

In: *The Cambridge Handbook of Computational Psychology*,
Ron Sun, (ed.), Cambridge University Press (2008), pages 101-126

CHAPTER 4

Dynamical Systems Approaches to Cognition

Gregor Schöner

1. Introduction

Think of a little boy playing in the playground, climbing up on ladders, balancing on beams, jumping, running, catching other kids. Or think of a girl who prepares to draw a picture, finding and setting up her painting utensils, dipping the brush in water, carefully wiping it off, whipping up the water paint of the selected color with small circular movements, the brush just touching the pad of paint. When she actually paints, she makes a sequence of brush strokes to sketch a house. Clearly, both scenes involve lots of cognition. The ongoing, complex behaviors of the two children are certainly not simple reflexes, nor fixed action patterns elicited by key stimuli, nor strictly dictated by stimulus–response relationships. Hallmarks of cognition are visible: selection, sequence generation, working memory. And yet, what makes these daily life activities most intriguing is how seamlessly the fine and gross motor control is tuned to the environment; how sensory information is actively sought by looking around, searching, establishing reference; and how seam-

lessly the flow of activities moves forward. No artificial system has ever achieved even remotely comparable behavior. Although computer programs may play chess at grand master level, their ability to generate smooth flows of motor actions in natural environments remains extremely limited.

Clearly, cognition takes place when organisms with bodies and sensory systems are situated in structured environments, to which they bring their individual behavioral history and to which they quickly adjust. There is a strong theoretical tension in cognitive science about the extent to which cognition can be studied while abstracting from embodiment, situatedness and the structure of the nervous systems that control cognitive processes in organisms. This chapter argues that in making such abstractions, important concepts are missed, including most importantly the concepts of stability and instability.

The embodied view of cognition emphasizes the close link of cognition to the sensory and motor surfaces and the structured environments in which these are immersed. The dynamical systems approach

to cognition is the theoretical framework within which this embodied view of cognition can be formalized. This chapter reviews the core concepts of the dynamical systems approach and illustrates them through a set of experimentally accessible examples. Particular attention will be given to how cognition can be understood in terms that are compatible with principles of neural function, most prominently, with the space–time continuity of neural processes.

2. Embodiment, Situatedness, and Dynamical Systems

Cognition is embodied in the obvious sense that natural cognitive systems are housed in a physical and physiological body, and that cognitive processes take place within the organism's nervous system. Cognition is situated in the similarly obvious sense that this body acts in a specific, structured environment from which it receives sensory information and on which it may have an effect. Body and nervous system are adapted to natural environments on many time scales, from evolution to development and learning. Any individual organism brings its particular history of behavior and stimulation to any situation in which cognition is acted out.

In another sense, embodiment is a scientific stance, in which researchers aim to understand cognition in ways that do not neglect the linkage between cognitive processes and the sensory and motor surfaces, do not neglect the structured environments in which cognition takes place, are mindful of the potential role of individual experience in cognition, and are careful when abstracting from the concrete neuronal processes that are the basis for the behavior of organisms.

Taking that stance does not prevent researchers from building artificial cognitive systems or from constructing abstract mathematical models of cognitive processes. But in each case, the potential link to a body, to an environment, and to a stream of behavior must be considered. Whether the theo-

retical constructs employed are compatible with organizational principles of the nervous system must be examined. Models of cognition that take the embodied stance must be process models that can capture, at least as a matter of principle, the unfolding in time of cognition and the associated sensory and motor processes.

Often, building a robotic demonstration of a process model is a useful test of the extent to which the principles of embodiment have been respected. Many classical production system modelers of cognition for instance, face serious, sometimes insurmountable problems when they try to feed their systems from real sensors in the real world and let their systems control real bodies. The limited success of the artificial intelligence approach to autonomous robotics reflects these difficulties (Brooks, 1991).

Dynamical systems theory provides the language in which the embodied and situated stance can be developed into a scientific theory of cognition. To examine this claim, we need to clarify what dynamical systems theory refers to. There is, of course, the field of mathematics that concerns itself with dynamical systems (Perko, 1991). The mathematical concepts capture the property of many natural systems in which a sufficiently rich characterization of the present state of the system enables prediction of future states. Scientific theories based on this mathematical framework have been extraordinarily successful in physics and many branches of science connected to physics. In each case, this required forming scientific concepts based on the mathematical framework, concepts that had to prove their power by capturing laws of nature, properties, and constraints of systems. The mathematics alone did not do that job. By analogy, developing an understanding of cognition within the mathematical framework of dynamical systems requires that concepts are defined that bring the mathematics to bear on the subject matter.

One level at which this has been done with considerable success is that of metaphor. Dynamical systems as a metaphor promote thinking about underlying

“forces” (vector-fields), from which the observed pattern or behavior emerges. The solutions of nonlinear dynamical systems may change qualitatively, even as the underlying vector-fields change only in a graded way. This fact may serve as a metaphor for how qualitatively distinct states or events may emerge from continuous processes, for how there may be multiple possible causes for the emergence of such qualities, and for how all contributions to a system may matter, not only the ones most specifically linked to the new quality. This image dates back, perhaps, to the notion of the Gestalt field in Gestalt psychology (Köhler, 1920/1939) and has been a source of fresh thinking in developmental psychology (Thelen & Smith, 1994). The strongest impact of dynamical systems as a metaphor may be in heuristics, that is, as a source of new questions and new view points.

