

CHAPTER 8

Dynamical Systems Thinking

From Metaphor to Neural Theory

Gregor Schöner

The Fundamental Tension in Development between Stability and Flexibility

As parents, we experience our children as unique individuals. When we have more than one child, we marvel at how each child is special from the very first day. On the other hand, our entire educational attitude is based on the assumption that children's experiences shape their development and that the environment to which they are exposed may make a difference in how they will develop.

As developmental scientists we have seen, time and again, how infants, toddlers, and children go through recognizable stages of development, formalized by developmental milestones. These stages invite accounts in which development resembles a program from which competences unfold in a fixed sequence. Growth processes in developmental biology provide the metaphor for how molecular processes could lead to maturation of the nervous system that may drive behavioral and cognitive development.

At the same time, there can be no doubt that the environment to which children are exposed matters and that their individual behavioral history influences the process of development. This observation invites accounts based on learning, in which the nervous system of the developing child is shaped by the activity induced by ongoing behavior and stimulation.

The Self-Organization Metaphor

One way developmental scientists have addressed this tension between the apparent stability of the developmental process and its flexibility in response to variable

environments and experiential histories is by invoking the concept of self-organization. In a self-organizing system, the flexible and individual processes of development may emerge from the confluence of various forces, while the inherent organizational principles would hold the developmental process on track. Exactly how that may work requires closer examination.

This idea has been promoted as a metaphor for development that has inspired a research program into possible mechanisms of emergence, of how environmental factors contribute to development, how individual differences persist, and how the ensemble of these influences jointly control development (“soft causation”) (Elman et al., 1997; Spencer et al., 2006; Thelen & Smith, 1994). This chapter reviews this metaphor, criticizes it, and proposes a specific direction in which this metaphor can be transformed into an operational theory of the development of behavior and cognition.

The perspective of self-organization was developed, in part, in opposition to an apparent alternative; that is, to an account of development that is based on maturation and innateness. In such an account the molecular machinery of growth guarantees the reproducible patterns of development and also explains how individual differences arise from different starting conditions, largely determined by genetic factors, and remain stable over the course of development (Fodor, 1981; for critiques, see Molenaar, 1986; Quartz, 1993).

Interestingly, the self-organization metaphor connects to older theoretical ideas about growth processes. In fact, within developmental biology an analogous tension exists between two metaphors, one based on information processing and programming, the other on self-organization and emergence. The fine-grained molecular machinery of growth engages gene expression, which may be thought of as the core driver of morphogenesis, which is the formation of macroscopic shape from microscopic processes. At this molecular level, information-processing metaphors such as *program*, *lock and key*, and *reading or writing* of genetic information, etc., are commonly invoked. These metaphors have their own interesting conceptual history (Fox Keller, 2002). Molecular biologists resonated with these concepts as these seemed to fit to their experimental tools.

In an older view, morphogenesis had been thought of as a form of pattern formation. This older view dates back to Waddington’s famous epigenetic landscape (Waddington, 1953; see also Ho, Chapter 5, and Newell & Liu, Chapter 12, this volume and Figure 8.1), often invoked but somewhat inconsequential to modern developmental biology. The metaphor suggests that “forces” shape the form of an organism, which then emerges from a process of equilibration of these forces. Over development, the landscape of forces and the associated equilibria become increasingly complex, leading to a more and more differentiated organism. This metaphor resonates with the famous mathematical model of pattern formation of Alan Turing (1952), in which the interaction of diffusion and chemical reaction kinetics generates concentration patterns. The structure of these patterns is encoded in the parameters of the chemical and diffusion dynamics that are hypothesized to drive growth. That and similar models remained essentially metaphorical too, as they did not make contact with experimentally accessible problems in development.

This changed somewhat in the 1980s, when experimental evidence was obtained that the substrate of such pattern formation, promoters of the growth of cell membranes, are at work in development. Mathematical models of pattern formation in morphogenesis (Gierer & Meinhardt, 1972; Murray, 2002) were based on the diffusion and chemical reaction kinetics of an activator and inhibitor of growth. The formation of chemical concentration patterns then creates a skeleton for growth processes. The theory was based on the mathematical framework of nonlinear dynamical systems, typically in the form of partial differential equations that undergo instabilities (Haken, 1983; Murray, 2002). The theory can account for features of morphogenesis, such as the qualitative form of misformation. For instance, an individual may have an extra finger, but the overall shape of the fingers and the hand is preserved. In the mathematical models, this is a property of the self-organization processes on which growth is based. The boundary conditions for a piece of tissue only allow for a discrete set of possible patterns. Among these, one typically emerges as the stable solution, but a perturbation or change of conditions may induce a neighboring pattern to “win” the competition (Murray, 2002). In fact, more generally, the study of malformations may provide provocative insight into development (Blumberg, 2010).

So, in a sense, in the domain of morphogenesis, the metaphor of Waddington has been shaped into mathematical models that have some traction empirically and are even, to some extent, reductionistic in nature. Is this the sort of transition from metaphor to theory at which we aim for an understanding of the development of behavior and cognition? Perhaps not. Even after the discovery of promoters, the mathematical theory of morphogenesis has had very little impact on the research program of developmental biologists, largely because the methods of molecular biology that have been most successful experimentally do not match the kinds of questions the theory promotes (Fox Keller, 2002). (This may begin to change as computational biology brings the mathematics of dynamical self-organization to the molecular level, see Harrison, 2011, for a survey).

Conversely, conceptually challenging and interesting questions that flow out of the metaphor of self-organization are not addressed by the mathematical theory. In particular, the metaphor of emergence suggests that growth processes are open to multiple forces that may literally help shape the organism. This idea has been at the core of the research program of developmental systems theory (Gottlieb, 2001; see also Griffiths & Gray, 2005, for a history of thought), which focuses on nonobvious interactions between the emerging function of a developing organism and the growth processes taking place through the environment, and modulating and mediating developmental processes.

Although the mathematics of self-organization may, in principle, be able to address such issues, models have not really played a critical role in this field. This is probably again due to a mismatch in levels of description, with the substrate of the mathematical theory being quite removed from the level of observation at which developmental systems theory seeks evidence for environmental influences.

From this history of ideas in developmental biology and morphogenesis, an important lesson can be derived for how the metaphor of self-organization can be

transformed into theory. Clearly, that transformation requires more than mere mathematization. Theoretical concepts must relate to the level of description at which development is characterized experimentally and must be able to articulate the role of the various factors found to impact on developmental processes.

The Dynamical Systems Metaphor for the Development of Behavior and Cognition

The metaphor of pattern formation and the associated mathematical concepts of dynamical systems theory have been used to talk about how behavior and cognition may emerge from learning. Thelen and Smith (1994), for instance, interpreted Waddington's epigenetic landscape in terms of emergent behavioral and cognitive competences. At any particular stage of development, a landscape is a potential function, whose minima represent behaviors that can be realized at that particular stage of development (Figure 8.1). Because the minima are thought to arise from the different forces that shape the behavioral landscape, they are not fixed competences, but emerge as a function of environmental conditions and of an individual's inner state; they are "soft-assembled" in the moment.

Over development, the landscape is postulated to evolve, becoming increasingly differentiated and complex to reflect the increasing number of adaptive behaviors that




FIGURE 8.1. The epigenetic landscape visualizes the emergence of behaviors within a dimension (horizontal) as a potential landscape (vertical) evolves over time (back to front). The system (black ball) moves to minima of the potential on the fast timescale of behavior (short arrow). These minima may change and new minima may appear on the slow timescale of development (long arrow).

are the new competences acquired over development. In the metaphor, the behaviors expressed by the system contribute to the developmental processes. One way to visualize this in the metaphor is to think of the system as moving within the landscape and carving out new minima. In this picture, individually different routes in development may emerge from the way ongoing behavior drives the developmental process.

This metaphor has important strengths. First, it promotes a process account of development by directing attention to how behavior unfolds in time under the influence of various factors. This perspective stands in contrast to the emphasis on description and normalization that a maturational account suggests by decoupling development from the moment-to-moment experience of an individual. Process accounts are inherently stronger and more naturally open to intervention.

Along the same line, the metaphor emphasizes that environmental conditions must be taken into account to understand how competences emerge during development. When environmental conditions are favorable, a child may be able to perform at a more mature level than when conditions are more difficult. Thus, during development, the environmental constraints required for a child to display a particular competence may relax. This is an important insight.

