Developmental Science 6:4 (2003), pp 392–412

Bridging the representational gap in the dynamic systems approach to development

John P. Spencer[1] and Gregor Schöner[2]

1. *Department of Psychology, University of Iowa, USA*
2. *Institut für Neuroinformatik, Ruhr-Universität Bochum, Germany*

Abstract

We describe the relationship between the dynamic systems approach to development and a recent approach to the dynamics of representational states – the dynamic field approach. Both approaches share an emphasis on the concepts of stability (attractor states), instability (especially bifurcations), soft-assembly and flexibility. But the dynamic field approach adds the concept of 'activation' to capture the strength with which behaviorally relevant information is specified. By explicitly linking these dynamic systems approaches, we allow for more direct comparisons between dynamic systems theory and connectionism. We note three current differences between these two approaches to development: (1) the notion of stability is central to how representational states are conceptualized in the dynamic field approach; (2) the dynamic field approach is more directly concerned with the sensorimotor origins of cognition; and (3) the dynamic approach is less advanced with regard to learning. We conclude that proponents of the two approaches can learn from the respective strengths of each approach. We suspect these differences will largely disappear in the next 20 years.

Introduction

Dynamic systems theory and connectionism have had a major impact in the last two decades on our understanding of development. These two approaches have led to new accounts of children's performance in classic Piagetian tasks (e.g. McClelland & Jenkins, 1991; Munakata, 1998; Smith, Thelen, Titzer & McLin, 1999; Thelen, Schöner, Scheier & Smith, 2001; van Geert, 1998), new explorations of social-emotional development (e.g. Fogel, Nwokah, Dedo & Messinger, 1992; Lewis, 2000; Lewis, Lamey & Douglas, 1999), revolutionary ideas about motor development (e.g. Thelen, 1995, 2000; Thelen & Ulrich, 1991) and fundamental progress in the understanding of how children acquire language (e.g. Bates & Elman, 2000; Elman, 2001). Given these impressive advances in overlapping domains, the question of this special issue becomes pressing: What is the relationship between these two theoretical approaches?

Although this question is timely, it is difficult to answer for a variety of reasons. Connectionism is a broad approach with many different flavors (e.g. Arbib, 1995; Bullock & Grossberg, 1988; Mareschal, Plunkett & Harris, 1999; McLeod, Plunkett & Rolls, 1998; O'Reilly & Munakata, 2000; Shultz, 1998). Similarly, within the dynamic systems approach, different variants coexist (e.g. Bidell & Fischer, 2000; Hartelman, van der Maas & Molenaar, 1998; Newell & Molenaar, 1998; Thelen *et al.*, 2001; Thelen & Smith, 1994; van der Maas & Molenaar, 1992; van Geert, 1998). A difficulty, therefore, is to select exemplary research programs and models to represent this diversity. Even after exemplary models have been selected, however, other problems wait in the wings. For instance, along which dimensions should one compare exemplary models (e.g. see Thelen & Bates, this issue)? Are such comparisons even possible in some cases? How, for example, does one compare the Munakata–McClelland flavor of connectionism (see this issue) which typically deals with issues of cognition and representation with one of the more well-known dynamic systems approaches – the approach pioneered by Thelen, Smith and colleagues (e.g. Thelen & Smith, 1994) – which typically deals with issues of motor control and development?

Faced with these daunting challenges, we opted to focus on a specific (yet still challenging!) goal: to build a 'representational bridge' from the dynamic systems approach to motor control and development toward a dynamic systems approach that includes the dynamics of representational states – the very issue many connectionist models treat as central. This serves two main purposes. First, the 'dynamic approach' as we perceive it is changing. Readers are likely to be more familiar with the

Address for correspondence: John P. Spencer, Department of Psychology, 11 Seashore Hall E., University of Iowa, Iowa City, IA 52242-1407, USA; e-mail: john-spencer@uiowa.edu

now classic 'motor approach' (e.g. Thelen & Smith, 1994; Thelen & Ulrich, 1991) than with our more recent dynamic field approach that incorporates the dynamics of representational states (e.g. Erlhagen & Schöner, 2002; Schöner, Dose & Engels, 1995; Schöner, Kopecz & Erlhagen, 1997; Schutte & Spencer, 2002; Thelen *et al.*, 2001). Thus, this paper provides an overview of the new dynamic approach and how this approach fits into the broader purview of dynamic systems theory. Second, we contend that building a representational bridge allows for more direct comparisons between the two approaches. In particular, we highlight three issues that currently differentiate dynamic systems and connectionist approaches: (1) within the dynamic field approach, the notion of *stability* is central to how representational states are conceptualized; (2) the bridge between the motor and dynamic field approaches allows representational states to emerge from sensorimotor origins; and (3) in contrast to connectionism, the dynamic approach described here has not yet formalized a well-defined theory of learning. In light of these differences, we conclude that the two approaches can learn from their respective strengths and weaknesses. Ultimately, we suspect these differences will largely disappear in the next 20 years.

Given our stated goal, it is worth noting two things. First, readers familiar with dynamic systems approaches might be having heart palpitations right now: a dynamic systems approach to representation!! Heresy! At face value, such a thing seems counter to important arguments against symbolic representation (see, e.g. Barsalou, 1999; van Gelder, 1998). Although we discuss what we mean by representational states later in the paper, we highlight two characteristics of our use of the 'R-word' now. First, throughout this manuscript, we use the term 'representational states' rather than 'representation'. While some readers might consider this window dressing, it is a reminder to us that we are talking about a time-dependent state of the nervous system, rather than a static thing sitting in the head somewhere. Second, our use of the term reflects a pre-cognitive revolution view of representation as meaning, literally, re-presentation. Thus, a representational state in our approach is a time-dependent state in which a particular pattern of neural activation that reflects, for instance, some event in the world is re-presented to the nervous system in the absence of the input that specified that event. Note that this view of re-presentation is related to recent ideas that the brain runs 'simulations' of past events during many cognitive tasks (see, e.g. Damasio & Damasio, 1994).

A second note of caution: this is not a review of dynamic systems approaches to development *en masse* (although some review is certainly included). Readers interested in such topics are referred to other recently published reviews (e.g. Fischer & Bidell, 1998; Lewis, 2000; Thelen & Smith, 1998). Moreover, we make no claims that our particular flavor of dynamics is completely representative of other dynamic systems approaches to development (e.g. Fischer & Bidell, 1998; van Geert, 1998; van der Maas & Molenaar, 1992). We made an effort to highlight points where our views converge with proposals made by other researchers. And we have included a discussion of differences among dynamic systems approaches in the final section. Nevertheless, we had to sacrifice breadth to achieve our stated goal. Similarly, we tended to emphasize some flavors of connectionism over others. As a basis for selection, we focused on those approaches that appeared to be most prominent in the developmental literature.

In the section that follows, we give a brief review of the dynamic systems approach to motor control and development (i.e. the motor approach). Next, we describe our dynamic field approach and how this approach incorporates central aspects of dynamic systems theory – notions of attractors, stability and so on. In the final section, we get to the business at hand: to evaluate how the new dynamic approach – which now includes the dynamics of representational states – differs from connectionist approaches to development.

The dynamic systems approach to motor control and development

Stability is necessary

Watching infants learn new motor skills such as how to reach is an often torturous experience. Early in the first year, their arms flail about in seemingly random ways as they stare intently, drooling, at the toy (Jones, 1996; Thelen, Corbetta, Kamm, Spencer, Schneider & Zernicke, 1993). Even when they appear to be 'getting it', small environmental events can perturb their efforts to grasp a toy that is within reach – a change in the supportive characteristics of an infant seat (Spencer, Veriejken, Diedrich & Thelen, 2000) or a change in the visual scene (Lee & Aronson, 1974). Similarly, various processes within the child's nervous system can potentially interfere with successful reaching. Early on, infants have difficulty activating the right sequence of muscles to move their hands in the vicinity of the toy (Spencer & Thelen, 2000). Later, more subtle effects can be observed such as when motor habits built up over the last few movements 'pull' the current reach onto the paths of previous reaches (Diedrich, Thelen, Smith & Corbetta, 2000).

Why is something as simple as a reaching movement such a complicated thing to learn? One reason is the complexity of the environments in which infants behave and their own inner complexity – the multitude of different interconnected subsystems and elements that make up nervous systems. These types of complexities ensure that infants are regularly bombarded by both external and internal 'noise' – unpredictable changes in the world or unpredicted changes in the internal elements of the nervous system. The trick is to maintain some goal state (a reaching plan) amidst this blooming, buzzing confusion. That is, the child must, at least for some small window of time, achieve a stable state – a state that resists perturbation. This is a fundamental challenge for the developing child. And understanding how such states arise over development is a fundamental challenge for developmental psychology.

In dynamic systems terms, stability is defined as the persistence of behavioral or neural states in the face of systematic or random perturbations. Stability may arise through a variety of mechanisms. In the simplest case, the physics of the body itself can provide stability. For example, the elastic and viscous properties of muscles can keep joints stable against perturbing forces (e.g. Latash, 1993). Often, however, the nervous system generates stability by, for instance, constantly monitoring and updating movements using sensory feedback. Although these mechanisms of stability differ, the resultant stable states can be usefully characterized using the concepts of dynamic systems theory (e.g. Braun, 1994). In this framework, the space of possible states of a system is spanned by state variables (or 'behavioral variables', see Schöner & Kelso, 1988a). For every possible state (or value of the state variables), a vector predicts in which direction and at which rate the system's state will evolve. Stable states are then values of the state variables at which the rate of change is zero and to which the system converges from nearby values (see Figure 1).

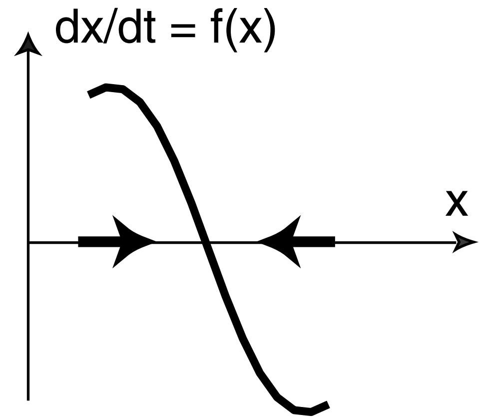

Figure 1 *Plot of a hypothetical differential equation (see solid line) showing the rate of change* (dx/dt) *of a state variable,* x. *Arrows indicate the direction in which the state will change when the system is positioned near the value of* x *specified by the start of each arrow. The value of* x *where* dx/dt *= 0 is called a fixed point attractor: the system will return to this value even when pushed to other nearby values of* x.

The need for instability

Although achieving behavioral stability is a fundamental part of developmental change, there is a down side to stability: it limits flexibility. Maintaining a stable behavior means, after all, that change to a qualitatively different behavior is prevented. Thus, behavioral flexibility requires that the stability of a particular state be dissolved; that the state be released from stability to allow for a new behavior. In dynamic systems terms, a change of a system that leads a particular state to become unstable is referred to as an instability (see Braun, 1994, Chapter 4).