This chapter, however, reviews efforts to form concepts based on the mathematical theory of dynamical systems into a rigorous scientific approach toward cognition that embraces the embodied and situated stance. The argument will be laid out that the concepts of attractor states with their stability properties, the loss of stability when such states go through bifurcations, and the emergence of new attractor states from instabilities are necessary ingredients of an account of embodied and situated cognition. No physical realization of cognition is possible without addressing the problems to which these concepts provide solutions. Not covered in this review is recent discussion about dynamical systems and embodiment within philosophy (see, for instance, Van Gelder, 1998; Juarrero, 1999; Keijzer, 2001).

3. Dynamical Systems Thinking: Uniquely Instantiated Dynamics

Control systems provide an interesting metaphor for the notion that meaningful function may emerge from simple, embodied mechanisms. A highly illustrative example comes from the orientation behaviors of the common house fly (Reichardt & Poggio,

1976; Poggio & Reichardt, 1976). Flies orient toward moving objects, which they chase as part of their mating behavior. Detailed analysis revealed that the circuitry underlying this behavior forms a simple controller: A motion detection system fed by luminance changes on the fly’s facet eye drives the flight motor, generating an amount of torque that is a function of where on the sensory surface motion was detected. If the speck of motion is detected on the right, a torque to the right is generated. If the speck is detected on the left, a torque to the left is generated. The level of torque passes through zero when the speck is right ahead. The torque changes the flight direction of the fly, which in turn changes the location on the facet eye at which the moving stimulus is detected. Given the aerodynamics of flies, the torque and its on-line updating generate an orientation behavior, in which the insect orients its flight into the direction in which a moving stimulus is detected.

That meaningful behavior emerges as a stable state, an attractor, from the neural circuitry linking the sensory surface to the flight motor, which together with the physics of flight establish a dynamical system (Figure 4.1). In the lab, the behavior can be elicited by imposing a motion signal on the fly’s facet eye. In the fly’s natural environment, the sensory signal typically comes from other flies. In fact, the system is tuned such that pursuit of another fly works amazingly well, probably the outcome of evolutionary optimization.

There is an irony in the scientific history of this analysis. Reichardt and colleagues (Reichardt & Poggio, 1976; Poggio & Reichardt, 1976) had opened the loop by fixing the fly to a torque meter, so that the amount of torque generated by the flight motor could be measured. This was done as a function of the location in the visual array, at which a moving stimulus was presented. From the measured torque, these authors predicted the closed loop behavior. Measuring closed loop behavior still required fixing the fly to the torque meter, but now the visual surround was moved as a function of the measured torque to imitate natural

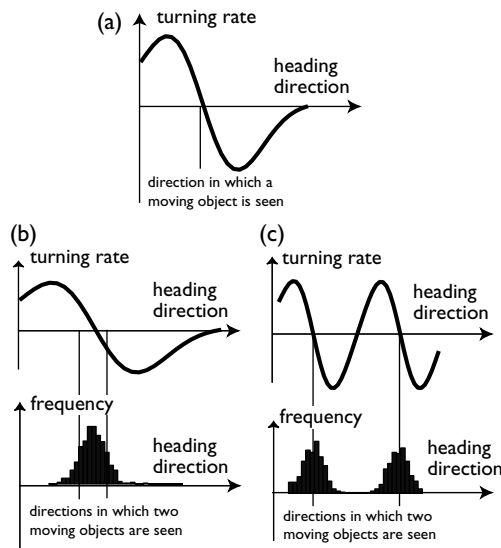


Figure 4.1. Schematic illustration of how the dynamics of heading direction accounts for how flies select and orient to visual targets. (a) The torque exerted by the fly's flight motor generates a turning rate, which is a function of the fly's heading direction relative to a visual target, an object that makes small oscillatory movements around the indicated position. This dependence effectively defines a dynamics of heading direction, which has an attractor (zero-crossing with negative slope) at the direction in which the target lies. (b) When two objects are metrically close (see the two vertical lines) they generate a fused attractor. Over time, the fly's average heading lies between the two targets. (c) When two objects are metrically far (vertical lines), a bistable dynamics results. The fly switches randomly between the two heading directions, generating a bimodal histogram (bottom) of flying directions over time.

flight. When imperfections of this early form of virtual reality were removed, researchers found that the prediction did not hold up (Heisenberg & Wolf, 1988). Apparently, the fly's simple nervous system computes an expected visual motion from its own motor commands (reafference) and treats detected visual motion matching that prediction differently than extraneous motion signals related to movement of an object relative to the fly. So even this very simple control system provides hints that uncovering the dynamics from which behavior emerges re-

quires more than an input-output analysis in open loop.

Even so, there is no explicit representation of the speck in motion, nor reasoning about goals and plans to reach those goals. The behavior emerges when a neural system linked to the sensory and motor surfaces is immersed in an environment to which it is adapted. The complete system, including neural network and coupling through the body and environment, can be understood as a dynamical system. Its attractor solution is the emergent behavior. Although in some abstract sense one could say that the neurons controlling the flight motor "estimate" the direction in which the target lies, their firing does not "re-present" this value because the estimate is implicit in the control circuitry driving the flight system and cannot be forwarded to be used in any other context.