The metaphor draws attention to learning as the core process of development. How experience matters for development is the central research question that the metaphor proposes for developmental science. The shaping of development by experience amplifies the role of environmental conditions: Because the environment may support, to varying degrees, the emergence of skilled behavior, it influences what experiences are accessible to the child and thus impact on the developmental process. Scaffolding, in which caregivers structure the environment to facilitate the emergence of new skills, makes immediate sense in this perspective (Lipscomb, Swanson, & West, 2004). Educators and parents are effective by enabling a child to gather its own experience.

Finally, the metaphor emphasizes individual differences and views them not as fixed traits, but as the result of the developmental process. Because experience shapes development, differences in experience have long-lasting impact and may lead to individual paths of development. This explains how individual differences may persist without attributing them necessarily to fixed structural differences. In fact, the metaphor suggests ways in which individual differences may emerge and then be amplified during development. Again, the implications for educators and parents are broad.

On a closer look, however, the dynamical systems metaphor is unclear and fuzzy in many ways. Consider the image of the epigenetic landscape again. Implicit in this image is a notion of a timescale: the slower timescale of development and the faster timescale of behavior. At a given moment during development, the different minima of the landscape represent the behavioral repertoire accessible at this time. How are these behaviors described? Are they embedded within a single space (illustrated along one dimension, but surely spanning multiple dimensions)? Are all behaviors equally accessible or does the distance between them matter? What are the processes of selecting a particular behavior? The metaphor is quite fuzzy about how behavior unfolds at any given moment during development.

The metaphor emphasizes that behavior is generated by an embodied and situated organism. However, how behavior is responsive to environmental conditions is obscured in the metaphor. Possible behaviors do not sit somewhere as minima, waiting for the state of the system to “fall” into the well. Instead, possible behaviors emerge “on the fly” in response to environmental conditions and the inner state of the system. For example, *looking* emerges in response to interesting things to look at and is strongly cued by salient visual transients. *Reaching* emerges when there are attractive objects within reach. Previous reaches affect new reaches.

The metaphor derives strengths from the idea that behaviors are shaped by the joint action of forces, the very notion of dynamics, from which stable states emerge. The landscape illustrates this notion through the minima to which the system is attracted. The time needed to move to the nearest minimum is really the timescale of behavior (Newell, Liu, & Meyer-Kress, 2009). But how can we conceive of these forces? How does the influence of the environment and inner states impact on the stability of a behavior? Behaviors may vary in stability. When environmental conditions change, behaviors may lose stability (Schöner & Kelso, 1988). So clearly, the shape of the minima vary with context. The dependence of the stability of behaviors on many factors, including the environment, intention, and attention, is an important dimension of development that needs to be made more explicit.

Relatedly, selecting a particular behavior involves active choice or decision processes. An infant may look, grasp, look away, or persist. This flexibility is not visualized in the metaphor, and the underlying process of decision making is another dimension of behavior that has its own dynamics and timescale (Newell et al., 2009). This point is critical, because in many instances it is this very flexibility that develops. The *A-not-B* error is a signature of a lack of flexibility when the goal of a reaching act is changed. Developmentally, this flexibility is an achievement. Similarly, younger infants tend toward stereotypy—the inflexible persistence of behavioral patterns—and over development become increasingly flexible. Looking behavior is a well-studied example, in which very long fixations may occur early on, whereas gaze shifts become increasingly nimble over development (Colombo, 2001).

Finally, the metaphor is suggestive of overt motor behavior as the primary mark of competence, realized while the system is sitting in a minimum of the landscape. Cognition entails inner states that affect overt motor behavior. Through working memory, for instance, perceptual or motor processes at one point in time impact on decisions at a later time. This is one factor in explaining how older infants succeed in overcoming perseveration in the *A-not-B* paradigm. Older infants are capable of stabilizing a motor intention to reach for the *B* location during the delay. In the simple form suggested by the potential landscape, the dynamical systems metaphor does not provide a substrate for such inner states. Such inner states may have their own dynamics—for instance, in the process of memory formation—that live at an intermediate timescale, longer than behavior in the here and now, and shorter than the timescale over which development takes place.

So let us turn to the slower timescale of development on which the landscape itself evolves. The core hypothesis is that the developmental process is driven by experience,

that is, by the behaviors (and inner states) that are realized on the faster timescale of behavior. How this happens remains vague in the metaphor. This is the major frontier, of course, for a process account of development. Here are a few questions that illustrate how much remains to be explored.

- What is the substrate for learning and development? Inner states are not obvious in the simple form of the dynamical systems metaphor. Connectionism proposes that change occurs in how the inner states are connected to the sensory and motor surfaces. But there may also be change in just the (recurrent) connectivity internal to the neural networks that drive behavior.
- What information is used to drive the learning process?
- Is learning largely driven by the statistics of afferent input, to which unsupervised learning rules are sensitive?
- Does it require a sensory feedback signal or reinforcement?
- Is the system open to learning in any of its subsystems, or is there an inherent organization of learning, in which one system is the focus of learning for some time and that focus is then shifted to the next system when a given level of performance has been reached?
- How does the system maintain a coherent learning process across the different episodes of experiencing a particular aspect of behavior or of stimulation?
- How does it autonomously control its learning processes?
- How do the learning processes actually lead to differentiation?
- Why does learning a new skill not unlearn a former skill? Is there a need for relearning old skills to stabilize them?

Clearly, at this point the self-organization metaphor has only scratched the surface of what a true process account of learning and development must account for. By way of summary, here are the three main challenges for a transition from the self-organization metaphor to dynamical systems theory of development.

- First, we need a much better, concrete account for what happens at the timescale of behavior. This account would include a systematic method to characterize behavior as the performance of a competence. This method must address how the environment—how stimuli and perception—shapes behavior and how ongoing behavior may be coupled with perception, but may also be conditioned by inner states, a particular mode of behavior that classically was called “task” and is all the more challenging to define for infants and children. And this method must characterize cognitive processes as a form of behavior. Not every behavioral state leads to immediate and overt motor action.

- Second, we need to account for flexibility and cognition that includes a systematic process understanding of how decisions are made in the here and now, and

how those decisions are linked to sensory information and may generate motor consequences. This also requires a substrate for representations and an account of the processes through which sensory and motor processes interface with representations and create, update, or remove memories.

- Third, we need a much more profound approach toward the processes of learning that are hypothesized to take place during development. Connectionism and dynamical systems thinking have only given first hints at the autonomous learning processes in which infants, toddlers, children, and even adults are permanently engaged.

From the Dynamical Systems Metaphor to Dynamic Field Theory as a Theoretical Framework for Understanding Development

What does the transition from metaphor to theory entail? This section provides a rough sketch and introduces the main concepts of dynamic field theory (DFT; for reviews, see Schöner, 2008). In the next section, we consider the same ideas at a more pedestrian pace by using two worked examples to make things concrete and simple.

We begin with the issue of how to characterize behavior in a way that is open to sensory and motor processes, but may also reach into cognition. The answer, perhaps not surprisingly, comes from the neural level of description. We know that neuronal activity in the central nervous system is determined in part by the pattern of connectivity from the sensory surfaces to the neurons in question. Their tuning curves to parameters of sensory stimulation delineate the extent of this influence. In this sense, a neuron “stands for” the stimulus condition to which it is sensitive. Conversely, at the output level the forward connectivity from a neuron to the motor surface determines the motor “meaning” of that neuron’s activity. Tuning to movement parameters estimates the extent of the motoric specificity of neural activation.

We also know that tuning curves throughout the higher nervous system, in the cortex, but also in such subcortical structures as the colliculus or the lateral geniculate nucleus, are relatively broad, so that when any single stimulus is presented or any individual motor act prepared, a considerable portion of the relevant neural populations is activated. Evidence for such population coding of perceptual and motor parameters comes from various sources that include the broad tuning curves of cortical neurons to behavioral parameters (Georgopoulos, Schwartz, & Kettner, 1986; Young & Yamane, 1992), the contribution that all activated neurons make to behavioral decisions (e.g., Groh, Bron, & Newsome, 1997; Lee, Rohrer, & Sparks, 1988), and the patterns of correlation between neurons and behavior (see contributions to Kriegeskorte & Kreima, 2012).

These facts support the proposal that DFT makes for the substrate of behavior and cognition: The activation of neuronal populations stands for particular perceptual states or motoric outcomes. These activation patterns need not be traced back to how they are distributed within the neural networks of the brain. Instead, we can think

of these patterns as being defined directly in terms of the sensory and motor states to which the populations are sensitive. Thus, for instance, a distribution of activation may be defined over retinal visual space, over color space, or over the space of movement direction of visual objects. Similarly, a distribution of activation may be defined over the space of possible movement directions, movement extent, or amount of force generated. These distributions are the activation fields on which DFT is based.