Under natural conditions, switches from one stable state to another are often quite rapid. Adults, for instance, will quickly and efficiently switch from a walking to a running pattern at the appearance of an oncoming car. In the laboratory, however, the release from stability may be studied in detail by varying environmental conditions gradually while monitoring stability. This was the objective of the classic work on 'finger twiddling' (Kelso, 1984; Kelso, Scholz & Schöner, 1986; Schöner & Kelso, 1988a). In these studies, adults reliably performed rhythmic in-phase and phase-alternating finger movements with low variable error when moving at a comfortable speed. When movement frequency was gradually increased, however, adults' production of the alternating pattern became increasingly variable, an indicator of reduced stability. Indeed, at high frequencies, adults were no longer able to maintain the alternating pattern and spontaneously switched to an in-phase pattern.

These data demonstrate that instabilities are central to behavioral switches from one pattern to another. Instabilities also occur when a new pattern is learned. For instance, Zanone and Kelso (Zanone & Kelso, 1992, 1994) asked adults to produce finger movements at a relative phase of 90°, a pattern that was not initially stable. Many participants had difficulty learning the new pattern, in part, because they would often slip into either an in-phase (0°) or a phase-alternating (180°) pattern. After learning the new pattern, however, participants were conversely attracted to 90° when attempting to perform other patterns, even, for some participants, when attempting to perform the formerly stable in-phase and phase-alternating patterns. Thus, learning not only stabilized the new pattern, but also destabilized previously stable states.

Flexibility and levels

The necessary balance between stability and instability makes the infant's rise out of the blooming, buzzing confusion of development even more impressive. By this view, development is not just a matter of achieving stability. Rather, the child must also acquire behavioral flexibility – the ability to flexibly destabilize one stable state and enter another as the situation demands.

Generating both stability and flexibility is a major challenge because nervous systems do not contain two separate systems dedicated to these two processes. To clarify this point, consider how a roboticist designs the motor control system for a non-autonomous robot. To achieve stability in the robot's behavior, the roboticist programs a feedback, error-correction mechanism. To make the robot flexible, however, the roboticist must write separate programs for each new task. In nervous systems, there is, of course, no programmer writing new routines for each new encounter with the environment. Rather, flexibility and stability must arise from the same, densely interconnected neural system.

Although there are not separate stability and flexibility levels in nervous systems, the notion of levels of control can still be used to understand how these two characteristics emerge from one complex system. Goal-directed movements, such as reaching for an object, involve at least three levels of control (Schöner, 1995). First, the person must specify global characteristics of the reach, for instance, its direction and amplitude (see Figure 2). Second, the person must determine timing – how fast or slowly he or she moves. Third, the mover must produce the precise forces needed to move the segments of the arm. Importantly, these levels of control are mutually coupled and interactive (see bidirectional arrows in Figure 2), rather than being linked hierarchically where one level controls another. Moreover, all three levels evolve in time and are continuously linked to sensory information (see 'world' in Figure 2). Thus, although we can measure movement direction, timing or muscular forces in the laboratory, these aspects of a reach are never generated in isolation, but are always part of a multiply-connected and interactive system.

In this multi-level framework, a stable reach is the resultant interplay of all the interlinked levels. The reach is stable because it has highly redundant levels of control. For example, an adult can achieve stable timing from multiple and distributed processes by several means (see Schöner, 2002): by controlling a timing pulse directly, by adjusting the amplitude of the reach, or by making adjustments at the level of forces (e.g. Gracco & Abbs, 1988; Kelso, Putnam & Goodman, 1983; Kelso, Tuller, Vatikiotis-Bateson & Fowler, 1984; Sternad, Collins &

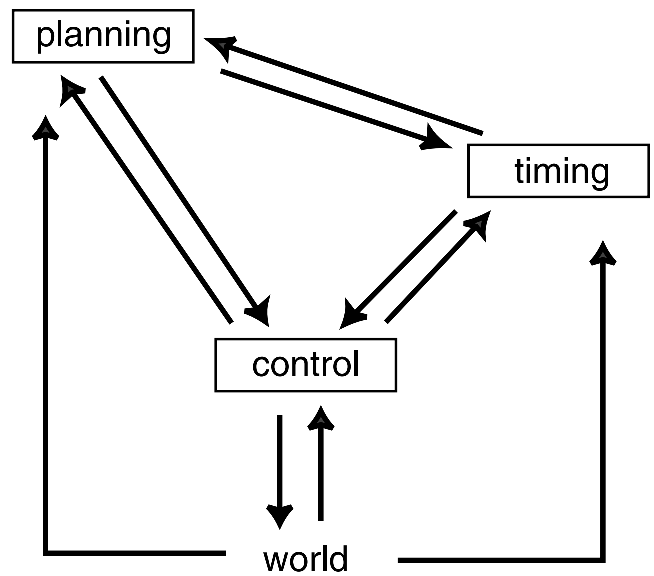

Figure 2 *A multi-leveled, dynamic systems view of how goal-directed actions are controlled. Each box represents a dynamic system responsible for a particular aspect of action control. Planning = selection of global movement parameters; Timing = timing of signals to the motor system; Control = coordination and control of force production at the effector level. Each level is mutually coupled to the other levels, and each level receives sensory feedback from the world.*

Turvey, 1995). Moreover, the whole system can be 'softly-assembled' to meet the demands of a given context. When people catch a ball, for instance, they adjust the timing of their limbs to their visual estimation of the ball's contact. When people jog on a treadmill, however, the timing of their limbs is tuned to the speed of the treadmill and visual cues are much less important.

It is precisely this same redundancy that gives dynamic systems flexibility. Because such systems are multiply-determined, there are many ways for patterns to become destabilized and reorganize in new ways. Consequently, there is no need for a central executive. Reorganization can occur when new perceptual information specifies a new state, but may also result from changes within the system itself. Thus, adults can change the phase relationships between two fingers in response to pacing by a metronome (e.g. Tuller & Kelso, 1989), but changes in timing may also arise due to muscular constraints when one limb moves a heavier weight than another (e.g. Sternad *et al.*, 1995). Similarly, infants reorganize their movements as support from the environment changes, but also when they grow or develop the ability to assume new postural configurations. Flexibility arises from the same system properties that generate stability.

Changes in stability and flexibility over development

The fundamental question of development is how new behavioral forms are created. The multi-leveled dynamic systems view discussed above provides a framework within which this question can be addressed. One challenge for the developing child is to identify which behavioral dimensions (i.e. which levels) are more or less related to the stability of a particular behavior. Moreover, the child must determine how the levels interact, that is, what types of coupling allow information to be exchanged among levels in such a way that the levels don't perturb one another. Finally, the child must learn to switch from one stable behavior to another as the context changes. Development is therefore not merely the evolution toward greater stability of particular behavioral states (akin to getting better over learning), but also the evolution toward an ability to release states from stability in targeted and meaningful ways (for related ideas, see Fischer, Rotenberg, Bullock & Raya, 1993).

The proposal that stability emerges from an ensemble of processes at multiple levels is consistent with a recent longitudinal study of the development of reaching (e.g. Corbetta & Thelen, 1996; Thelen *et al.*, 1993; Thelen, Corbetta & Spencer, 1996). Early in the first year, the infants in this study moved their arms in seemingly random ways through many different regions of the task space using a variety of muscle contraction patterns. At this point in development, infants had difficulty stabilizing force production around the joints of the arm to move the hand in a controlled manner through space. The week when infants first began reaching, however, they learned to stabilize force production in a clever, but rather brute force way: infants began using the muscles of the upper arm to get the hand moving in the right direction and then co-contracted many arm muscles to keep the hand in the vicinity of the toy (Spencer & Thelen, 2000; for related postural changes, see Spencer *et al.*, 2000).

With this hurdle cleared, infants began to refine the relationship between the spatio-temporal aspects of their reaching movements and force control. That is, they began to improve the stability and coupling between timing and force production levels (Thelen *et al.*, 1996). During the weeks after the onset of reaching, infants would occasionally produce straight, accurate reaches. More typical, however, were torturous reaches containing several changes in direction and speed. Such reaches were particularly likely when infants moved quickly, indicating that timing processes could perturb the stable production of forces at another level of control. Around 30–36 weeks of age after several months of experience reaching for toys, infants solved this problem: they were able to produce relatively smooth and straight reaches, even when moving quickly.

More recent studies of infants' reaching skill have revealed that, although infants generally produce smooth and straight reaches by 8 to 9 months, their ability to stabilize a plan for *where* to reach is still developing. For instance, infants will reliably make the classic Piagetian 'A-not-B' error at this point in development: after infants have repeatedly searched for a hidden toy at an A location, they will search back at this location, even after watching a toy being hidden at a nearby B location (e.g. Piaget, 1954; Smith *et al.*, 1999). Importantly, recent data demonstrate that the goal-directed processes that underlie such errors can be influenced by coupling to the timing and force levels. For instance, when infants' body posture is changed just before a toy is hidden at the B location, 8- to 10-month-old infants are less likely to erroneously reach to A (Smith *et al.*, 1999). Thus, the bias to reach toward A is not just about where A is located in space, but also about how the arm must be moved to arrive at this location (see Thelen *et al.*, 2001).

Beyond 12 months (the age at which infants typically succeed in the Piagetian A-not-B task), infants begin to show more behavioral flexibility. For instance, around 2 years of age, children's decision to reach with one hand or two becomes finely scaled to the size of the object (e.g. Fagard & Jacquet, 1996). Moreover, 2-year-olds are able to use external landmarks to remember the location of a toy hidden in a large sandbox, even when their body position is manipulated between hiding and search (Newcombe, Huttenlocher, Drummey & Wiley, 1998).

Taken together, these data suggest that the development of complex skills can be usefully studied within the stability framework outlined above. That is, development can be viewed as the emergence of stability at different levels of control, progressively more refined coupling among levels, and enhanced behavioral flexibility as system-wide organization becomes tuned to the details of the behavioral context.

Discussion: evaluation of the motor approach

Strengths

The dynamic approach to motor control and development has clear strengths. Of particular note, its ability to provide a detailed 'collective' picture of behavior at a relevant level, the level at which behavioral patterns arise, has made the approach attractive to experimentalists. Conceptual theory (such as notions of attractors and their disappearance through instabilities, the ideas of emergence and self-organization) as well as formal theory (based on the mathematics of stochastic differential equations) has had close ties to experimental work, both qualitatively and quantitatively. Moreover, novel predictions have been possible in some cases (e.g. Schöner, 1989; Schöner & Kelso, 1988b), and several formal models of phenomena have been proposed that are consistent with all known facts (e.g. Schöner, Haken & Kelso, 1986).

