. These changes result in relatively unstable peaks early in development, which are sensitive to input across a broad spatial range, and stable peaks later in development, which are only sensitive to input at narrow separations.

The developmental changes in neural interaction captured by SPH can explain several behavioral results with children. For instance, in a sandbox task over development, children show a reduction in effects that are similar to the A-not-B error. In particular, when children are asked to find an object repeatedly at location A in a sandbox, their responses on a subsequent B-trial are biased in the direction of A. This supports the *drift* effect shown in Figure 5.8C, although drift in this case is toward an LTM of location A. More importantly, for the present discussion, however, there is a change




Figure 5.13. Changes in strength of excitatory and inhibitory coupling over development

This figure provides an illustration of how changes in the strength of excitatory and inhibitory coupling in a three-layer field model lead to increased stability and spatial precision of peaks across development. In panel A, darker activation profiles depict the structure of self-sustaining peaks for younger subject models, while lighter profiles depict self-sustaining peak profiles for older subject models. Note the narrowing and strengthening of peaks over development. Panel B depicts the strength changes in parameters used to produce the shift in peak structure over development shown in panel A. See text for further details.

Source: Adapted from simulations by Schutte and Spencer (2008, submitted for publication).

in the spatial range across which such effects are observed between 2 and 4 years of age; with enhanced precision over development, A and B must be spatially close for the A-not-B error/bias to be observed (for details, see Schutte et al., 2003). The enhanced stability of working memory peaks 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 SPH captures a number of changes in performance across a range of ages, many of these developmental effects are quite intuitive (e.g., spatial memory improves over development). Does SPH capture something less obvious? For instance, one of the more dramatic developmental changes in spatial cognition is the shift in spatial recall biases near a perceived reference frame: young children show biases toward symmetry axes, while older children show biases away from such axes (Huttenlocher, Newcombe, & Sandberg, 1994; Schutte & Spencer, 2002; Spencer & Hund, 2003). Can we capture this *qualitative* developmental transition using the same hypothesis?

A recent paper shows that the answer is yes. Schutte and Spencer (2008, submitted for publication) showed that broader, weaker neural interactions in the three-layer model (as depicted in Fig. 5.8) actually lead to attraction toward the landmark, rather than repulsion away from the landmark. This occurs because neural interactions are too weak to support the active maintenance of two peaks simultaneously. Using Figure 5.8 as an example, interactions in this field are too weak to support both the *reference* peak in PF and the *target* peak in WM. In the absence of a reference peak in PF, the peak in WM *drifts* toward the weak input associated with the landmark.

In addition to capturing the general direction of drift in early versus late development, Schutte and Spencer (2008, submitted for publication) were able to capture the details of a complex developmental transition as children shifted month-by-month from an attraction

pattern to a repulsion pattern. Critically, these details were captured via a *quantitative* shift in model parameters. In particular, to shift from the *child* parameters to *adult* parameters, we scaled the inputs and the strength of the projections among the layers of the three-layer model as specified by SPH. The specific strength changes are depicted in Figure 5.13B. We increased the strength of local excitatory interactions in PF (see line *dev_cuu*), the strength of local excitatory interactions in WM (see line *dev_cww*), and the strength of the inhibitory projections from Inhib to both layers (see line *dev_cv*). Interestingly, these strength changes were sufficient to produce the *sharpening* of neural interaction evident in Figure 5.13A. That is, the enhancement in the precision of neural representations (shown in Fig. 5.13A) was, in fact, completely emergent from changes in the strength of neural interactions.

Results from this study are 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 (Smith & Thelen, 1993; Thelen & Smith, 1994; van der Maas & Molenaar, 1992; van Geert, 1998). But this account goes beyond other dynamical systems account in that SPH also has ties to the emerging literature on developmental cognitive neuroscience.