Peaks of activation are then instances of a perceptual or motor state (see Figure 8.2). Generating such a peak amounts to falling into one of the minima of the potential landscape (see Figure 8.1). Note, however, that the peaks are not all standing there and waiting to be selected. In fact, typically a single peak or a small number of peaks may be activated. And peaks are tunable: They can be localized at different positions along the behavioral dimensions for which they stand, representing different perceptual or motor outcomes. Peaks can directly drive motor behaviors, or they can drive other representations to build other peaks. The processes of peak formation, peak updating, and peak deletion are the neural dynamics of neuronal populations. Input plays an important role in these processes. Environmental condition may impact on activation fields by generating patterns of input that come through the sensory surfaces.

How the forward connectivity determines input patterns is the central topic of connectionist modeling, and DFT shares this principle with connectionism. In connectionism, cascades of layered input stages may lead to increasing abstraction in sensory representations. The motor side is much less well understood, in large part because it is not well described by the feedforward conception. In DFT, even perceptual events are not purely input driven. Peaks of activation are postulated to be stable states, attractors of the neural dynamics, so that they can be continuously coupled to sensory information. Recurrent connectivity within the neural population stabilizes such peaks against decay and against competing inputs. This is based on the principle of local excitation, according to which similarly tuned neurons are exciting each other, and global inhibition, according to which differently tuned neurons are effectively inhibiting each other (through interneurons).

A consequence of this principle of the stabilization of activation peaks is that such peaks are separated from nonpeak patterns of activation by an instability, in which input-driven patterns of activation reach a critical level of activation, at which the detection decision is made. The instability comes from the nonlinear properties of neural




FIGURE 8.2. An activation field defined over a perceptual or motor dimension has a peak whose location specifies a particular value along that dimension.

dynamics, in which only sufficiently activated neurons in a population contribute to neural interaction, as described by a sigmoidal function. The concept of self-stabilized peaks of activation, sustained by neural interaction, is a central tenet of DFT, in which it departs from standard connectionist thinking (although there are connectionist models that entail the same kind of instabilities, e.g., Usher & McClelland, 2001). We consider the detection instability again in the next section (see Figure 8.5).

Activation peaks are units of representation in DFT. In whichever way a peak was induced, it resists change and impacts on whatever downstream neural networks or motor systems it projects. When a peak is induced by a single localized input, the decision to detect this input resists change: It persists even if the input is weakened or noisy. When multiple localized inputs are presented, an activation field may make a selection decision by generating a single peak located over the input that arose first or was strongest. That selection decision resists change (up to a point) when, for instance, distractor inputs impinge on the system. In fact, under appropriate dynamic conditions for the activation dynamics, a peak may persist after all localized input has been removed. That persistence provides an account for working memory of metric information, represented by the location of the peak in the activation field.

The stability of activation peaks has important consequences for learning processes. The detection instability may, essentially, amplify small inhomogeneities in a neural population into a macroscopic decision represented by a peak. It may be sufficient, for instance, to just push global activation into a field, without any specific information about the location to activate; that is, no specific information about the desired state or behavior. As soon as a field location creates enough activation so that the activation engages the sigmoidal function and interaction sets in, an entire peak is generated, essentially pulling itself up through its own excitatory interaction. The locations at which this is possible are dictated by inhomogeneities—that is, by small, competitive advantages for field locations through prior activation or strengthened synaptic connections to any input structures.

Learning Processes Are Also a Form of Neural Dynamics in DFT

The concept of self-stabilized peaks lowers demands on learning process. In conventional, largely feedforward neural networks, learning must shape synaptic connectivity sufficiently to make a difference at the output level, for instance, to generate a new “winner.” In DFT, even small biases may be sufficient to lead to the formation of new peaks. As result, even very simple learning mechanisms, such as a memory trace of prior patterns of activation, may have a significant impact on a neural dynamics.

Such memory traces are, in a sense, a zeroth order learning mechanism because they are sensitive only to the history of activation, not to correlations between patterns of input and patterns of output. Hebbian learning, which is sensitive to such correlations, may, of course, also shape the field and determine which activation patterns are most easily induced.

Both the laying down of memory traces and the strengthening of connections according to a Hebbian rule can be understood as a form of dynamics at a slower timescale (Erlhagen & Schöner, 2002; Schöner, 2008). These dynamics converge toward the learned state, which is a stable state of the learning rules. The learning dynamics are driven by the ongoing activation in the system. This is how experience matters and shapes the future network.

Understanding the dynamics of learning deeply is the theoretical frontier in our understanding of development. In spite of much work by connectionist modelers and the demonstration of the power of the memory trace within DFT, most of the fundamental questions about learning as a process have not been addressed. One big question is how learning is organized autonomously, so that a reproducible sequence of learning phases emerges from the experience of infants and children. How is consistency created across the varying experiences over the course of a day or of weeks, with continuing progress on a particular front at a particular stage of development? What role does reward play, and what are intrinsic rewards?

Worked Examples of DFT

Neural Grounding of Fields and Neural Dynamics

DFT is based on the hypothesis that neural processes control behavior through the activity of populations of neurons. To be specific, consider the preparation of a movement of the hand toward an object or location. Recording from the motor and premotor cortical areas of macaque monkeys, researchers have found populations of neurons whose firing is modulated in time, reflecting the movement task (Georgopoulos et al., 1986). These neurons begin to fire as the movement is prepared, and then stop firing at some point early in the actual hand movement. Moreover, the firing rate of the neurons depends on the spatial direction of the hand movement. That dependence is captured by the tuning curve, depicted schematically in Figure 8.3, which represents the neural firing rate as a function of movement direction. The tuning curves of most cortical neurons are single-humped and broad and can be approximated by a cosine function centered on a “preferred movement direction” of each neuron. The cosine function is the broadest smooth period function of direction, so tuning is broad in an objective sense. By implication, whenever a single movement direction is specified, a large ensemble of neurons fires—all those whose preferred direction lies within a range of about 180 degrees from the specified direction. Approximately half of the entire population of neurons tuned to movement direction is active when any specific movement is being prepared. Similarly broad patterns of activation have been observed in many other cortical and subcortical areas.

The activation of a population of neurons that is tuned to a task can be visualized as the formation of a peak in a field of neural activation. In some cortical areas such as the visual and other primary sensory cortices, neighboring neurons tend to have similar tuning curves. In such cases, neurons activated in the task will tend to be

neighbors. The peak of activation can be thought of as sitting on the cortical surface, over which the neural population is distributed. In other areas, such as the motor and premotor cortex, neighboring neurons may have quite different tuning curves. The neurons activated in a task may be broadly distributed over the cortical area. We can still visualize the activation pattern as the formation of a peak, however, if we reorder the neurons in a virtual (or functional) space, in which neurons are put next to each other if they have similar tuning curves. The concept of a distribution of population activation (DPA) does this job. Each neuron contributes its entire tuning curve to a distribution of activation defined over the task space (Figure 8.3). In the example, the activation is thought of as distributed over the space of movement directions of the hand. When a specific movement is prepared, a peak in this activation field is positioned over that movement direction, which is currently specified (Figure 8.4).

As a movement is prepared, a peak arises in response to inputs that specify the upcoming motor act. In the laboratory, visual cues to the movement target may be provided at some point, but a participant may also anticipate and preactivate movements that have been frequently elicited in given task setting (Erlhagen & Schöner, 2002). Such prior activation can be induced by the layout of the environment, in which a few graspable or reachable objects may be visible in the work space. It may also arise out of learned motor habits, in which previous reaching acts leave memory traces near the movement directions of those previous reaches. In each individual reaching incident, one particular movement direction must be selected from among the preactivated locations in the field.

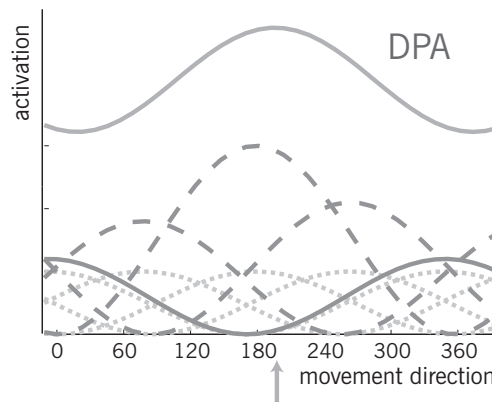


FIGURE 8.3. From tuning curves (schematically shown as dotted lines) of neurons in motor and premotor cortex to the direction of a hand movement, the distribution of population activation (DPA; schematically shown as a solid line) can be constructed by weighting the tuning curves with the firing rate currently observed (dashed line). In the illustration, the movement indicated by the arrow was prepared, leading to a current firing pattern in which neurons tuned to directions close to that value, firing more and thus weighing more than neurons tuned to other directions. As a result, the DPA develops a single peak centered on the currently prepared movement direction. Redrawn from Bastian, Schöner, and Riehle (2003).




