Contributions of dynamic systems theory to cognitive development

John P. Spencer [a,*], Andrew Austin [a], Anne R. Schutte [b]

[a] *Department of Psychology and Delta Center, University of Iowa, Iowa City, IA, United States*
[b] *Department of Psychology, University of Nebraska–Lincoln, Lincoln, NE, United States*

ARTICLE INFO

ABSTRACT

We examine the contributions of dynamic systems theory to the field of cognitive development, focusing on modeling using dynamic neural fields. After introducing central concepts of dynamic field theory (DFT), we probe empirical predictions and findings around two examples—the DFT of infant perseverative reaching that explains Piaget's A-not-B error and the DFT of spatial memory that explain changes in spatial cognition in early development. Review of the literature around these examples reveals that computational modeling is having an impact on empirical research in cognitive development; however, this impact does not extend to neural and clinical research. Moreover, there is a tendency for researchers to interpret models narrowly, anchoring them to specific tasks. We conclude on an optimistic note, encouraging both theoreticians and experimentalists to work toward a more theory-driven future.

Mathematical modeling of human behavior has a long history dating back to the early 19th century (Fechner, 1860; Weber, 1842–1853). The history of formal modeling in developmental science is much shorter. Thus, this special issue offers a welcome opportunity to evaluate the contributions of computational modeling to developmental science in its infancy, when prospects for the future are just beginning to come into focus.

Our article emphasizes a particular type of computational modeling using dynamic neural fields (DNFs) that has emerged from the broader framework of dynamic systems theory (DST). Our goal

* Corresponding author at: E11 Seashore Hall, Department of Psychology, University of Iowa, Iowa City, IA 52242, United States. Tel.: +1 319 335 2482; fax: +1 319 335 0191.
 E-mail address: John-spencer@uiowa.edu (J.P. Spencer).

is to highlight how dynamic field theory (DFT) has been useful in understanding cognitive development and generating new empirical predictions and findings. We focus on two examples—the DFT of infant perseverative reaching (Thelen, Schöner, Scheier, & Smith, 2001) proposed to explain the classic Piagetian A-not-B error, and the DFT of spatial memory used to explain changes in spatial cognition in early development (Spencer, Simmering, Schutte, & Schöner, 2007). These examples are ideal in the context of this special issue because in each case there are alternative formal theories. This allows us to discuss the impact of DFT in particular as well as the impact of formal modeling more generally. Analyses of the literature in each domain reveal that computational modeling is making inroads into mainstream cognitive development; however, there is much to be done to fully integrate formal approaches into mainstream cognitive development. Doing so will require effort from both theoreticians and experimentalists.

1. Dynamic systems theory overview

DST emerged within developmental science within the last 20 years. It is based on advances in physics, mathematics, biology, and chemistry that have changed our understanding of non-linear, complex systems (Prigogine & Stengers, 1984). The developmental concepts that underlie DST are based on pioneering work by Thelen and Smith (1994) as well as early work from other theoreticians such as Fischer (Fischer & Rose, 1996), Van Geert (1997, 1998), and Molenaar (Molenaar & Newell, 2010; Van der Maas & Molenaar, 1992).

DST has made major contributions to developmental science by formalizing multiple concepts central to a developmental systems perspective (Lerner, 2006). The first is that systems are self-organizing. Complex systems such as a developing child consist of many interacting elements that span multiple levels from the genetic to the neural, the behavioral, and the social. Interactions among elements within and across levels are nonlinear and time-dependent. Critically, such interactions have an intrinsic tendency to create pattern (Prigogine & Nicolis, 1971). Thus, there is no need to build pattern into the system ahead of time—developing systems are inherently creative, organizing themselves around special habitual states called "attractors."

The notion that human behavior is organized around habits dates back at least to James (1897). But DST helps formalize the more specific notion of an attractor, providing tools to characterize these special states (Van der Maas, 1993; Van der Maas & Molenaar, 1992). For instance, a typical way to characterize a habit is to simply measure how often the habitual state is visited. Importantly, DST has encouraged researchers to also measure how variable performance is around that state and whether the system stays in that state when actively perturbed. This is particularly revealing over learning and development because habits often become more stable—more resistant to perturbations—over time.

Within this context, DST also helps clarify the relation between two related concepts central to developmental science—qualitative and quantitative change (Spencer & Perone, 2008; Van Geert, 1998). Qualitative change occurs when there is a change in the number or type of attractors, for instance, going from one attractor state in a system to two. Such special changes—called bifurcations—can arise from gradual, quantitative changes in one aspect of the system. A simple example is the shift from walking to running. As speed quantitatively increases across this transition in behavior, a sudden and major reorganization of gait occurs having a qualitatively new arrangement of elements (Diedrich & Warren, 1995).

Gait changes are one of the classic examples first studied by researchers interested in applying the concepts of DST to human behavior. This early work led naturally to the use of dynamic systems concepts to explain transitions in motor skill both in real time and over learning and development (Adolph & Avolio, 2000; Fogel & Thelen, 1987; Thelen, 1995; Thelen, Corbetta, & Spencer, 1996; Whitall & Getschell, 1995). One conclusion from these studies is that the brain is not the "controller" of behavior. Rather, it is necessary to understand how the brain capitalizes on the dynamics of the body and how the body informs the brain in the construction of behavior (Thelen & Smith, 1994). This has led to an emphasis on embodied cognitive dynamics (Schöner, 2009; Spencer, Perone, & Johnson, 2009), that is, to a view of cognition in which brain and body are in continual dialogue. We return to this theme in our discussion of dynamic field theory.

Another DST concept that has been particularly salient in developmental science is the notion of "soft assembly." According to this concept, behavior is always assembled from multiple interacting components that can be freely combined from moment to moment on the basis of the context, task, and developmental history of the organism. Thelen characterized this as a form of improvisation in which components freely interact and assemble themselves in new, inventive ways, like musicians playing jazz, which gives behavior an intrinsic sense of exploration and flexibility (Spencer et al., 2006).

A final contribution of DST is the host of formal modeling tools that can capture and quantify the emergence and construction of behavior over development (such as growth models, oscillator models, dynamic neural field models), and statistical tools that can describe the patterns of behavior observed over development (Lewis, Lamey, & Douglas, 1999; Molenaar, Boomsma, & Dolan, 1993; Molenaar & Newell, 2010; Van der Maas et al., 2006). These tools have enabled researchers to move beyond characterizing what changes over development toward deeper understanding of how these changes occur.

1.1. Dynamic field theory: cognition and real-time neural dynamics

DST is very good at explaining the details of action, for instance, how infants transition from crawling to walking. Consequently, it has had a major impact in motor development. DST also provides a good fit with aspects of perception. For instance, there are elegant dynamic systems models of how the visual array changes as animals move through the environment that explain, for instance, when a gannet will pull in its wings when diving for a fish (Schöner, 1994).

But can DST explain something as complex as working memory, executive function, and language? This was a central challenge to the theory in the 1990s, following innovative studies applying DST to motor development. Several initial models captured cognition at a relatively abstract level of analysis. For instance, Van der Maas and Molenaar (1992) proposed a model that captured transitions in children's conservation behavior using a specific variant of DST called catastrophe theory. The model provided a quantitative analysis of stage-like transitions in thinking defined over abstract dimensions of cognitive level, perceptual salience, and cognitive capacity. Similarly, Van Geert (1998) proposed a model defined over the abstract dimension "developmental level" to reinterpret several classical concepts from Piaget's and Vygotsky's theories. Both approaches showed the promise of DST for offering new insights into classic questions—such as the nature of quantitative versus qualitative developmental change (Spencer & Perone, 2008)—and also highlighted the potential for integrating quantitative models and rich behavioral data sets.

A second group of dynamic systems models also moved into the foreground during the 1990s—connectionist models of development (Spencer, Thomas, McClelland, 2009). These models, highlighted in other contributions to this special issue, attempted to explain cognition at a less abstract level and interface with known properties of the brain. A third dynamic systems approach to cognition also emerged in the late 1990s—DFT. DFT represented an explicit effort to create an embodied approach to cognition building from and connecting to the dynamic systems concepts emerging in the fields of perceptual and motor development. Thus, DFT retains transparent ties to central dynamic systems concepts such as attractor states, bifurcations, and soft assembly. But it also offers a mechanistic-level understanding of how brains work with a well-specified perspective on how the brain and body work together to enable cognition and action in the world (Engels & Schöner, 1995; Lipinski, Sandamirskaya, & Schöner, 2009; Sandamirskaya & Schöner, 2010).