In development, the approach has often been criticized as being 'metaphorical'. Indeed, most dynamic systems research in development has stayed at the conceptual and empirical levels. Nevertheless, the dynamic approach has provided critical insights into the nature of developmental change. Dynamic systems theory has helped move the field away from notions of developmental programs and controlling processes. Rather, there is an emerging emphasis on the self-organizing tendencies of a spontaneously active, complex organism (see also, Fischer & Bidell, 1998; Lewis, 2000; Newell & Molenaar, 1998). This view is consistent with ideas proposed by other developmental theorists as well, including Jean Piaget (1952), Eleanor Gibson (1991) and Kurt Lewin (1936, 1946). Moreover, the nonlinear nature of dynamic systems has demonstrated that both quantitative and qualitative developmental change can be generated by the same system. Thus, it is not necessary to posit new control processes to account for qualitative shifts in behavior. Rather, such changes can occur via a reorganization of the system during the transition (e.g. Thelen & Smith, 1994; van der Maas & Hopkins, 1998; van der Maas & Molenaar, 1992; van Geert, 1998).

In addition to these insights, the concepts of multicausality and soft-assembly have focused attention away from the search for single, causal factors in development, toward the confluence of factors that create more or less stable behavioral patterns. For instance, rather than searching for the trigger that 'turns on' walking during infancy, Thelen and colleagues examined a host of component systems that might contribute to the development of this skill. This led to the striking discovery that infants will show alternating stepping patterns on a treadmill long before they begin walking independently (see Thelen & Ulrich, 1991). Examples such as these illustrate that development can be usefully studied by examining how infants and children assemble behaviors 'in the moment' based on their own past behavioral histories and the nature of the current behavioral context (Thelen & Smith, 1994).

The multi-causal nature of how stability is realized also helps explain the origin of the inherent variability and context-dependency of development (Siegler, 1994; Thelen, 1992; Thelen & Smith, 1994). Behavioral development is variable in this view, because stability requires coordination among a complex array of components. As such, it is unlikely that stable behavioral patterns will be realized in precisely the same way each time a child engages in a behavior. Rather, the way in which stability is achieved will depend on the details of the context and each individual's developmental history.

Behavior is context-dependent in this view, in part, because stability often requires strong input from the environment to suppress internal and external sources of perturbation. Thus, if contextual supports are not present, the behavioral state cannot be stabilized. A multileveled view of stability is also relevant here, because couplings established in one context may not produce behavioral stability in a different context. Such context-dependency may be particularly prevalent in early development when only a limited number of ways to stabilize a particular behavior have been discovered. For instance, Adolph (1997) reported that after several months of crawling experience, infants developed a relatively precise perceptual threshold for slopes they could safely descend. However, once these infants began walking, they plunged over all slopes as if they had never seen slopes before! This dramatic example of context-dependency may reflect context specific coupling between visual and motor processes. That is, perceptual processes clearly informative for stabilizing crawling movements were not effectively coupled to the motor processes involved in walking.

Limitations

One criticism of the motor approach is that dynamic systems concepts provide what is essentially a description of observed patterns, without helping us understand mechanism, that is, how behavioral patterns are actually generated by the nervous system. Are dynamic systems approaches purely descriptive?

A first answer is, yes, dynamic concepts are descriptive in that formal dynamic systems models require that the scientist select which variables to use to represent a particular behavioral state and distance from a goal state. Relative phase, in the rhythmic movement example, expresses such a choice. The characterization of rhythmic movements in terms of relative phase is descriptive relative to more specific levels such as the signals coming from particular sensors or going to particular effectors. A coupled oscillator picture expresses a different choice. Here the mechanical position of each limb is the level of description, so that the relative phase level appears as a more macroscopic, approximate description.

A closer look reveals, of course, that the description in terms of mechanical positions of limbs is rather macroscopic too, as there are many potential mechanisms that contribute to the mechanical position of a limb. Those include efference copy of motor commands, proprioceptive, spindle and other sensory information about limb position, and many others. In fact, with this critical eye we discover that we can move to finer and finer measurement scales, yet each scale still requires decisions by the scientist about which variables to use to represent some behavioral state. Thus, as with relative phase, the description of the motor system in terms of firing rates of motor neurons or in terms of their membrane potential is still a particular choice of level of description, and even these hide other complex microscopic dynamics. And so it continues down to the biochemistry of neurotransmitters or membrane gates.

This leads to a second answer to the question about description: no, formal dynamic concepts are not more descriptive and less explanatory than any other set of concepts currently in use, because all formal models require selection and specification of the variables deemed 'relevant' to the phenomenon in question. The question is less whether there is an absolute mechanistic level of description from which all behavior must be explained than what is an appropriate choice (or choices) of level of description for each particular behavioral phenomenon.

From a dynamic systems perspective, there are two arguments that help answer this question. First, an appropriate level of description must be closely linked to the *stability* of the behavioral pattern under study. This follows from the goal of a dynamic systems model – to capture how the stability of behavioral states changes over time. Thus, if we find, for instance, that membrane mechanisms work the same way irrespective of whether a particular behavioral pattern is currently stable or not, then we are not looking at a level of description specific to the phenomenon that we aim to understand.

The second argument that helps select an appropriate level of description is related to the issue of multicausality. Recall from our discussion of the multi-leveled view that behavioral states may arise under a wide set of circumstances and may be influenced by processes at multiple levels. For instance, a stable timing pattern might be influenced by constraints in the environment (e.g. the tick of a metronome), as well as by mechanical constraints that arise from the weighting of a limb. To effectively model the characteristics of timing patterns, then, a level of description is required that captures when timing is more or less stable, irrespective of which particular mechanism/level was the primary cause for the resultant timing pattern. This does not prevent us, of course, from also addressing that latter question, for which an adequate characterization of timing dynamics is a prerequisite.

A second criticism of dynamic systems approaches to development is that they have been largely driven by a conceptual or metaphorical understanding of dynamic systems theory without a strong connection to formal modeling. Although this criticism is generally accurate (with some notable exceptions), a conceptual approach has been quite generative to date. Thus, while it should certainly be a goal to move developmental research toward more formal theory building, it is important not to undersell the role played by the conceptual ideas that underlie dynamic systems approaches (see Thelen & Bates, this issue). Indeed, formal approaches are less useful unless the conceptual groundwork has been laid within a research domain.

A final limitation of the dynamic approach we have sketched to this point is its relationship to perception, cognition and, specifically, representation. The notion of dynamic state variables that evolve continuously toward stable states seems to require that at all times information about such states is available, well specified and changing gradually in time. This is not always the case with representational states. Consider three situations where a person decides to grab a coffee cup on a table. In one case, the cup is clearly visible and far from other objects; in another case, the cup is surrounded by identical 'distracting' cups; in a third case, vision of the cup is obstructed by a stack of journal articles. Assume, further, that the person makes an identical movement in these cases – an accurate, stable, efficient reach that successfully makes contact with the cup. Does the motor approach capture everything about this situation? The answer is 'no'. This approach fails to capture differences in the representational states underlying these movements. In particular, in one case, there was a high degree of certainty regarding where to move, in the second case, the decision of where to move was much less certain, and in the third case, the decision took on a new flavor – rather than being generated based on visible information, the decision was generated based on a longer-term memory of the cup's location. The dynamic approach must somehow deal with these characteristics of representational states. This is the issue to which we turn in the next section as we introduce the dynamic field approach.

The dynamic field approach and representational states

Stability is a necessary component of representational states

In the dynamic systems approach to motor control and development, behavioral variables are used to capture the collective structure of, for instance, overt motor behavior and which behaviors are more or less stable at particular moments in time. Although this approach has many strengths, the example above illustrates that something is missing. Can the dynamic approach capture the properties of representational states that seem necessary to deal with perception, cognition and the development of these systems?

Before describing our approach to representational states, it is important to clarify exactly what we think is missing from the motor approach. First, the motor approach fails to capture the *graded certainty* of representational states that underlie behavior. In the example above, the person set up a comparable dynamic system for the arm movement, even though there were important differences in the underlying representational states. What we need is a coupled dynamic system that captures the certainty with which movements are planned, information is integrated and decisions are made *and* captures the generation of movements as well.

The second limitation of the motor approach is more subtle, but equally important. With more classical dynamic systems, the state of the system is always clearly specified. This is appropriate because the hand (or any other effector) is always somewhere! By contrast, information about, for instance, the location of a target can sometimes be insufficient. In these cases, a well-defined movement plan may not be formed. At other times, in spite of remaining uncertainty about exactly which target to move to, a decision might be made to select a particular target and a clearly defined movement plan is formed. Thus, the dynamic approach must allow for both the presence and the absence of a clearly defined representational state in a manner that extends beyond simply mirroring the graded certainty of available information. To foreshadow our approach to this issue, we need a dynamic system that is bi-stable – it can stably express whether a clearly defined representational state has been established or not.

The third key feature of representational states is that they can be *discontinuous in content*. There are times (particularly with adults) when the content of a response doesn't appear to be systematically related to the content of previous responses. For instance, in a target detection task, the representational state that underlies the verbal response 'the far left target' does not obligatorily evolve through intervening locations en route to the response 'the far right target'. This is not to say that the processes underlying these two representational states were not continuous, nor does it imply that continuous change of representational state never occurs (e.g. continuous state change *does* occur in many mental rotation tasks). The central point is that the content of representational states does not necessarily take on this continuous character. This stands in contrast to the motor system which must always evolve continuously in 'content'. Thus, a pointing response to the 'left' must necessarily evolve through intervening locations en route to the 'right'.

To account for these three missing characteristics of representational states, we adopt the concept of activation as it has been used in the past in mathematical psychology, connectionism and theoretical neuroscience (Churchland & Sejnowski, 1992; Williams, 1986). In our dynamic field approach, this concept takes the form of an activation field, defined over the metric dimension represented. In the coffee cup example above, activation might be distributed across the dimension of reachable locations, a continuous metric dimension stretching from a far left location to a far right location. A localized peak of activation within this field indicates that a target object (a cup) has been detected at a particular location. Such a peak might be built up via perceptual input that specifies where the cup is located within reachable space.

How does 'activation' allow for the characteristics of representational states mentioned above? Graded certainty can be accommodated by allowing activation to take on a complex distributional form. In the presence of multiple inputs (i.e. multiple cups), for example, some locations might be more strongly 'represented' than others, leading to a multi-modal distribution of activation across reachable space. Alternatively, in the presence of no inputs, a homogeneous, low-level amount of activation might be distributed across all locations. In this case, the activation field (in conjunction with stability characteristics described below) captures the absence of a clearly defined representational state. Finally, if we consider the evolution of activation over time, discontinuous change in content can be effectively described. For instance, if a left target suddenly disappears and a right target appears, there can be a discontinuous shift in what is represented by activation that reflects this change. In particular, activation associated with the left target can decay while activation associated with the right target is built up. Importantly, this can occur even though activation never builds up at intermediate reachable locations.