FIGURE 8.4. The DPA is obtained as a function of time by using the current firing rate from about 100 neurons in the motor cortex when a movement direction is first cued at “PS” and then the “go” signal is given at “RS.” Data from Bastian, Schöner, and Riehle (2003). An optimal linear estimator was used to smooth the representation, as described in Erlhagen, Bastian, Jancke, Riehle, and Schöner (1999).

Selection decisions require competition within the neural population. All movement directions cannot become activated at the same time. Inhibition within the population is a neural mechanism that may bring about such competition. Inhibition is generic, of course, within cortical and subcortical neuronal networks. Excitatory coupling may also exist, and there is a typical gradient in which neurons with similar tuning curves tend to have net excitatory coupling, whereas neurons with dissimilar tuning curves tend to have net inhibitory coupling. In the network dynamics of neural populations this pattern of interaction stabilizes localized patterns of activation—peaks or blobs or bubbles depending on how many dimensions we envisage. The excitatory interaction among similarly tuned neurons stabilizes a localized pattern of activity against decay, as the neurons contributing to the localized patterns push each other up toward higher activation. The inhibitory interaction among more dissimilarly tuned neurons counterbalances this excitatory influence, which would tend to activate the neighborhood of a localized patterns of activation, leading to activation spreading away from the localized pattern. Inhibition thus stabilizes peaks against diffusive decay. Localized peaks of activation are thus the stable states, or attractors, of the neural dynamics that may emerge under appropriate circumstances.

In DFT we abstract from the details of neural population activity and tuning curves. We think of distributions of neural activation defined directly over the functionally relevant spaces, here, the space that consists of all possible directions in space of hand movement. Patterns of localized activation are then peaks that are localized in these functional spaces, and that location specifies values along these dimensions. A movement that has been prepared is thus a peak of activation positioned at the direction in which the movement will be made. Connectivity downstream from the neural population captured by the dynamic field brings about that movement. That is how the tuning curves of the neurons that contribute to the field come about. For sensory representations, analogously, connectivity from the sensory surface to the neurons that contribute to the dynamic field brings about the sensory event encoded by a localized peak of activation, again, as captured by the tuning curves of the contributing neurons.

The A-not-B Paradigm

We have seen how neural dynamics are grounded in the population activity of the brain. How do we use the concepts of DFT to understand how behavior and cognition unfold in the here and now? To answer this question in the concrete, let us have a look at what happens when an infant is induced to reach toward objects in Piaget's *A-not-B* paradigm (Figure 8.5). The infant would be seated in an infant chair and you, the experimenter, would interact with the baby.

At some point you retrieve an attractive toy from somewhere and show it to the infant by bringing into his or her field of view. Even very young infants will direct their gaze to the toy. Infants around 5–6 months of age have learned to reach. Such infants will often initiate a reach for the toy when you move the toy close enough so that it is reachable. How would we think of such a reaching act within DFT (Thelen, Schöner, Scheier, & Smith, 2001)? A reach is initiated when a peak of activation arises in a dynamic field that represents prepared movements. Because the location of the peak determines the orientation of the reach, the peak should be positioned over the movement direction in which the object lies.

Sensory information about this object comes from the visual system. Extracting from the visual system information about the direction of a reach to a visible object actually requires a whole sequence of neural events, which are not modeled in any detail in this simplest description. For instance, neural mechanisms generate a shift of attention to the visual location of salient stimuli, usually accompanied by a shift of gaze to foveate the object. This entails a perceptual process of figure–ground segmentation in which the visual system focuses on the object. Extracting movement parameters from visual information about the object further requires transforming from visual, head-centered coordinates to body-centered coordinates and taking the initial position of the hand into account. All this is summarized in a simplest DFT description of a reach by assuming that a localized pattern of input is provided by the visual system to the dynamic field representing the movement direction of a prepared hand movement.

So before the infant's attention has been captured by the object, no localized input is provided to the movement direction field. When the infant attends to the object, such localized input is present and promotes the building of an activation peak localized over the movement direction in which the object lies. In DFT, the decision to reach amounts to the building of an activation peak in the relevant dynamic field. As the peak is built, could other visual objects, shifts of attention, or other movement intentions weaken the input and prevent the peak from remaining stable? Yes, in principle, this could happen. In DFT the peak is stabilized, however, to some extent. It is stabilized against building competing peaks at other locations, in which other distractor objects lie, by inhibitory interaction. It is stabilized against decay by local excitatory interaction, so a wavering of attention will not necessarily lead to a decay of the peak. A decision to reach is stabilized in the sense that the sensory input required to keep a peak of activation stable is less than the sensory information required to first bring about a peak.

In DFT models the induction of a peak by localized input involves the detection of instability, in which a peak-less pattern of activation becomes unstable, and the system switches to a peaked pattern of activation that is stable. This is illustrated in Figure 8.5. What happens is that below a threshold (conventionally defined as "zero" activation), neurons in the population do not interact. This reflects that only sufficiently activated (i.e., spiking) neurons impact on their postsynaptic targets. So excitatory and inhibitory interaction do not come into play until activation somewhere in the field exceeds that threshold. When that happens, local excitatory interaction drives up




FIGURE 8.5. *Left:* An infant reaching to one of two cued locations, “A” or “B,” at which lids are placed on a box. On an “A” trial, the experimenter attracts the attention of the infant to the “A” location, for instance, by waving a lid in front of the infant and then setting that lid down on the “A” location. *Right:* An activation field defined over the movement parameter “direction” accounts for the motor plans that the infant is generating. The waving action leads to localized (“specific”) input to the field at the “A” location. The lid visible at the “B” location receives smaller input. Sufficiently strong, specific input at “A” induces a suprathreshold peak of activation. Local excitatory interaction leads to higher levels of activation within the peak than dictated by input. Global inhibitory interaction suppresses activation below resting level elsewhere in the field.

activation around the location where the threshold was first pierced. A suprathreshold peak grows. At the same time, inhibitory projections from that same region of the field begin to affect other regions, suppressing the field at other, alternative reaching directions. Once a peak has been generated, it is stabilized by the excitatory interaction within the peak region and inhibitory projection from the peak onto the rest of the field. Even when input is weakened, the peak may resist because it is, in a sense, self-stabilized by excitatory interaction.

Not every object in the visual array will automatically generate a reach. Infants are not little reaching machines that reach no matter what. For one thing, even as an infant attends to the object, he or she will often initiate a reach only when the object is pushed close enough to be in reaching space. Thus, other factors than the localized input that specifies the movement direction may influence whether or not a peak is generated. These factors can be captured as nonlocalized or as constant inputs to the field. For instance, seeing that the object is reachable, an encouragement by the parent to take the object or an endogenous intention to reach may provide a global boost to the dynamic field of movement direction. Such global input may push the activation field through the threshold somewhere, most likely where localized input is already present but was not sufficient by itself to induce a peak. In DFT we call this the *boost-driven detection instability*, and it is illustrated in Figure 8.6.

Another factor that may support initiating a reach is learning—that is, building the habit to reach. If we play a little game with the infant, presenting an object several times in a row, the infant often willingly reaches for the object. The infant may now need less of a cue to initiate a reach. This is exploited by experimenters, who often use an initial training regimen in which they push an object closer to the infant on the first few trials. In fact, in Piaget's *A-not-B* experiment, the lid on the "A" trough is often left a bit off kilter, closer to the infant on the first few "A" trials to encourage the infant to reach toward the "A" location. I explain this experiment a little better below.




FIGURE 8.6. Pushing the box with the two lids into the reaching space of the infant is modeled by applying a homogeneous excitatory input, a *boost*, to the entire field. This boost may push activation at the "A" or the "B" location through the detection instability and induce a peak there. The field is preactivated at these locations due to residual activation from a peak that decayed during the delay, due to the memory trace of earlier reaches to a location, or due to the visual structure of the scene with lids visible at the two locations.