Fig. 1 provides an overview of the neural concepts that underlie the basic computational unit in DFT—a dynamic neural field (DNF). Fig. 1A shows an arrangement of excitatory neurons on a cortical surface in, for instance, parietal cortex. Each excitatory neuron has a receptive field that is, in this case, tuned to spatial information such as the direction of a stimulus relative to the head. As an example, the third neuron on the cortical sheet fires maximally when a leftward stimulus is presented—this is the neuron's "preferred" direction indicated by the vector. The neuron also fires when a less preferred stimulus direction is presented (e.g., slightly down and leftward), but it will fire at baseline levels when a rightward stimulus direction is presented. These tuning properties are captured by the tuning curve shown for neuron 3. The curve shows the average activation level interpolated across 360° of stimulus space (Georgopoulos, Lurito, Petrides, Schwartz, & Massey, 1989). Fig. 1A also shows
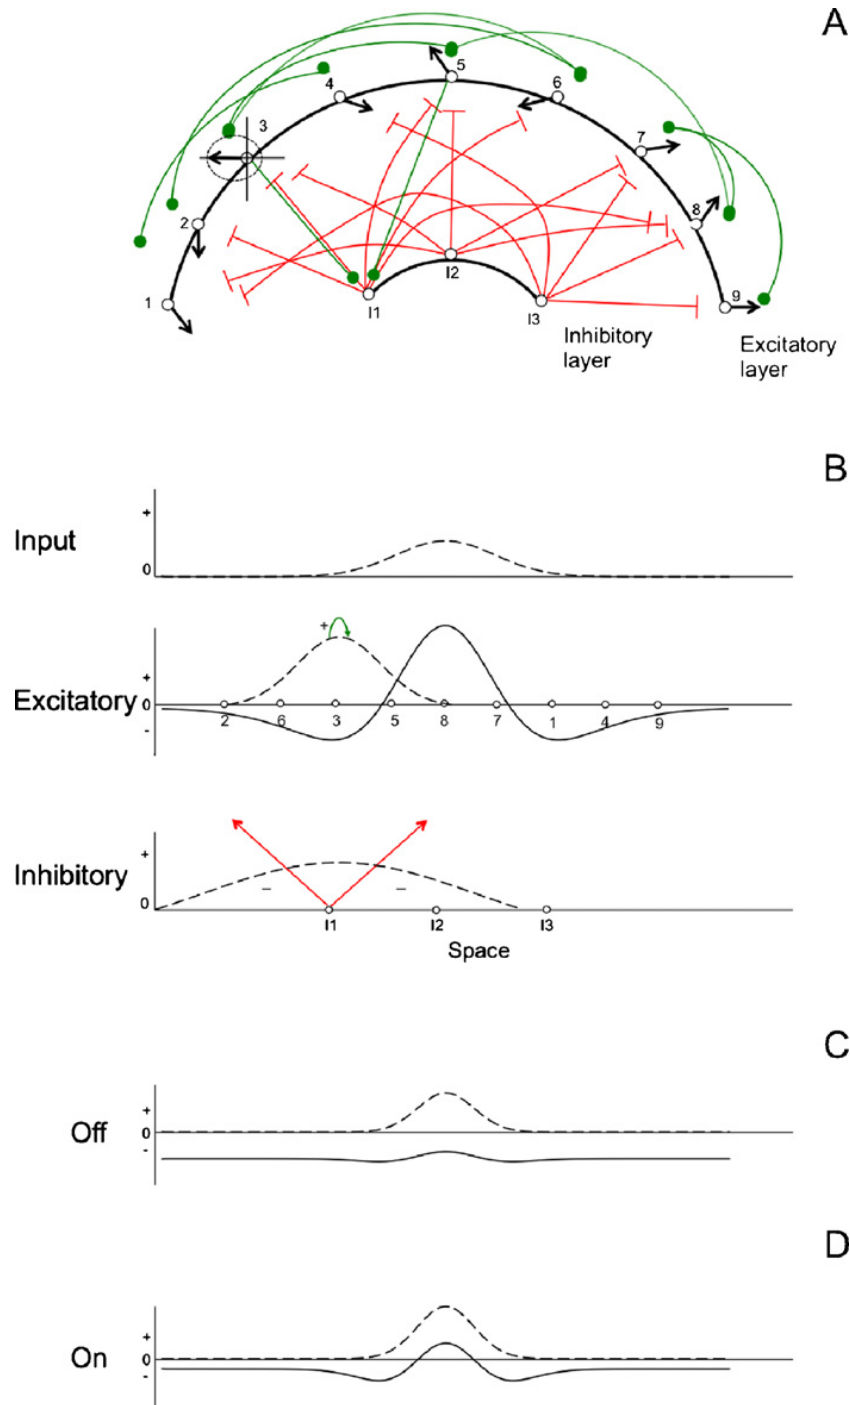
Fig. 1. Overview of concepts underlying dynamic neural fields (DNFs). (A) A layer of excitatory neurons coupled to a layer of inhibitory interneurons. Each excitatory neuron is tuned to a particular spatial direction, indicated by black arrows showing its preferred stimulus direction. Green connections are excitatory; red connections are inhibitory. (B) The system of cortical connections in (A) rearranged as a layered DNF architecture. Neurons are rearranged according to functional topography, such that neurons' preferred direction runs systematically left to right. The dashed line in the top panel of (B) shows input given to the model that forms a "peak" of activation in the excitatory layer. Dashed line in the excitatory layer shows the tuning curve of activation for neurons surrounding neuron 3. Dashed line in the bottom panel shows the broad projection of inhibition back into the excitatory layer. (C) The "off" state of the field with weak input. (D) The "on" state of the field. Here, slightly stronger input engages neural interactions forming a peak. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

"lateral" excitatory connections among the excitatory neurons in the cortical sheet (see green lines). For instance, neuron 3 excites neuron 5. This makes sense because both neurons prefer leftward stimulus directions.

The second layer in Fig. 1A shows a collection of inhibitory interneurons. These neurons are stimulated by neurons in the excitatory layer and inhibit other neurons in that layer. For instance, inhibitory neuron 1 is stimulated by neurons 3 and 5 in the excitatory layer (see excitatory synaptic connections in Fig. 1A; for simplicity, we omitted all remaining connections from the excitatory to the inhibitory layer). When inhibitory neuron 1 becomes active, it inhibits the firing of neurons 2, 6, 3, 5, and 8 in the excitatory layer. This keeps excitation from spreading uncontrollably in the excitatory layer. This is needed because stimulation of neuron 3 excites neuron 5, which excites neuron 8, which excites neuron 7, and so on.

Fig. 1B takes this cortical picture and reorganizes it to form a DNF. Rather than showing the neurons organized as they are on the cortical surface (Fig. 1A), Fig. 1B shows them organized by their functional connection properties (called a functional topographic map). Thus, neurons in the excitatory layer have preferred directions that run systematically from left to right. The interneurons are organized similarly. Next, each neuron's tuning curve is standardized to create the pattern of local excitation shown around neuron 3 (see dashed curve). This Gaussian distribution shows the strength of excitation passed between neuron 3 and its neighbors (e.g., neuron 5) once neuron 3 exceeds its firing threshold (i.e., its activation goes above 0). This occurs when, for instance, a leftward stimulus direction is presented. Neuron 3 also stimulates the first interneuron (see green arrow). Activation of inhibitory neuron 1 causes inhibition to be broadly spread to the excitatory layer (see red arrows) based on the tuning properties of this interneuron (see dashed curve).

The result of these volleys of locally excitatory and broadly inhibitory neural interactions is a "bump" (Edin, Macoveanu, Olesen, Tegnér, & Klingberg, 2007; Wang, 2001) or "peak" of activation—the basic unit of cognition in DFT. The top panel of Fig. 1B shows the input pattern (dotted line) generated when a stimulus is presented directly in front of the simulated neural system. This input stimulates neuron 5 a little, neuron 8 a lot, and neuron 7 a little. These neurons, in turn, excite one another and pass activation to the inhibitory layer. The inhibitory layer passes inhibition back. As these interactions play out over time, a peak forms (solid line in excitatory layer) that actively "represents" the stimulus direction at the level of the neural population. That is, the population of neurons "knows" at some moment in time that the stimulus is directly ahead.

Although this is an abstraction away from real neurons in the brain, all the steps noted can be reconstructed with real neural data. For instance, Bastian and colleagues (Bastian, Schöner, & Riehl, 2003) used multi-unit recordings in motor cortex to construct a dynamic neural field from neurophysiological data. Results were then used to quantitatively test predictions of a DFT of motor preparation (Erlhagen & Schöner, 2002). Thus, when asked whether we think the brain actually works like a dynamic neural field model, the answer is 'yes'—at least at the level of neural populations.

Fig. 1 follows from a population dynamics approach to neuroscience (Amari, 1977; Amari & Arbib, 1977; Wilson & Cowan, 1972); it also carries forward some of the key concepts of DST. For instance, in the lower panels of Fig. 1, the state of the excitatory layer follows two inputs differing very slightly in strength. In Fig. 1C, the input was sufficiently weak that strong local excitatory interactions were never engaged. We call this the "off" state because the activation pattern will relax back to the neuronal resting state given sufficient time. Fig. 1D shows what happens when the input is increased slightly—now a robust peak turns "on" and remains stable as long as the input remains present. Fig. 1C and D are qualitatively different—they are formally different attractor states—yet they arise from a small quantitative difference in the strength of neural input. This is one example of how DFT sheds light on the quantitative versus qualitative distinction in development: Small differences—a slightly more salient input—can lead to a non-linear or qualitative shift in responding, actively encoding the stimulus versus missing it altogether.