As an example, it is likely that development of the prefrontal cortex (PFC) is related to the changes in spatial recall discussed above. Data from several studies demonstrate that the prefrontal cortex plays a central role in SWM (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 prefrontal cortex continues into the post-adolescent years (Gogtay et al., 2004; Sowell, Thompson, Tessner, & Toga, 2001). Synaptic density in the PFC reaches its peak at about 2 years of age in humans (Huttenlocher, 1979) and then declines until it finally reaches adult levels

at about 16 years of age (Huttenlocher, 1990). Myelination is also still occurring in the frontal lobe between 2 and 6 years of age (Sampaio & Truwit, 2001). Thus, there is dramatic brain development happening during the period we see dramatic changes in spatial cognition. At this point, however, there is no strong neurophysiological support for the type of changes captured by 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 SPH.

Do these biological changes point toward a maturational picture? Perhaps, but they are equally consistent with an experience-dependent view of development. In particular, the types of change specified by SPH could arise via simple Hebbian processes that strengthen neural connections among layers within cortical fields. We are currently examining whether the specific strength changes depicted in Figure 5.13B can arise via Hebbian learning, that is, in an experience-driven manner.

To summarize, the spatial precision hypothesis effectively captures a broad array of developmental changes in spatial cognition including changes in task performance (such as the A-not-B error) (Schutte et al., 2003), changes in the stability of SWM (Schutte, 2004, unpublished; Spencer & Schöner, 2003), developmental differences in spatial recall versus spatial recognition performance (Spencer, Simmering, & Schutte, 2006), and the qualitative developmental shift in spatial recall biases discussed here (Schutte, 2004, unpublished). Thus, this relatively simple, intuitive hypothesis about how neural interactions change systematically over the longer timescale of development has been quite generative. This is no accident; in our view, the richly structured real-time and learning-time dynamics of DFT can operate in concert with relatively simple developmental changes to yield dramatic, qualitative changes in behavior over long timescales. Importantly, we contend that the relatively modest view of development captured by SPH brings an understanding of developmental mechanisms within reach.

Although SPH has been an impressive workhorse thus far, we readily admit that this hypothesis captures only one aspect of developmental change. Another key issue is the emergence of behavioral flexibility over development (see discussions in Schöner & Dineva, 2007; Spencer & Schöner, 2003). We discussed a dynamic field approach to flexibility in the previous section. With regard to development, a central issue is how the neural system begins to gain active control over the mechanisms described above, that is, control over the resting level of neural populations such that cortical fields can be moved into and out of a peak state. Although we have clearly just scratched the surface on this topic, we contend that DFT has the potential to advance our understanding of this aspect of development.

CONCLUSION

The goal of this chapter was to shed new light on the central concepts of dynamic field theory. Toward this end, we highlighted three neural attractor states that have been foundational in our approach, as well as many characteristics of activation peaks that appear to *come for free* within our theoretical framework. Next, we used a multilayered architecture to show how neural fields can integrate perceptual and working memory processes in real time, and highlighted some emergent features of the particular three-layer architecture we have developed.

We also focused on a different type of integration: integrating processes across multiple timescales using long-term memory dynamics. We showed how a type of Hebbian process in dynamic field theory can lead to category formation; however, this mechanism can also lead to less adaptive behaviors such as perseverative responding. This pointed toward another innovation of our framework—a novel account of behavioral flexibility. Here, we painted a picture of a neural system going stably into attractor states, but then quickly moving out of these states as the situation demands. We framed some of this discussion in the context of work on autonomous robotics, a field which has both demonstrated the embodied nature of dynamic

field theory and highlighted the importance of stability for the generation of real actions in the context of real sensors.

In the final section, we highlighted aspects of our approach to development. Although this work is still in progress, our discussion focused on several novel contributions of our spatial precision hypothesis and pointed toward future work exploring the emergence of behavioral flexibility over development.