Some have argued that such emergence of a behavior from a closed control could form the core of a potential dynamical systems refoundation of cognition (Gelder & Port, 1995). But is control already a form of cognition? Would that not imply that every heater with a thermostat is already a cognitive system? One dimension along which systems maybe be distinguished is flexibility. One could argue that the threshold of cognition has not been reached as long as an emergent behavior is uniquely determined by sensory inputs. By contrast, when a control system makes decisions, selects among a range of inputs, and generates behavior based on its own inner state, then this might represent the most elementary form of cognition. That implies a form of flexibility, in which the emergent outcome is no longer dictated by the sensed outer world but is instead, at least to some minimal extent, based on choices generated from within the system.

The flight control system of house flies is capable of such flexibility. When confronted with several patches of visual motion on its facet eye, the fly selects one of the patches and directs its flight in pursuit of that visual object (Poggio & Reichardt, 1976). This capacity to select emerges from the control dynamics. Superposing the torque patterns

generated by each motion patch, two attractors emerge (Fig. 4.1c). One of these is selected depending on the initial orientation and on chance. When the two patches are close to each other, the two attractors merge, and the fly flies in an averaged direction (as shown in Fig. 4.1b).

This example illustrates ideas that generalize well beyond the framework of control theory. A summary of the principles of a dynamic approach to behavioral patterns can be formulated as follows (Schöner & Kelso, 1988a). (1) Patterns of behavior are characterized by inner states, which determine the persistence of the patterns over time and under changing conditions. A state can be characterized by variables with specific values corresponding to specific states. States and associated variables are not limited to those of sensorimotor loops. State variables may originate, for instance, from within the neural networks that control behavior. (2) The evolution in time of these state variables is generated by neural networks linked to sensory and motor surfaces that can be modeled as dynamical systems. Many factors may contribute to the effective dynamics of such systems, including the physical dynamics and material properties of the body and of the environment. Sensory inputs, including internal sensory feedback, also act as forces on this dynamics. (3) Asymptotically stable states structure the solutions of this dynamical system. Over the long run, only attractor solutions are robust and likely to be observed. The nervous system is extensively interconnected, so that for any particular circuit and any particular pattern, other connections act effectively as perturbations, as do variable sensory inputs and the complex and temporally variable natural environment. (4) As a corollary, only when states are released from stability does behavioral flexibility arise. Release from stability takes the form of instabilities (bifurcations) in which the restoring forces around an attractor become too weak to resist change. New solutions may be reached or even created from instabilities. The full complexity of behavior may ultimately be generated from the complex structure of stable dynamical

states and their instabilities in a nonlinear, strongly interlinked dynamical system.

As an illustration of these principles, consider the coordination of rhythmic movement, a domain in which dynamical systems ideas have been developed and evaluated in detail (Schöner & Kelso, 1988a; Kelso, 1995). Patterns of coordination underlying such behaviors as the gaits of locomotion, speech articulatory movements, or the playing of a musical instrument can be characterized through measures of the relative timing of components, such as the relative phase, ϕ . Their evolution reflects the coupling between the neural networks that control the components (Grossberg, Pribe, & Cohen, 1997) as well as, in some instances, mechanical coupling (Turvey, 1990). The temporal evolution and stability of the coordination patterns can be described by an effective dynamical system governing the measures of relative timing, which can be modeled as a relative phase dynamics (Schöner, Haken, & Kelso, 1986). Stable states (attractors) of the dynamics correspond to stable patterns of coordination. The coordination of homologous limbs, for instance, occurs generally in two symmetric patterns, the limbs either moving in-phase or in phase alternation ("anti-phase"). These patterns stay invariant under a variety of conditions, including changes in the frequency of the rhythmic movement. Their stability does not stay invariant, however. The anti-phase pattern of coordination typically becomes less stable at higher movement frequencies. This manifests itself in associated changes of stability measures, such as an increase in the amount of fluctuation of relative phase and an increase in the amount of time needed to recover from a perturbation of the coordination pattern (Schöner et al., 1986). Stability is thus an essential property of coordination patterns. Without stability, patterns do not persist. In fact, at sufficiently high movement frequencies, an involuntary switch out of the anti-phase into an in-phase pattern of coordination occurs. An understanding of coordination thus requires more than an account of the information processing needed to compute

the pattern. It also requires an account for how the pattern is reached from all kinds of perturbed states. In dynamical systems thinking, both the specification of the state and the mechanisms for its stabilization emerge from the same underlying dynamical system.