We can take this example a step further by noting multiple types of "on" states (i.e., different attractor states) within DFT. Fig. 2 shows two qualitatively different types of peak states that have played a central role in the examples we discuss here. The first example was generated by the same model shown in Fig. 1. The only difference here is that we show the model in action through time. In Fig. 2A, an input is presented 100 ms into the simulation. This builds a robust peak of activation
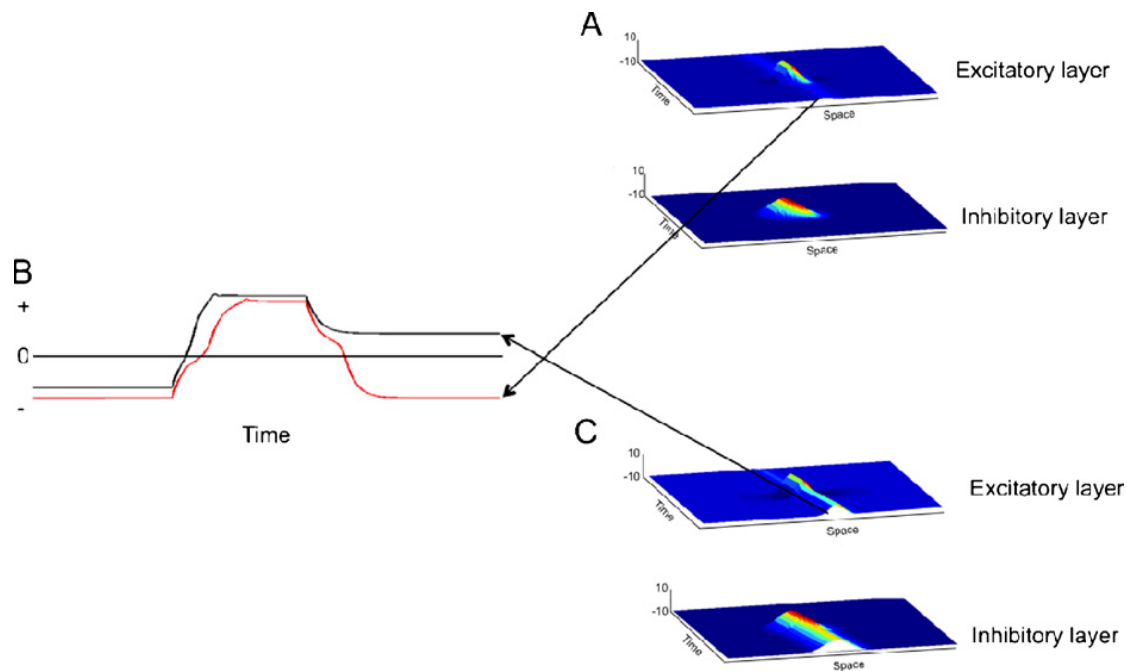
Fig. 2. Two qualitatively different states of dynamic neural fields: self-stabilized (A) and self-sustained (C) representations. The red line in (B) shows how activation at the center field site in (A) returns to baseline after the removal of the stimulus. The black line in (B) shows how activation at the center field site in (C) is sustained throughout the simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

in the excitatory layer. Critically, the peak becomes unstable once the input is removed and the field returns to its resting state (red line in Fig. 2B). This state is the self-stabilized state—the peak is only maintained when stabilized by an input; it mimics the properties of "encoding".

Fig. 2C shows a similar run of the model. Now, however, we have boosted the strength of locally excitatory interactions. As seen, this small change in excitation leads to a big difference in performance—the peak is maintained throughout the simulation, although the input was removed as in Fig. 2A (black line in Fig. 2B). This state is the self-sustaining state—the peak sustains itself in the absence of input; this mimics the properties of "working" or "active" memory.

Returning to a key issue raised earlier, in what sense is DFT "embodied?" That is, how can dynamic neural fields be integrated with perceptual and motor systems? Schöner and colleagues (Sandamirskaya & Schöner, 2010) have done extensive work on this topic and have developed a formal approach to embodiment that specifies the link between DNFs and motor control. For instance, Bicho and Schöner (1997), Bicho and Schöner (1998) have shown how DNFs can interface continuously with a motor system to enable an autonomous robot to navigate in real-world contexts. Although a formal treatment of this work is beyond our scope here (for a review, see Schöner, 2009), it highlights that DFT offers a rigorous approach to cognitive dynamics that is grounded in both neurophysiology and more classic approaches to motor control and development.

1.2. Dynamic field theory: learning and development

How does DFT theory address learning and development? Fig. 3A shows a variant of the DNF model from Fig. 1 with one layer added. We call this a Hebbian layer (HL) because it implements a form of Hebbian learning. Neurons in the excitatory layer are connected one-to-one with neurons in the HL. Consequently, when there is robust activation in the excitatory layer (i.e., neuron 8 > 0), activation begins to build at the associated site in the HL. This, in turn, projects activation back onto the excitatory layer (see green arrows from HL to the excitatory layer). Critically, activation in HL grows quite slowly, that is, over a learning timescale rather than over the timescale of, say, encoding or working memory maintenance. Moreover, when activation grows at some sites in HL, it decays at all other un-stimulated sites. Thus, this form of Hebbian learning is competitive. Note that sites in the HL operate more like
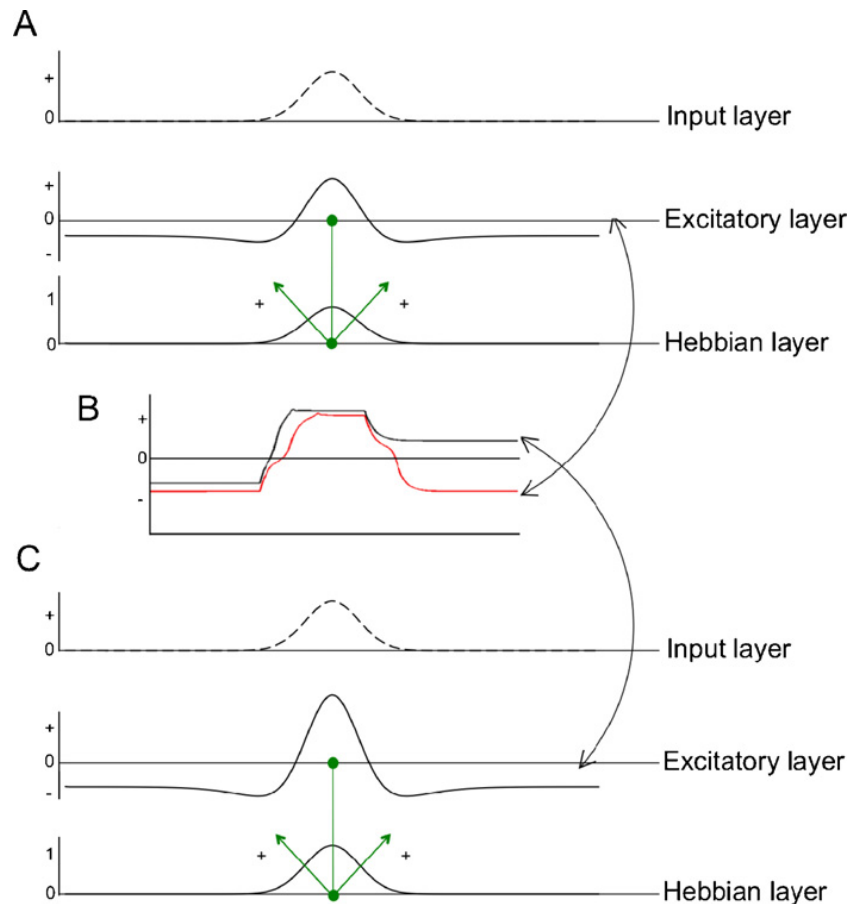
Fig. 3. Variation of the three-layer DNF architecture showing input to the model, a layer of excitatory neurons, and a Hebbian layer. Neurons in the excitatory layer are connected one-to-one to neurons in the Hebbian layer (green connections); these sites project activation back to the excitatory layer. When the excitatory layer is given a relatively weak input, activation from the Hebbian layer can help create a self-stabilized peak (A and red line in B). With slightly more Hebbian activation, the weak input can create a self-sustaining peak in the excitatory layer (C and black line in B). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

synapses than neurons, that is, they grow or shrink in activation continuously over a slow timescale rather than actively firing in real-time. In this sense, the HL is like a weight matrix in a connectionist network (Faubel & Schöner, 2008; Spencer, Dineva & Schöner, 2009).

What is the effect of activation in the HL? As in a connectionist network, stronger weights lead to stronger excitatory interactions among a local group of connected neurons. For instance, the pattern shown in Fig. 3A enhances processing of the input, leading to a robust peak. Indeed, in this example, we used the weak input from Fig. 1C. Recall that this weaker input failed to build a robust peak and the DNF stayed in the "off" state. Now, after a bit of learning, the DNF model can build a robust peak even when the input is less salient (red line in Fig. 3B).

Fig. 3C takes this learning one step further. The DNF model from Fig. 1 operated in the self-stabilized or "encoding" state. Consequently, when the input was removed, the peak returned to the resting level. After some learning, however, this same DNF model enters the self-sustaining state: With the extra excitation provided by a strong memory trace in the HL, the field is able to form a working memory for the stimulus and maintain this memory after the stimulus is removed (black line in Fig. 3B). Thus, learning can qualitatively alter the type of peak present in a neural field (Spencer & Perone, 2008).

Most of our efforts understanding learning within DFT have focused on properties of the DNF model shown in Fig. 3 (see, e.g., Lipinski, Simmering, Johnson, & Spencer, 2010; Lipinski, Spencer, & Samuelson, 2010). But what about development? Can DFT capture patterns of change that happen over weeks, months, and years? The first avenue we have explored on this front involves developing the feed-forward mapping from the input layer into the neural field. This type of slow learning—akin to perceptual learning— has been important in projects modeling infants' and young children's

performance in quantitative detail (Perone & Spencer, in press; Schutte & Spencer, 2009). In these cases, we have implemented this form of developmental change by changing the precision (i.e., the width) and strength of the input pattern. A second avenue involves changing the connection pattern within and across layers according to a spatial precision hypothesis (SPH) which posits that locally excitatory and laterally inhibitory interactions become stronger over development. Increasing the strength of neural interactions in excitatory and inhibitory layers (Perone & Spencer, in press; Schutte & Spencer, 2009) effectively captures both quantitative and qualitative developmental change.