A central issue in the context of this chapter is whether the dynamic field theory basics we discussed solve the challenges we mentioned at the outset. For instance, one central challenge within cognitive science has been to develop a framework that can handle the real-time integration of perception, action, and cognition. Several examples from this chapter demonstrate that dynamic field theory can effectively tackle this problem head-on. The most direct example was work on autonomous robotics; however, the examples of our three-layer architecture also highlighted the integration of perceptual and working memory processes. A key insight in this work has been that stability in real time is necessary to realize a coherent interface between brain and behavior.

In addition to real-time integration, we also discussed the challenges inherent in dealing with lower level perception/action on one hand and higher level cognition on the other. Is dynamic field theory up to this challenge? More specifically, are the concepts of dynamic field theory relevant for higher level cognition? In our view, the answer is yes, and recent efforts to examine executive control within a dynamic field framework point in this direction (Buss & Spencer, 2008). In addition, Faubel and Schöner (2008) have recently demonstrated that dynamic field concepts can be used in an autonomous robotics context to support the fast learning of object labels (for related work, see chapter 4 and Samuelson et al., in press). This work uses coupled two-dimensional fields to show how higher dimensional categories of objects can emerge from a lower dimensional representation of features in cortical fields. We have used a similar approach to understand how the brain might solve the *binding problem* in vision (see Johnson et al., 2008).

We want to conclude by acknowledging that there are certainly many areas of overlap between our approach and connectionist approaches to development. Although we did not say much about connectionist approaches in the present chapter (for a more extensive discussion, see chapter 18), we concur with many of the authors of chapters in this volume that dynamic systems theory and connectionism offer complementary views of developmental processes. In our effort to clarify the concepts of dynamic field theory, we have presented material that will be central to a more substantive evaluation of the two approaches. We look forward to engaging in that evaluative process over the coming decade.

ACKNOWLEDGMENTS

Preparation of this manuscript was supported by NSF HSD 0527698 awarded to John P. Spencer, as well as support from the Obermann Center for Advanced Studies at the University of Iowa. We thank Gregor Schöner and Michael Thomas for helpful comments on an earlier version. Correspondence concerning this article should be addressed to John P. Spencer, Department of Psychology, University of Iowa, E11 Seashore Hall, Iowa City, Iowa 52242. E-mail: john-spencer@uiowa.edu.

REFERENCES

- Amari, S.-I. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2), 77–87.
- Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7(3), 237–252.
- 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(5), 433–437.
- Baddeley, A. D. (1986). *Working memory*. New York: Oxford University Press.
- Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking* (pp. 129–163). New York: Cambridge University Press.