Are stability properties perhaps generated at a lower level of control, whereas relative timing per se is planned at a more abstract, disembodied level? Information processing models of timing have invoked “clocks” that generate time signals representing more or less complex patterns of coordination, which are then handed down to a “motor” system that handles the control (Vorberg & Wing, 1996). A first response is that the clocks themselves must have stability properties if they are to account for coordination, and this makes them dynamical systems as well (Schöner, 2002). Abstracting from the dynamic, embodied properties of timing means, however, missing out on important constraints for higher cognitive function. How people switch intentionally from one pattern of coordination to another, for instance, is constrained by stability (Scholz, Kelso, & Schöner, 1988). First, switching from a more stable to a less stable pattern takes longer than vice versa. Second, the intention to switch to a pattern increases that pattern’s stability so that it is possible to switch to a pattern that is unstable under less specific intentional constraints. The experimental results were quantitatively accounted for by a model in which the intention to switch to a coordination pattern is a force in the coordination dynamics that increases the stability of the intended pattern (Schöner & Kelso, 1988b). Another study had participants learn a new pattern of coordination that initially was not in the behavioral repertoire (Schöner, Zanone, & Kelso, 1992). The process of learning amounted to increasing the stability of the target pattern. That the underlying coordination dynamics was changed could be shown when participants were asked at different points during their training to produce patterns of coordination near the target pattern. Before learning performance was biased toward the in-

trinsic patterns of in-phase and anti-phase. After learning a new bias toward the learned pattern was observable. The bias could be attributed to the underlying dynamics of relative timing, which changed during learning, with new forces stabilizing the learned pattern (Schöner et al., 1992). Related work established that perceived or memorized constraints for relative timing could likewise be understood as contributions to the dynamics of coordination (Schöner & Kelso, 1988c).

The picture that emerges from this exemplary system is that movement coordination emerges as stable states from a nonlinear, potentially multistable dynamics, realized by neural networks coupled to the body in a structured environment. Cognitive aspects of motor control, such as intentional movement goals, motor memory, or skill learning, are all mediated through this dynamics of coordination. Its graded change may lead to both graded and categorical change of movement behavior.

Beyond motor control, nonlinear dynamics has been invoked as a general framework for cognition, in which the concept of representation is unneeded (Gelder & Port, 1995). This has been viewed as a strength by some, as a weakness by others (Markman & Dietrich, 2000). Extending dynamical systems ideas beyond motor control, we run into a conceptual limitation, however. Take the coordination of rhythmic movement we just discussed, for example. What value does relative phase have when the movement is stopped? When movement is restarted, does the coordination system start up at the last value that relative phase had? Describing the state of a motor system by a variable such as relative phase requires that variable to have a unique value at all times. That value must evolve continuously in time, cannot jump, cannot split into two values, or disappear and have no value. The dynamical system description of coordination by relative phase is thus “uniquely instantiated.”

For another example, consider the biomechanics of the human arm, which has a well-defined physical state at all times,

characterized by the spatial positions of its segments or the ensemble of its joint angles. That physical state changes only continuously, that is, the arm obviously does not disappear in one position and reappear at another. The biomechanics of the arm are characterized by the equations of motion, a set of differential equations for the joint angles that generate continuous joint angle trajectories. The variables in this dynamical description are uniquely instantiated. Every joint angle has exactly one value at each time, and that value changes continuously in time.

Now think about planning a movement. A movement plan may exist before a movement is initiated. This is revealed by the movement starting out in the direction of the target, by the latency between a movement command and the initiation of the movement being shorter when the movement goal is known ahead of time, or by observing specific neuronal activity prior to movement initiation. A movement plan may be described in the same terms as the state of the arm, for example, as a desired configuration of the joint angles. But are those variables uniquely instantiated? After having made a movement, is the movement plan still around? When a new movement is prepared, do the joint angle variables containing the planned arm configuration evolve continuously from the values of the previous movement plan to the required new values? Clearly, that does not make sense. In a first approximation, the preparation of a new movement does not depend on the previous motor act. Also, early during movement planning, movement parameters may have multiple values (Ghez et al., 1997; Wilimzig, Schneider, & Schöner, 2006).

Is it possible that the planning of movements does not fall into the domain of dynamical systems thinking? The answer is no, because there are clear indications of dynamics at the level of movement preparation (Erlhagen & Schöner, 2002). Movement plans evolve continuously in time (Ghez et al., 1997) and are updated at any time during movement preparation when sensory information changes (Goodale, Pélisson, &

Prablanc, 1986). The neural substrate reflects both neuronal properties (Georgopoulos et al., 1989; Scott, 2004).

Similarly, perception has signatures both of dynamics and of a lack of unique instantiation. That percepts cannot be described by uniquely instantiated variables is intuitive. When we watch a slide show, each slide induces a new percept. It does not seem to make sense to say that the new percept induced by the next slide emerges from the percept of what was on the previous slide by continuous transformation. Evidence for a new percept, depending on what has just previously been perceived, comes, however, from multistable stimuli. The motion quartet (Hock, Kelso, & Schöner, 1993) is a particularly clear example, illustrated in Figure 4.2. Spots at the four corners of an imaginary rectangle have luminance levels above background. Two spots lying on one diagonal are much brighter, the two spots on the other diagonal are only a little bit brighter than background. If the two diagonals switch from frame to frame, then one of two motion patterns is clearly seen: Either the bright spots move horizontally (panel c) or vertically (panel d), but never both at the same time. Different stimulus geometries favor either perceptual outcome: Flat rectangles (panel e) generate predominantly vertical motion percepts, tall rectangles (panel f) generate predominantly horizontal motion percepts. When the stimulus geometry is changed continuously in time, the perceptual state tends to persist, leading to perceptual hysteresis (panel g). This is a very robust finding, immune to intentional or semantic influences and to eye movements, and not caused by response bias (Hock et al., 2005). Hysteresis is evidence for continuity of the underlying perceptual state and supports a dynamical systems account for perception (Hock, Schöner, & Giese, 2003).

