

# Moving Toward a New Grand Theory While Valuing the Importance of the Initial Conditions

John P. Spencer<sup>1,2</sup>, Evelina Dineva<sup>1,2</sup> & Gregor Schöner<sup>3</sup>

<sup>1</sup>Department of Psychology, University of Iowa

<sup>2</sup>Iowa Center for Developmental and Learning Sciences

<sup>3</sup>Institute for Neuroinformatics, Ruhr-University, Bochum, Germany

Dynamical systems and connectionist approaches to development are fundamentally complementary (Smith & Samuelson, 2003). The two communities that have grown up around these approaches share a basic outlook on development, share a conception of what understanding development entails, and are committed to joint principles. These principles include a firm grounding in neural function, an awareness of the role structured environments play during development, a mistrust of static architectural accounts which both communities view as incompatible with learning and development, and a deep appreciation for the step-by-step, emergent, and self-organizing nature of developmental processes.

Our impression of convergence between these two approaches was not seriously destabilized at the conference from whence this volume originated, even though disagreements were openly and sometimes controversially discussed. Indeed, in several instances, we felt the disagreements that did arise were based on misperceptions rather than deep differences. That said, several “real” differences did arise, and both groups articulated their positions on issues that ultimately require clarity and resolution. We also discovered that there are sometimes subtle differences in outlook that are not readily perceived—differences in the role that theory and models play that, in the extreme case, can lead people to use the same mathematical equations to quite different ends. Such differences reside at a metatheoretical level and have practical consequences for what we communicate to the field of developmental science more generally.

We left the conference inspired and with the sense that the stuff of “real” development is well within grasp as our two communities work to understand development, to understand each other, and to ultimately capitalize on our complementary strengths. As we look toward this integrative future, it will be critical to establish a deep appreciation for the initial conditions, that is, to articulate and value the current strengths of each approach. This will foster a deep sense of integration that will push both ways of thinking to new levels. In this spirit, our “reactions” chapter works to articulate and clarify cases of perceived differences between approaches that, in our view, are *not* real, as well as cases of perceived differences that are real. We conclude by discussing the implications of efforts to integrate the two approaches for developmental science more generally. In our view, clarifying similarities and differences between approaches offers far more than just technical clarity for “modeler types”—it truly offers a vision of a new, integrative, developmental theory.

## **Perceived differences that are not real**

In this section, we discuss two perceived differences that, in our estimation, are not real differences between dynamical systems theory (DST) and the connectionist (CN) approach: the neural commitment of each approach and the degree to which each approach addresses the processes that underlie learning.

### *How neural are CN and DST?*

Connectionism is explicitly based on a neuronal framework. The very definition of the connectionist program refers to this: parallel distributed processing, neural units that integrate inputs across synaptic weights, architectures and connectivity patterns that determine neural function, and so on.

Less clear is whether DST shares this neural commitment. This is understandable given that several flavors of DST have emphasized the “collective” behavior of organisms, as in the cases of movement coordination (Schöner & Kelso, 1988), motor development (Spencer et al., 2006), early language development (van Geert, 1991, 1993), and educational theory (Fischer, in press). Such efforts to understand the dynamics of collective behavior often over a developmental time scale have resulted in abstract dynamical models that capture, for instance, the rate of change in relative phase as people oscillate their fingers (Schöner, Haken, & Kelso, 1986) or the complex behavior of systems consisting of multiple “growers” that change in complex ways over developmental time (Van Geert, 1998).

Although it is clear that not all flavors of dynamic systems theory have been firmly grounded in neural principles, some have. Our Dynamic Field Theory (DFT) is an example. The dynamic field framework was originally developed to capture the dynamics of neural activation in visual cortex (Amari, 1977). More recently, this framework has been extended to account for the processes that underlie saccadic eye movements (Kopecz & Schöner, 1995; Wilimzig, Schneider, & Schöner, 2006), motor planning (Erlhagen & Schöner, 2002; Schutte & Spencer, 2007), the dynamics of neural activation in visual and premotor cortex (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schöner, & Riehle, 2003), the behavior of autonomous robots (Bicho, Mallet, & Schöner, 2000; Iossifidis & Schöner, 2006; Steinhage & Schöner, 1998), infants’ performance in Piaget’s A-not-B task (Thelen, Schöner, Scheier, & Smith, 2001), and developmental changes in visuo-spatial cognition (Simmering, Schutte, & Spencer, in press; Spencer, Simmering, Schutte, & Schöner, 2007).

Reviews of the dynamic field approach can be found in the chapters by Schöner (chapter 2) and Spencer (chapter 5). For the purposes of the present issue—the neural commitment of DST—we focus on one particular example.

Erlhagen and Schöner (2002) proposed a dynamic field theory of movement preparation. One of the central concepts in this theory is that movement preparation can be “preshaped” by either perceptual pre-cues or by a memory trace of past motor decisions (a concept which also underlies performance in other tasks in early development, see Spencer & Hund, 2002, 2003; Thelen et al., 2001). For instance, in the case of perceptual pre-cues, the theory predicted that patterns of neural activation at the population level would be centered around pre-cued locations. If the pre-cued pattern was relatively precise, this would form a sharp distribution of activation associated with the average, pre-cued location. When the target cue was given, this sharp distribution could quickly shift in the direction of the target, leading to fast RTs and accurate performance with sufficient planning time. Distributed pre-cues, however, would lead to a distributed activation pattern with consequently slower RTs and lower accuracy under certain testing conditions (for successful tests of these claims, see Schutte & Spencer, 2007).

Although Erlhagen and Schöner (2002) demonstrated that the DFT could capture behavioral patterns from a variety of paradigms, more direct support for the “preshape” concept came from neurophysiological evidence. In particular, time-dependent changes in neural activation in the DFT were compared to single-unit neural activity in the premotor cortex of

Rhesus monkeys measured in a pre-cue paradigm using population coding techniques (Bastian et al., 1998; Bastian et al., 2003; Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999). The first step in making this comparison was to map the responses of neurons to basic stimuli and create a continuous field by ordering the neurons based on their “preferred” stimulus directions. This was followed by a behavioral pre-cuing task that probed whether neural activation was “preshaped” as predicted by the DFT (Bastian et al., 1998; Bastian et al., 2003; Erlhagen & Schöner, 2002). Results confirmed this prediction, showing a robust relationship between the dynamics of neural activation in real and simulated cortical fields.

Tests of the DFT of movement preparation clearly show that dynamic systems theories can be tightly linked to neural data. And there are many such examples from our own group (Bastian et al., 2003; Erlhagen et al., 1999; Jancke et al., 1999) as well as groups using related theoretical concepts (Pouget, Dayan, & Zemel, 2000; Rolls, Stringer, & Trappenberg, 2002). We also note that links between dynamic fields and neural data are not confined to single- or multi-unit neurophysiology. For instance, the DFT of movement preparation has been tested using ERP techniques (McDowell, Jeka, Schöner, & Hatfield, 2002). Moreover, recent ties between our dynamic field approach and more biophysical approaches to working memory suggest that neural measurements can also be used to *test theories of developmental change* (see Simmering et al., in press).

The dynamic field theory shows that dynamical systems approaches can have a strong neural commitment. Given this, where does DFT sit in the broad collection of dynamic systems approaches? Put differently, what is the link between DFT and other dynamical systems approaches that emphasize the “collective” behavior of organisms?