- Bastian, A., Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information preshapes the population representation of movement direction in motor cortex. *NeuroReport*, 9(2), 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(7), 2047–2058.
- Bicho, E., Mallet, P., & Schöner, G. (2000). Target representation on an autonomous vehicle with low-level sensors. *International Journal of Robotics Research*, 19(5), 424–447.
- Braitenberg, V., & Schüz, A. (1991). *Anatomy of the cortex*. New York: Springer.
- Braun, M. (1994). *Differential equations and their applications: An introduction to applied mathematics* (4 ed.). New York: Springer.
- Bridgeman, B., Gemmer, A., Forsman, T., & Huemer, V. (2000). Processing spatial information in the sensorimotor branch of the visual system. *Vision Research*, 40(25), 3539–3552.
- Brooks, R. A. (1991). New approaches to robotics. *Science*, 253(5025), 1227–1232.
- Buss, A., & Spencer, J. P. (2008, July). *The emergence of rule-use: A dynamic neural field model of the DCCS*. Paper presented at the 30th Annual Conference of the Cognitive Science Society, Washington, DC. Retrieved October 17, 2008, from <http://www.cogsci.rpi.edu/csjar/chive/proceedings/2008/pdfs/p463.pdf>
- Camperi, M., & Wang, X.-J. (1998). A model of visuospatial working memory in prefrontal cortex: Recurrent network and cellular bistability. *Journal of Computational Neuroscience*, 5(4), 383–405.
- Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press–Bradford Books.
- 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(9), 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(2), 1352–1355.
- Corbetta, D., & Thelen, E. (1996). The developmental origins of bimanual coordination: A dynamic perspective. *Journal of Experimental Psychology: Human Perception and Performance*, 22(2), 502–522.
- Damasio, A. R., & Damasio, H. (1994). Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In C. Koch & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 61–74). Cambridge, MA: MIT Press–Bradford Books.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.
- D'Esposito, M. (2007). From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 761–772.
- Diedrich, F. J., Highlands, T. M., Spahr, K. A., Thelen, E., & Smith, L. B. (2001). The role of target distinctiveness in infant perseverative reaching. *Journal of Experimental Child Psychology*, 78(3), 263–290.
- Diedrich, F. J., Thelen, E., Smith, L. B., & Corbetta, D. (2000). Motor memory is a factor in infant perseverative errors. *Developmental Science*, 3(4), 479–494.
- Dineva, E., Faubel, C., & Schöner, G. (2007, July). *Dynamic field theory and embodiment: Robotic demonstration of infant perseveration*. Paper presented at the 14th International Conference on Perception and Action, Yokohama, Japan.
- Dineva, E., Schöner, G., & Thelen, E. (2008). Behavioral history matters: A dynamic field theory of spontaneous and perseverative errors in infant reaching. Manuscript in preparation.
- 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* (4th ed., pp. 459–509). New York: Oxford University Press.
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience*, 3(Suppl.), 1184–1191.
- Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection*. New York: Basic Books.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109(3), 545–572.
- Faubel, C., & Schöner, G. (2008). Learning to recognize objects on the fly: A neurally based dynamic field approach. *Neural Networks*, 21(4), 562–576.

- Finkel, L. H. (2000). Neuroengineering models of brain disease. *Annual Review Biomedical Engineering*, 2, 577–606.
- Fuster, J. M. (1973). Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *Journal of Neurophysiology*, 36(1), 61–78.
- 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., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental rotation of the neuronal population vector. *Science*, 243(4888), 234–236.
- Georgopoulos, A. P., & Massey, J. T. (1988). Cognitive spatial-motor processes. 2. Information transmitted by the direction of two-dimensional arm movements and by neuronal populations in primate motor cortex and area 5. *Experimental Brain Research*, 69(21), 315–326.
- Georgopoulos, A. P., Taira, M., & Lukashin, A. (1993). Cognitive neurophysiology of the motor cortex. *Science*, 259(5104), 47–52.
- Glenberg, A. M. (1997). What memory is for: Creating meaning in the service of action. *Behavioral and Brain Sciences*, 20(1), 1–19, commentary 20–40, author's response 41–55.
- 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, USA*, 101(21), 8174–8179.
- Goldman-Rakic, P. S. (1995). Cellular bases of working memory. *Neuron*, 14(3), 477–485.
- Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, 52, 217–257.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806.
- 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(2), 141–174.
- Hock, H. S., Schöner, G., & Giese, M. A. (2003). The dynamical foundations of motion pattern formation: Stability, selective adaptation, and perceptual continuity. *Perception and Psychophysics*, 65(3), 429–457.
- Hund, A. M., & Spencer, J. P. (2003). Developmental changes in the relative weighting of geometric and experience-dependent location cues. *Journal of Cognition and Development*, 4(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(3), 352–376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129(2), 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(2), 195–205.
- Huttenlocher, P. R. (1990). Morphometric study of human cerebral cortex development. *Neuropsychologia*, 28(6), 517–527.
- Iossifidis, I., & Schöner, G. (2006, October). *Dynamical systems approach for the autonomous avoidance of obstacles and joint-limits for an redundant robot arm*. Paper presented at the IEEE International Conference on Intelligent Robots and Systems (IROS), Beijing, China.
- Jackendoff, R. (1996). The architecture of the linguistic-spatial interface. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garret (Eds.), *Language and space* (pp. 1–30). Cambridge, MA: MIT Press–Bradford Books.
- Jancke, D., Erhlagen, W., Schöner, G., & Dinse, H. R. (2004). Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex. *Journal of Physiology (Paris)*, 556(3), 971–982.
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (in press). A dynamic neural field model of visual working memory and change detection. *Psychological Science*.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (2006). *A dynamic neural field theory of multi-item visual working memory and change detection*.