Although the introduction of 'activation' moves toward a dynamic approach to representational states, a key characteristic of dynamic systems has been left out – stability! To illustrate the role of stability in our dynamic field approach, we turn to a new, but related example: the detection of reachable targets in the classic Piagetian A-not-B situation. Infants in this situation are typically shown a box with two hiding wells (e.g. Smith *et al.*, 1999). Placed on top of these wells are lids that are often the same color as the box. Infants are shown an attractive toy, the toy is hidden in one of the wells and covered up, there is a short memory delay, and infants are allowed to reach. What do infants represent about the state of the world that would usefully guide action in this task? For instance, are desired targets present or absent? Is the location of the target specified by perceptual cues (lids)? How salient are these cues?

We can think about the characteristics of infants' representational states in this situation in terms of the same activation field described above: an activation field defined over the behavioral dimension of reachable locations. A localized peak of activation within this field indicates that a target object has been detected at a particular location. For instance, localized sensory input driven by the hiding event at a particular location can build a localized peak of activation at sites in the activation field associated with this location in reachable space.

Does such an activation field achieve stable representational states? Or have we simply described a way to transform sensory input into neural output? If we want the child to behave stably and robustly, our activation field must do more than simply reflect localized sensory input. This becomes clear when one considers what infants do in the A-not-B situation. After a target has been hidden, infants are still wiggling around, making noise, looking from side to side. All of these things may disrupt the continued detection of the toy's location. If our activation field merely transformed sensory input into neural output, infants would be stopping and starting all the time – starting to reach for a detected target, aborting this action in the face of new sensory input, redetecting the target location, starting again, aborting the action, and so on.

What we need is an activation field that goes through an instability when a target is detected and forms a stable attractor state, in this case, a stable, localized peak of location-specific activation. A dynamic system of the activation field – what we refer to as a *dynamic field* – generates such an instability and stable attractor state through a well-known brain-like interactive mechanism: local excitation and lateral inhibition. This leads to an emergent property critical to performance in the A-not-B situation: when an object is hidden, dynamic fields can enter a state in which a localized peak of activation built up by sensory input (the hiding event) remains stable even when this input is removed. Thus, dynamic fields can actively and stably retain a memory of the hiding event, a form of 'working' memory.

The A-not-B error: capturing early changes in the stability of representational states

To explore how dynamic fields stably represent information, we turn to the dynamic field model proposed by Thelen and colleagues (2001) to capture infants' performance in the A-not-B situation. Figure 3 shows how this model behaves on the first B trial in the A-not-B task (i.e. the first trial after pre-training to the A location and two A trials). The top panel shows a simulation of an 8- to 10-month-old infant's performance, the age range in which infants typically make the A-not-B error (e.g. Diamond, Cruttenden & Neiderman, 1994; Diamond & Doar, 1989; Smith *et al.*, 1999). The bottom panel shows a simulation of a 12-month-old infant's performance, the age at which infants typically search correctly in this task.

The figures in the left column of each panel show three inputs to the model that capture events in the task, while the figures to the right in each panel (Figures 3d, 3h) show dynamic fields. In each field, the range of possible reaching locations is captured along the x-axis, time from the start of a trial (0 s) to the end of a trial (10 s) is on the y-axis, and activation is on the z-axis. Note that we refer to the figures in 3d and 3h as 'working memory fields' rather than 'motor planning fields' (see Thelen *et al.*, 2001) to link up with recent extensions of the dynamic field model to studies of spatial working memory beyond infancy (see below). Nevertheless, our notion of working memory remains closely tied to sensorimotor phenomena, a point we return to later.

The inputs depicted in the left column of the upper panel capture the different events that occur in the A-not-B task. Figure 3a shows the first input, the *task input*. In the canonical A-not-B task, there are two hiding locations covered by identical lids. This is captured in the figure by the two peaks of input activation centered over the A and B locations. The *specific input* is shown in Figure 3b. This input captures the hiding event. At the start of the trial when the toy is not visible, the specific input is zero. Next, the experimenter waves the toy near the B location and hides it under the B lid. This event is captured by the strong input activation at the B location between 2 and 4 s. After the hiding event is over (after 4 s), the specific input is once again zero (i.e. the toy is not visible). The final input, the *memory input*, captures the infant's longer-term memory of previous trials (see Figure 3c). The memory input has activation centered at the A location. Recall that Figure 3 shows how the model behaves on the first B trial. Thus, the stronger activation at A in the memory input reflects the infant's longer-term memory of the previous trials to A.

These inputs are integrated within the spatial working memory field shown in Figure 3d. At the start of the trial, there is stronger activation in the field at A than at B. This is due to the stronger memory input at A. From 2 to 4 s, the experimenter waves the toy at the B location. The strong specific input that captures this event builds

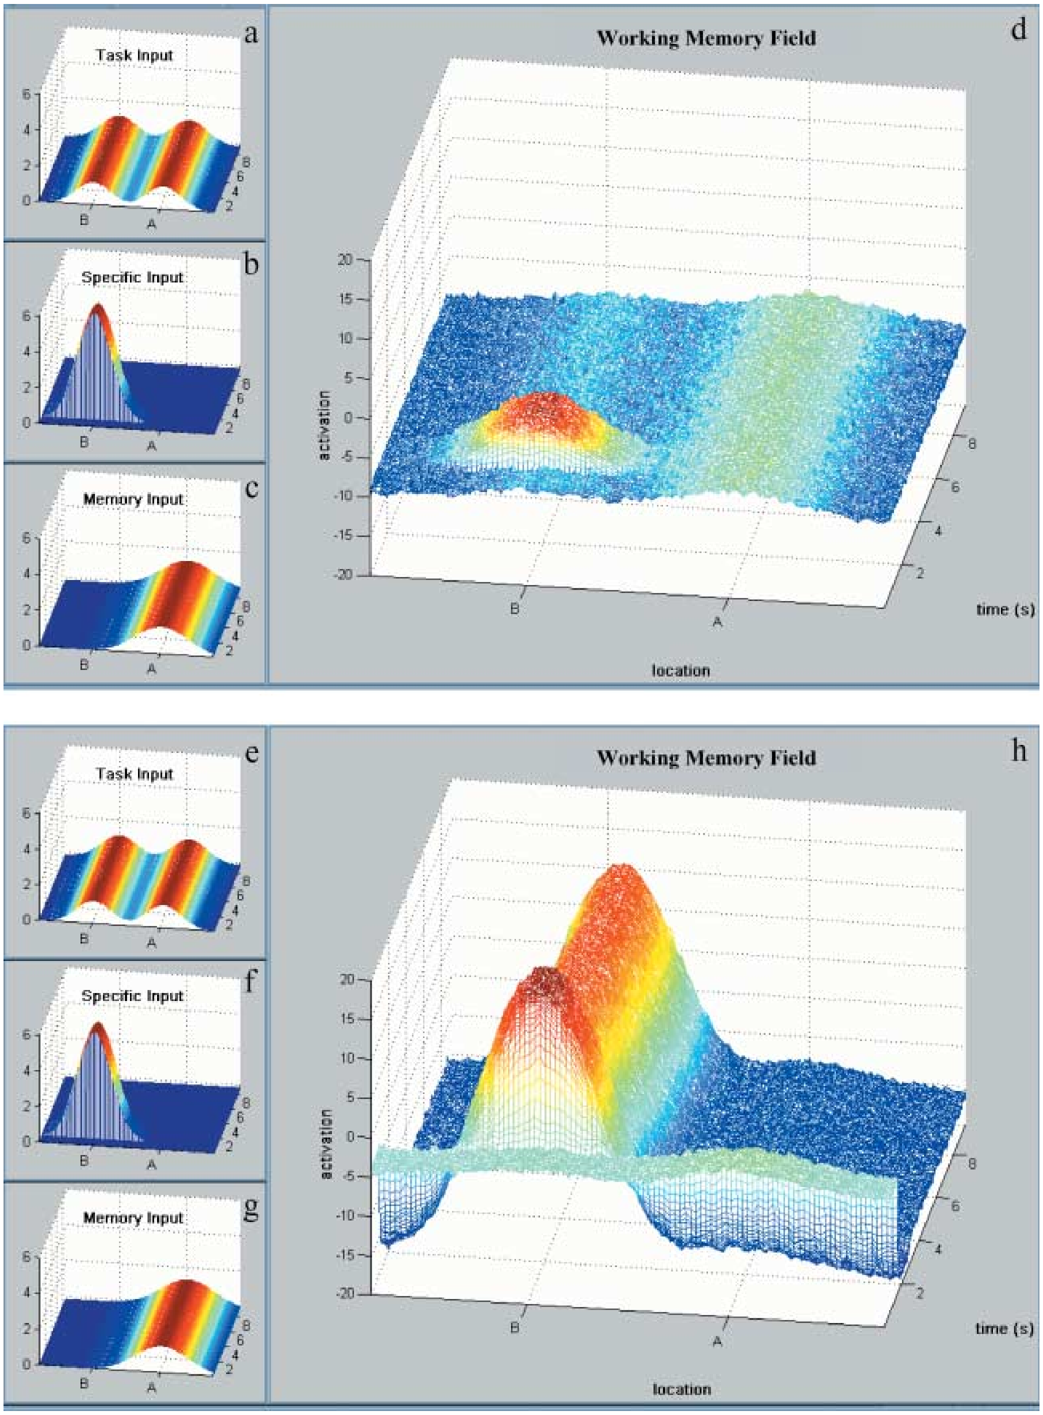

Figure 3 *Dynamic field theory of the A-not-B error. Top panel shows time-dependent changes in a spatial working memory field (d) in the context of three inputs (a, b, c) during the first B trial for an 8–10-month-old infant. Bottom panel shows changes in spatial working memory (h) in the context of the three inputs (e, f, g) during the first B trial for older infants. See text for further details.*

a peak of activation in the field at the B location (see Figure 3d). At 4 s, the toy is hidden. In the absence of strong specific input, the activation peak at B in the field *decays*. Consequently, by the time the infant is allowed to reach (at 10 s), activation at A – driven by the memory input – dominates and the infant reaches to A. That is, the infant makes the A-not-B error.

The bottom panel of Figure 3 shows a simulation of the model that captures the performance of older infants (e.g. 12 months). Notice, first, that the inputs in the left column (Figures 3e, 3f, 3g) are identical to the inputs shown in the top panel. However, the working memory field behaves quite differently (see Figure 3h). When the toy is held up and hidden at the B location, the strong specific input at the B location is amplified in the field. Critically, the resultant activation peak in the field is *maintained during the delay*, even after the toy is hidden and the specific input goes to zero (see Figure 3f). Thus, even in the absence of strong input, the field retains a memory of the hiding event at B. Consequently, after the delay (at 10 s), the older infant reaches correctly to the B location.