2. Applications of dynamic field theory: the A-not-B error

We have highlighted the link between the psychological construct of "encoding" and self-stabilized peaks. Here we carry this idea forward in the context of a specific example—the Piagetian A-not-B error. In so doing we highlight how DFT has been useful in both understanding cognitive development and generating new empirical predictions and findings.

In *The construction of reality in the child*, Piaget (1954) described a hiding and finding game he played with his infants that has become a signature task in the study of infant cognition: the A-not-B task. An adult hides a toy repeatedly at location A. After some initial training, an 8–10-month-old infant will retrieve the toy from A when presented with two possible hiding locations following a brief delay (e.g., 3 s). After several successful hiding and finding events at A, the adult hides the toy at a nearby location B. At 8 months, infants reach for the toy at A; by 12 months they successfully search at B.

The A-not-B task has a long history, with over a hundred different experimental variants, providing fertile ground for theory development. Thelen and Smith (1994) proposed an account based on the dynamic systems concept of an attractor and a consideration of changes in reaching skill during the first year. They tested several implications of this dynamic systems account (see Smith, Thelen, Titzer, & McClin, 1999). Most radically, they predicted that infants would make the A-not-B error even when no toys were hidden at A or B. This prediction ran counter to the dominant explanations of this error at the time that centered on infants' concepts of objects and object permanence. As predicted, infants continued to make the A-not-B error although all objects were in plain sight.

Thelen et al. (2001) formalized this perspective in the DFT of infant perseverative reaching. This model consisted of the basic architecture shown in Fig. 3—an input layer, a dynamic neural field, and a Hebbian layer. The 8–10-month-old model had weak neural interactions such that peaks in the field—that is, decisions to reach to an "A" to the left or a "B" to the right—were self-stabilized: they relaxed to the neural resting level when input was removed. The older infant model, by contrast, had stronger neural interactions such that peaks were more likely to enter the self-sustaining state and be actively maintained during the delay. The model required only a small quantitative change to capture the dramatic shift in infants' perseverative tendencies (much like the small quantitative change required in Fig. 2 to shift from the "encoding" mode to the "working memory" mode).

The DFT of perseverative reaching successfully captured many effects from the literature including manipulations to the stable perceptual cues in the task (the appearance of the lids and hiding box), the salience of the cuing event (whether the adult drew attention to one side of the box or the other), and the influence of long-term memory built up from trial to trial. The model also generated several novel predictions. For instance, Clearfield, Dineva, Smith, Diedrich, & Thelen, 2009 predicted a set of interactions between delay and the salience of the cuing event in the A-not-B task. Most notably, the DNF model predicted that infants would show perseveration after *no* delay with a *weak* cue. This prediction contrasts with previous accounts in which delay is critical because the error is centrally about representation of the toy or hiding location. As predicted by the DFT, infants did perseverate in a no delay condition, but only with a weak cue. With a stronger cue and no delay, infants were accurate as reported in previous studies (Diamond, 1985; Wellman, Cross, & Bartsch, 1986). Thus, the DFT predicted a suite of effects in quantitative detail that involved the interaction of two factors. Such predictions provide a strong test of the theory.
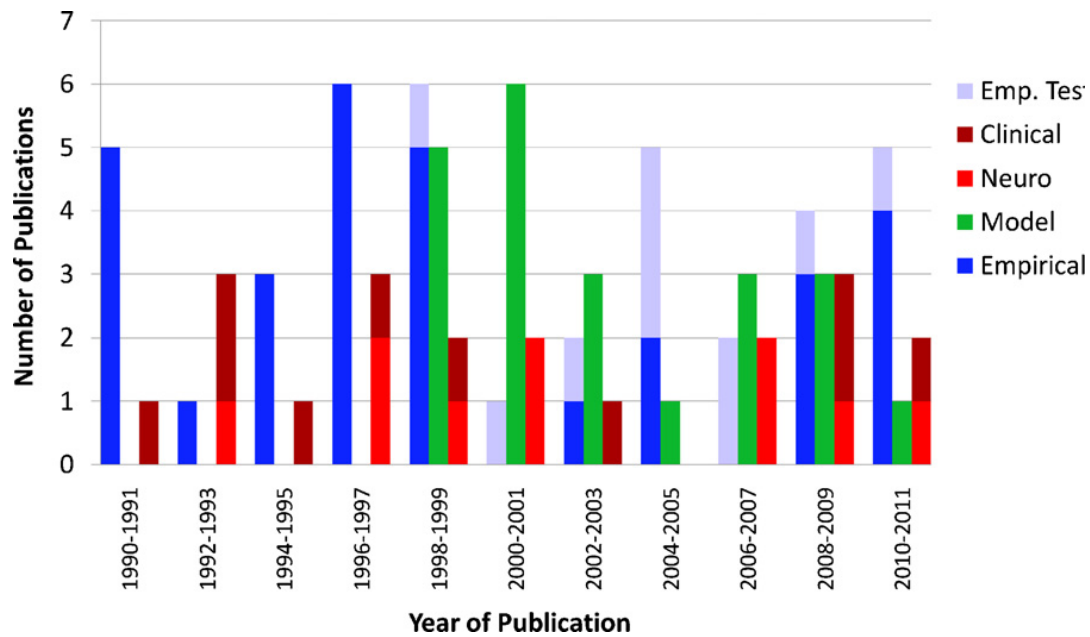
Fig. 4. Analysis of the A-not-B literature from 1990 to 2011.

2.1. Formal models of A-not-B: what do they contribute?

The previous example highlights one way to evaluate the usefulness of a theory. Does the theory explain a broad range of findings? In the case of the DFT of perseverative reaching, the answer is certainly 'yes'. One can also ask whether the theory leads to novel predictions. Again, in the case of the DFT, the answer is 'yes.' All of the studies noted, however, came from one large group of researchers. What has the DFT contributed beyond this group? Has the theory been useful to others? And beyond the DFT, are formal theories of the A-not-B error having an impact? We conducted a systematic review of the A-not-B literature during a two-decade span from 1990 to 2011. All articles were examined to ensure they were developmental studies (empirical or theoretical) that did more than tangentially mention the concept. The final sample included 82 articles. (A list and details of the search procedure are available from the authors.)

We classified these articles into five categories: modeling, neurophysiological/biological, clinical, empirical-only, and empirical tests of formal theories. The results are shown in Fig. 4. We classified articles as "modeling" if they included a formal mathematical model or were from one of the modeling groups and tested a key aspect of the proposed theoretical account. Included were articles examining Munakata's (1998) Parallel Distributed Processing model, as well as a connectionist model by Mareschal, Plunkett, and Harris (1999). Also included were Thelen et al.'s (2001) DFT of perseverative reaching and several articles by Zelazo, Marcovitch, and colleagues related to a hierarchical competing systems model (Marcovitch & Zelazo, 2006), as well as two developmental robotics papers. A model by Changeux and Dehaene (1989) was excluded because it fell outside of the temporal window we examined.

As shown in Fig. 4, publication of research employing formal models increased in 1998 with the publication of Munakata's PDP model. The rate of modeling publications has been steady since then and rivals or exceeds that of empirical research in some years. But have these formal models had an impact beyond the researchers who employ them? The neurophysiological/biological category comprised mostly electrophysiological research (Bell, 2001; Bell & Fox, 1992, 1997), as well as comparative (Diamond, 1990a, 1990b) and biological studies (Dettmer, Novak, Novak, Meyer, & Suomi, 2009). None of these used formal models to significantly inform or generate hypotheses. Indeed, none of the articles in this category after 1998 even cited a formal model. The clinical category consisted of research on atypically developing populations, with only a single article (Mauerberg-deCastro et al.,

2009) employing a formal model. Across the neurophysiological and clinical categories, then, formal models have had little impact.

The lack of contact between modeling work and these specialized areas of research is unfortunate because models may be particularly useful in these areas. For instance, there is clearly an opportunity to tie neural and biological studies of behavior to neurally-grounded models of behavior such as the DFT (Bastian et al., 2003; McDowell, Jeka, Schöner, & Hatfield, 2002). The clinical literature also represents a missed opportunity, as formal models can be useful tools to help evaluate the mechanisms thought to underlie psychopathology (Harm & Seidenburg, 1999; Joanisse & Seidenberg, 2003; Lewis & Elman, 2008; McMurray, Samelson, Lee, & Tomblin, 2010; Plaut, McClelland, Seidenburg, & Patterson 1996; Thomas & Karmiloff-Smith, 2003).

Another category of research we considered included all other empirical contributions, that is, research without neural/biological measures that assessed the performance of non-clinical populations by authors not directly associated with a modeling group. We categorized these by whether or not they used formal models to significantly inform or generate hypotheses. A study Berger (2004), for example, did so. It examined the effect of cognitive load on infant perseveration in a locomotor A-not-B task, while controlling for motor habit, which is central to the DFT. Similarly, Ruffman, Slade, Sandino, & Fletcher (2004) designed a series of experiments to differentiate between hypotheses generated from formal theories.