- Paper presented at the 28th Annual Conference of the Cognitive Science Society, Vancouver, British Columbia, Canada. Retrieved October 21, 2008, from <ftp://ftp.neuroinformatik.rub.de/pub/manuscripts/articles/JohnsonSpencerSchoner2006.pdf>
- Johnson, J. S., Spencer, J. P., & Schöner, G. (2008). Moving to a higher ground: The dynamic field theory and the dynamics of visual cognition. In F. Garzón, A. Laasko, & T. Gomila (Eds.), *Dynamics and psychology* [Special issue]. *New Ideas in Psychology*, 26(2), 227–251.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press–Bradford Books.
- Kelso, J. A. S., & Tuller, B. (1984). A dynamical basis for action systems. In M. S. Gazzaniga (Ed.), *Handbook of cognitive neuroscience* (pp. 321–356). New York: Plenum Press.
- 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(2), 101–113.
- Kopecz, S., Engels, C., & Schöner, G. (1993). Dynamic field approach to target selection in gaze control. In S. Gielen & B. Kappen (Eds.), *Proceedings of the International Conference on Artificial Neural Networks, Amsterdam* (pp. 96–101). Berlin, Germany: Springer.
- Lipinski, J., Spencer, J. P., & Samuelson, L. K. (in press). It's in the eye of the beholder: Spatial language and spatial memory use the same perceptual reference frames. In L. B. Smith, M. Gasser, & K. Mix (Eds.), *The spatial foundations of language*. New York: Oxford University Press.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–284.
- Mareschal, D., Johnson, M., Sirois, S., Spratling, M., Thomas, M. S. C., & Westermann, G. (2007). *Neuroconstructivism: Vol. 1. How the brain constructs cognition*. New York: Oxford University Press.
- Morton, J. B., & Munakata, Y. (2002). Active versus latent representations: A neural network model of perseveration, dissociation, and decalage. *Developmental Psychobiology*, 40(3), 255–265.
- 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(1), 109–116.
- Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM*, 19(3), 113–126.
- Palmer, S. E., & Kimchi, R. (1986). The information processing approach to cognition. In T. J. Knapp & L. C. Robertson (Eds.), *Approaches to cognition: Contrasts and controversies* (pp. 37–77). Hillsdale, NJ: Erlbaum.
- Plumert, J. M., & Hund, A. M. (2001). The development of memory for location: What role do spatial prototypes play? *Child Development*, 72(2), 370–384.
- Port, R. F., & van Gelder, T. (Eds.). (1995). *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press–Bradford Books.
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science*, 276(5313), 821–824.
- Rumelhart, D. E., McClelland, J. L., & the PDR Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. Foundations*. Cambridge, MA: MIT Press–Bradford Books.
- Salinas, E. (2003). Self-sustaining activity in networks of gain-modulated neurons. *Neurocomputing*, 52–54, 913–918.
- 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–Bradford Books.
- Samuelson, L. K., Schutte, A. R., & Horst, J. S. (in press). The dynamic nature of knowledge: Insights from a dynamic field model of children's novel noun generalizations. *Cognition*.
- Schöner, G. (1995). Recent developments and problems in human movement science and their conceptual implications. *Ecological Psychology*, 7(4), 291–314.
- Schöner, G., & Dineva, E. (2007). Dynamic instabilities as mechanisms for emergence. *Developmental Science*, 10(1), 69–74.
- Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychology Review*, 113(2), 273–299.
- Schutte, A. R. (2004). A developmental transition in spatial working memory. Unpublished doctoral dissertation, University of Iowa, Iowa City.
- Schutte, A. R., & Spencer, J. P. (2002). Generalizing the dynamic field theory of the A-not-B error beyond infancy: Three-year-olds' delay- and