Central to DFT are the attractor states in which patterns of neural activation live. Much of our understanding of these attractor states comes from initial work by Amari (1977) who provided a quasi-formal analysis of the attractor layout of dynamic neural fields. Our work on spatial cognition, for example, has utilized Amari’s analysis, emphasizing three particular attractors: a resting state, a self-sustaining state, and a self-stabilized state where peaks of activation remain stable only in the continued presence of input (see Spencer & Schöner, 2003; Spencer et al., 2007). We have a firm grasp on the parameters of the model that move it into and out-of these states; we understand the bifurcations that arise when one state is de-stabilized and another state is created; and we have demonstrated how these attractor states give rise to observable behavioral and neural properties.

Importantly, this emphasis on attractors, bifurcations, stability/instability, and so on, has real consequences for the behavior of the systems we develop. As an example, we contend that stability is a fundamental requirement for *any autonomous, real-time behavioral system*. If this is the case, then understanding the attractor states of a neural network becomes a fundamental requirement for any theory of behavior. We discuss this claim in greater detail below.

Our emphasis on dynamical systems concepts also connects DFT at a deep level with other dynamical systems approaches that emphasize these concepts at other levels of analysis and at longer time scales (Newell & Molenaar, 1998; van der Maas & Molenaar, 1992; Van Geert, 1998). In this context, we emphasize that not all approaches—DST, CN, or otherwise—need to be tightly linked to the neural level of analysis. Rather, the level of analysis should depend on the questions being asked and the phenomena under study (see van Geert & Fischer, chapter 16).

Critically, however, dynamical systems concepts provide a link among these different behavioral and measurement levels. Our work on the Piagetian A-not-B error and related

phenomena illustrates how this link can be quite generative. The dynamical systems approach to the A-not-B error was motivated largely from a conceptual analysis of infants' performance in the classic Piagetian task using the concepts of DST (see Thelen & Smith, 1994). This conceptual analysis led to a host of novel behavioral predictions (Smith, Thelen, Titzer, & McLin, 1999) which were collected and formalized in the dynamic field theory of infant perseverative reaching (Dineva, 2005; Thelen et al., 2001). This theory, in turn, enabled strong ties between work in infancy and work in later development on children's spatial working memory abilities (Simmering & Spencer, in press). And there may be ways to test the developmental hypotheses that have emerged from this body of work in the near future using fMRI (see Edin, Macoveanu, Olesen, Tegnér, & Klingberg, 2007).

What, ultimately, will this move toward neural data buy us? In our view, moving to this level provides more constraints for theory and novel ways to test theoretical claims. Work by Bastian and colleagues (Bastian et al., 1998; Bastian et al., 2003; Erlhagen et al., 1999) described above provides an example of this use of neurophysiology—this measurement level provided direct evidence for the preshape concept, testing in precise detail whether our population dynamics account offered an accurate view of the brain's role in the behavioral pattern under study.

Thus far, we have described how dynamical systems theory can have a strong neural commitment. In this context, we can also ask how committed connectionism is to neural principles. Real neurons in the higher nervous system typically have 10,000 synapses, not the limited number that most connectionist models never exceed. Moreover, real neural networks are overwhelmingly recurrent. Braitenberg (Braitenberg & Schüz, 1991) has estimated that more than 95% of all neuronal connections are parts of recurrent loops rather than of feed-forward paths. The simplifications commonly used in connectionist networks—a small number of neurons with 50 or so connections—seem like reasonable simplifications of this complexity until one considers that well known learning algorithms do not generalize to more complex networks with high degrees of interconnectivity (Braitenberg & Schüz, 1991; Douglas & Martin, 2004). The status of these simplifications is important in determining what CN tells us about how brains really work.

Furthermore, in contrast to the tight demonstrated links between neurophysiological recording and activation in dynamic neural fields (Bastian et al., 1998; Bastian et al., 2003; Erlhagen et al., 1999), connectionist networks have been linked to neural data primarily at the macroscopic level, where a few units of computation stand for activation in entire brain areas. The connectivity of the network then reflects the macro-connectivity between such areas (Cohen, Braver, & O'Reilly, 1996; Cohen & Servan-Schreiber, 1992; O'Reilly, Braver, & Cohen, 1999). This is a useful approach and a timely one, given the enormous amount of data emerging from neuroimaging labs. But it highlights that connectionist networks may provide more abstract descriptions of neurophysiology than neuronally-grounded variants of dynamical systems thinking.

In summary, although there is a sense in the literature that CN has a stronger neural commitment than DST, this is not the case. There are differences in how both approaches have been linked to neural data, but both approaches have a shared commitment to neural function. It will be useful to clarify the differential strengths/weaknesses of each approach in this regard as we collectively try to integrate brains, behavior, and development.

Critically, however, we must place this future work in the right context. Such work is geared toward refinement of theoretical concepts, tighter and richer links to neural data, and so

on. It is not about determining whether one group is “more” committed to neural function than the other. We feel that the commitment to a grounding of perceptual, motor, and cognitive function in principles of neuronal function is a shared and deep conviction between CN and DS approaches. We are much closer here than it first appears. In fact, when contrasted with the still largely dominant information processing approach and good-old-fashioned cognitive development (see chapter 13), our differences appear negligible.

#### *How much do CN and DST address learning?*

As with claims about neural grounding, there exist misconceptions regarding learning. In particular, DS thinking is viewed as relatively mute on learning while CN models are assumed to confront learning head on.

We certainly agree that the CN approach has led to major advances in our understanding of particular classes of learning. For instance, the CN approach has fostered a revolution in thinking about statistical learning and the ability of neural systems to extract statistical patterns from complex and rich inputs. This work has had a particularly strong impact within the domain of language (e.g., Altmann, 2002; Elman, 1990; MacDonald & Christiansen, 2002), and stands in opposition to nativist claims about language acquisition (Elman et al., 1997; Spencer et al., in press).

But this is only one aspect of learning. As we highlight below, dynamical systems approaches to learning have much to offer that is complementary to insights revealed by the CN approach. In particular, DST offers a path toward process models of learning that explain how learning is constrained by an organism’s own activity and ability to determine its own stimulation. By this view, learning is not solely about the statistics of input; learning is also about the statistics of experience, of internal states, of the systems’ own autonomous decisions, all occurring within the context of the systems’ intrinsic dynamics.

Dynamical systems thinkers have examined learning within a variety of domains. We highlight two here: motor learning and spatial cognition. Within the domain of motor learning, there have been at least two dynamical systems approaches to learning. The first is summarized by Newell and colleagues in chapter 6; thus, we focus solely on a second DS approach to learning that emerged from the study of coordination dynamics by Schöner, Kelso, and Zanone (1992).

Schöner and colleagues examined learning within the “finger twiddling” paradigm where participants move their fingers at a particular relative phase (e.g., in-phase [ $0^\circ$ ] where the index fingers move toward the palm in synchrony, or anti-phase [ $180^\circ$ ] where one finger moves toward the palm while the other moves away). A series of empirical and theoretical analyses of these coordination patterns revealed that they showed key “signatures” of a dynamical system with attractors at  $0^\circ$  and  $180^\circ$  (Scholz & Kelso, 1989; Schöner & Kelso, 1988). Within this context, the learning question was simple: what happens as people learn a new pattern of coordination, for instance,  $90^\circ$  relative phase?

To examine this question, Schöner, Kelso, and Zanone (1992) first had to specify what it is that is being learned as people acquire a new coordination pattern. Schöner (1989) proposed that learning can be conceptualized as a dynamic process that operates on the coordination dynamics but at a slower time scale. This means the people should be learning about the *dynamics* of the coordination pattern and not simply about the specific, target coordination pattern ( $90^\circ$  relative phase). Moreover, because the dynamics at the learning time scale have their

own stability properties, there can be dramatic, qualitative transitions over learning which lead to large fluctuations at specific times during the learning process.