Because there were few formal computational models in the developmental A-not-B literature prior to 1998 (but see Changeux & Dehaene, 1989), we focused on research published from 2000 to 2011 (to give the modeling articles some time to impact the empirical literature). Within this period, 19 articles were classified in an empirical category. Of these, 9 (47%) were empirical tests of a formal model. These data are encouraging—models are clearly having an impact.

3. Application of dynamic field theory 2.0: the development of spatial cognition

We now examine how we have moved from the dynamic field model of the A-not-B task to a more general theory of the development of spatial cognition (Schutte & Spencer, 2009; Simmering, Schutte, & Spencer, 2008; Spencer et al., 2007). Note that this work has also been extended into the field of visual cognition, including perceptual discrimination and working memory for features (Johnson, Spencer, Luck & Schöner, 2009; Johnson, Spencer, & Schöner, 2009; Simmering & Spencer, 2008).

Achieving something as simple as remembering the location of a favorite toy can require a remarkable degree of sophistication. First, you need to perceive the location of the toy the first time you play with it. You can easily do this retinally, but that is not very useful because our eyes move around so much. Thus, children and adults commonly encode locations relative to a body-centered or world-centered frame of reference (e.g., a few inches to the right of the table's edge). Achieving this is a real trick because everything needs to be actively coordinated. For instance, to establish a neural representation in world-centered coordinates requires that you update the relationship between your body and the world every time you move (Pouget, Deneve, & Duhamel, 2002; Schneegans & Schöner, 2012). But once you encode the location in a world-centered frame, you need to remember the location during, for instance, a 20 s delay when the toy is occluded by your brother who has come over to interfere with your play time. This requires some form of active or working memory, and this memory has to be updatable as you move around the world and as objects move. Next, you need to turn that working memory into an action that will, for instance, get your hand to the toy after pushing your brother out of the way. Finally, you need to use encoding and memory abilities to do something longer-term—remember where the toy is located days and weeks later.

In what follows, we present an overview of a DFT of spatial cognition that can overcome many of these challenges. More detailed discussions are provided by Spencer et al. (2007), Schutte and Spencer (2009), and Lipinski, Schneegans, Sandamirskaya, Spencer, and Schöner (2011).

3.1. Real-time cognitive dynamics

As a first step in understanding how people remember the locations of objects, we used a simple task developed by Huttenlocher, Newcombe, and Sandberg (1994). It involves having children remember
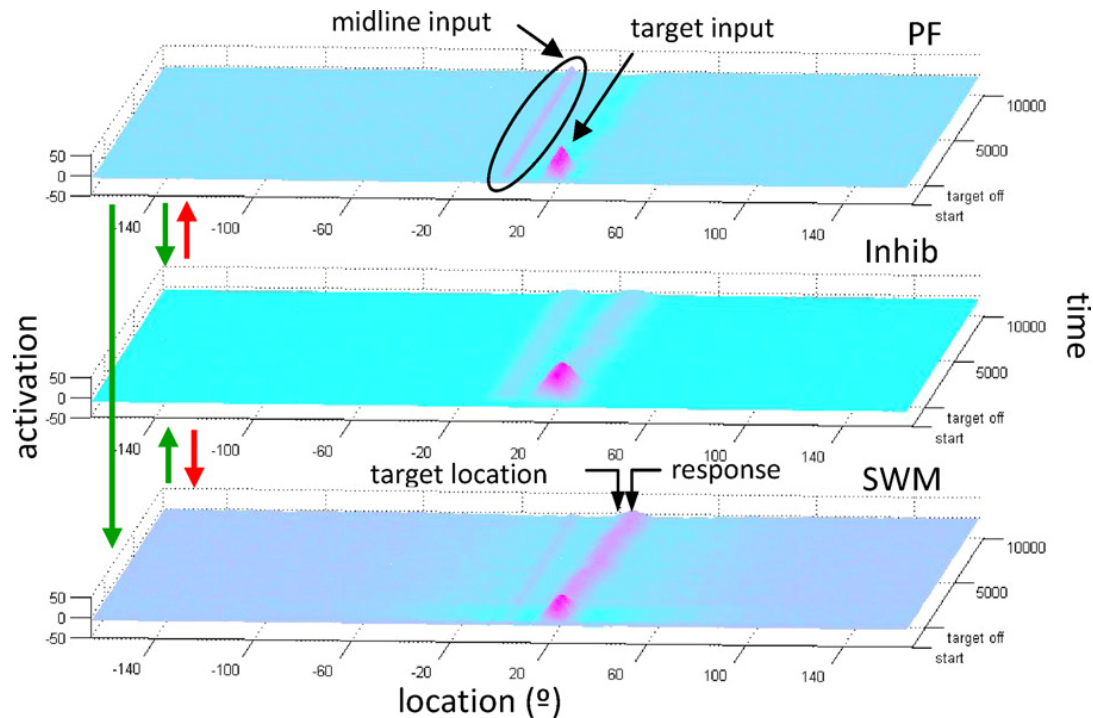
Fig. 5. Simulation of the 3-layer DNF model. Panels represent: perceptual field [PF]; inhibitory field [Inhib]; working memory field [SWM]. Arrows show connections between fields. Green arrows represent excitatory connections and red arrows represent inhibitory connections. In each field, location is represented along the *x*-axis (with midline at location 0), activation along the *y*-axis, and time along the *z*-axis. The trial begins at the front of the figure and moves toward the back. See text for additional details. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

the location of a toy buried in a long, narrow sandbox. The task is simple enough for 18-month-olds, yet it is empirically rich. To remember a location in a sandbox, one anchors memory to the available reference frame, the edges and symmetry axes of the box. This anchoring results in biases away from the reference axes and toward the center of, for instance, left and right spatial categories (Huttenlocher, Hedges, & Duncan, 1991; Spencer & Hund, 2002). When the location is aligned with a reference axis, errors are small and the variance across trials is small as well (Schutte & Spencer, 2002; Spencer & Hund, 2002).

The DFT of spatial cognition offers a detailed, mechanistic explanation of how people encode a location in a frame of reference, actively remember the location in working memory, and keep themselves in register with the perceptual surrounds such that they can reach for the hidden object a short time later (Lipinski et al., 2011; Spencer et al., 2007). For simplicity, we focus on the perceptual and working memory demands in the sandbox task using the three-layer architecture shown in Fig. 5. This DNF model consists of an excitatory perceptual field (PF; first layer of Fig. 5) that codes perceptual structure in the task space; an excitatory spatial working memory field (SWM; third layer of Fig. 5) that receives excitatory input from the perceptual field (green arrows) and maintains the memory of the target location; and an inhibitory field (Inhib; second layer of Fig. 5) that receives input from both the perceptual and SWM fields (green arrows) and sends inhibition back to both fields (red arrows). For each layer, location is represented along the *x*-axis, activation along the *y*-axis, and time along the *z*-axis.

Fig. 5 shows the model's performance in a single trial with a target 20° from the midline (0°) of the task space. When the target appears ("target input" in Fig. 5), peaks build in the excitatory fields at the target location. During the memory delay, the target peak in the perceptual field dies out; however, the peak in the SWM field self-sustains due to the stronger neural interactions in this layer. Note that a second peak is maintained in the perceptual field during the delay at the location of the nearest reference axis—the symmetry axis aligned with the center of the sandbox ("midline input"). This peak helps the model maintain neural activation in a world-centered reference frame (Spencer et al., 2007). Spatial memory biases in the sandbox task emerge from the interaction of the midline peak in the

perceptual field and the target peak in SWM during the delay. In particular, inhibition associated with the reference frame pushes or "repels" the WM peak away from the midline axis (arrows in the lower panel of Fig. 5 mark the target location and the response location).

The DNF model predicts that there should be a time-dependent "drift" of spatial memory during the delay, that is, memory should be more and more strongly repelled from midline as the delay increases. Spencer and Hund confirmed this prediction with both children (2003) and adults (2002). In addition, the model predicts that response variability should increase over the memory delay. This has also been empirically confirmed (Spencer & Hund, 2002, 2003). Finally, the model explains why there is no memory drift and low response variability when the target is aligned with the reference frame. In this case, excitatory input from the midline peak in the perceptual field attracts the WM peak, keeping it accurately positioned during the delay. The low response variability when a target is aligned with a perceived reference frame is theoretically important because a second account of spatial memory—the Category Adjustment Model (CAM; Huttenlocher et al., 1991; Huttenlocher et al., 1994)—does not capture this finding.

3.2. Spatial memory and learning

The model in Fig. 5 captures the real-time neural dynamics that underlie memory for a single location in a world-centered reference frame. But how do you learn, for instance, where your favorite toy typically is? For this, we need to add learning and long-term memory formation to the model in the form of a Hebbian layer, creating a "3-layer+" model. This brings the 3-layer model into register with the type of perseverative phenomena studied in the A-not-B task. For instance, Simmering et al. (2008) showed that the 3-layer+ model will make the A-not-B error, and Schutte and Spencer (2009) showed that the 3-layer+ model captures A-not-B-type biases in the sandbox task (Spencer, Smith, & Thelen, 2001).

To probe whether this addition to the DNF model effectively captures longer-term spatial learning, we conducted several studies with adults in another task pioneered by Huttenlocher et al. (1991). In this task, adults remember the location of a dot inside a circle on a computer screen. After a brief delay, they are asked to reproduce the location. An initial probe of spatial learning in this task revealed null results. Huttenlocher, Hedges, Corrigan, and Crawford (2004) gave two groups different dot distributions over trials. In one—the "+" distribution—dots were clustered near the horizontal and vertical symmetry axes. In a second—the "X" distribution—dots were clustered near the diagonals. If participants build up a long-term memory of the target distribution, one might expect opposite spatial memory biases in these conditions—toward the vertical and horizontal axes with the + distribution and toward the diagonal axes with the X distribution. Huttenlocher et al. (2004) found no systematic differences across conditions. They concluded that—consistent with the CAM—participants did not remember the target distributions in detail but rather used a prototype representation at the center of each quadrant.