When, on the other hand, the perception of motion is briefly interrupted while the stimulus series is presented, hysteresis is abolished (Hock & Ploeger, 2006; Nichols, Hock, & Schöner, 2006). A sophisticated way of demonstrating that fact employs the background relative luminance contrast

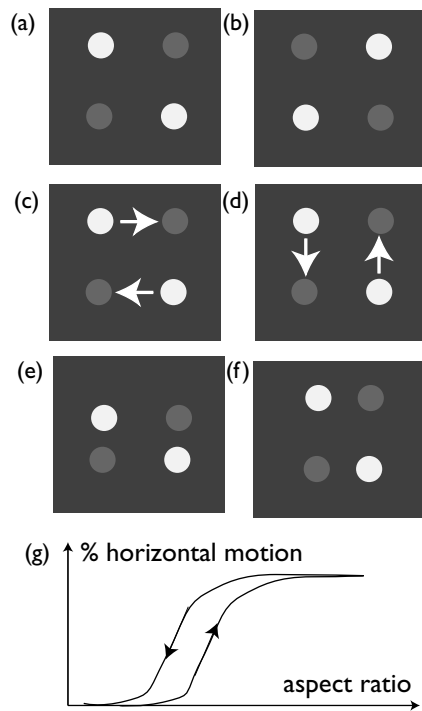


Figure 4.2. The two frames (a and b) of the motion quartet stimulus and its two possible percepts, horizontal (c) and vertical (d) motion, indicated by white arrows. Motion is seen from locations at which luminance decreases to locations at which luminance increases. If the stimulus is periodically repeated, the direction of motion is inverted on subsequent frame changes, but the pattern of horizontal versus vertical tends to persist. Low aspect ratios (e) favor vertical, high aspect ratios (f) favor horizontal motion. When the aspect ratio is continuously varied from frame to frame (g), the initially established motion direction tends to persist, leading to hysteresis, the dependence of perceptual outcome on the direction of stimulus change (arrows).

(BRLC; see Figure 4.3). This is the amount of change in luminance from frame to frame in relation to how much the spots' average luminance is above background. A BRLC of 2 (panel a) provides the strongest motion signal. This is when the change of luminance goes all the way down to background luminance. BRLCs below two have considerable contrast in both frames, but more contrast in one frame than in the other. Near a BRLC of zero (panel b), there is hardly any lumi-

nance change between frames. When BRLC is varied, the probability of motion being perceived varies between these limit cases (panel c); Hock, Gilroy, & Harnett, 2002). (This transition from nonmotion to motion perception is likewise hysteretic, a fact we shall return to later.) At intermediate levels of BRLC, motion perception is bistable: Sometimes, motion will be perceived, sometimes not. Thus, if during a stimulus series that changes the geometry from, say, flat to tall, the BRLC level is abruptly lowered for just a single frame, then on a certain percentage of trials, the perception of motion will stop altogether, whereas on the other trials, motion will continue to be perceived through the perturbation. Hock and Ploeger (2006) found that hysteresis was abolished on those trials, on which motion had stopped, but not on those on which motion had continued. Thus, whether or not the direction of motion is preserved over

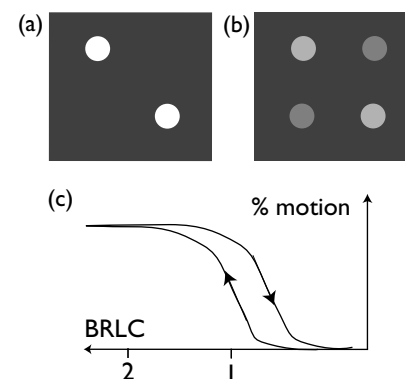


Figure 4.3. The background relative luminance contrast (BRLC) is the amount luminance changes from frame to frame divided by the distance between average luminance and background. (a) A BRLC of two means luminance changes between a high level and background. (b) A small BRLC means luminance changes little from frame to frame, so that all locations in the motion quartet have similar contrast relative to background in both frames (c). The probability of perceiving motion increases with increasing BRLC. The transition between motion and nonmotion shows hysteresis, that is, depends on the direction of BRLC change (arrows).




Figure 4.4. Activation fields are defined over the metric space, x , relative to which information is being represented, for example, space, feature, or movement parameters. Activation, u , itself encodes the amount of information about specific values, for example, the readiness to activate a specified action or the certainty of a sensory estimate.

time depends on whether or not motion remained active.

The continuity captured by the dynamical variables of uniquely instantiated descriptions is thus not sufficient to move from motor control toward movement planning or to perception. What is missing is a dynamic representation of the presence or absence of plans or percepts. The classical concept of activation can play that role (Spencer & Schöner, 2003).

4. Dynamical Field Theory

4.1. *Activation Fields*

To represent metric information in terms of dynamical state variables, we need two dimensions (Figure 4.4). One is the metric dimension along which information is specified. Each location in the field corresponds to a particular value along the metric dimension. For each such value, the activation level of the field spans the second dimension that encodes a measure of the amount of information about that value.