Empirical observations confirmed these insights (Zanone & Kelso, 1992, 1997). As people learned the new coordination pattern (90° relative phase), the entire landscape of coordination patterns changed: the anti-phase pattern (180°) became unstable and the attractor landscape showed a qualitative shift toward the new pattern (90°). This qualitative shift was accompanied by large fluctuations that were distinct from fluctuations at the faster time scale of real-time behavior. Thus, learning is about changing the dynamics of the system rather than learning the details of each individual pattern; as a new attractor is created, others can be destabilized (see also Schöner & Dineva, 2007).

A second dynamical systems approach to learning has emerged within the framework of Dynamic Field Theory. This approach to learning emerged from a fundamental observation. Organisms determine their own stimulation by exploring their surrounds, looking at different things, moving through the world, and so on. These activities strongly constrain the external stimulation they receive. But the organisms' own *internal* activity—the second-to-second decisions they make about what to attend to, what is task relevant, what must be remembered—also places strong constraints on what is learned.

The DFT approach captures this sense of active learning as our work on infants' perseverative reaching illustrates. Based on an analysis of data from hundreds of infants who participated in the Piagetian A-not-B task, Dineva and colleagues (Dineva, 2005; Dineva, Schöner, & Thelen, 2008) have demonstrated that the trial-by-trial decisions infants make in this task have a profound impact on their performance on the critical B trial. In particular, infants who spontaneously reach to the B location on an A trial have a significantly weaker tendency to reach back to A on a subsequent B trial. Thus, infants' own spontaneous decisions—and not just the statistics of input—structure what is learned as well as infants' ability to make subsequent decisions.

Within the dynamic field framework, we have captured these findings using a simple form of Hebbian learning implemented using memory traces. In particular, when a peak of activation emerges in, say, a motor planning field that makes decisions about which lid to pick up in the A-not-B task, it leaves a trace of activation in a “preshape” or long-term memory (LTM) field (Dineva et al., 2008; Simmering & Spencer, in press). Such traces build up over a slower time scale than the activation dynamics within the planning field, and can become quite focused as similarly positioned peaks form from trial to trial. Thus, if the system builds a peak associated with an A location over and over, a long-term memory trace associated with A will also form.

Critically, activation in LTM fields is coupled back into the field of origin. In this way, long-term traces can bias which decisions or peaks are formed in the motor planning field on subsequent trials. This is the mechanism that underlies infants' perseverative reaching in the A-not-B task. As 8- to 10-month-old infants repeatedly reach to the A location, they build up stronger memory traces which make it more likely that they will reach to A again on subsequent trials, that is, they form a habit. When the hiding location is switched to B, however, the strong traces associated with A can lead to perseverative reaching if infants are not able to sustain activation associated with responding to B. Note that this is less likely to occur if the memory traces associated with A are weaker and there is some trace associated with B—precisely what occurs when infants make spontaneous errors (i.e., reaches to B on an A trial) in the A-not-B task.

The tendency to perseveratively reach to A in the canonical A-not-B task disappears around 12 months as infants develop the ability to sustain an active or “working” memory for B in the context of the perceptual cues available in the task. Nevertheless, the same long-term memory mechanism leads to systematic trial-to-trial distortions between 2 and 6 years when toys are hidden in a sandbox instead of beneath visibly marked hiding locations (Schutte, Spencer, & Schöner, 2003; Spencer, Smith, & Thelen, 2001). In this case, children’s reaches to a B location in a sandbox are biased in the direction of a nearby A location after a short delay.

Given that our LTM fields implement a form of Hebbian learning, what is the link between this learning mechanism and Hebbian learning in connectionist networks? In many cases, these different forms of Hebbian learning lead to the same outcomes. For instance, just as strengthening synapses in a connectionist network makes it more likely that a stimulated neuron will excite other neurons to which it is strongly connected, LTM traces in the DFT boost activation within a local region of a field, making it more likely that a local group of neurons will excite one another to, for instance, form a stabilized activation peak.

Similarly, just as Hebbian learning can lead to the extraction of statistical patterns from input, memory traces in the DFT can reflect the statistics of past decisions (i.e., past peaks). For instance, we have shown that such traces produce a bias toward the average responded-to location in spatial planning and memory tasks (Schutte & Spencer, 2007; Spencer & Hund, 2002; Spencer et al., 2007). Traces are sensitive to frequency information as well. If, for instance, one location in a set of three close locations is more frequent than the others, spatial recall responses will not be distorted toward the average location; rather, recall responses will be pulled toward the frequent location (Spencer & Hund, 2002). Note that there are some differences between our implementation of Hebbian learning and Hebbian learning in connectionist networks. These differences are discussed in the section that follows.

In conclusion, the examples of DS approaches to learning highlighted in this section illustrate two points. First, we believe there is a strong, shared commitment to understanding learning in the CN and DS approaches. Second, we think both approaches offer unique insights into learning. With regard to the DS approach, dynamical systems concepts offer insights into the process of learning, usefully confronting classes of learning that extend beyond the boundaries of statistical learning. We think such efforts are needed to capture aspects of *active learning* that most certainly play an important role over development.

## **Perceived differences that are real**

Although there is a shared commitment to the issues of neural grounding and learning, DST and the CN approach do differ in important ways. We discuss six “real” differences below, highlighting issues that we think must be addressed head-on as we collectively build on the differential strengths of the two approaches.

### *Learning the statistics of decisions versus inputs*

Although the learning mechanism we use in DFT implements a form of Hebbian learning, because this learning is a direct result of field dynamics—the formation of an activation peak—it is not identical to the type of Hebbian learning implemented in most connectionist networks. Put simply, learning in DFT is a reflection of *past decisions* rather than *past inputs*. This is not a trivial distinction.

Consider the dynamic field model of the A-not-B error described above. There are two central points of contrast between this model and the PDP model proposed by Munakata (Munakata, 1998) with regard to learning. First, the way the output of the PDP model is mapped to infants' performance in the A-not-B situation is to interpret the level of activation of each "response" node as a probability of responding. Thus, the PDP model does not make a response on each trial per se. Second, the "responses" generated on each trial do not feedback and affect synaptic weights in the network. Because of these features of the model, it cannot account for spontaneous errors and the effect such errors have on perseveration.

This contrasts with the dynamic field model proposed by Dineva et al. (Dineva, 2005; Dineva et al., 2008) where an active decision—a stabilized peak of activation—is generated at the end of each trial. This difference in "read-out" has consequences: because the DFT forms a stabilized peak of activation corresponding to the A or B location at the end of each trial, the model generates a robust trace of each decision. Such traces are necessary to explain why spontaneous errors in the A-not-B task have such a dramatic effect on infants' perseverative reaching. The alternative—accumulating small fluctuations in the activation of response nodes from trial-to-trial—does not create enough of a bias toward B to account for the data and, in some cases, to override the statistics of input (see Dineva et al., 2008).