A challenge in studying long-term memory is that effects often build up slowly over learning. Critically, participants in the Huttenlocher et al. (2004) study only responded once to each target. Lipinski, Simmering, et al. (2010) replicated key aspects of this study but had participants respond multiple times to each target. Results showed a systematic bias toward a long-term memory of each target distribution, that is, memory biases were pulled in opposite directions across conditions. These data were captured in quantitative detail by the 3-layer+ model. Because the CAM has no memory beyond a summary prototypical representation for each category, it fails to capture these results.

3.3. The development of spatial memory

Thus far, we have described how the DFT captures the integration of perception and working memory in real-time as well as over learning. But what about development? Much of our developmental work has focused on a transition in spatial category biases first documented by Huttenlocher et al. (1994), who found that memory was biased toward rather than away from midline in the sandbox task.

To explain this developmental change, Spencer, Schutte, and colleagues (Schutte & Spencer, 2009; Schutte, Spencer, & Schöner, 2003; Simmering et al., 2008; Spencer et al., 2007) proposed the spatial precision hypothesis which posits that the neural interactions that underlie spatial cognition become stronger over development; that is, locally excitatory interactions and laterally inhibitory interactions strengthen. When this hypothesis is implemented in the 3-layer model, the model captures a suite of developmental changes in spatial memory biases between 3 and 5 years (Simmering, Spencer, & Schöner, 2006). For instance, one consequence of stronger neural interactions is that peaks can become more precise. There thus should be systematic changes in precision in perceiving and remembering locations relative to a reference frame. Recent data confirm that perception of symmetry axes becomes more precise in early development during the same developmental period when there is a shift in spatial memory biases (Ortmann & Schutte, 2010).

Most surprisingly, simulations of the DFT that implemented the spatial precision hypothesis predicted that the developmental shift in spatial memory—from biases toward midline to biases away from midline—do not happen in an all-or-none fashion (Schutte & Spencer, 2009). Rather, the model predicted a complex pattern of change between 3 and 5 years from gradual narrowing of the bias toward midline, to a period of intermediate responding where targets very close and very far from midline are accurate but targets in an intermediate spatial zone show bias away from midline, to finally an expansion of the repulsion effect across a broad spatial range. Empirical data confirmed this pattern. Three-year-olds showed a bias toward midline across a broad range of targets (Schutte & Spencer, 2002). At 3 years 8 months, children's responses were biased toward midline only at the hiding location closest to midline; other locations were accurate (Schutte & Spencer, 2009). At 4 years 4 months, responses were biased away from midline only at an intermediate location and the other locations were accurate. Finally, by 5 years 4 months, responses were biased away from midline across a broad range that became broader still at 6 years (Spencer & Hund, 2002). In follow-up simulations, Schutte and Spencer (2009) showed that the 3-layer model can quantitatively simulate this complex pattern. At present, no other model explains this suite of developmental effects.

Schutte and Spencer (2010) tested an additional prediction that making the midline symmetry axis more salient would shift the transition in spatial memory biases to an earlier point in development. To test this, they added two lines to the midline axis of the task space. As predicted, the added perceptual structure switched the memory responses of children 3 years 8 months of age from bias toward midline to bias away from midline. This influence of perceptual structure on the developmental transition in memory bias provides further support for the specific type of integration of perception and working memory captured by the DFT.

3.4. Formal models of spatial memory: what do they contribute?

What has the DFT has contributed to the field of spatial cognition? First, it explains a broad range of findings and integrates the A-not-B error (Simmering et al., 2008), changes in position discrimination in early development (Simmering et al., 2006), and a suite of changes in spatial memory biases all within a single framework. Second, it has led to novel predictions that have been tested empirically. These include observations difficult to explain otherwise, such as the complex pattern of change in spatial memory biases between 3 and 5 years. Finally, the model effectively integrates perceptual, working memory, and long-term memory processes in a way that is predictive in real time (Spencer & Hund, 2003), over learning (Lipinski, Spencer, et al., 2010), and over development (Schutte & Spencer, 2009). This was our initial goal as we seek to develop a broad, integrative theory of spatial cognitive development.

Has this work had an impact on the field more generally? We conducted a second review of the literature covering a three-decade span from 1980 to 2011 (details available from the authors). We chose to include a wider temporal window because some of the early work on the CAM was published in the 1980s. The final sample included 168 articles that were classified into the same five categories used previously: modeling, neurophysiological/biological, clinical, empirical only, and empirical tests of formal theories. We classified articles as "modeling" if they included a formal mathematical model as part of the research or if they were from one of the modeling groups and probed issues related to the theoretical account. The modeling category included articles related to the CAM, as well as those
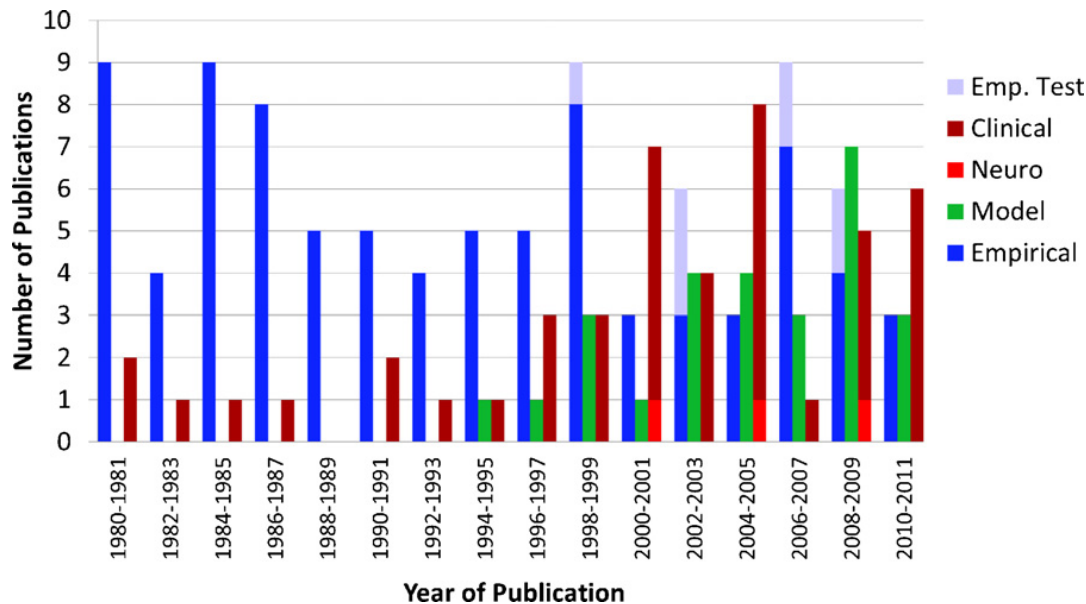
Fig. 6. Analysis of the spatial categorization/spatial memory literature from 1980 to 2011.

related to DFT. One modeling paper by Rotzer et al. (2009) used a neural network model of spatial working memory. As shown in Fig. 6, publication of research employing formal models began with the first application of the Category Adjustment Model to the study of development (Huttenlocher et al., 1994). Modeling work remained relatively dormant for several years until the early 2000s with the publication of the DFT account of spatial memory. Since then, the rate of modeling publications rivals that of empirical research and in some years exceeds it.

What about the non-modeling categories? For neural research, models have once again had relatively little impact. The neurophysiological/biological category comprised mostly fMRI studies (Schweinsburg, Nagel, & Tapert, 2005; Vuontela et al., 2009). None used formal models of spatial memory to inform or generate hypotheses. The clinical category consisted of research using atypically developing and other specialized populations (including some cross-cultural work). Once again, none made significant contact with formal models. However, many of them used tasks not obviously addressed by either the CAM or DFT (e.g., the Corsi block task).

The final empirical category, included papers that used formal models to inform or generate hypotheses. For instance, Hund and Plumert (2003) investigated the role of object categorization on spatial learning and categorization, drawing inspiration from the DFT. In another example, Nardini, Thomas, Knowland, Braddick, and Atkinson (2009) tested the effects of different frames of reference in an experiment that built on the CAM model (Huttenlocher et al., 1991; Huttenlocher et al., 1994).

As with our analysis of the A-not-B literature, we investigated the effects of models on the empirical literature by focusing on the 12-year period from 2000 to 2011. Of 30 empirical reports published during this period, 7 were classified as "empirical tests" (23%). Thus, models appear to be having a lesser impact on studies of spatial cognitive development than on the A-not-B literature. That said, it is important to place this result in the context of several observations. First, Fig. 6 shows a slight *increase* in papers in the modeling category toward the end of the 12-year period. Given that most of these studies directly test formal models with empirical data, the decline in "empirical tests" is somewhat misleading. Second, many of the "empirical" papers in the 12-year period focused on the topic of spatial reorientation (Learmonth, Newcombe, Sheridan, & Jones, 2008). This literature examines spatial memory in the context of navigation, inspired, in part, by research examining the role of hippocampus in navigation (Burgess, Maguire, & O'Keefe, 2002). It is not transparent how models of spatial memory developed to explain location memory in a small-scale space can be extended to capture findings in navigation and reorientation (although recent DNF models of navigation take a step in this direction; Bicho & Schöner, 1997). This observation resonates with our evaluation of the clinical literature in that clinical studies often use tasks that have not been directly captured by formal

models (e.g., the Corsi block task; Pagulayan, Busch, Medina, Bartok, & Krikorian, 2006). Thus, results from Fig. 6 appear to reflect a compartmentalization of the spatial memory literature into modeling papers, empirical probes of spatial memory in the context of navigation, and clinical studies using standardized assessment tasks.