A habit to reach toward a specific location can be described in DFT as coming from an additional source of localized input that does not originate from the visual array, but from a memory trace of past reaches. Classically, habits have been viewed as powerful ways to reduce the amount of sensory specification required to bring about an action. William James (1890) famously speculated about the pervasiveness and power of habits, which he described as “ease of action” and considered the most common and basic form of learning. The neural mechanisms underlying habit formation are well studied in biological psychology (Yin & Knowlton, 2006), but little is known about how human infants acquire habits. In any case, in the *A-not-B* paradigm, a very small number of reaches is sufficient to induce a reaching habit, so this is a fast and robust form of learning. It is, perhaps, unclear at this time if the emergence of a memory trace of previous reaches over a few trials is truly akin to the formation of behavioral habits, as observed in paradigms of instrumental conditioning.

How about a hidden object? This is already quite a complex scenario. Presumably, the infant “knows” about the hidden object because earlier he or she saw the object, until it was hidden in a trough in a box, say, as in the typical *A-not-B* experiment. The trough was then covered with a lid, so the toy is hidden under the lid. When the object was shown, did the infant prepare a movement to reach for it? Perhaps not if the project was outside its reaching range. But there was some localized input that induced some localized activation, which is now decaying because the input is no longer available. How do we elicit the reach to the hidden toy? By moving the hiding location into the reach space of the infant. The lid is an object, so the reach may be induced just by that—by the visual stimulus of a reachable object. The hidden toy is discovered, becomes visible, and may elicit further object-oriented action. So it isn’t necessarily the case that reaching for a hidden toy requires a true memory of the toy.

In the *A-not-B* experiment (Piaget, 1954; Wellman, Cross, & Bartsch, 1986) there are two locations, typically a box with two troughs, each covered with a lid. When the toy is first hidden under the “A” lid, the infant does select the “A” lid for his or her reach in most cases. This means that the previous action of attracting attention to the “A” location—of hiding an attractive toy there and putting down the lid over that location—left some activation trace that biases the competition among the two reaching locations, the two lids, toward the “A” location.

In many *A-not-B* experiments, the two lids are not equally spaced on the first few training trials: as noted previously, the “A” lid is left closer to the infant than the “B” lid, inviting a reach toward “A.” This may bias the decision independently of any remaining activation from previous events.

Once the reach to “A” has been elicited several times over a few “A” trials, a memory trace solidly biases the motor decision toward the “A” location each time the box is pushed into the infant’s reaching space. By the time the hiding location is switched to the “B” location, the bias toward “A” from the memory trace may be strong enough to overcome any remaining activation trace from the stimulation at “B” when the toy is hidden. The *A-not-B* perseverative error ensues.

The hallmark of the *A-not-B* error is the critical role of the delay. When the box is pushed toward the infant right after the toy is hidden at “B,” the infant is likely

to reach correctly toward “B.” Young infants make the error even at short delays, whereas older infants tolerate increasingly longer delays. The delay directly controls the tradeoff between the decaying activation induced by the stimulus at “B” when the toy is hidden there and the bias to “A” that the motor habit, the memory trace at “B,” imposes.

Presumably, the decay of activation at “B” is slower in older infants because they may engage neural interaction more easily, better stabilizing the activation at “B” against decay. (The possibility that older infants may sustain a peak at “B” over long delays as a form of working memory for a planned action requires an account for why they still reach only when the box is pushed into their reaching space. It is likely that understanding the initiation of reaching requires a more refined account, in any case.)

Clearly, the *A-not-B* paradigm does not probe reaching to hidden objects per se. Rather, it probes how flexible infants are in switching to a new pattern of activation that must be stabilized against decay. The memory trace is a slow and inflexible form of stabilization of motor behavior, whereas the neural interaction within a neural field is a faster form of stabilization that may thus flexibly respond to changed demands.

Is the toy then actually needed? A toyless variant of the paradigm demonstrates that perseverative reaching emerges even in a purely sensory–motor version of the task (Smith, Thelen, Titzer, & McLin, 1999). So this is at heart a sensory–motor decision task in which different factors that bias the selection decision are varied. The developmental trajectory is one in which increasingly, the system is less dependent on current sensory input to bring about the selection decision.

Are the lids really needed? In a lidless version of the task (Spencer, Smith, & Thelen, 2001), the toys are hidden in sand, leaving no visible trace of the hiding location. Perseverative errors now show up as metric errors in which the searching location is biased metrically toward the “A” location. This can be understood in a different picture, in which a peak is sustained during the delay, but drifts under the influence of the memory trace. This metric bias occurs in much older children. Toddlers of 4 years, for instance, show a strong metric bias consistent with perseveration. The localized perceptual input at the locations marked by the lids in the conventional *A-not-B* paradigm locks peaks of activation in place. The absence of such visible perceptual markers makes the field more sensitive to other inputs or to internal inhomogeneities. This is how the memory trace of past reaches may have observable consequences in 4-year-olds in the sandbox, whereas infants older than about a year no longer show perseveration in the conventional *A-not-B* paradigm.

The sandbox reaches into spatial working memory (Schutte & Spencer, 2009). The hiding location must be remembered as just that, a location along a metric continuum. In DFT, working memory for metric dimensions arises from the capacity of neural interaction to sustain peaks of activation once the inducing localized input has been removed. This is the most common neural account for working memory and is consistent with the observation of sustained firing of neurons in many cortical structures (Fuster, 1995). Excitatory and inhibitory connectivity within a neural population is critical to enable self-sustaining neural activation patterns. Input from the sensory surface alone is not sufficient. The stronger such intrapopulation connectivity, the

stabler the sustained activation patterns. Stable patterns of sustained activation may resist distractor inputs more effectively and are less easily biased by such inputs.

Across a broad set of experimental paradigms, the developmental trajectory has been found to shift from the dominant role of feedforward sensory inputs into neural representations toward a dominant role of intrapopulation recurrent connectivity (Schutte & Spencer, 2009; Schutte, Spencer, & Schöner, 2003; Simmering, Schutte, & Spencer, 2008). Activation patterns in the brains of younger infants seem to be more strongly dictated by inputs from the sensory surfaces and by fixed habits (modeled as inputs from memory traces). During development, activation patterns are increasingly controlled by input from within the neural population, whereas the role of external inputs is reduced. This developmental trend has been called the *spatial precision hypothesis* (Schutte, Spencer, & Schöner, 2003) because it manifests in increasing precision with which spatial (or other metric) information can be retained and processed. This same tradeoff between the influence of inputs, including a form of memory (latent memory implicit in connections), and the influence of recurrent connectivity (active memory) is at the core of a connectionist account of perseverative reaching (Munakata, McClelland, Johnson, & Siegler, 1997) that is largely equivalent to DFT.

The Habituation Paradigm

How about perception? The neural concepts at work are the same as for the motor domain: The space of possible percepts is spanned by feature values such as retinal location, visual orientation, movement direction, or color (see Figure 8.7 for an example). Cortical neurons are sensitive to such feature dimensions as indicated by their tuning curves. Any given stimulus evokes a distribution of population activation that may be construed as a peak localized within the feature spaces. Such a peak is then the neural representation of a percept, whose location represents an estimate of the feature value while the amplitude may reflect other stimulus properties such as contrast.

Perceptual peaks arise as a sensory stimulus impinges on the sensory surface. The forward connectivity from the sensory surface to the neural representation is largely responsible for the feature extraction, which determines where a peak is induced. The peak itself is a stable state of the perceptual field, stabilized by interaction. The detection instability, at which the peak arises, separates the preperceptual from the perceptual state. Typically, when the sensory array is sufficiently complex, a perceptual detection decision entails at the same time a selection decision in which one particular perceptual object or one particular perceptual dimension is brought into the foreground. Featural salience influences which object is thus selected. Infants are particularly sensitive to movement, for instance. In DFT, this selection decision amounts to generating a peak localized over one stimulated location, while inhibiting activation over other stimulated locations within a feature dimension.

This is a simple picture, in which percepts are attentional blobs in feature space, with a small number of estimated feature values. Even this simple-minded picture may account for much of what is known about infant perception and its development. In particular, the metrics along the feature dimensions matter for the perceptual




FIGURE 8.7. A perceptual activation field defined over a feature dimension, in this case, the direction of visual motion. An object moving in different directions is represented by a peak localized over the corresponding location in the field (*top*: solid line for diagonal motion, dashed line for horizontal motion). The luminance and complexity of the stimulus are reflected in the amplitude of the induced activation peak, higher for a bright stimulus with inner structure (*bottom*: dashed line) than for a plain stimulus (*bottom*: solid line).

experience of infants: Metric differences along feature dimensions separate familiar from novel stimuli, and development manifests in the metric differences required to shift from familiar to novel.