A second example from spatial cognition highlights another way in which Hebbian learning in DFT and CN models differ. Spencer and colleagues have shown that spatial recall responses are biased away from perceived reference frames during memory delays. If, for instance, people are asked to remember a location to the left of a visible line, they exaggerate the "leftness" of the location after a short memory delay (see Simmering & Spencer, 2007). Critically, several lines of evidence suggest that these spatial recall biases reflect repulsion away from perceived reference axes (rather than attraction toward spatial prototypes as proposed by Huttenlocher, Hedges, & Duncan, 1991; see Spencer et al., 2007 for discussion). This type of bias creates an interesting case study because working memory is *drifting in a direction with no clear input*. Importantly, results have shown that these drift effects are superimposed on long-term memory biases. For instance, biases over learning in spatial recall tasks are not precisely toward the average remembered location. Rather, biases over learning can be systematically distorted away from reference axes in the task space (Schutte et al., 2003). This provides empirical evidence that the long-term memory traces learned by participants are a reflection of distortions that occur in working memory rather than the statistics of the input distribution.

A final example supports another aspect of our account of learning—that memory traces can directly affect the dynamics of decision-making in dynamic fields, leading to non-linear changes over learning. As memory traces build in the DFT, they effectively elevate the overall level of excitation within a local region of the field. Such increases in activation can make it more likely to form self-sustaining activation peaks, that is, to move from one state—an input-driven state—to a qualitatively different state—the self-sustaining or "working memory" state. This has been empirically demonstrated recently by Dineva and Smith (2008) who used distributed memory traces to make it more likely that infants would be able to sustain a decision to reach to B in the A-not-B task. The result: with distributed practice, 8- to 10-month-olds were able to sustain a memory of the B location during a 3-5 s delay, that is, during delays that typically produce perseverative responses. Note that this example parallels work on learning and coordination dynamics by (Schöner, 1989) in that learning in this case directly affects the dynamics of the system (see also, Schöner & Dineva, 2007).

### *Autonomy and stability*

One generally perceived difference between the DST and CN approaches is that dynamicists have focused more on the real-time behavior of organisms in time varying and structured environments. We think that this perception is right.

Organisms do not sit around waiting to receive a stimulus to give a response. They actively generate behavior and, thus, control their own sensory stimulation. The resulting continuous flow of behavior, which we like to think of as an expression of the fundamental autonomy of organisms, is not simply understandable as a form of stimulus-response chaining. Instead, the current behavioral state and recent behavioral history both matter and impact on upcoming behavioral decisions. The dynamical systems approach is centrally concerned with such autonomy and the associated real-time and closed-loop coupling of an organism to its environment (for related concerns, see Turvey, 1990; Warren, 2006).

Any system that controls its own sensory information must have stability properties that ensure that small changes in sensory information lead to small changes in behavior and vice versa. We know this in the technical domain since the inception of control theory (Mees, 1981). Nevertheless, this fundamental constraint is often overlooked in models that address brain-behavior relations.

An insightful example comes from our own work. In the original dynamic field model of perseverative reaching in young infants (Thelen et al., 2001), activation evolved under the influence of various inputs (e.g., the presentation of a toy, “vision” of lids in the task space, etc.). To generate a response at the end of the delay, the location in the field with maximal activation was mapped onto the reaching target that was closest to that field location. This is a standard “read-out” procedure used in many other approaches that involve piercing a threshold of activation to make a response (Schall, 2004).

When we embodied this model on an autonomous robot, however (something pushed for by the ever-insightful Esther Thelen), we discovered that our read-out procedure was missing something fundamental—stability! Because the robot was continuously linked to perceptual input and motor output, the representation of the movement direction after the delay—that is, the maximal activation value in the field that specified whether the robot should turn left or right to “acquire” the target—was pushed around by all sorts of things. The result: the robot would change course erratically, wiggling its way somewhere in-between the two targets, and not reaching either.

The solution was simple: we needed to stay true to the dynamical systems concepts we had proposed. In particular, rather than reading out the maximal activation in the field after the delay—something that the robotic implementation showed was intrinsically unstable—we needed to have the model generate a *stable* decision to reach to A or B. This led to the next iteration of our account of infant perseverative reaching (Dineva et al., 2008). As described previously, this new model generates a stable decision—a stabilized activation peak—at the end of each trial due to a boost in the inputs (which corresponds to pushing the box toward the infant in the A-not-B task at the end of each trial). And recent work with an autonomous robot shows that this works: our robotic implementation now effectively captures the details of decision-making in the A-not-B situation (Dineva, Faubel, & Schöner, 2007).

Note that there was an important corollary of this work with regard to learning: stable decisions also enable organisms to build up different behavioral histories depending on the choices they make. As described previously, small fluctuations in the activation of response nodes do not generate sufficient differences across trials to explain how infants’ own decisions

impact subsequent behavioral events. Stable decisions, by contrast, are an effective way to structure trial-to-trial learning, and simulations of the new DFT explain how the behavioral history of individuals in the A not B task predicts the extent to which they persevere (for related results, see Clearfield, Smith, Diedrich, & Thelen, 2006; Dineva et al., 2008; Schöner & Dineva, 2007).

### *Embodiment*

A prerequisite for autonomy is that cognitive systems can be interfaced with the sensory and motor surfaces of a body and that cognitive processes can be acted out by such an embodied system in a real structured environment. We believe that DST and CN differ in the extent to which they enable, and are committed to, embodiment.

The best way to prove that a model may be embodied and situated is to implement it on an actual body—a robotic device—with real-time sensory and motor systems, and to have the robotic system generate behavior in a natural environment. This has been done extensively within the realm of DST. In fact, an approach to autonomous robotics inspired by the principles of DST has been developed and demonstrated in robotic navigation tasks (Bicho et al., 2000; Schöner & Dose, 1992; Schöner, Dose, & Engels, 1995; Steinhage & Schöner, 1997), tasks involving robotic cooperation (Large, Christensen, & Bajcsy, 1999; Soares, Bicho, Machado, & Erlhagen, 2007), in tasks involving autonomous reaching by robot arms (Iossifidis & Schöner, 2006), and in tasks that involve the fast learning and recognition of objects in the work space of a robotic arm (Faubel & Schöner, in press). This approach literally employs the same attractor dynamics and their instabilities used to model the behavior of human infants, children, and adults to generate robotic behavior acted out using conventional robotic hardware in real-world situations.

Connectionist models have proved more difficult to implement in robotic devices (but see Schlesinger, chapter 9; Lovett & Scassellati, 2004; Sirois, 2005). One reason is that connectionist models commonly are interfaced with sensory and motor systems in relatively abstract ways. For instance, Munakata and Morton (chapter 7) use a discrete set of input neurons to represent the details of lids and toys in their model of the A-not-B task. Similarly, they use a discrete set of input neurons to represent objects such as trucks or flowers in their model of the Wisconsin Card Sorting Task. As another example, the categorization model of McClelland and Vallabha (chapter 1) has input units representing plants or animals. These interfaces essentially presuppose that an object recognition system extracts the appropriate categorical information from raw sensory input to represent these objects. The interface between model and sensory surface, thus, presupposes that fundamental problems of vision such as segmentation and invariant object recognition are somehow solved. By contrast, we have demonstrated that dynamic neural field models can work with relatively “raw” sensory inputs, even in cases where some form of “higher-level” learning must take place such as learning the names for objects (Faubel & Schöner, in press)

At the other end, single neurons in CN models sometimes represent planned actions. When the activity of these neurons exceeds a threshold, the response is assumed to be generated. But as our robotic implementation of the A-not-B task showed, this form of read-out falls far short of the temporal and spatial dynamics required to control an actual effector system. To realize the real-world control of an effector, one must tackle the problems of initiating action, stabilizing decisions even when the sensory information that led to the decision is lost during movement, controlling the effector, and updating motion plans in closed loop.