4. Conclusions: what role will formal theories play in the future?

Our goal has been to highlight how DFT has been useful in understanding cognitive development using two literatures as illustration. These examples showed how DFT has been an effective framework for doing what theories are supposed to do—integrate empirical findings, explain links among diverse results, and generate novel predictions. We also noted cases where DNF models went above and beyond what can be accomplished with verbal concepts alone, generating predictions that were hard to explain otherwise.

The two empirical literatures we examined were also useful in that there are multiple formal models of the respective phenomena. This allowed us to conduct two systematic analyses of the literature to evaluate the contributions of formal models more generally. These revealed a steady rate of modeling publications in the last decade. More critically, the analyses also revealed several epochs during which empiricists have engaged formal models in a significant way. These data show that models are having an impact in cognitive development beyond the group of researchers using these tools directly. From our vantage point, this is an exciting result.

Yet these analyses also reveal an important missed opportunity in that neurophysiological and clinical studies rarely interface with formal models. It is unfortunate that neural studies do not engage formal models like DFT that have an established tie to neurophysiology, because models offer a way to test and generate neural hypotheses and link them to behavior. Similarly, models can help organize and explain data from complex networks in the brain. Clinical studies should also be a target for further integration with formal models because models can help probe hypotheses about the mechanisms that underlie atypical behaviors. For instance, models have played an instrumental role in several studies examining the processes that underlie dyslexia (Harm & Seidenburg, 1999; Plaut et al., 1996) and specific language impairment (Joanisse & Seidenberg, 2003; McMurray et al., 2010).

What does the future hold? In a recent paper, Simmering and colleagues (Simmering, Triesch, Deák, & Spencer, 2011) argued that we need more communication between experimentalists and theoreticians and, critically, that both need to be valued. Our analysis shows positive trends in this regard, but we need to strive to do better in the future. How? One idea is to train graduate students to at least be familiar with different classes of models. Consistent with this, there are a growing number of pre-conference meetings and summer courses to give students exposure to modeling techniques. Another idea is to engage in more formal discussions in which we try to bring experimentalists and modelers together around specific topics. Too often, detailed discussion happens in separate "camps" and there is little resolution or substantive movement forward.

Regardless of the mechanisms that move us forward as a field, our view is that complex theories and formal models are here to stay. Why? Because development is the most complex of topics—more complex than major topics in biology, chemistry and physics. Moreover, we have to do the job of neuroscientists and cognitive psychologists *and* understand how the cognitive and behavioral system changes through time. We cannot just think our way out of this degree of complexity using verbal concepts alone. The sooner we embrace this view and more fully integrate empirical and theoretical work, the faster we will become a mature, cumulative, and groundbreaking science.

Acknowledgements

References

Adolph, K., & Avolio, A. M. (2000). Walking infants adapt locomotion to changing body dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, *26*(3), 1148–1166.

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

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.

Bastian, A., Schöner, G., & Riehl, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, *18*(7), 2047–2058.

Bell, M. A. (2001). Brain electrical activity associated with cognitive processing during a looking version of the A-Not-B task. *Infancy*, *2*(3), 311–330.

Bell, M. A., & Fox, N. A. (1992). The relations between frontal brain electrical activity and cognitive development during infancy. *Child Development*, *63*(5), 1142–1163.

Bell, M. A., & Fox, N. A. (1997). Individual differences in object permanence performance at 8 months: Locomotor experience and brain electrical activity. *Developmental Psychobiology*, *31*(4), 287–297.

Berger, S. E. (2004). Demands on finite cognitive capacity cause infants' perseverative errors. *Infancy*, *5*(2), 217–238.

Bicho, E., & Schöner, G. (1997). The dynamic approach to autonomous robotics demonstrated on low-level vehicle platform. *Robotics & Autonomous Systems*, *21*(1), 23–35.

Bicho, E., & Schöner, G. (1998). Robot target position estimation. *Cahiers de Psychologie Cognitive*, *17*(4–5), 1044–1045.

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*(4), 625–641.

Changeux, J.-P., & Dehaene, S. (1989). Neuronal models of cognitive functions. *Cognition*, *33*, 63–109.

Clearfield, M. W., Dineva, E., Smith, L. B., Diedrich, F. J., & Thelen, E. (2009). Cue salience and infant perseverative reaching: Tests of the dynamic field theory. *Developmental Science*, *12*(1), 26–40.

Dettmer, A. M., Novak, M. F. S. X., Novak, M. A., Meyer, J. S., & Suomi, S. J. (2009). Hair cortisol predicts object permanence performance in infant *Rhesus macaques* (Macacamulatta). *Developmental Psychobiology*, *51*(8), 706–713.

Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on A-not-B. *Child Development*, *56*, 868–883.

Diamond, A. (1990a). Developmental time course in human infants and infant monkeys, and the neural bases of inhibitory control in reaching. *Annals of the New York Academy of Sciences*, *608*, 637–676.

Diamond, A. (1990b). The development and neural bases of memory functions as indexed by the AB and delayed response tasks in human infants and infant monkeys. *Annals of the New York Academy of Sciences*, *608*, 267–317.

Diedrich, F., & Warren, W. H. (1995). Why change gaits? Dynamics of the walk-run transition. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 183–202.

Edin, F., Macoveanu, J., Olesen, P., Tegnér, J., & Klingberg, T. (2007). Stronger synaptic connectivity as a mechanism of working memory-related brain activity during childhood. *Journal of Cognitive Neuroscience*, *19*(5), 750–760.

Engels, C., & Schöner, G. (1995). Dynamic fields endow behavior-based robots with representations. *Robotics & Autonomous Systems*, *14*(1), 55–77.

Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*(3), 545–572.

Faubel, C., & Schöner, G. (2008). Learning to recognize objects on the fly: A neurally-based dynamic field approach. *Neural Networks*, *21*(4), 562–576.

Fechner, G. T. (1860). *Elementederpsychophysik*. Leipzig: Breitkopf, Hartel.

Fischer, K. W., & Rose, S. P. (1996). Dynamic growth cycles of brain and cognitive development. In R. W. Thatcher, G. R. Lyon, J. Rumsey, & N. Krasnegor (Eds.), *Developmental neuroimaging: Mapping the development of brain and behavior* (pp. 263–279). New York: Academic Press.

Fogel, A., & Thelen, E. (1987). Development of early expression and communicative action: Reinterpreting the evidence from a dynamic systems perspective. *Developmental Psychology*, *23*(6), 747–761.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental rotation of the neuronal population vector. *Science*, *243*, 234–236.

Harm, M. W., & Seidenburg, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, *106*(3), 491–528.

Hund, A. M., & Plumert, J. M. (2003). Does information about what things are influence children's memory for where things are? *Developmental Psychology*, *39*(6), 939–948.

Huttenlocher, J., Hedges, L. V., Corrigan, B., & Crawford, L. E. (2004). Spatial categories and the estimation of location. *Cognition*, *93*(2), 75–97.

Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*(3), 352–376.

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

James, W. (1897/1894). In B. Wilshire (Ed.), *The essential writings*. New York, NY: State University of New York Press.

Joanisse, M. F., & Seidenberg, M. S. (2003). Phonology and syntax in specific language impairment: Evidence from a connectionist model. *Brain Language*, *86*(1), 40–56.

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science*, *20*(5), 568–577.

Johnson, J. S., Spencer, J. P., & Schöner, G. (2009). A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain Research*, *1299*, 17–32.

Learmonth, A. E., Newcombe, N. S., Sheridan, N., & Jones, M. (2008). Why size counts: Children's spatial reorientation in large and small enclosures. *Developmental Science*, *11*(3), 414–426.

Lerner, R. M. (2006). Developmental science, developmental systems, and contemporary theories of human development. In W. Damon, & R. M. Lerner (Eds.), *Handbook of child psychology. Vol 1: Theoretical models of human development* (6th ed., pp. 1–17). Hoboken, NJ: Wiley.

Lewis, J. D., & Elman, J. L. (2008). Growth-related neural reorganization and the autism phenotype: A test of the hypothesis that altered brain growth leads to altered connectivity. *Developmental Science*, 11, 135–155.

Lewis, M. D., Lamey, A. V., & Douglas, L. (1999). A new dynamic systems method for the analysis of early socioemotional development. *Developmental Science*, 2(4), 457–475.

Lipinski, J., Sandamirskaya, Y., & Schöner, G. (2009). Swing it to the left, swing it to the right: Enacting flexible spatial language using a neurodynamic framework. *Cognitive Neurodynamics*, 3(4), 373–400.

Lipinski, J., Schneegans, S., Sandamirskaya, Y., Spencer, J. P., & Schöner, G. (2011). A neurobehavioral model of flexible spatial language behaviors. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, http://dx.doi.org/10.1037/a0022643

Lipinski, J., Simmering, V. R., Johnson, J. S., & Spencer, J. P. (2010). The role of experience in location estimation: Target distributions shift location memory biases. *Cognition*, 115, 147–153.