- experience-dependent location memory biases. *Child Development*, 73(2), 377–404.
- Schutte, A. R., & Spencer, J. P. (2008). Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. Manuscript submitted for publication.
- 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., Johnson, J. S., & Spencer, J. P. (2008). Does change detection underestimate capacity? Insights from a dynamic field theory of visual working memory. Manuscript in preparation.
- Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. In S. Becker (Ed.), *Computational cognitive neuroscience* [Special issue]. *Brain Research*, 1202(C), 68–86.
- Simmering, V. R., & Spencer, J. P. (2008). Generality with specificity: The dynamic field theory generalizes across tasks and time scales. *Developmental Science*, 11(4), 541–555.
- Simmering, V. R., Spencer, J. P., & Schöner, G. (2006). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception and Psychophysics*, 68(6), 1027–1046.
- Skarda, C. A., & Freeman, W. J. (1987). How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10(2), 161–173, commentary 174–183, authors' response 183–192.
- Smith, L. B., & Thelen, E. (Eds.). (1993). *A dynamic systems approach to development: Applications*. Cambridge, MA: MIT Press–Bradford Books.
- 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(2), 235–260.
- 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(22), 8819–8829.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, 131(1), 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., & Perone, S. (2008). Defending qualitative change: The view from dynamical systems theory. *Child Development*, 79, 1639–1647.
- Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamical systems approach to development [Special section/issue]. *Developmental Science*, 6(4), 392–412.
- 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 and Performance*, 32(2), 473–490.
- Spencer, J. P., Simmering, V. R., Schutte, A. R., & Schöner, G. (2007). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In J. M. Plumert & J. P. Spencer (Eds.), *The emerging spatial mind* (pp. 320–361). New York: Oxford University Press.
- Spivey, M. J. (2007). *The continuity of mind*. New York: Oxford University Press.
- Steinhage, A., & Schöner, G. (1997). Self-calibration based on invariant view recognition: Dynamic approach to navigation. *Robotics and Autonomous Systems*, 20(2–4), 133–156.
- Steinhage, A., & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. In P. S. Schenker & G. T. McKee (Eds.), *Proceedings of SPIE: Vol. 3523. Sensor fusion and decentralized control in robot systems* (pp. 160–180). Bellingham, WA: SPIE Publishing.
- Thelen, E. (1995). Motor development: A new synthesis. *American Psychologist*, 50(2), 79–95.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24(1), 1–34, commentary 34–70, authors' response 70–80.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press–Bradford Books.
- Thelen, E., & Ulrich, B. D. (1991). Hidden skills: A dynamic systems analysis of treadmill stepping

- during the first year. *Monographs of the Society for Research in Child Development*, No. 223, 56(1), 1–98, discussion 99.
- Townsend, J. T., & Busemeyer, J. (1995). Dynamic representation of decision making. In R. F. Port & T. van Gelder (Eds.), *Mind in motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.
- van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stages of cognitive development: An application of catastrophe theory. *Psychological Review*, 99(3), 395–417.
- van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, Vygotsky, and beyond. *Psychological Review*, 105(4), 634–677.
- van Gelder, T., Bechtel, W., Beer, R. D., Braisby, N., Cooper, R., Franks, B., et al. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628, commentary 629–654, author's response 654–661.
- Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12(1), 1–24.
- Zelazo, P. D. (2006). The dimensional change card sort (DCCS): A method of assessing executive function in children. *Nature Protocols*, 1, 297–301.
- Zelazo, P. D., Müller, U., Frye, D., Marcovitch, S., Argitis, G., Boseovski, J., et al. (2003). The development of executive function in early childhood. *Monographs of the Society for Research in Child Development*, 68(3), 1–137.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235.