Although simple input and output neurons might be a reasonable simplification for an abstract model of behavior, one cannot overlook such problems if the demonstration of embodiment is a goal (which, to be clear, is not an explicit goal of many CN approaches). Interestingly, when connectionist models have been implemented in robotic systems, they start to look remarkably like dynamical systems models (see, e.g., Sirois, 2005). For instance, Schlesinger and colleagues (Schlesinger, Amso, & Johnson, 2007) proposed a model for visual selective attention in infants that was applied to real world images. The raw image stream was first filtered to produce feature maps (intensity, oriented edges, motion, and opponent-color contrast). In each feature dimension, excitatory short-range and inhibitory long-range connections worked for several iteration steps. Then, the sum of feature intensities was used to form a saliency map. For read-out, a stochastic process mapped saliency to a probability distribution for fixation. Given that neural representations in this model operate over continuous feature dimensions via short-range excitation and long-range inhibition, these models are similar to the types of dynamic neural fields used in our own work. Thus, this example highlights how there can be clear convergence between the two approaches when there is a shared commitment to embodiment.

Although convergence is possible, embodiment has certainly been more of a focus within DST than in the CN approach. The critical question is whether this difference matters. Put differently, is it acceptable to use abstract input-output nodes for the sake of simplification, or does this simplification miss something critical? We suspect that this simplification does hide some important aspects of cognitive and neural systems, even when applied to issues in higher-level cognition. This suspicion comes from our work developing a formal account of “embodied” cognition within dynamic field theory.

Our use of stabilized activation peaks as a form of “working” memory provides one example (Spencer & Plumert, 2007). This form of working memory differs in several ways from CN instantiations of working or “active” memory. Active memory traces are transient in nature, slowly returning to a resting level. Such transients have limited stability and will tend to blend inputs presented to the network provided that sufficient time elapses. Working memory peaks in DFT, by contrast, are in a stable attractor state capable of suppressing distracting inputs via lateral inhibition or maintaining separate peaks simultaneously in working memory. This mechanism has behavioral consequences that can be measured in experiment (e.g., Johnson & Spencer, 2008). For instance, we recently developed a dynamic neural field model of multi-item working memory and change detection (Johnson, Spencer, Luck, & Schöner, 2008). This model made the counterintuitive prediction that visual working memory for similar non-spatial features (e.g., two similar colors; two similar orientations) would be enhanced relative to memory for dissimilar items. This effect arises in the model because close peaks share lateral inhibition, effectively sharpening locally excitatory interactions. Most theories of working memory predict the opposite—that there should be interference when similar items are held in memory. Behavioral results confirmed that adults are, in fact, better at detecting changes when they hold similar items in visual working memory.

In what sense is this prediction linked to embodiment? We have explicitly linked the “working memory” construct to self-sustaining peaks in DFT based on our past work implementing neural fields on autonomous robots. This work led to the insight that working memory must be stable to effectively drive real behavior. Importantly, there is a conceptual bridge between working memory in DFT and classic ideas discussed by Baddeley (1986). Baddeley defines working memory as a type of memory used in the service of other cognitive

operations—it does some *work* for the cognitive system. In our view, such work requires stability as a pre-requisite. Just as stability is needed when a movement decision is required to drive a robotic arm in real time, we suspect stability is needed when working memory must help drive a cognitive operation in real time.

We believe that the pressure to examine the embodiment and situatedness of cognition will ultimately be a force of convergence between the two approaches. As models by Schlesinger and colleagues (Schlesinger et al., 2007; see also Schlesinger & Parisi, 2007) and Sirois (2005) attest, embodiment is not outside the purview of the CN approach. But we do encourage connectionists to take a closer look at this issue.

### *Higher-level cognition*

The previous two sections highlighted how the DS approach has been used to examine the issues of autonomy and embodiment, issues that have been explored less thoroughly by connectionists. The state of affairs is mirror reversed in the case of higher-level cognition: the CN approach has been much more thoroughly applied to problems in higher-level cognition. Multiple examples in this book bear this out including work on the acquisition of grammar (Tabor, chapter 9), analogy (Mareschal et al., chapter 11), and language development (McMurray et al., chapter 12).

Although DST has been less thoroughly applied to issues in higher-level cognition, there are some notable efforts in this direction. van Geert and colleagues (van Geert & van Dijk, 2002; see also van Geert & Fischer, chapter 16), for example, has used dynamical concepts to examine individual trajectories in language development. Similarly, van der Maas and colleagues (Jansen & van der Maas, 2002; see also van der Maas & Raijmakers, chapter 7) have used dynamical concepts to understand developmental changes in performance on the balance scale task. And more recent work within the dynamic field framework has examined several issues in early word learning (see Smith, chapter 4; Samuelson, Schutte, & Horst, 2007), including an embodied approach to the fast acquisition of novel words in an autonomous robotics setting (Faubel & Schoner, in press).

Considered in the context of the issues of autonomy and embodiment, one can readily see the complementary nature of the CN and DS approaches: each approach has its respective strengths. A fundamental goal for the future is to clarify these strengths and work to build on them to achieve systems that can develop higher-level forms of cognition yet stay grounded in the real-time details of sensori-motor systems and real-world environments.

### *DST and CN have different scientific histories*

Although we have described the CN and DS approaches as complementary, the presence of differences can be confusing, particularly given overlap in the concepts and mathematical equations we use. In a formal sense, for instance, all recurrent neural networks are discrete time dynamical systems. Similarly, there are instances of connectionist models which are formulated as continuous time dynamical systems all the way through, sometimes using the same equations that dynamicists claim for themselves (e.g. Usher & McClelland, 2001; McClelland & Vallabha, chapter 1). Conversely, dynamic neural fields can be viewed as a special class of recurrent neuronal networks (Amari, 1977; Wilson & Cowan, 1973). If the equations are the same, then why are there differences between these approaches? One reason: due to our unique scientific histories, models play a different role in the two lines of thinking.

Connectionism emerged within the tradition of mathematical psychology and the information processing approach to cognition and development. This tradition places a heavy emphasis on models of cognition and is concerned with the identification of the best-fitting model that is compatible with the broadest set of experimental facts. Note, however, that connectionism made a fundamental contribution to mathematical psychology in that it expanded the range of relevant evidence to include neuroscientific observations at multiple levels.

To illustrate this scientific tradition, consider the model of perceptual decision making proposed by Usher and McClelland (2001; see chapter 1). The mathematics of this model are, essentially, a special case of the generic dynamic field model of Amari (1977) in which only two locations of a field are modeled—the two associated with one perceptual choice or another. The perceptual decision is conceptualized as an attractor. And the stability of the attractor is varied depending on, for instance, the strength of available sensory input.

To test this account, the model has been set in competition with a set of alternative models, including diffusion models. The central question here is which model fits reaction time distributions more effectively and how closely each model maps onto neural observations. A somewhat biased analysis by Smith and Ratcliff (2004) provides an example of this style of thinking and arrives at the conclusion that a diffusion account fits more facts than the competitive accumulator model of Usher and McClelland (2001).

Although this is a valid way to test competing theories, the tradition within DST is to take a more qualitative view of theoretical comparisons. For instance, the competitive accumulator model has some obvious conceptual strengths that are unmatched by the diffusion account. Unlike the decision boundaries invoked in diffusion models, it is easy to see how choices in the competitive accumulator model could reflect learned neuronal representations. Moreover, the diffusion model is inherently non-autonomous (e.g., how is the diffusion process reset to prepare for the next decision? what neuronal processes amplify small changes at a decision boundary into a macroscopic behavioral act?) The competitive accumulator model, in contrast, could easily be extended to address the dynamics of continuous, online behavior.