What accounts for the qualitatively different patterns of activation in the two working memory fields shown in Figure 3? This shift in behavior reflects a difference in the attractor states in which the two models operate caused by differences in how activation spreads among sites in the field. Underlying the performance of each field is a local excitation/lateral inhibition function. According to this function, an activated site in the field will increase the activation of neighboring sites (local excitation) and suppress the activity at far away sites (lateral inhibition). The field in Figure 3d is only weakly interactive: sites do not have a strong effect on one another. Instead, patterns of activation through time are largely determined by input. By contrast, the field in Figure 3h is strongly interactive. That is, sites in the field interact so strongly that patterns of activation can take on a life of their own – even after input has disappeared, sites can continue to stably excite one another within the local region initially stimulated by input. Thus, underlying the shift in infants' performance captured by these simulations is a change in how stably information is represented.

It is important to note two characteristics of the developmental switch from weak to strong interaction shown in Figure 3. First, it only takes a small, quantitative change in the interaction function to switch dynamic fields from operating in the weakly interactive mode (Figure 3d) to the strongly interactive mode (Figure 3h). Thus, as is the case with many dynamic systems accounts of developmental phenomena, a small change in the parameters of the model can lead to qualitatively different behaviors over development (e.g. van der Maas & Molenaar, 1992; van Geert, 1998). Second, although changes in interaction can be realized with a small parameter change, there are also other ways to create self-sustaining local excitation. For instance, self-sustaining activation can be created by very strong input. Thus, the developmental change depicted here should not be considered an all-or-none developmental switch. Rather, this type of developmental change is likely context- and experience-dependent (Smith *et al.*, 1999; Thelen *et al.*, 2001; Thelen & Smith, 1994).

Beyond infancy: a dynamic field theory of the development of spatial working memory

An important aspect of Thelen *et al.*'s account of infants' behavior in the A-not-B situation is that such behavior reflects the operation of general processes that make goal-directed actions to remembered locations. As such, the principles captured by the dynamic field model should not be specific to a particular task or to a particular period in development. Recent work by Spencer and colleagues (e.g. Schutte & Spencer, 2002; Spencer, Smith & Thelen, 2001) demonstrates that this is indeed the case. Moreover, this body of evidence suggests that the same developmental insights that capture changes in infants' performance in the A-not-B task can account for both qualitative and quantitative changes in the stability of spatial working memory processes later in development.

As a first step toward moving beyond infancy, Spencer, Smith and Thelen (2001) examined 2-year-olds' responses in an A-not-B version of a sandbox task. In this task, children watch as a toy is buried somewhere in a long, rectangular sandbox (see also, Huttenlocher, Newcombe & Sandberg, 1994). The toy is covered up, there is a short delay, and the child is allowed to search for the toy. Spencer *et al.* hid toys repeatedly at an A location. These trials were followed by several trials to a nearby B location. Results showed that 2-year-olds – an age at which children no longer reliably show the A-not-B error – were biased to point in the direction of the A location on the B trials. Moreover, as is the case with 8- to 10-month-old infants (Smith *et al.*, 1999), the magnitude of this bias depended on the number of trials to the A location, that is, on the strength of the longer-term memory of A. More recent work has demonstrated that these A-not-B-type biases occur with children as old as 6 years of age (Schutte, Spencer & Schöner, in press).

How might the dynamic field model account for these biases? More specifically, why do 2-year-olds show A-not-B-type errors in the sandbox task, but not in the canonical Piagetian task? Shouldn't 2-year-olds, like

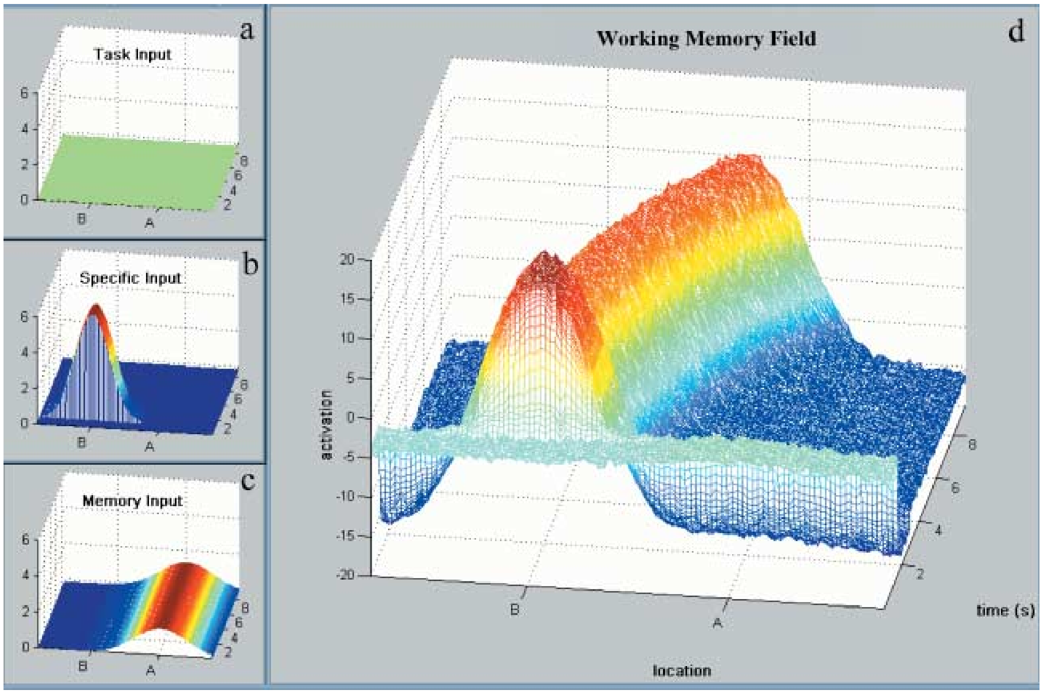

Figure 4 *Dynamic field theory with homogeneous task input. Figure details are identical to Figure 3 with one exception – the task input (a) was set to zero. See text for further details.*

12-month-olds in Thelen *et al.*'s account, be able to sustain target-specific activation in working memory in the sandbox task? The solution to these puzzles lies in an emergent feature of dynamic fields that operate in the self-sustaining state: self-sustaining activation peaks can show *systematic spatial drift during memory delays* (Schutte & Spencer, 2002; Spencer & Schöner, 2000).

Figure 4 shows a simulation of the dynamic field model with the same parameter setting used in the lower panel of the previous simulations (see Figure 3h). The only difference here is that we have modified the task input: we have set this input to zero since there are no salient location cues in the sandbox (see Figure 4a). Figure 4d shows how this change in the task input modifies the behavior of the working memory field on the first B trial (following several trials to the A location). At the start of the trial, the toy is hidden at the B location. The strong target input generates a self-sustaining peak of activation in working memory centered at the target location. During the memory delay, however, the self-sustaining peak begins to drift toward the A location. This occurs because the self-sustaining peak 'feels' stronger excitation on its A-ward side due to the longer-term memory input at A. This extra bit of excitatory input recruits new sites into the self-sustaining interaction and the peak drifts toward the A location. This type of drift did not occur in Figure 3h because the location of the self-sustaining peak was stabilized by the task input at the B location. This explains why 2-year-olds do not make the A-not-B error in the canonical Piagetian task with marked hiding locations but do show biases toward A in the sandbox task. And, importantly, these task-specific differences in performance emerge out of the same general processes.

The delay-dependent nature of the spatial drift in Figure 4 suggests that biases toward previously responded-to locations should increase systematically as the memory delay is varied. This is indeed the case with 3-year-olds (Schutte & Spencer, 2002), 6- and 11-year-olds (Hund & Spencer, 2003; Spencer & Hund, in press), and adults (Spencer & Hund, 2002). Importantly, however, the magnitude of delay-dependent drift becomes quantitatively smaller over development. For instance, in a task similar to the sandbox task with adjacent targets separated by 20°, 3-year-olds erred roughly 10° toward A on the B trials following a 10 s delay (Schutte & Spencer, 2002). Six- and 11-year-old children showed comparable effects in the same task; however, their biases were typically 2°–4° (Spencer & Hund, in press). Finally, adults' errors were quite small – biases toward previously responded-to targets tended to be roughly 1°–2° (Spencer & Hund, 2002).

These quantitative changes raise the question of what is changing over development to create a more accurate memory of locations. Spencer, Schutte and Schöner (Schutte, Spencer & Schöner, in press; Spencer & Schöner, 2000; Spencer & Schöner, 2003) proposed that such effects could be accounted for by narrowing the local excitation/lateral inhibition function that underlies site-to-site interactions in the field. A reduction in A-not-B-type biases then occurs for two reasons. First, with narrower interaction, it is less likely that self-sustaining peaks will overlap with activation associated with the longer-term memory input. Consequently, activation peaks in working memory will not feel the effects of such input. Second, dynamic fields with narrower local excitation and, conversely, more extensive lateral inhibition have more stable self-sustaining peaks (for related effects, see Compte, Brunel, Goldman-Rakic & Wang, 2000). More stable peaks are more resistant to delay-dependent drift. Consistent with this second point, there is a systematic reduction in the variability of children's and adults' responses from 6 years into adulthood (Hund & Spencer, 2003; Spencer & Hund, in press, 2002).

A final body of work has demonstrated that the narrower interaction view of development can also account for a qualitative shift in spatial categorization performance between 3 and 6 years. Several studies have shown that young children's (e.g. 3-year-olds') responses in spatial recall tasks are biased toward the midline of the task space (e.g. toward the midline symmetry axis of the sandbox) (Huttenlocher *et al.*, 1994; Schutte & Spencer, 2002; Spencer *et al.*, 2001). Older children (e.g. 6-year-olds), by contrast, are biased away from the midline of the task space (Hund & Spencer, 2003; Huttenlocher *et al.*, 1994; Sandberg, Huttenlocher & Newcombe, 1996; Spencer & Hund, in press). These children show a type of categorization behavior, where they appear to group locations into left and right regions. A recently-developed version of the dynamic field theory can account for this qualitative shift in performance (Spencer *et al.*, 2002). Although the details of this new model are beyond the scope of the current paper, it is relevant here because the qualitative change in categorization behavior by the model occurs over development as site-to-site interactions are narrowed. Thus, the same mechanism in the model that produces quantitative changes in A-not-B-type effects and variable errors over development can also account for qualitative changes in spatial categorization.

Discussion: evaluation of the dynamic field approach

Connections to the motor approach

The dynamic field approach is a departure from dynamic models of motor control and motor development in that dynamic fields represent information as attractor patterns of activation. This overcomes a major limitation of the motor approach. Nevertheless, the concepts of the motor approach remain at the core of the dynamic field perspective. Dynamic fields are, after all, dynamic systems.