Perception is much harder to observe in infants than overt motor action, of course. The tools of modern psychophysics are not readily applied. One common way of indirectly assessing the state of visual perceptual representations is to register patterns of looking (Cohen, 1972). This method is based, of course, on the well-documented principles of perceptual habituation, in which the probability of continuing to fixate on an unchanged stimulus decreases as a function of the total looking time (Colombo & Mitchell, 2009). This more indirect assessment of the state of the perceptual system requires some extra machinery in DFT to link activation peaks in feature spaces to looking behavior.

Does decreased looking reflect a decrease or an increase in perceptual activation? Any dependence of “dwell time” on prior perceptual experience is used as an index of perceptual processing (Colombo, 2001). *Familiarity preference* is the tendency to look more at a similar or identical stimulus to the one previously looked at; *novelty*

preference is the tendency to look more at a dissimilar stimulus. An account for how perceptual activation is linked to looking within DFT (Schöner & Thelen, 2006) is illustrated in Figure 8.8. A first assumption is that fixation is strengthened by perceptual activation. Thus, whenever a visual stimulus is perceived, the induced activation keeps the infant fixating on it. Perceptual activation thus promotes continued looking, accounting for the familiarity preference early during looking (Roder, Bushnell, & Sasseville, 2000). The excitatory link from visual activation to fixation is consistent with prolonged saccadic response times when a visual stimulus continues to be visible at fixation (overlap) than when it is removed (gap) in adults (Kopecz, 1985).

A second assumption is that perceptual activation drives the buildup of inhibition that is represented in a separate field in the model. Increasing inhibition, in turn, weakens perceptual activation and thus promotes looking away, accounting for novelty preference later during looking. Such inhibition is a well-known form adult visual perception, wherein it is observed as an increase in perceptual threshold after prolonged experience of a particular percept, a phenomenon known as *selective adaptation* (see Hock, Schöner, & Hochstein, 1996, for the subtle issue of whether conscious perception is required to induce selective adaptation).

Together, the two assumptions lead to an account of the typical pattern of looking in infants exposed to a visual stimulus—an early increase followed by a decrease in looking time that is referred to as *visual habituation* (Colombo, Frick, & Gorman, 1997). The early increase is not always observed for reasons that can be understood. The DFT model also predicts how habituation is modulated by stimulus strength and prior levels of activation, which account for the difference between “fast” and “slow” habituators (Schöner & Thelen, 2006).

The DFT model also explains the pattern of dishabituation routinely used to assess infants’ visual perception and cognition (Kaplan & Werner, 1986). Figure 8.8 illustrates a typical paradigm. In a first phase the infant is exposed repeatedly to a habituation stimulus, here a figure moving either horizontally or diagonally, until the infant has *habituated*, that is, until the time the infant looks at the stimulus has fallen from the initial level to a criterion level. In the following test phase, the infant is presented with new stimuli that differ in some way from the habituation stimulus. In the example the motion direction may change or the moving object may change color and have an additional smiley face added to it. If the change in stimulus induces renewed looking, called *dishabituation*, then this is evidence that the infant detected the change of stimulus and thus discriminated the test from the habituation stimulus.

The stimuli shown in Figure 8.9 were used in an unpublished experiment performed in Esther Thelen’s lab by Connie Jing Feng, Melissa Clearfield, and myself. Overall, ninety-six 16-month-old infants participated, 24 in each of the four stimulus conditions shown. The typical stimulus-controlled procedure was used, in which an attention-grabbing stimulus (chime with a concentric moving pattern) attracted the infant’s gaze to the display, which was then switched to the current stimulus until the child looked away for more than 2 seconds or the trial lasted more than 120 seconds. The habituation phase lasted at least 4 trials and a maximum of 20 trials. It was aborted when the mean looking time for three consecutive trials fell below 50% of the

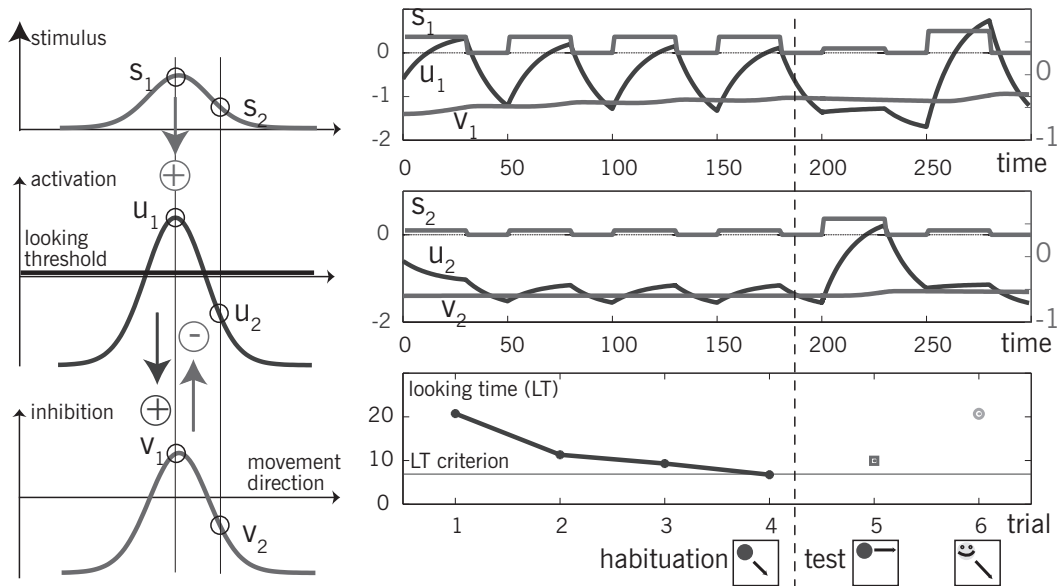


FIGURE 8.8. *Left:* In a model of visual habituation (Schöner & Thelen, 2006), an activation field (*middle*) defined over a relevant feature dimension (here, the direction of visual motion) receives input from the sensory surface (stimulus on top). The activation field drives an inhibitory field (*bottom*), which conversely inhibits the activation field. Levels of activation above a looking threshold promote fixation on the stimulus. *Right:* The neural dynamics of the activation and inhibition fields are simulated at two locations along the feature dimensions corresponding to diagonal (index 1) and horizontal motion (index 2). The moving objects are shown repeatedly (illustrated by the boxcar stimulus trace for S_1 and S_2). Whenever a matching stimulus is presented, activation at the corresponding location is driven up. For instance, while a diagonally moving object is presented (habituation phase), activation at location 1 (*top panel*) is driven up, leading to looking at the moving object (*bottom panel*). Activation decays when the stimulus is removed. Across stimulus presentations, inhibition increases whenever there is sufficient activation, leading to a reduction of activation and looking time during the habituation phase. In the test phase, dishabituation occurs when a new stimulus (horizontal motion, *second from right*) is shown or when the stimulus receives added strength (diagonal motion with brighter and more complex object, *far right*).

looking time during the first three trials. Figure 8.10 shows the mean looking times across conditions together with error bars and simulation results.

The pattern of looking in these experiments shows dishabituation when the movement direction changes from diagonal to horizontal (e.g., in condition 1A for the first test stimulus) or from horizontal to diagonal (e.g, in condition 2A for the first test stimulus). In the DFT model, such dishabituation arises when the test stimulus differs sufficiently along the perceptual dimension from the habituation stimulus such that it activates a different location in the perceptual field. That new location is not yet inhibited by previous perceptual experience and may therefore build up a strong activation

peak that supports renewed looking. This account for dishabituation is symmetrical: Either direction of change along the perceptual dimension leads to dishabituation as long as the distance along that dimension is sufficiently large. This symmetrical form of dishabituation is thus truly a signature of discrimination.

The results also show dishabituation that is asymmetrical. A change from the simple to the bright and complex object moving in the same direction leads to renewed looking (e.g., in condition 1B, first test). The reverse change from a bright, complex to the simple object (e.g., in condition 4B, first test) does not lead to dishabituation. In the DFT model, such asymmetrical dishabituation reflects a change in input strength that leads to renewed looking through higher perceptual activation only when input strength is increased, not when it is decreased. In fact, the DFT model predicts shorter looking time than during habituation—the only mismatch between the model and this experiment (on condition 4). (In the experiment such a negative difference is not possible, because the attention-grabbing stimulus is presented until the infant looks, so looking times cannot fall below a lower bound. The attention-grabbing stimulus is not modeled).




FIGURE 8.9. Stimuli used in a habituation experiment. The stimuli involved a moving schematic figure that differed between the habituation and test phases by movement direction (either horizontal or vertical) and complexity (either plain blue or bright yellow with a smiley face).