Lipinski, J., Spencer, J. P., & Samuelson, L. K. (2010). Biased feedback in spatial recall yields a violation of delta rule learning. *Psychonomic Bulletin and Review*, 17, 581–588.

Marcovitch, S., & Zelazo, P. D. (2006). The influence of number of A trials of 2-year-olds' behavior in two competing A-not-B-type search tasks: a test of the Hierarchical Competing Systems Model. *Journal of Cognition and Development*, 7(4), 477–501.

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

Mauerberg-deCastro, E., Cozzani, M. V., Polanczyk, S. D., dePaila, A. I., Lucena, C. S., & Moraes, R. (2009). Motor perseveration during an "A not B" task in children with intellectual disabilities. *Human Movement Science*, 28(6), 818–832.

McDowell, K., Jeka, J. J., Schöner, G., & Hatfield, B. (2002). Behavioral and electrocortical evidence of an interaction between probability and task metrics in movement preparation. *Experimental Brain Research*, 144, 303–313.

McMurray, B., Samelson, V. M., Lee, S. H., & Tomblin, J. B. (2010). Individual differences in online spoken word recognition: Implications for SLI. *Cognitive Psychology*, 60(1), 1–39.

Molenaar, P. C. M., Boomsma, D. I., & Dolan, C. V. (1993). A third source of developmental differences. *Behavior Genetics*, 23(6), 519–524.

Molenaar, P. C. M., & Newell, K. M. (Eds.). (2010). *Individual pathways of change: Statistical models for analyzing learning and development.*. Washington, DC: American Psychological Association.

Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP model of the AB task. *Developmental Science*, 1(2), 161–184.

Nardini, M., Thomas, R. L., Knowland, V. C. P., Braddick, O. J., & Atkinson, J. (2009). A viewpoint-independent process for spatial reorientation. *Cognition*, 112(2), 241–248.

Ortmann, M. R., & Schutte, A. R. (2010). The relationship between the perception of axes of symmetry and spatial memory during early childhood. *Journal of Experimental Child Psychology*, 107(3), 368–377.

Pagulayan, K. F., Busch, R. M., Medina, K. L., Bartok, J. A., & Krikorian, R. (2006). Developmental normative data for the Corsi block-tapping task. *Journal of Clinical and Experimental Neuropsychology*, 28(6), 1043–1054.

Perone, S., & Spencer, J. P. (in press). Autonomy in action: Linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive Science*.

Piaget, J. (1954). *The construction of reality in the child*. Abingdon, Oxon: Routledge Press.

Plaut, D. C., McClelland, J. L., Seidenburg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in semi-regular domains. *Psychological Review*, 103(1), 56–115.

Pouget, A., Deneve, S., & Duhamel, J.-R. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience*, 3, 741–747.

Prigogine, I., & Nicolis, G. (1971). Biological orders, structure, and instabilities. *Quarterly Review of Biophysics*, 4(2–3), 107–148.

Prigogine, I., & Stengers, I. (1984). *Order out of chaos: Man's new dialogue with nature*. New York, NY: Bantam Books.

Rotzer, S., Loenneker, T., Kucian, K., Martin, E., Klaver, P., & von Aster, M. (2009). Dysfunctional neural network of spatial working memory contributes to developmental dyscalculia. *Neuropsychologia*, 47(13), 2859–2865.

Ruffman, T., Slade, L., Sandino, J. C., & Fletcher, A. (2004). Are A-not-B errors caused by a belief about object location? *Child Development*, 76(1), 122–136.

Sandamirskaya, Y., & Schöner, G. (2010). An embodied account of serial order: How instabilities drive sequence generation. *Neural Networks*, 23(10), 1164–1179.

Schneegans, S., & Schöner, G. (2012). A neural mechanism for coordinate transformation predicts pre-saccadic remapping. *Biological Cybernetics*, 106(2), 89–109.

Schöner, G. (1994). Dynamic theory of action-perception patterns: The time-before-contact paradigm. *Human Movement Science*, 13, 415–439.

Schöner, G. (2009). Dynamic systems approaches to cognition. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 101–126). New York, NY: Cambridge University Press.

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. (2009). Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 1698–1725.

Schutte, A. R., & Spencer, J. P. (2010). Filling the gap on developmental change: Tests of a dynamic field theory of spatial cognition. *Journal of Cognition and Development*, 11, 1–27.

Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74, 1393–1417.

Schweinsburg, A. D., Nagel, B. J., & Tapert, S. F. (2005). fMRI reveals alteration of spatial working memory networks across adolescence. *Journal of the International Neuropsychological Society*, 11(5), 631–644.

Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental timescales. *Brain Research*, *1202*, 68–86.

Simmering, V. R., & Spencer, J. P. (2008). Generality to specificity: The dynamic field theory generalizes across tasks and time scales. *Developmental Science*, *11*(4), 541–555.

Simmering, V. R., Spencer, J. P., & Schöner, G. (2006). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception and Psychophysics*, *68*, 1027–1046.

Simmering, V. R., Triesch, J., Deák, G. O., & Spencer, J. P. (2011). To model or not to model? A dialogue on the role of computational modeling in developmental science. *Child Development Perspectives*, *4*, 152–158.

Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review*, *106*(2), 235–326.

Spencer, J. P., Clearfield, M., Corbetta, D., Ulrich, B., Buchanan, P., & Schöner, G. (2006). Moving toward a grand theory of development: In memory of Esther Thelen. *Child Development*, *77*, 1521–1538.

Spencer, J. P., Dineva, E., & Schöner, G. (2009). Moving toward a unified theory while valuing the importance of the initial conditions. In J. P. Spencer, M. S. Thomas, & J. L. McClelland (Eds.), *Toward a unified theory of development: Connectionism and dynamic systems theory re-considered* (pp. 354–372). New York: Oxford University Press.

Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience- dependent spatial categories. *Journal of Experimental Psychology: General*, *131*, 16–37.

Spencer, J. P., & Hund, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*, *47*, 432–480.

Spencer, J. P., & Perone, S. (2008). Defending qualitative change: The view from dynamic systems theory. *Child Development*, *79*(6), 1639–1647.

Spencer, J. P., Perone, S., & Johnson, J. S. (2009). Dynamic field theory and embodied cognitive dynamics. In J. P. Spencer, M. S. C. Thomas, & J. McClelland (Eds.), *Toward a unified theory of development: Connectionism and dynamic systems theory re-considered* (pp. 86–118). New York, NY: Oxford University Press.

Spencer, J. P., Simmering, V. R., Schutte, A. R., & Schöner, G. (2007). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In J. M. Plumert, & J. P. Spencer (Eds.), *The emerging spatial mind* (pp. 320–361). New York, NY: Oxford University Press.

Spencer, J. P., Smith, L. B., & Thelen, E. (2001). Tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of 2-year-olds. *Child Development*, *72*, 1327–1346.

Spencer, J. P., Thomas, M. S. C., & McClelland, J. (Eds.). (2009). *Toward a unified theory of development: Connectionism and dynamic systems theory re-considered*. New York, NY: Oxford University Press.

Thelen, E. (1995). Motor development: A new synthesis. *American Psychologist*, *50*(2), 79–95.

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

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

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

Thomas, M. S. C., & Karmiloff-Smith, A. (2003). Connectionist models of development, developmental disorders and individual differences. In R. J. Sternberg, J. Lautrey, & T. Lubart (Eds.), *Models of intelligence: International perspectives* (pp. 133–150). American Psychological Association.

Van der Maas, H. L. J. (1993). *Catastrophe analysis of stagewise cognitive development.* Amsterdam: Universiteit van Amsterdam.

Van der Maas, H. L. J., Dolan, C. V., Grassman, R. P. P. P., Wicherts, J. M., Huizenga, H. M., & Raijmakers, M. E. J. (2006). A dynamical model of general intelligence: The positive manifold of intelligence by mutualism. *Psychological Review*, *113*(4), 842–861.

Van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stagewise cognitive development: An application of catastrophe theory. *Psychological Review*, *99*(3), 395–417.

Van Geert, P. (1997). Variability and fluctuation: A dynamic view. In E. Amsel, & K. A. Renninger (Eds.), *Change and development: Issues of theory, method and application* (pp. 193–212). Mahwah, NJ: Erlbaum.

Van Geert, P. (1998). A dynamic systems of basic developmental mechanisms: Piaget, Vygotsky, & Beyond. *Psychological Review*, *105*(4), 634–677.

Vuontela, V., Steenari, M.-R., Aronen, E. T., Korvenoja, A., Aronen, H. J., & Carlson, S. (2009). Brain activation and deactivation during location and color working memory tasks in 11–13-year-old children. *Brain and Cognition*, *69*(1), 56–64.

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neuroscience*, *24*, 455–463.

Weber, E. (1842–1853). Dertastsinn und das gemeingefuhl. In R. Wagner (Ed.), *Handworterbach derphysiologie*. Braunschweig: Vieweg.

Wellman, H. M., Cross, D., & Bartsch, K. (1986). Infant search and object permanence: A meta- analysis of the A-not-B error. *Monographs of the Society for Research in Child Development*, *51*(1–51), 62–67.

Whitall, J., & Getschell, N. (1995). From walking to running: Applying a dynamic systems approach to the development of locomotor skills. *Child Development*, *66*(5), 1541–1553.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*, 1–24.