Importantly, this emphasis on qualitative comparisons between classes of models has deep roots in the application of dynamical systems concepts within psychology. In dynamical systems thinking, the identification of a specific model is a relevant goal only in a limited sense. Providing quantitative fits between a dynamical model and a set of experimental data is primarily an existence proof, showing that there is at least one model that is consistent with all known facts. Typically, however, such fits do not provide a way of uniquely identifying a single model as the only plausible formulation for a given system. Why? The problem is that dynamical systems are strongly non-linear and multi-causal. A whole equivalence class of models may bring about a particular state, and different members of the equivalence class may be used to invariantly realize some function depending on the details of the environmental conditions or behavioral history. This makes it difficult (perhaps impossible) to uniquely identify the system in question.

This is one of the reasons why DST has emphasized theoretical concepts rather than models (see, e.g., Schönér, chapter 2). Indeed, when we think of the “New Grand Theory” mentioned in the title of this volume, we think of the shared theoretical concepts that have emerged from the CN and DS approaches rather than an effort to develop a “New Grand Model” of development.

This is also one of the reasons why much of the early work within the DST framework was explicitly aimed against information processing and the identification of the modules and

stages of processing that underlie cognition (Port & van Gelder, 1995). In contrast, emphasis was placed on evaluating, for instance, the concept of stability and loss of stability as a necessary ingredient of understanding behavioral patterns (van der Maas et al., 2006; van der Maas & Molenaar, 1992). DST predicts that signatures of instability should be observed such as enhanced fluctuations and critical slowing down. These signatures have been empirically measured. For instance, data suggest that switches in coordination patterns are driven by instabilities because prior to coordination changes there are large fluctuations in the behavioral pattern (Schöner & Kelso, 1988). Although these data have been modeled quantitatively (Schöner et al., 1986), it is also true that *any* model consistent with the qualitative theory of dynamical systems leads to these predictions. The predictions are, therefore, universal within the dynamical systems framework (van der Maas et al., 2006; van der Maas & Molenaar, 1992). This is a strength in that it provides evidence for the framework. This also means, however, that the observation of instabilities cannot be used to identify any particular model of the coordination system.

Given that DST and CN have emerged from different scientific traditions, are they ultimately incompatible? We, of course, do not think this is the case. For instance, the mathematical psychology tradition can be quite informative. We can ask, for instance, whether the dynamic field model of infants' perseverative reaching provides a better quantitative fit to empirical observations than Munakata's PDP model (Munakata, 1998). Similarly, we can ask which framework has been generalized to other observations more effectively—the PDP model to different forms of perseveration (e.g., perseveration in the Wisconsin Card Sorting Task; see Munakata & Morton, chapter 8) or the DFT which has been generalized to the development of visuo-spatial working memory (Spencer et al., 2007) and language development (Lipinski, Spencer, & Samuelson, in press; Samuelson et al., 2007). Finally, we can ask which set of concepts map more effectively onto neural reality (Cohen et al., 1996; O'Reilly et al., 1999; Simmering et al., in press).

Although these are all valid questions, we also need to be cautious about how seriously we take these comparisons and how deeply we allow ourselves to fall into the information processing attractor, so to speak. In our view, the non-linear concepts we are collectively working with are simply not compatible with strong forms of information processing that embrace modules, serial stages, and so on. Thus, we must stay true to our scientific traditions, but we must also look toward the future and acknowledge that we are pursuing qualitatively new ways of thinking.

### *Is learning the same thing as development?*

We end this section with a topic that is more food for thought than a well-articulated case of a “real” difference between approaches. As we discussed previously, both CN and DST have a shared commitment to understanding the processes that underlie learning, although the CN approach has taken this issue as central. Interestingly, however, the issue of learning raises a second fundamental issue: is learning the same thing as development?

Although there is not a clear answer to this question at present, our sense is that dynamicists tend to answer this question ‘no’, while connectionists tend to answer ‘yes’. If learning and development are not the same thing, then what is development? One place dynamicists have looked for development is in patterns of neural reorganization that reliably occur at different times across the lifespan (Fischer, 1980; Fischer & Bidell, 2006). Dynamicists have also looked for development in the body—changes in motor coordination patterns related

to, for instance, muscle growth and physical development (Thelen, 1985; Thelen, Fisher, & Ridley-Johnson, 1984). And in our dynamic field approach, we have made a distinction between learning that operates over a trial-to-trial time scale and more general changes at the level of the tuning of cortical maps that occurs over a slower time scale (see also Kohonen, 1982; Miikkulainen, 1991; Sirosh & Miikkulainen, 1994). Although changes at these two levels can both arise due to Hebbian processes (Abbott & Regehr, 2004), our empirical evidence suggests that these processes occur over at least two different time scales (Spencer et al., 2007).

Clearly the issue of whether learning and development are the same thing is fundamental to our science. And we look forward to future discussions about this topic. Critically, however, we contend that the only way to examine this issue effectively is to use the full power of computational modeling and formal theory. We take up this issue in greater detail in our concluding section below.

### **Implications for development: Where next and why should developmentalists care?**

Some of the issues we have raised in this “reactions” chapter are quite technical in nature: how does one evaluate different types of models? What do different forms of Hebbian learning produce? In this context, why should developmentalists not firmly rooted in the CN or DST “camp” care about either of these two approaches? Is this book merely about two groups squabbling over technical issues? Or do the issues raised in this volume pose deeper problems that require decisions (and actions!) on the part of those who don't model? In our view, the issues at stake here are important, not just for those of us enmeshed in the world of modeling and formal theories. These issues matter for all developmental scientists. Here's why.

First off, real thinking about learning and development is happening in both approaches, and both approaches have adopted a process view of development at their core. This view contrasts with talk about innateness, modules, and encapsulated, evolutionary adaptations. Thus, in the grand scheme, there is a paradigm shift happening. We contend that it is finally time to leave the old framing of our science behind. And, critically, this argument is not based on belief in one set of concepts over another. It is based on empirical facts and contemporary developmental theory (see, e.g., Spencer et al., in press). We had a tangible sense at the conference in 2005 that the group assembled was close to a deep understanding of developmental process that can make this paradigm shift complete. The discussions were thoughtful, the work was innovative, and we left more energized than from any other scientific gathering.

Second, the conceptual issues that have emerged from the in-depth exploration of DST and CN approaches are important to all researchers in development. Indeed, we contend that a process-based account of development must address *all* of the central themes we have raised in this chapter and others that form the backbone of this volume: the autonomy of behavior; the processes that underlie statistical learning as well as other classes of learning; the embodiment of cognitive systems that are situated in real environments; the emergent, multi-causal nature of development that yields individual pathways over the lifespan; the emergence and soft-assembly of behavioral functionality; the integration of processes over multiple time scales; the real-time and developmental relation between “lower” and “higher” level forms of cognition; and the interface between brain and behavior as we seek to understand the processes that underlie development from neurons to neighborhoods. As examples from this volume illustrate, DST and CN have begun to address all of these challenges, but we are only just beginning to integrate the

insights from our diverse endeavors. Both approaches must make steps to achieve more. As they do, both will change and, we contend, may ultimately fuse.

Are these concepts unique to DST and CN? No. They have been invoked in other forms within other theoretical traditions (e.g., Piagetian theory; Vygotsky's theory; the Gibsonian tradition; Good-old-fashioned Cognitive Development, see chapter 13). That said, CN and DST have taken these concepts to new levels of clarity. Indeed, we would go so far as to state that *developmental science cannot make progress without the type of computational modeling and formal theories pioneered within the DS and CN approaches*. In our view, the problems are simply too hard. This means that the differences between CN and DST are not merely technical. They are fundamental to our science.