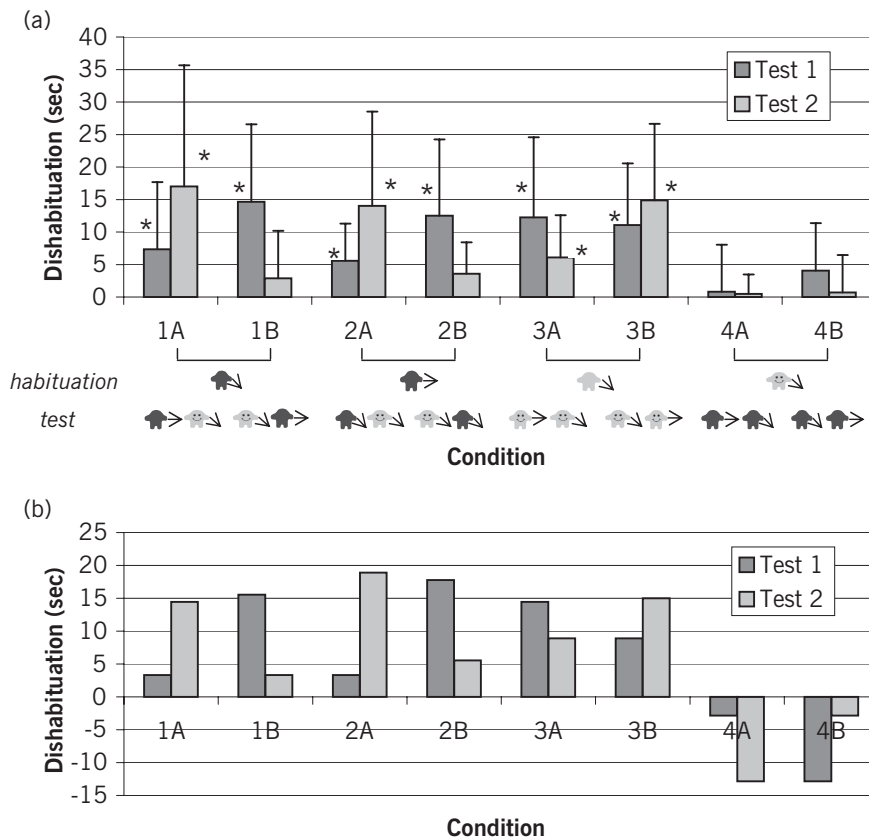


FIGURE 8.10. Mean looking time during the test phase for the four experimental conditions shown in Figure 8.9. Experimental results are shown on top (a). Significant dishabituation is marked by an asterisk (*). Results from model simulations are shown on bottom (b).

Over development, perceptual discrimination generally improves (Kaplan & Werner, 1986). In DFT this improvement is consistent with a generalization of the spatial precision hypothesis, according to which activation peaks become sharper and stabler over development. As a result, they will overlap less. Shorter displacements of the test stimulus along the feature dimension are required to induce new activation peaks that do not share inhibition with the habituation stimulus.

A variant of the habituation paradigm has been extensively used to study infant cognition. In this variant, hypotheses are formed about infant “knowledge.” For instance, in the visual *drawbridge paradigm* (Baillargeon, 1987a), the hypothesis is tested that infants “know” that solid bodies cannot occupy the same location in space. Infants are habituated to a visual stimulus in which a wooden panel (the “drawbridge”) rotates from horizontal back to horizontal 180 degrees toward and away from the infant (Figure 8.11). Then one out of two new stimuli is presented. One is “novel”: It

differs from the habituation stimulus more than the other new stimulus. In the novel condition, a wooden block is visible to the infant, and the drawbridge is lifted until it occludes and touches that block, which leads to less rotation, about 120 degrees. This stimulus “respects” the infant’s knowledge that solids cannot move through each other, because the drawbridge’s motion is stopped by the block, or so it appears. It is the “possible” stimulus in light of the infant’s hypothesized knowledge. The other “familiar” condition is similar to the habituation stimulus, in that the drawbridge makes the same 180 degree rotation away and toward the infant. This condition is “impossible,” because it violates the infant’s knowledge about solids: Initially, a block is visible that must lie in the path of the drawbridge, but the drawbridge continues motion through the block (in reality, the block is dropped down through an opening in the table).

The logic of these types of paradigms is this: Infants are expected to dishabituate more to the novel than to the familiar stimulus. If they look longer at the familiar but impossible stimulus, then this is evidence that there is something surprising about that stimulus—the violation of infant expectation derived from their knowledge. Looking longer at impossible than at possible stimuli is thus considered evidence in support of the hypothesis regarding infant knowledge that the impossible stimulus challenges.

DFT suggests that the pattern of looking in this kind of paradigm can be accounted for without any reference to knowledge or violated expectation. In DFT, the metrics of perceptual experience alone can account for the pattern of results. To illustrate the idea, we analyzed the metrics of the drawbridge stimulus (see Figure 8.11). In this paradigm, the block is a new element that appears in both test stimuli, but not in the habituation stimulus. So the presence of the block alone leads to some dishabituation through increased input strength. Now the spatial structure of the stimulus is changed by the block: The drawbridge moves through space covering a particular area that overlaps with the block. This overlap is more extensive when the drawbridge moves through the spatial location occupied by the block than when the drawbridge stops short of that location. That difference accounts for why there is stronger perceptual input for the “impossible” stimulus, leading to more looking.

Similar analyses can be made for many of the classical habituation scenarios that are hypothesized to probe infant knowledge. The DFT model also makes specific predictions that support the account (Schöner & Thelen, 2006). In particular, the temporal order in which test stimuli are presented is shown to matter: When the familiar–impossible is presented first on the test, its stronger overlap with the habituation stimulus makes that activation start from a higher level of activation, leading to more dishabituation than when the novel–possible stimulus is presented first. In experiments, this interaction between the order of presentation and the preference for the impossible is routinely observed. In fact, in the original experiment (Baillargeon, 1987a), the looking advantage of the impossible stimulus comes entirely from this order of presentation.

Similarly, individual differences are understood in the DFT account (Schöner & Thelen, 2006). Infants with a higher initial level of activation will have a long initial look, leading to faster and deeper habituation and a larger difference between the two




FIGURE 8.11. The stimuli used in the drawbridge experiment illustrated by seeing the arrangement from the side. The infant sees a wooden panel (“drawbridge”) shown here at three moments in time (dashed and solid straight lines). The rectangle illustrates a wooden block that sits in the path of the drawbridge for the test stimuli.

test stimuli. Again, in the original experiment, the difference between the two test stimuli came entirely, statistically speaking, from the fast habituators (Baillargeon, 1987a).

This DFT account is only a first sketch of all that is entailed in understanding infant perception and visual habituation. More recently, Perone and Spencer (2013) have substantially extended this account into a comprehensive theory of visual habituation and visual cognition. In their model, the control that infants exert over their own visual stimulation is modeled in more depth than was done in Schöner and Thelen (2006) by adding a dynamic fixation system constrained by what is understood from work on adult visual fixation (Kopocz, 1995). More importantly, this work links the familiarity-to-novelty transition to the emergence of working memory. The (spatial) precision hypothesis, shown in earlier work to account for the increase in the capacity of working memory in infancy (Perone, Simmering, & Spencer, 2011) is thus being confirmed as a pervasive principle of development.

Conclusion

The two worked examples have underscored how far we have moved from the self-organization metaphor that was based on analogies with physics to neurally mechanistic accounts of infants’ and childrens’ behavior and cognition. The strong neural emphasis in this transition from metaphor to theory may have come as a surprise to some readers. But it is consistent with the need to ground the concepts of self-organization in empirically valid mechanisms: that is, the patterns of neural interaction. The critical question is if this neural grounding preserves key features of the

metaphor. Does the neural theory address the notion of the emergence of cognitive and motor competences under favorable environmental conditions? Does the theory articulate how these competences may emerge from experience and account for individual routes of development?

Emergence is, in fact, an inherent property of DFT. Cognitive competences, such as the ability to stabilize decisions and keep them in working memory, are captured by attractor states of the neural dynamics: stable peaks of activation within neural populations. The stability of these states emerges from the confluence of different contributions. For instance, the neural dynamics of an older infant may stabilize a peak that keeps a reaching goal in working memory just based on its more strongly developed neural interaction. A younger infant may achieve the same with weaker neural interaction if he or she receives additional broad input, for instance, from an environment with rich perceptual structure (Schöner & Dineva, 2007). Thus, whether a hidden object is “out of sight, out of mind” for an infant depends not just on age, but on environmental conditions. The neural dynamics of younger infants make stronger demands on the environment to achieve the cognitive capacity that older infants achieve more easily. Perone and colleagues (2011) have provided a similar analysis of how the capacity of visual working memory emerges depending on the task settings. The conceptual implications are lucidly discussed in Simmering and Perone (2013), who review 50 studies of visual working memory, highlighting how the confluence of different factors impacts on how infants perform in different circumstances.