For a sensory representation, the first dimension may entail parameters whose values must be estimated from sensory information. Such parameters may be, for instance, the location in visual space of a stimulus, perceptual features such as orientation, spatial frequency, and pitch, or even relatively high-level parameters, such as the

estimation of the pose of a visual object. The body scheme is another example, which may involve estimating joint angles from proprioceptive information. For such sensory representations, low levels of activation at a particular field location indicate that the value of the represented dimension associated with that location is not a likely estimate. High levels of activation mark field locations that contribute substantially to the current estimate of sensory information.

For motor representations, metric dimensions may be spanned by movement parameters like movement direction or movement extent, level of force, or movement time. Low levels of activation at a field location indicate that the associated movement is not a likely movement plan. High levels of activation indicate that the movement represented at that location is close to being initiated, and activation from that field site will be handed down to the motor control system.

In this picture, localized peaks of activation are units of representation. The location of the peaks in the field encodes metric information about the underlying dimension. The activation level of the peaks is the strength of representation, which may variably encode the certainty of an estimate, the closeness to execution of a plan, or the physical intensity of a stimulus. A flat distribution of activation, by contrast, represents the absence of specific information.

The limit case of uniquely instantiated dynamics is modeled whenever a single positive peak moves continuously along the metric dimension. Its motion may be described by an instantiated dynamics in which the peak location is the dynamical state variable, whose time course is generated through an ordinary dynamical system. By contrast, in the more general conception of a dynamic field, it is the activation level at each field location that acts as state variable. Thus, dynamic fields are infinite dimensional dynamical systems, and activation levels rather than peak locations evolve continuously in time. Peaks may be suppressed and created. For instance, harking back to the problem of coordination, a peak over the dimension “relative phase” would indicate that relative

phase has a well-defined value. If the movement is stopped, that peak decays. When movement resumes, a peak could be generated at a new location, so that relative phase could start up at a new value. Similarly, peaks of activation in a dynamic field defined over the direction and location of perceived motion signify the perception of a particular movement pattern. When the peaks decay (e.g., because BRLC was lowered), the motion percept is lost. When the stimulus is restored in strength, peaks may come up at new locations, restoring the percept of motion but potentially in a new direction.

4.2. Field Dynamics

The dynamical system from which the temporal evolution of activation fields is generated is constrained by the postulate that localized peaks of activation are stable objects, or, in mathematical terms, fixed point attractors. Such a field dynamics has the generic form

$$\begin{aligned} \tau \dot{u}(x, t) = & -u(x, t) + \text{resting level} \\ & + \text{input} + \text{interaction} \end{aligned} \quad (4.1)$$

where $u(x, t)$ is the activation field defined over the metric dimension, x , and time, t . The first three terms define an input-driven regime, in which attractor solutions have the form $u(x, t) = \text{resting level} + \text{input}$. The rate of relaxation is determined by the time scale parameter, τ . The interaction stabilizes localized peaks of activation against decay by local excitatory interaction and against diffusion by global inhibitory interaction (Figure 4.5). In Amari's formulation (Amari, 1977) the mathematical form is specified as

$$\begin{aligned} \tau \dot{u}(x, t) = & -u(x, t) + h + S(x, t) \\ & + \int dx' w(x - x') \sigma(u(x', t)). \end{aligned} \quad (4.2)$$

Here, $h < 0$ is a constant resting level, $S(x, t)$ is spatially and temporally variable input function, $w(\Delta x)$ is an interaction ker-




Figure 4.5. Local excitatory interaction helps sustain localized peaks of activation, whereas long-range inhibitory interaction prevents diffusion of peaks and stabilizes against competing inputs.

nel, and $\sigma(u)$ is a sigmoidal nonlinear threshold function (Figure 4.6). The interaction term collects input from all those field sites, x' , at which activation is sufficiently large. The interaction kernel determines if inputs from those sites are positive, driving up activation (excitatory), or negative, driving down activation (inhibitory). Excitatory input from nearby location and inhibitory




Figure 4.6. The interaction kernel, $w(\Delta x)$, in Amari's neural field dynamics depends only on the distance, Δx , between the field locations as illustrated on top. The kernel depicted here is excitatory only over small distances, whereas over larger distances, inhibitory interaction dominates. Only sufficiently activated field sites contribute to interaction. This is modeled by sigmoidal threshold functions, such as the one illustrated on bottom, $\sigma(u) = 1/(1 + \exp(-\beta u))$.




Figure 4.7. The dynamic activation field (solid line) in response to input distributions (dashed line). (a) Localized input is applied to the activation field that is initially at the negative resting level. This induces a subthreshold peak, which does not yet engage interaction. (b) When input is slightly increased, excitatory interaction pulls the activation peak up, which in turn inhibits the field elsewhere.

input from all field locations generically stabilize localized peaks of activation. For this class of dynamics, detailed analytical results provide a framework for the inverse dynamics task facing the modeler, determining a dynamical system that has the appropriate attractor solutions.

A literature on neuronally more realistic or detailed models includes multilayer field dynamics (in which excitation and inhibition are separated, as in real neuronal systems; Wilson & Cowan, 1972) and models of spiking neurons (Gerstner & Kistler, 2002). The qualitative dynamics of the generic Amari formulation are shared features of this entire class of neuronal dynamics, however. In particular, two functionally relevant kinds of attractor solutions arise. The input-driven attractor is a largely subthreshold pattern of activation in which the contribution of the neuronal interaction is negligible (Figure 4.7a). Self-excited attractors, by contrast, are localized patterns of activation with levels sufficient to engage neuronal interaction (Figure 4.7b). In this state, local excitatory interaction lifts activation within the peak beyond levels induced by input, whereas global inhibitory interaction suppresses levels elsewhere below the levels justified by the resting level or inputs.