Central to the dynamic field approach is the concept of stability. Dynamic fields provide a theoretical language that captures how a network of neurons can solve the stability problem, that is, how a network of neurons can stably maintain a pattern of activation in the face of perturbations. Formally, this occurs when activation in the field goes from a stable resting state through an instability (bifurcation) into a new attractor state – the self-sustaining state. Dynamic fields also embody the concepts of multi-causality and self-organization. Attractor states in dynamic fields emerge from the confluence of inputs and the intrinsic excitatory/inhibitory properties of intra-field interactions. Thus, as noted above, a dynamic field can enter the self-sustaining state in many ways. This can occur because a strong input is presented (e.g. a toy), or because multiple weak cues combine (e.g. weak perceptual cues plus a weak longer-term memory).

Importantly, input-related effects always depend on the current state of activation in the field and the nature of intra-field interactions. For instance, when information about a B location is being actively maintained, weak perceptual inputs at A may be suppressed via lateral inhibition – a form of decision-making (choosing B over A). This may not be the case earlier in the trial, however, when a weak cue at A might 'win' over a weak cue at B. Such multi-causal, self-organizing properties of dynamic fields create the potential for an impressive amount of behavioral flexibility, flexibility realized in our extensions of the dynamic field approach beyond infancy.

Additional possibilities for behavioral flexibility emerge within the context of the multi-level view of stability considered previously. Now, rather than thinking about multiple levels of control solely in the context of motor control and motor development, one can consider multiple levels of representation by coupling multiple dynamic fields together. Steinhage and Schöner (1998), for instance, have used dynamic fields to organize the behaviors of autonomous robots. Such robots are able to acquire targets in the world while avoiding obstacles, make decisions about which targets to acquire, and plan paths to acquire multiple targets. This impressive behavioral flexibility originates from a multi-leveled dynamic system containing multiple dynamic fields as well as more classic dynamic systems involved in, for instance, the control of the robot's effectors.

These ideas about the nature of cognitive processes and the relationship between cognition and action have fundamental implications for cognitive development.

First, they provide a framework for thinking about a classic Piagetian issue: how cognition can emerge from sensorimotor origins over development (see Thelen *et al.*, 2001). Because representational states in the dynamic field approach are closely tied to sensory input and motor output, it is possible for sensorimotor phenomena to help create such states. This is the case, for example, when strong perceptual cues move a dynamic field from the input-driven mode to the self-sustaining mode in the context of the A-not-B situation. Second, the dynamic field view of cognition moves the focus away from what children 'know' at different points in development toward an understanding of which cognitive states children reliably enter and how those states come about via the integration of information in the world, the child's past history in the task, and so on.

Beyond the motor approach

Dynamic fields extend the motor approach through the inclusion of representational states that have close ties to neurophysiology. Representation within dynamic fields emerges from the self-organizing properties of a population of neurons whose activation is time- and context-dependent. Importantly, this population activity does not merely form an imprint of sensory events or transform sensory input into motor output. Rather, representational states can take on their own intrinsic flavor. In this sense, representation in the dynamic field approach is consistent with other dynamic approaches to representational states in cortex (e.g. Skarda & Freeman, 1987).

Recent approaches have demonstrated, moreover, that the state of dynamic fields can be directly estimated from firing rates of populations of cortical neurons using population coding ideas (Erlhagen, Bastian, Jancke, Riehle & Schöner, 1999). For example, the activation of neurons in motor cortex (e.g. Georgopoulos, Kettner & Schwartz, 1988; Georgopoulos, Taira & Lukashin, 1993), premotor cortex (e.g. di Pellegrino & Wise, 1993), and prefrontal cortex (e.g. di Pellegrino & Wise, 1993; Graziano, Hu & Gross, 1997; Wilson, Scalaidhe & Goldman-Rakic, 1993) is broadly tuned such that neurons respond maximally to stimulation at a 'preferred' location, and less vigorously as stimulation is moved away from the preferred location. A population representation can be constructed by lining these neurons up, not according to their cortical locations, but according to their preferred spatial locations. Then, the activation of the newly aligned neurons in some task can be plotted through time and the resulting distributions of activation can be compared to activation profiles predicted by dynamic fields. These techniques have been used to directly observe, for instance, the representation of movement direction in motor and premotor cortex, providing evidence for the presence of preactivation when precues are given (Bastian, Riehle, Erlhagen & Schöner, 1998). In addition to this population coding approach, the general principles of dynamic fields can be usefully integrated with a more biophysical approach that attempts to incorporate the details of neurotransmitter action, timing properties of neurons, etc. (Compte *et al.*, 2000).

Although these connections to neural processes are exciting, it is important to note that dynamic field models are descriptive in the same sense we described previously. That is, these models still embody decisions by the experimenter about which behavioral dimensions are relevant to the behavior in question. Thus, while dynamic fields certainly interface more clearly with neurophysiology, this does not obviate the need for tough theoretical decisions about the appropriate level of description for a given behavioral system.

A third limitation of the motor approach discussed previously was that dynamic systems ideas are typically applied to development in a purely conceptual or metaphorical way. The dynamic field approach clearly moves toward a formal theory of infant perseverative reaching and spatial working memory. Nevertheless, it is important not to undersell the role conceptual thinking played in the development of this formal theory. This point is dramatically demonstrated by the evolution of the A-not-B model. Smith, Thelen and colleagues (Smith *et al.*, 1999; Thelen & Smith, 1994) published a re-thinking of the A-not-B error based solely on the conceptual application of dynamic systems concepts to this error. This conceptual analysis was then formalized by Thelen *et al.* (2001) using dynamic field concepts. And both the conceptual and formal theories have generated a variety of novel experiments (e.g. Diedrich, Highlands, Thelen & Smith, 2001; Diedrich *et al.*, 2000; Smith *et al.*, 1999). Thus, the evolution of the A-not-B story can be usefully compared to the now classic story that emerged from the study of finger twiddling and relative coordination. In both cases, there was a tight interplay among empirical work, conceptual theory and formal theory.

Connections to connectionism

Shared concerns between connectionism and the redefined dynamic systems approach

The dynamic systems approach to development pioneered by Thelen, Smith and colleagues developed out of the area of motor control. Consequently, this view began with a strong anti-representational stance largely directed against unreasonable symbolic motor programming accounts of movement preparation. Moreover, this early approach was abstract, attempting to capture global patterns in behavior, rather than how such patterns were realized by different physiological systems. The newer dynamic field approach embraces the notion of representational states and the dynamics of cognition in a way that interfaces with neurophysiology. In this regard, comparisons with connectionism are facilitated.

The major shared concern of both approaches is the drive to rethink the basis of cognition, taking into account from the start the constraints of the nervous system, the effector and sensor structure that this nervous system is connected to, and the structure of the specific environment in which this nervous system acts, learns and develops. The dynamic field concepts add a neurophysiologically plausible, sub-symbolic form of representation to the dynamic approach that is tightly linked to the sensory and motor 'surfaces'. The dynamic field representation is sub-symbolic because it does not solve 'matching' problems of mapping current sensory input onto a particular representation. Instead, that matching is achieved as an emergent property: when input is sufficiently close to current activational states in the field, matching or blending of activations occurs. In this, there are close ties to connectionism whose 'distributed representations' are likewise sub-symbolic, graded and relatively close to sensory and motor processes. The field approach makes explicit the metric properties of the underlying behavioral dimensions, which is not always the case in connectionist models. However, this is certainly not excluded within the connectionist approach (see, e.g. Bullock & Grossberg, 1988; Bullock, Grossberg & Guenther, 1993).

Furthermore, we argue that in both approaches there is a notion of 'emergence' of higher cognitive functioning. In the dynamic systems framework, new modes of behavior and representation emerge multi-causally through instabilities. In connectionist models, learning may lead to the extraction of invariances that are contained in the environment or the system–environment interaction. Emergence is also used in connectionist thinking in a slightly different sense, such as when epochs of different effective rates of learning emerge from a time-invariant learning rule. Because learning has not played a comparable role in the dynamic systems approach to motor control and development, emergent properties of learning dynamics have not been investigated, although instabilities probably form a common mathematical basis (see, e.g. Amari, 1989).

With regard to development, there is dramatic conceptual overlap (see Bates & Thelen, this issue). Both approaches view development as a step-by-step emergent process. There is also a strong emphasis on understanding the details of children's performance in a variety of situations, rather than stressing the competencies of children at different ages. In this way, these developmental approaches have tried to simplify the representational demands necessary for complex behavior, because not all knowledge must be stored in the brain. Finally, both approaches stress the importance of nonlinear processes in development. Nonlinear systems provide fundamental insights into how development can be both gradual and sudden at the same time.

Differences between the dynamic systems and connectionist approaches

The dynamic field approach carried over from the earlier motor approach the idea that instabilities mark qualitative changes then of behavior, now of representational states. Instabilities disrupt the one-to-one, input-output mapping and make dynamic fields non-standard neural networks. Dynamic fields are, however, a special case of competitive dynamic neural networks (e.g. Amari, 1977; Amari & Arbib, 1977; Hopfield, 1982) with a special emphasis on instabilities and metrics. Within this context, therefore, the evolution of activation in connectionist networks is described in essentially the same way as in the dynamic field approach, although missing a good understanding of instabilities.

The lack of emphasis on stability in connectionism has often led to oversimplified views of sensorimotor activity (overly simple forms of 'input'), and an impoverished treatment of the real-time coupling among perception, action and cognition (a simplification of real time via 'steps' in connectionist nets with no well-specified time scale). The challenge to scale up connectionist models from simple 'toy' simulations to simulations that would work with real-world sensors and actuators is probably surmountable, but not trivial (see Bullock & Grossberg, 1988; Bullock *et al.*, 1993 for examples that move in this direction). The dynamic field approach has already surmounted that challenge (e.g. Bicho, Mallet & Schöner, 2000). In this sense, the field approach is closer to addressing how cognition emerges from sensorimotor origins over development (see Thelen *et al.*, 2001).

The differences in emphasis between the two approaches have, in part, led researchers to gravitate toward different problems. We have focused on 'lower' levels of cognition with an eye toward the idea that higher forms of cognition might emerge from lower forms over development. By contrast, connectionist modelers have, in some cases, tried to deal with 'higher' level forms of cognition. In this case, extraction of relevant information about the world is already assumed. Indeed, some connectionist approaches to higher level cognition have built amodal, arbitrary symbols into the network architecture (for examples, see Rumelhart, McClelland & the PDP Research Group, 1986; see Barsalou, 1999, for a critique of such approaches). Such networks, though certainly not representative of connectionism *en masse*, clearly differ from the approach discussed here which is grounded in sensorimotor phenomena. This grounding is important, not only to circumvent the symbol grounding problem (see Barsalou, 1999), but also to account for the many reports where cognition takes on a sensorimotor character (e.g. Freyd, 1983).