What is the way forward? We must first appreciate our strengths. But we must also appreciate our differences, and the initial conditions that brought us to this point. Bringing these issues into the foreground reveals our different commitments as scientists, our different points of emphasis, our different styles. We hope this volume has taken an important first step in this direction. If we stay grounded in the concepts of our science and the shared theoretical ideas that are the foundation for our different approaches, we will ultimately be stronger and more cohesive. We agree with L. B. Smith and Samuelson (2003; see also, Thomas et al., chapter 17) that diversity is good, complementarity is good. There is no rush to fuse our approaches. But we do think a new grand theory of development is in the offing.

Critically, however, an ungrounded grand theory will do us absolutely no good. Just as we need strong theory to tackle the complexity of development, we need tight theory-experiment links to understand the processes that operate at multiple levels and across multiple time scales. In this context, *we need to forge strong ties to experimentalists*. This is no small feat. There is a real mistrust of modeling (sometimes, rightly so!) by some researchers in the developmental community, often expressed by claiming that "anything" could be modeled. The reality of modeling is, of course, quite the opposite. It is generally quite difficult to come up with a single model that is consistent with all constraints known experimentally. Thus, we need to find ways to bridge the gap between theoreticians and experimentalists, to forge a shared sense of trust, to deeply value both approaches (see Simmering, Spencer, Deák, & Triesch, 2008). Clearly, this is possible as this volume—rich with theory and data alike—attests.

Surveying the content at the conference in 2005 and surveying the content of this book, it is clear to us that there is no other community making more progress on discovering the reality of developing systems. We are eager to push forward and do the hard work of understanding what William James, Kurt Lewin, Jean Piaget, Lev Vygotsky, Elenor Gibson, Elizabeth Bates, Esther Thelen, and countless other scholars all saw as the path to understanding the details of developing systems. Let us take up their charge, move forward smartly, and do the hard work of understanding development.

## References

- Abbott, L. F., & Regehr, W. G. (2004). Synaptic computation. *Nature*, 796-803.
- Altmann, G. T. M. (2002). Learning and development in neural networks - the importance of prior experience. *Cognition*, 85(2), B43-B50.
- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77-87.
- Baddeley, A. D. (1986). *Working memory*. Oxford: Oxford 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, 315-319.
- Bastian, A., Schöner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, 18, 2047-2058.
- Bicho, E., Mallet, P., & Schöner, G. (2000). Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research*, 19, 424-447.
- Braitenberg, V., & Schüz, A. (1991). *Anatomy of the Cortex*. Berlin: Springer Verlag.
- Clearfield, M. W., Smith, L. B., Diedrich, F. J., & Thelen, E. (2006). Young infants reach correctly on the A-not-B task: On the development of stability and perseveration. *Infant Behavior & Development*, 29(3), 435-444.
- Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society London, B: Biology*, 351, 1515-1527.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, 45-77.
- Dineva, E. (2005). *Dynamical field theory of infant reaching and its dependence on behavioral history and context*. Unpublished PhD, Ruhr-Universität-Bochum, Bochum, Germany.
- Dineva, E., Faubel, C., & Schöner, G. (2007). *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*.
- Dineva, E., & Smith, L. B. (2008). Dynamical field approach to memory guided reaching: The role of behavioral history for stabilizing decisions. *Manuscript in preparation*.
- Douglas, R. J., & Martin, K. A. C. (2004). Neural circuits of the neocortex. *Annual Review of Neuroscience*, 27, 419-451.
- Edin, F., Macoveanu, J., Olesen, P., Tegnér, J., & Klingberg, T. (2007). Stronger synaptic connectivity as a mechanism behind development of working memory-related brain activity during childhood. *Journal of Cognitive Neuroscience*, 19(5), 750-760.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1997). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

- Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schöner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, *94*, 53-66.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545-572.
- Faubel, C., & Schöner, G. (in press). Learning to recognize objects on the fly: A neurally based dynamic field approach. *Neural Networks*.
- Fischer, K. W. (1980). A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review*, *87*, 477-531.
- Fischer, K. W. (in press). Dynamic cycles of cognitive and brain development: Measuring growth in mind, brain, and education. In A. M. Battro, K. W. Fischer & P. Léna (Eds.), *The Educated Brain*. Cambridge U.K.: Cambridge University Press.
- Fischer, K. W., & Bidell, T. R. (2006). Dynamic development of action, thought, and emotion. In W. Damon & R. M. Lerner (Eds.), *Handbook of child psychology: Vol. 1. Theoretical models of human development* (6 ed., pp. 313-399). New York: Wiley.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352-376.
- Iossifidis, I., & Schöner, G. (2006). *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.
- Jancke, D., Erlhagen, W., Dinse, H. R., Akhavan, A. C., Giese, M., Steinhage, A., et al. (1999). Parametric population representation of retinal location: neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience*, *19*, 9016-9028.
- Jansen, B. R. J., & van der Maas, H. L. J. (2002). The development of children's rule use on the balance scale task. *Journal of Experimental Child Psychology*, *81*, 383-416.
- Johnson, J. S., & Spencer, J. P. (2008). Working memory for location is attracted toward the focus of attention. *Manuscript submitted for publication*.
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2008). A dynamic neural field model of visual working memory and change detection. *Manuscript submitted for publication*.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*, 59-69.
- Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural, dynamic fields. *Biological Cybernetics*, *73*, 49-60.
- Large, E. W., Christensen, H. I., & Bajcsy, R. (1999). Scaling the dynamical approach to path planning and control: Competition among behavioral constraints. *The International Journal of Robotics Research*, *18*(1), 37-58.
- 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*. Oxford: Oxford University Press.
- Lovett, A., & Scassellati, B. (2004). Using a robot to reexamine looking time experiments, *Third International Conference on Development and Learning*. San Diego, CA, USA.
- MacDonald, M. C., & Christiansen, M. H. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, *109*(1), 35-54.

- McDowell, K., Jeka, J. J., Schöner, G., & Hatfield, B. D. (2002). Behavioral and electrocortical evidence of an interaction between probability and task metrics in movement preparation. *Experimental Brain Research*, *144*, 303-313.
- Mees, A. I. (1981). *Dynamics of Feedback Systems*. Hoboken, NJ: John Wiley and Sons, Ltd.
- Miikkulainen, R. (1991). Self-organizing process based on lateral inhibition and synaptic resource redistribution. In T. Kohonen, K. Mäkisara, O. Simula & J. Kangas (Eds.), *Artificial Neural Networks*. North-Holland: Elsevier Science Publishers B.V.
- Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP model of the AB task. *Developmental Science*, *1*(2), 161-184.
- Newell, K. M., & Molenaar, P. C. M. (Eds.). (1998). *Applications of nonlinear dynamics to developmental process modeling*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. Cambridge: Cambridge University Press.
- Port, R. F., & van Gelder, T. (Eds.). (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.
- Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*, 125-132.
- Rolls, E. T., Stringer, S. M., & Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proceedings of Biological Science*, *269*(1496), 1087-1093.
- Samuelson, L. K., Schutte, A. R., & Horst, J. S. (2007). The dynamic nature of knowledge: Insights from a dynamic field model of children's novel noun generalizations. *Manuscript submitted for publication*.
- Schall, J. (2004). On building a bridge between brain and behavior. *Annual Reviews of Psychology*, *55*, 23-50.
- Schlesinger, M., Amso, D., & Johnson, S. P. (2007). The neural basis for visual selective attention in young infants: A computational account. *Adaptive Behavior*, *15*, 135-148.
- Schlesinger, M., & Parisi, D. (2007). Connectionism in an Artificial Life perspective: Simulating motor, cognitive, and language development. In S. S. D. Mareschal, G. Westermann, & M.H. Johnson (Ed.), *Neuroconstructivism: Vol. 2. Prospectives and prospects* (pp. 129-158). Oxford, UK: Oxford University Press.
- Scholz, J. P., & Kelso, J. A. S. (1989). A quantitative approach to understanding the formation and change of coordinated movement patterns. *Journal of Motor Behavior*, *21*, 122-144.
- Schöner, G. (1989). Learning and recall in a dynamic theory of coordination patterns. *Biological Cybernetics*, *62*, 39-54.
- Schöner, G., & Dineva, E. (2007). Dynamic instabilities as mechanisms for emergence. *Developmental Science*, *10*, 69-74.
- Schöner, G., & Dose, M. (1992). A dynamical system approach to task-level system integration used to plan and control autonomous vehicle motion. *Robots and Autonomous Systems*, *10*, 253-267.
- Schöner, G., Dose, M., & Engels, C. (1995). Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, *16*, 213-245.
- Schöner, G., Haken, H., & Kelso, J. A. S. (1986). A stochastic theory of phase transitions in human hand movement. *Biological Cybernetics*, *53*, 247-257.