That these two kinds of attractors are qualitatively different states can be seen from the fact that there is a dynamical instability separating them, the detection insta-

bility (see Bicho, Mallet, & Schöner, 2000, for discussion). This instability can be observed, for instance, if the amplitude of a single localized input is increased. Below a critical point, this leads to a subthreshold input-driven solution (Figure 4.7a). When input strength reaches a threshold, this solution becomes unstable and disappears. The system relaxes to a peaked solution, which coexists bistably with the (input-driven solution). As a result, the detection decision is stabilized against small changes of input: When the input level drops again, the peak is sustained within a range of input strengths. This leads to hysteresis, that is, dependence of the observed state on the direction of change.

Next to detection decisions, selection among multiple inputs is another elementary form of cognition. This function emerges from a second instability, illustrated in Figure 4.8. When inputs are sufficiently strong and metrically close, the detection instability leads to a peak positioned over an averaged location. For broad input distributions, averaging may occur in the input stream, although the excitatory interaction may bring about averaging even when input is bimodal (as in Figure 4.8a). When the metric distance between inputs is larger, however, the dynamic field is bistable, instead. A single peak emerges from the detection decision, localized either at one or at




Figure 4.8. (a) The dynamic activation field (solid line) may generate a peak at an averaged position in response to bimodal input (dashed line) when the input peaks are within the range of excitatory interaction. (b) At larger distance, inhibitory interaction dominates, and the field dynamics becomes bistable: Either a peak positioned over one input mode may be generated (black solid line) or a peak positioned over the other (gray solid line), but not both at the same time.




Figure 4.9. In the left column, a self-stabilized activation peak generated in response to a localized input is only stable while that input is present (a). When the input is removed (c), the peak decays to a subthreshold distribution. In the right column, the input-induced peak (b) is sustained when input is removed (d). The difference between the two columns is the slightly larger resting level (dashed horizontal lines) on the right.

the other location of input. Asymmetries in input, fluctuations, or prior activation history may favor one over the other state, but competition prevents simultaneous activation at both locations.

When activation levels are sufficiently high in the field, many locations have activation levels above threshold and contribute to interaction. This may enable the dynamic field to sustain a peak of activation, even when the localized input that first generated the peak (through a detection instability) is removed (Figure 4.9). Sustained activation of localized peaks of activation provides a neural mechanism for metric working memory (Schutte, Spencer, & Schöner, 2003). Metric information about past stimulation is thus preserved over much longer time scales than the dynamic time scale of individual neurons or field sites (Amit, 1994).

This is true, of course, only to the extent to which there are not other localized inputs that would attract sustained peaks. Such inputs may be small enough to not be able to push the dynamic field through the

detection instability. Even so, they may induce drift of sustained peaks (Schutte et al., 2003; Spencer et al., 2007). Relatedly, small inputs may be sufficient to trap peaks that are induced by broad or even homogenous inputs to the field, which push the field broadly through a detection instability. This may lead to categorical representations of metrically graded inputs (Spencer, Simmering, & Schutte, 2006).

Instabilities may amplify small differences. A peak of activation can be induced, for instance, by providing a homogeneous boost to the field. Where the peak comes up then depends on small subthreshold patterns of localized input or any other inhomogeneities in the field. Such inhomogeneities may arise because input connections have slightly different strengths as a result of Hebbian strengthening of those inputs that have successfully induced peaks. Another simple form of learning is the laying down of a memory trace at those field sites at which peaks have been generated. In either case, the history of activation may influence ongoing decision making. This is illustrated in the following section by showing how habits may be formed and how the behavioral history biases decisions.

4.3. Behavioral Signatures of Dynamic Fields

How may dynamic fields, their instabilities, and their functional modes help understand the emergence of elementary forms of cognition? We will answer this question in the context of a few exemplary case studies and show, at the same time, how behavioral signatures of the neural field dynamics may provide evidence for the Dynamical Field Theory (DFT) account of cognition.

Most psychophysics makes use of detection decisions in one form or another. Are these decisions related to the detection instability of DFT (Figure 4.7)? Hock, Kogan, and Espinoza (1997) observed that a psychophysical detection decision was self-stabilizing when the perceptual alternative to a detection was perceptually clear. They asked participants to indicate whether they

saw apparent motion or flicker between two spots of changing luminance. The parameter they varied was the BRLC discussed earlier (Figure 4.3). Changing BRLC continuously in time led to perceptual hysteresis, the initially established percept persisting into the regime, in which both motion and nonmotion percepts were possible (illustrated schematically in Figure 4.3c). The authors argued that perception is always based on the activation of ensembles of neurons. Within such ensembles, interaction supports self-stabilization of patterns of activation, and this could account for the observed stabilization of the detection decision in the manner of the detection instability (Figure 4.7). In a follow-up study, Hock et al. (2004) exposed participants to motion stimuli of constant BRLC within the bistable region for a variable amount of time. They asked participants to report when percepts switched from motion to nonmotion and vice versa. The resultant switching rates revealed that the switching probability increased over time both for switches to motion as well as for switches to nonmotion perception. This would typically be interpreted in terms of selective adaptation, indicating that both the motion and the nonmotion percepts are embodied by neural populations that build up inhibition while they are activated. (See Hock et al., 2003, for how adaptation relates to dynamical systems ideas in perception.)