In the context of symbolic representation and symbol grounding, it is also important to acknowledge that the dynamic systems and connectionist approaches to development share different histories. Although there are many important historical similarities described in the Thelen and Bates article (this issue), generally speaking, the dynamic approach described here is closely allied with Gibsonian approaches, while connectionism has historically been connected with Information Processing. To the extent that dynamic systems and connectionist researchers still strongly share such allegiances, important differences will remain. We, for instance, disagree with several core assumptions of the Information Processing approach, including the emphasis on 'computation' and recursive decomposition (see Palmer & Kimchi, 1986). Nevertheless, we have moved away from Gibsonian approaches here by acknowledging the importance of representational states. In this sense, we are explicitly trying to tackle some of the ground typically occupied by the Information Processing approach. In general, we don't see these historical roots as precluding unification of dynamic systems and connectionist approaches. Rather, we see the historical context as a point of caution: as the two camps move forward, we should not always assume that we are using concepts in the same way.

A more practical issue resulting from the differential emphasis on stability/instability in the two approaches is how measures of variability are used to constrain formal models. Within the dynamic systems approach, variability provides an important index of behavioral stability. For instance, variable errors played a central role in decisions about how to modify our dynamic field model of spatial working memory over development. The fact that both constant and variable errors decreased suggested that development might best be captured by changing the stability properties of working memory, rather than other characteristics (e.g. coupling between working memory and longer-term memory). Our perception is that variable errors have been less informative for connectionist approaches.

A related weakness of some forms of connectionist modeling is the lack of analytical understanding of connectionist networks, which has sometimes obscured where constraints come from and what aspects of models are insightful (but see, Smolensky, Mozer & Rumelhart, 1996, for a mathematical treatment of neural networks). For instance, the worst abuse of connectionist networks is when they are essentially universal approximators, so that they may learn any input–output relationship. The fact that they may learn the one learned by nervous systems then means rigorously nothing (as they could have learned any other one too, including those relationships nervous systems cannot learn).

The primary weakness of the dynamic systems framework discussed here is the limited amount of serious work on the processes of learning and of development. Although it has been stated that there are dynamics at many time scales, there is no coherent framework at the moment for how the slower scales of development can be characterized and identified in experiment. We believe this can be addressed (and have presented some first examples), but connectionism appears to be more developed here. Moreover, there has been little serious work within the dynamic approach on the processes of adaptation and selection. In their book, Thelen and Smith (1994) provide a useful dynamic framework that incorporates these processes, drawing heavily from Edelman's ideas about neuronal selection (e.g. Edelman, 1987). Nevertheless, such ideas have not, as of yet, been formally handled within the dynamic approach sketched here.

Dynamic systems and connectionism beyond the redefined dynamic approach

It is on the front of learning and development that we also see the largest point of departure between the dynamic approach described here and other approaches to dynamic systems theory and development. Several researchers have applied dynamic concepts to understanding the macrostructure of development, including developmental transitions/stages (see, e.g. Fischer & Bidell, 1998; van Geert, 1998). For instance, van der Maas and Molenaar (1992) have used catastrophe theory to examine the characteristics of developmental stages. These researchers have proposed a set of 'flags' that should be empirically present if a particular developmental transition is due to a qualitative shift in the attractor structure underlying behavior. Similarly, van Geert (1998) has used logistic growth equations to model qualitative transitions in development.

We share with these approaches the idea that developmental transitions can be thought of as global changes in the attractor structure of a system. Our work on working memory fits with this view in that qualitative changes in the stability of representational states in working memory differentiate some transitions in development (for related ideas, see Case, 1998). However, we have focused more on the real-time, microstructure of behavior than on this macrostructure of development. In this sense, these other dynamic systems approaches share a common interest with connectionism. As such, these approaches might provide useful insights as dynamic systems theorists and connectionist researchers seek to integrate micro and macro views of development.

With an eye toward this goal, we see two important limitations of other dynamic systems approaches that must be considered thoughtfully as we look to the future. First, these other approaches have sometimes taken on an analogical flavor, without strong ties to empirical work. Van Geert (1998), for example, has demonstrated that at a macro level, developmental stages look like transitions in logistic growth curves. Although this analogy has been useful in highlighting how qualitative and quantitative changes over development relate, this approach has been only loosely tied to empirical data. The 'flags' approach described above has fared better on this front. This approach has been used to search for empirical markers of particular classes of developmental transitions, and, in some cases, empirical data have been linked up with formal models of transitions (see Hartelman *et al.*, 1998; van der Maas & Hopkins, 1998). However, it is not always clear what the empirical agenda is after one has identified a particular type of transition. Some recent work has tried to manipulate potential 'control' parameters that move children through particular transitions (see Hartelman *et al.*, 1998). Although this work is promising, our sense is that, because the flags approach is grounded in the macro structure of development, it can be difficult to move easily from models to new experiments. Again, we emphasize that these approaches have contributed to an understanding of the macro structure of development. We raise these limitations simply to highlight the challenges ahead as we seek to link micro and macro levels.

Another important issue as we look to the future is the manner in which representations have been used in other approaches to dynamic systems and development. Several dynamic systems researchers have adopted the Piagetian idea that cognition can be conceptualized in terms of formal representational structures (e.g. Fischer & Bidell, 1998; van Geert, 1998). Such structural representations often take on a disembodied, non-metric flavor that appears to differ from the sense of representation described here. This is not always the case, however. For instance, Fischer and colleagues (1993) have embedded structural representations into a connectionist architecture that relies on graded representational states. Thus, we see the issue of representation more as a future challenge than as something that precludes bringing micro and macro approaches together.

Conclusion

The central goal of this manuscript was to use the dynamic field theory to build a representational bridge from the dynamic systems approach to motor control and development to our redefined dynamic approach. In the context of the special issue, this served to facilitate comparisons between our dynamic systems approach and connectionism, comparisons which have been difficult to make in the past because the approaches typically examine very different aspects of development. By tracing the shared emphases of the motor and dynamic field approaches, we highlighted how stability plays a central role in the dynamic approach. Connectionism has emphasized this concept less strongly. Moreover, we highlighted how the redefined dynamic approach integrates perception, action and cognition. By staying close to the sensorimotor world, the dynamic approach offers insights into how cognition might emerge from perception-action over development. Again, connectionism has emphasized this issue less strongly. Finally, we pointed toward a fundamental limitation of our dynamic approach at this time: we have not proposed a coherent account of the dynamics of learning and development. Connectionism and other dynamic systems approaches appear to be more advanced on this front.

We conclude by emphasizing, again, the deep conceptual overlap between dynamic systems and connectionist approaches. This overlap provides the foundation for future convergence, as researchers working within the different camps learn from the strengths and limitations of each approach. Our dynamic approach must address more profoundly learning and development. Connectionism must address the metrics of representations, as well as the central importance of stability and instability (and here embrace dynamic systems models). In 20 years, we don't think these will be two distinguishable approaches.

Acknowledgements

We would like to thank Esther Thelen, Kurt Fischer and Marc Lewis for helpful comments on earlier versions of this manuscript. In addition, Linda Smith and Larissa Samuelson provided a useful castigation for our use of the 'R-word' in an earlier (and current?) version. This research was supported by NIMH RO1 MH62480 and NSF BCS 00-91757 awarded to John P. Spencer.

References

Adolph, K.E. (1997). Learning in the development of infant locomotion. *Monographs of the Society for Research in Child Development*, **62**, 1–140.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, **27**, 77–87.

Amari, S. (1989). Dynamical stability of formation of cortical maps. In M.A. Arbib & S. Amari (Eds.), *Dynamic interactions in neural networks: Models and data* (pp. 15–34). New York: Springer.

Amari, S., & Arbib, M.A. (1977). Competition and cooperation in neural nets. In J. Metzler (Ed.), *Systems neuroscience* (pp. 119–165). New York: Academic Press.

Arbib, M.A. (Ed.) (1995). *Handbook of brain theory and neural networks*. Cambridge, MA: MIT Press.

Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral & Brain Sciences*, **22**, 577–660.

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.

Bates, E., & Elman, J. (2000). The ontogeny and phylogeny of language: a neural network perspective. In S.T. Parker & J. Langer (Eds.), *Biology, brains, and behavior: The evolution of human development. School of American Research advanced seminar series* (pp. 89–130). Santa Fe, NM: School of American Research Press.

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.

Bidell, T.R., & Fischer, K.W. (2000). The role of cognitive structure in the development of behavioral control: a dynamic skills approach. In W.J. Perrig & A. Grob (Eds.), *Control of human behavior, mental processes, and consciousness: Essays in honor of the 60th birthday of August Flammer* (pp. 183–201). Mahwah, NJ: Lawrence Erlbaum Associates.

Braun, M. (1994). *Differential equations and their applications* (4 edn.). New York: Springer Verlag.

Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological Review*, **95**, 49–90.

Bullock, D., Grossberg, S., & Guenther, F.H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, **5**, 408–435.

Case, R. (1998). The development of conceptual structures. In W. Damon (Editor-in-Chief), D. Kuhn & R.S. Siegler (Eds.), *Handbook of child psychology (fifth edition) Volume 2: Cognition, perception, and language* (pp. 745–800). New York: Wiley.

Churchland, P.S., & Sejnowski, T.J. (1992). *The computational brain*. Cambridge, MA: Bradford Book/MIT Press.

Clark, A. (1997). The dynamical challenge. *Cognitive Science*, **21**, 461–481.

Compte, A., Brunel, N., Goldman-Rakic, P.S., & Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, **10**, 910–923.

Corbetta, D., & Thelen, E. (1996). The developmental origins of bimanual coordination: a dynamic perspective. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 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). Boston: MIT Press.

Diamond, A., Cruttenden, L., & Neiderman, D. (1994). AB with multiple wells: 1. Why are multiple wells sometimes easier than two wells? 2. Memory or memory + inhibition? *Developmental Psychology*, **30**, 192–205.

Diamond, A., & Doar, B. (1989). The performance of human infants on a measure of frontal cortex function, the delayed response task. *Developmental Psychobiology*, **22**, 271–294.

Diedrich, F.J., Highlands, T., Thelen, E., & Smith, L.B. (2001). The role of target distinctiveness in infant perseverative reaching errors. *Journal of Experimental Child Psychology*, **78**, 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**, 479–494.

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.

Edelman, G.M. (1987). *Neural Darwinism*. New York: Basic Books.

Elman, J.L. (2001). Connectionism and language acquisition. In M. Tomasello & E. Bates (Eds.), *Language development: The essential readings* (pp. 295–306). Malden, MA: Blackwell Publishers.

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.

Fagard, J., & Jacquet, A.Y. (1996). Changes in reaching and grasping objects of different sizes between 7 and 13 months. *British Journal of Developmental Psychology*, **14**, 65–78.

Fischer, K.W., & Bidell, T.R. (1998). Dynamic development of psychological structures in action and thought. In W. Damon (Editor-in-Chief) & R.M. Lerner (Eds.), *Handbook of child psychology (fifth edition) Volume 1: Theoretical models of human development* (pp. 467–561). New York: Wiley.