- Schöner, G., & Kelso, J. A. S. (1988). Dynamic pattern generation in behavioral and neural systems. *Science*, *239*, 1513-1520.
- Schöner, G., Zanone, P. G., & Kelso, J. A. S. (1992). Learning as change of coordination dynamics: Theory and Experiment. *Journal of Motor Behavior*, *24*, 29-48.
- Schutte, A. R., & Spencer, J. P. (2007). Planning 'discrete' movements using a continuous system: Insights from a dynamic field theory of movement preparation. *Motor Control*, *11*, 166-208.
- 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., Schutte, A. R., & Spencer, J. P. (in press). 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*.
- Simmering, V. R., & Spencer, J. P. (2007). Carving up space at imaginary joints: Can people mentally impose spatial category boundaries? *Journal of Experimental Psychology: Human Perception & Performance*, *33*, 871-894.
- Simmering, V. R., & Spencer, J. P. (in press). Generality with specificity: The dynamic field theory generalizes across tasks and time scales. *Developmental Science*.
- Simmering, V. R., Spencer, J. P., Deàk, G., & Triesch, J. (2008). Why should developmentalists care about computational modeling? *Manuscript submitted for publication*.
- Sirois, S. (2005). *Hebbian motor control in a robot-embedded model of habituation*. Paper presented at the International Joint Conference on Neural Networks (IJCNN).
- Sirosh, J., & Miikkulainen, R. (1994). Cooperative self-organization of afferent and lateral connections. *Biological Cybernetics*, *71*, 65-78.
- Smith, L. B., & Samuelson, L. K. (2003). Different is good: connectionism and dynamic systems theory are complementary emergentist approaches to development. *Developmental Science*, *6*(4), 434-439.
- Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review*, *106*, 235-260.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *TRENDS in Neurosciences*, *24*(3), 161-168.
- Soares, R., Bicho, E., Machado, T., & Erhagen, W. (2007). Object transportation by multiple mobile robots controlled by attractor dynamics: theory and implementation, *International Conference on Intelligent Robots and Systems IROS 2007, IEEE/RSJ* (pp. 937-944). San Diego, CA, USA.
- Spencer, J. P., Blumberg, M. S., McMurray, B., Robinson, S. R., Samuelson, L. K., & Tomblin, J. B. (in press). Short arms and talking eggs: Why we should no longer abide the nativist-empiricist debate. *Child Development Perspectives*.
- Spencer, J. P., Clearfield, M. W., Corbetta, D., Ulrich, B. D., Buchanan, P., & Schöner, G. (2006). Moving toward a grand theory of development: In memory of Esther Thelen. *Child Development*, *77*, 1521-1538.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, *131*, 16-37.
- Spencer, J. P., & Hund, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*, *47*(4), 432-480.

- Spencer, J. P., & Plumert, J. M. (2007). What makes thinking about development so hard? In J. M. Plumert & J. P. Spencer (Eds.), *The emerging spatial mind* (pp. 375-386): Oxford University Press.
- Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamical systems approach to development. *Developmental Science*, *6*, 392-412.
- Spencer, J. P., 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, NY: Oxford University Press.
- Spencer, J. P., Smith, L. B., & Thelen, E. (2001). Tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of 2-year-olds. *Child Development*, *72*, 1327-1346.
- Steinhage, A., & Schöner, G. (1997). Self-calibration based on invariant view recognition: Dynamic approach to navigation. *Robotics and Autonomous Systems*, *20*, 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.), *Sensor Fusion and Decentralized Control in Robot Systems: Proceedings of SPIE* (Vol. 3523, pp. 160-180): SPIE Publishing.
- Thelen, E. (1985). Developmental origins of motor coordination: Leg movements in human infants. *Developmental Psychobiology*, *18*(1), 1-22.
- Thelen, E., Fisher, D. M., & Ridley-Johnson, R. (1984). The relationship between physical growth and a newborn reflex. *Infant Behavior & Development*, *7*, 479-493.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral & Brain Sciences*, *24*, 1-86.
- Thelen, E., & Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge: MIT Press.
- Turvey, M. T. (1990). Coordination. *American Psychologist*, *45*(8), 938-953.
- Usher, M., & McClelland, J. L. (2001). On the Time Course of Perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550-592.
- van der Maas, H. L. J., Dolan, C. V., Grasman, R. P. P., Wicherts, J. M., Huizenga, H. M., & Raijmakers, M. E. J. (2006). A dynamical model of general intelligence: the positive manifold of intelligence by mutualism. *Psychological Review*, *113*(4), 842-861.
- van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stagemwise cognitive development: An application of catastrophe theory. *Psychological Review*, *99*, 395-417.
- van Geert, P. (1991). A dynamic systems model of cognitive and language growth. *Psychological Review*, *98* 3-53.
- van Geert, P. (1993). A dynamic systems model of cognitive growth: Competition and support under limited resource conditions. In L. B. Smith & E. Thelen (Eds.), *A dynamic systems approach to development: Applications* (pp. 265-331). Cambridge, MA: MIT Press.
- Van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, Vygotsky, and beyond. *Psychological Review*, *105*, 634-677.
- van Geert, P., & van Dijk, M. (2002). Focus on variability: New tools to study intra-individual variability in developmental data. *Infant Behavior and Development*, *25*(4), 340-374.
- Warren, W. H. (2006). The Dynamics of Perception and Action. *Psychological Review*, *113*(2), 358-389.

- Wilimzig, C., Schneider, S., & Schöner, G. (2006). The time course of saccadic decision making: Dynamic field theory. *Neural Networks, 19*, 1059-1074.
- Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik, 13*, 55-80.
- Zanone, P. G., & Kelso, J. A. S. (1992). Evolution of behavioral attractors with learning: Nonequilibrium phase transitions. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 403-421.
- Zanone, P. G., & Kelso, J. A. S. (1997). Coordination dynamics of learning and transfer: collective and component levels. *Journal of Experimental Psychology: Human Perception and Performance, 23*(5), 1454-1480.

## Acknowledgements

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 awarded to John P. Spencer and Gregor Schöner. 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](mailto:john-spencer@uiowa.edu).