Individual differences are captured through differences in the neural dynamics formalized as different parameter values. This accounts for covariation of behavioral signatures of the dynamics. For instance, in the DFT account of infant habituation (Schöner & Thelen, 2006), looking patterns during the habituation phase are predictive of the amount of dishabituation on test. Perone and Spencer (2013) have pushed the account of DFT for individual differences to a new level. They show how individual differences may emerge from learning processes. In simulations of long looking episodes in a visual habituation task, fluctuations in the looking behavior early during a trial may lead the system to build visual memories more efficiently, leading to characteristic signatures of a more mature memory system.

This chapter has illustrated the transition from metaphor to neural theory based on elementary forms of cognition, close to cognition’s sensory–motor origin. Social interaction is, of course, central to development. Can the neural theory be put to work on language and social interaction? The answer is yes, although moving to higher cognition is one of the frontiers of dynamical systems thinking. In a trailblazing study, Samuelson, Smith, Perry, and Spencer (2011) show how the reference problem that is key to how children learn words while interacting with parents may be solved by activation fields over visual space. Their DFT model accounts for experimental data on word learning and leads to new, tested predictions, suggesting that the neural theory is as productive as the metaphor was.

The neural grounding of cognitive development is a strategy that dynamical systems thinking shares with connectionist and neuroconstructivist approaches to development. The convergence of these currents of theoretical thinking was the topic of a

conference and book (Spencer, Thomas, & McClelland, 2009). The joint frontier of these approaches was clearly recognized as the problem of understanding autonomous development. How do infants and children shape their own behavior to provide the stimulation and experience through which they learn and develop? And how may the reproducible pattern of development emerge through autonomous learning?

REFERENCES

- Baillargeon, R. (1987a). Object permanence in 3.5 and 4.5 month-old infants. *Developmental Psychology*, 23, 655–664.
- Baillargeon, R. (1987b). Young infants' reasoning about the physical and spatial characteristics of a hidden object. *Cognitive Development*, 2, 179–200.
- Bastian, A., Schöner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, 18, 2047–2058.
- Blumberg, M. (2010). *Freaks of nature: What anomalies tell us about development and evolution*. New York: Oxford University Press.
- Cohen, L. (1972). Attention-getting and attention-holding processes of infant visual preferences. *Child Development*, 43, 869–879.
- Colombo, J. (2001). The development of visual attention in infancy. *Annual Reviews of Psychology*, 52, 337–367.
- Colombo, J., Frick, J. E., & Gorman, S. A. (1997). Sensitization during visual habituation sequences: Procedural effects and individual differences. *Journal of Experimental Child Psychology*, 67, 223–235.
- Colombo, J., & Mitchell, D. (2009). Infant visual habituation. *Neurobiology of Learning and Memory*, 92, 225–234.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1997). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schöner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, 94, 53–66.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109, 545–572.
- Fodor, J. (1981). Fixation of belief and concept acquisition. In M. Piatelli-Palmani (Ed.), *Language and learning: The debate between Piaget and Chomsky* (pp. 143–149). London: Routledge.
- Fox Keller, E. (2002). *Making sense of life: Explaining biological development with models, metaphors, and machines*. Cambridge, MA: Harvard University Press.
- Fuster, J. M. (1995). *Memory in the cerebral cortex: An empirical approach to neural networks in the human and nonhuman primate*. Cambridge, MA: MIT Press.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neural population coding of movement direction. *Science*, 233, 1416–1419.
- Gierer, A., & Meinhardt, H. (1972). A theory of biological pattern formation. *Kybernetik*, 12, 30–39.
- Gottlieb, G. (2001). A developmental psychobiological systems view: Early formulation and

- current status. In S. Oyama, P. S. Griffiths, & R. D. Gray (Eds.), *Cycles of contingency: Developmental systems and evolution* (pp. 41–54). Cambridge, MA: MIT Press.
- Griffiths, P. E., & Gray, R. D. (2005). Discussion: Three ways to misunderstand developmental systems theory. *Biology and Philosophy*, 20, 417–425.
- Groh, J. M., Born, R. T., & Newsome, W. T. (1997). How is a sensory map read out?: Effects of microstimulation in visual area MT on saccades and smooth pursuit eye movements. *Journal of Neuroscience*, 17, 4312–4330.
- Haken, H. (1983). *Synergetics: An introduction* (3rd ed.). Berlin: Springer Verlag.
- Harrison, L. G. (2011). *The shaping of life*. Cambridge, UK: Cambridge University Press.
- Hock, H. S., Schöner, G., & Hochstein, S. (1996). Perceptual stability and the selective adaptation of perceived and unperceived motion directions. *Vision Research*, 36, 3311–3323.
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Kaplan, P., & Werner, J. (1986). Habituation, response to novelty, and dishabituation: Tests of a dual-process theory. *Journal of Experimental Child Psychology*, 42, 199–217.
- Kopecz, K. (1995). Saccadic motor planning by integrating visual information and preinformation on neural dynamic fields. *Biological Cybernetics*, 73, 49–60.
- Kriegeskorte, N., & Kreima, G. (Eds.). (2012). *Visual population codes: Toward a common multivariate framework for cell recording and functional imaging*. Cambridge, MA: MIT Press.
- Lee, C., Rohrer, W., & Sparks, D. (1988). Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature*, 332, 357–360.
- Lipscomb, L., Swanson, J., & West, A. (2004). Scaffolding. In M. Orey (Ed.), *Emerging perspectives on learning, teaching, and technology*. Retrieved March 18, 2013, from <http://projects.coe.uga.edu/epltt>.
- Molenaar, P. C. M. (1986). On the impossibility of acquiring more powerful structures: A neglected alternative. *Human Development*, 29, 245–251.
- Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104, 686–719.
- Murray, J. D. (2002). *Mathematical biology II* (3rd ed.). New York: Springer Verlag.
- Newell, K. M., Liu, Y.-T., & Mayer-Kress, G. (2009). Timescales of change in connectionist and dynamical systems approaches to learning and development. In J. P. Spencer, M. S. C. Thomas, & J. L. McClelland (Eds.), *Toward a unified theory of development* (pp. 119–138). New York: Oxford University Press.
- Perone, S., Simmering, V. R., & Spencer, J. P. (2011). Stronger neural dynamics capture changes in infants' visual working memory capacity over development. *Developmental Science*, 14, 1379–1392.
- Perone, S., & Spencer, J. P. (2013). Autonomy in action: Linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive Science*, 37, 1–60.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Quartz, S. R. (1993). Neural networks, nativism, and the plausibility of constructivism. *Cognition*, 48, 223–242.
- Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. *PLoS One*, 6(12), e28095.
- Schöner, G. (2008). Dynamical systems approaches to cognition. In R. Sun (Ed.), *Cambridge handbook of computational cognitive modeling* (pp. 101–126). Cambridge, UK: Cambridge University Press.

- Schöner, G., & Dineva, E. (2007). Dynamic instabilities as mechanisms for emergence. *Developmental Science*, *10*, 69–74.
- Schöner, G., & Kelso, J. A. S. (1988). Dynamic pattern generation in behavioral and neural systems. *Science*, *239*, 1513–1520.
- Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychological Review*, *113*, 273–299.
- 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., & Schöner, G. (2003). Testing the Dynamic Field Theory: Working Memory for Locations Becomes More Spatially Precise Over Development. *Child Development*, *74*, 1393–1417.
- Simmering, V. R., & Perone, S. (2012). Working memory capacity as a dynamic process. *Frontiers in Psychology*, *3*, 567.
- 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.
- 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.
- Spencer, J. P., Clearfield, M., Corbetta, C., 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., 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. L. (Eds.). (2009). *Toward a unified theory of development*. New York: Oxford University Press.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Brain and Behavioral Sciences*, *24*, 1–33.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, Massachusetts: MIT Press.
- Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London*, *B237*, 37–72.
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550–592.
- Waddington, C. H. (1953). *The epigenetics of birds*. Cambridge, UK: Cambridge University Press.
- 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* (No. 214), *51*, 1–67.
- Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*, 464–476.
- Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science*, *256*, 1327–1331.