Fischer, K.W., Rotenberg, E.J., Bullock, D.H., & Raya, P. (1993). The dynamics of competence: how context contributes directly to skill. In R.H. Wozniak & K.W. Fischer

(Eds.), *Development in context: Acting and thinking in specific environments* (pp. 93–117). Hillsdale, NJ: Erlbaum.

Fogel, A., Nwokah, E., Dedo, J.Y., & Messinger, D. (1992). Social process theory of emotion: a dynamic systems approach. *Social Development*, **1**, 122–142.

Freyd, J. (1983). The mental representation of movement when static stimuli are viewed. *Perception and Psychophysics*, **33**, 575–581.

Georgopoulos, A.P., Kettner, R.E., & Schwartz, A.B. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, **8** (8), 2928–2937.

Georgopoulos, A.P., Taira, M., & Lukashin, A.V. (1993). Cognitive neurophysiology of the motor cortex. *Science*, **260**, 47–52.

Gibson, E.J. (1991). *An odyssey in learning and perception*. Cambridge, MA: MIT Press.

Gracco, V.L., & Abbs, J.H. (1988). Central patterning of speech movements. *Experimental Brain Research*, **71**, 515–526.

Graziano, M.S.A., Hu, X.T., & Gross, C.G. (1997). Coding the locations of objects in the dark. *Science*, **277**, 239–241.

Hartelman, P., van der Maas, H.L.J., & Molenaar, P.C.M. (1998). Detecting and modelling developmental transitions. *British Journal of Developmental Psychology*, **16**, 97–122.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, **79**, 2554–2558.

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**, 3–38.

Huttenlocher, J., Newcombe, N., & Sandberg, E.H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, **27**, 115–147.

Jones, S.S. (1996). Imitation or exploration? Young infants' matching of adults' oral gestures. *Child Development*, **67**, 1952–1969.

Kelso, J.A.S. (1984). Phase transitions and critical behavior in human bimanual coordination. *American Journal of Physiology: Regulatory, Integrative, and Comparative Physiology*, **15**, 1000–1004.

Kelso, J.A.S., Ding, M., & Schöner, G. (1993). Dynamic pattern formation: a tutorial. In L.B. Smith & E. Thelen (Eds.), *A dynamic systems approach to development: Applications* (pp. 13–50). Boston, MA: MIT Press.

Kelso, J.A.S., Putnam, C.A., & Goodman, D. (1983). On the space–time structure of human interlimb coordination. *Quarterly Journal of Experimental Psychology*, **35A**, 347–375.

Kelso, J.A.S., Scholz, J.P., & Schöner, G. (1986). Non-equilibrium phase transitions in coordinated biological motion: critical fluctuations. *Physics Letters A*, **118**, 279–284.

Kelso, J.A.S., Tuller, B.T., Vatikiotis-Bateson, E., & Fowler, C.A. (1984). Functionally specific articulatory cooperation following jaw perturbations during speech: evidence for coordinative structures. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 812–832.

Latash, M. (1993). *Control of human movement*. Chicago, IL: Human Kinetics.

Lee, D.N., & Aronson, E. (1974). Visual proprioceptive control of standing in human infants. *Perception and Psychophysics*, **15**, 529–532.

Lewin, K. (1936). *Principles of topological psychology*. New York: McGraw-Hill.

Lewin, K. (1946). Behavior and development as a function of the total situation. In L. Carmichael (Ed.), *Manual of child psychology* (pp. 791–844). New York: Wiley.

Lewis, M.D. (2000). The promise of dynamic systems approaches for an integrated account of human development. *Child Development*, **71**, 36–43.

Lewis, M.D., Lamey, A.V., & Douglas, L. (1999). A new dynamic systems method for the analysis of early socioemotional development. *Developmental Science*, **2**, 457–475.

McClelland, J.L., & Jenkins, E. (1991). Nature, nurture, and connections: implications of connectionist models for cognitive development. In K. VanLehn (Ed.), *Architectures for intelligence: The twenty-second Carnegie Mellon symposium on cognition* (pp. 41–73). Hillsdale, NJ: Lawrence Erlbaum Associates.

McLeod, P., Plunkett, K., & Rolls, E.T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford: Oxford University Press.

Mareschal, D., Plunkett, K., & Harris, P.L. (1999). A computational and neuropsychological account of object oriented behaviors in infancy. *Developmental Science*, **2**, 306–317.

Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: a PDP model of the AB task. *Developmental Science*, **1**, 161–184.

Newcombe, N., Huttenlocher, J., Drummey, A., & Wiley, J. (1998). The development of spatial location coding: place learning and dead reckoning in the second and third years. *Cognitive Development*, **13**, 185–200.

Newell, K.M., & Molenaar, P.C.M. (Eds.) (1998). *Applications of nonlinear dynamics to developmental process modeling*. Mahwah, NJ: Lawrence Erlbaum Associates.

O'Reilly, R.C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

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: Lawrence Erlbaum Associates.

Piaget, J. (1952). *The origins of intelligence in children*. Madison, CN: International Universities Press.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.

Rumelhart, D.E., McClelland, J.L., & The PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.

Sandberg, E.H., Huttenlocher, J., & Newcombe, N. (1996). The development of hierarchical representation of two-dimensional space. *Child Development*, **67**, 721–739.

Schöner, G. (1989). Learning and recall in a dynamic theory of coordination patterns. *Biological Cybernetics*, **62**, 39–54.

Schöner, G. (1995). Recent developments and problems in human movement science and their conceptual implications. *Ecological Psychology*, **7** (4), 291–314.

Schöner, G. (2002). Timing, clocks, and dynamical systems. *Brain and Cognition*, **48**, 31–51.

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. (1988a). Dynamic pattern generation in behavioral and neural systems. *Science*, **239**, 1513–1520.

Schöner, G., & Kelso, J.A.S. (1988b). A dynamic theory of behavioral change. *Journal of Theoretical Biology*, **135**, 501–524.

Schöner, G., Kopecz, K., & Erlhagen, W. (1997). The dynamic neural field theory of motor programming: arm and eye movements. In P.G. Morasso & V. Sanguineti (Eds.), *Self-organization, computational maps and motor control* (Vol. 119, pp. 271–310). Amsterdam: Elsevier-North Holland.

Schutte, A.R., & Spencer, J.P. (2002). Generalizing the dynamic field theory of the A-not-B error beyond infancy: three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, **73**, 377–404.

Schutte, A.R., Spencer, J.P., & Schöner, G. (in press). Testing the dynamic field theory: working memory for locations becomes more spatially precise over development. *Child Development*.

Shultz, T.R. (1998). A computational analysis of conservation. *Developmental Science*, **1**, 103–126.

Siegler, R.S. (1994). Cognitive variability: a key to understanding cognitive development. *Current Directions in Psychological Science*, **3**, 1–5.

Skarda, C.A., & Freeman, W.J. (1987). How brains make chaos in order to make sense of the world. *Behavioral & Brain Sciences*, **10** (2), 161–195.

Smith, L.B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: the task dynamics of the A-not-B error. *Psychological Review*, **106**, 235–260.

Smolensky, P., Mozer, M.C., & Rumelhart, D.E. (1996). *Mathematical perspectives on neural networks*. Mahwah, NJ: Lawrence Erlbaum.

Spencer, J.P., & Hund, A.M. (in press). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*.

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., & Schöner, G. (2000). A dynamic field model of location memory. Paper presented at the 22nd Annual Conference of the Cognitive Science Society, Philadelphia, PA.

Spencer, J.P., & Schöner, G. (2003). A dynamic field theory of spatial working memory and its development. Manuscript in preparation.

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.

Spencer, J.P., & Thelen, E. (2000). Spatially-specific changes in infants' muscle coactivity as they learn to reach. *Infancy*, **1**, 275–302.

Spencer, J.P., Veriejken, B., Diedrich, F., & Thelen, E. (2000). Posture and the emergence of manual skills. *Developmental Science*, **3**, 216–233.

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). Bellingham, WA: SPIE Publishing.

Sternad, D., Collins, D., & Turvey, M.T. (1995). The detuning factor in the dynamics of interlimb rhythmic coordination. *Biological Cybernetics*, **73** (1), 27–35.

Thelen, E. (1992). Development as a dynamic system. *Current Directions in Psychological Science*, **1**, 189–193.

Thelen, E. (1995). Motor development: a new synthesis. *American Psychologist*, **50**, 79–95.

Thelen, E. (2000). Motor development as foundation and future of developmental psychology. *International Journal of Behavioral Development*, **24**, 385–397.

Thelen, E., Corbetta, D., Kamm, K., Spencer, J.P., Schneider, K., & Zernicke, R.F. (1993). The transition to reaching: mapping intention and intrinsic dynamics. *Child Development*, **64**, 1058–1098.

Thelen, E., Corbetta, D., & Spencer, J.P. (1996). Development of reaching during the first year: role of movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, **22** (5), 1059–1076.

Thelen, E., Schöner, G., Scheier, C., & Smith, L.B. (2001). The dynamics of embodiment: a dynamic field theory of infant perseverative reaching errors. *Behavioral and Brain Sciences*, **24**, 1–86.

Thelen, E., & Smith, L.B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.

Thelen, E., & Smith, L.B. (1998). Dynamic systems theories. In W. Damon (Editor-in-Chief) & R.M. Lerner (Eds.), *Handbook of child psychology (fifth edition) Volume 1: Theoretical models of human development* (pp. 563–634). New York: Wiley.

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*, **56**, 1–104.

Tuller, B.T., & Kelso, J.A.S. (1989). Environmentally-specified patterns of movement coordination in normal and split brain subjects. *Experimental Brain Research*, **75**, 306–316.

Van der Maas, H.L.J., & Hopkins, B. (1998). Developmental transitions: so what's new? *British Journal of Developmental Psychology*, **16**, 1–13.

Van der Maas, H.L.J., & Molenaar, P.C.M. (1992). Stagewise cognitive development: an application of catastrophe theory. *Psychological Review*, **99**, 395–417.

Van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, Vygotsky, and beyond. *Psychological Review*, **105**, 634–677.

Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral & Brain Sciences*, **21**, 615–665.

Williams, R.J. (1986). The logic of activation functions. In D.E. Rumelhart, J.L. McClelland, & The PDP Research Group (Eds.), *Parallel distributed processing*. Volume 1 (pp. 423–443). Cambridge, MA: MIT Press.

Wilson, F.A.W., Scalaidhe, S.P., & Goldman-Rakic, P.S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science*, **260**, 1955–1958.

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. (1994). The coordination dynamics of learning: theoretical structure and experimental agenda. In S.P. Swinnen & H. Heuer (Eds.), *Interlimb coordination: Neural, dynamical, and cognitive constraints* (pp. 461–490). San Diego, CA: Academic Press.