Thus, even though the mean behavior may be perfectly well described in the classical threshold manner, psychophysical detection decisions may involve the stabilization of decisions through a bistable regime around threshold. The decisions underlying selection are less commonly studied, in part, because tasks in which the stimulus does not uniquely specify the required choice tend to be fragile experimentally (e.g., by being subject to cognitive penetration, response strategies, and response bias). The gaze control system frequently performs selection decisions and is relatively immune to these problems. Our gaze is reoriented to new visual targets at a high rate of approximately three times a second. Select-

ing the next visual target is thus one of the most basic sensorimotor decision problems solved by the human central nervous systems. Empirically, a transition is observed from averaging for visual targets that can be fovealized simultaneously to selection when targets are metrically too far from each other for that to be possible (Ottes, van, Gisbergen, & Eggermont, 1984). DFT has provided an account for this transition that captures a range of experimental details (Kopocz & Schöner, 1995; Trappenberg et al., 2001). The most recent addition to that range is an account for the time course of selection, with fast saccades tending more toward averaging than slow saccades, because the competitive inhibition required for selection needs more time to become effective (Wilimzig et al., 2006).

Development offers another, wonderful window into the study of selection. Infants are not at risk of adopting dodgy cognitive schemes when confronted with a range of choices and no stimulus that disambiguates the selection. Instead, they select movement targets reliably, such as in the classical paradigm of Jean Piaget (Piaget, 1954) in which two locations, A and B, are perceptually marked on a box. In the classical version, a toy is hidden at the A location, covered by a lid, and after a delay, the box is pushed toward the infant, who reaches for the A lid and may also retrieve the toy. After four to six such A trials, the toy is hidden at the B location. If a delay continues to be imposed, young infants below about 12 months of age are likely to make the A-not-B error, that is, they persevere and reach to the A location rather than the cued B location. Older infants do not make the perseverative error, nor do young infants when the delay is short. Smith et al. (1999) have demonstrated that a toyless version works just as well: The cue consists of waving the A or the B lid and attracting the infant's attention to the corresponding location. Thus, sensorimotor decision making is a critical component of this paradigm.

In the dynamic field account of perseverative reaching (Thelen et al., 2001; Schöner & Dineva, 2006; Dineva & Schöner, 2007), an




Figure 4.10. A Dynamical Field Theory model of perseverative reaching in Piaget's A not B paradigm represents the planned movement by an activation field defined over a movement parameter. A peak over the A location (here at the left-most lid) represents a reach to A. The dynamic activation field evolves under the influence of specific input (attracting attention to the A location by waving the lid), task input (the two visibly distinct lids), boost (the box with lids is pushed toward the infant), and the memory trace left by earlier reaches to A.

activation field represents the range of possible reaching movements, described parametrically, for instance, by the direction of end-effector movement in space (Figure 4.10). A suprathreshold activation peak centered over the direction in which either of the two lids are located represents an intention to move the hand in the corresponding direction. Various forms of input drive the generation of such a peak. The two locations are perceptually marked by the lids, so both associated directions receive permanent "task" input representing the layout of the reaching space. The cuing action that attracts the infant's attention to one of the two locations generates transient "specific" input to the cued location only. Finally, when the box is pushed into the reaching space of the infant, all reaching locations on the box receive a homogeneous "boost" input. On later trials, an accumulated "memory trace" of previous activation patterns also acts as input, preactivating the locations to which earlier reaches were directed.

The mathematical model that formalizes this account is reviewed in the Appendix. Figure 4.11 illustrates the temporal evolu-

tion of the field and the memory trace over the course of an "A not B" experiment. On the initial A trials, the specific input generates some activation at the cued A location, which decays again during the delay. When the boost is applied after the delay, this pushes the field through the detection instability, generating a peak at the A location. The small remnant activation advantage of A over B left over from the specific input is sufficient to bias the decision in favor of A. The peak at A signifies a reach to A and a memory trace is laid down at that location. This memory trace preactivates the field near A on subsequent trials, further biasing the selection toward the A location. The memory trace thus represents the motor habit formed during A trials.

The memory trace is sufficiently strong to tip the balance on the first B trial in favor of the A location. In that trial, the specific cue provided input to the B location, but the remnant activation after the delay is not strong enough to overcome the bias to A induced by the memory trace. Generally, when the delay is sufficiently long and when sufficiently many reaches to A have




Figure 4.11. Temporal evolution of activation in the reaching field (top) and the associated memory trace (bottom) during the time course of an A not B experiment performed on the Dynamical Field Theory model. Six A trials are followed by two B trials. In each trial, a specific cue is presented at A or B, inducing activation at the associated locations (small bumps in the middle of each trial in the top graph). During the ensuing delay, these activations decay. When the boost arrives (ridges along the ϕ -axis in top graph, a peak is induced at the location with most remaining activation. This peak generates a memory trace (bottom), which biases peak generation on B trials. In these simulations, the model perseverates by generating peaks at the A location on B trials.

built enough of a memory trace at A, the bias to A generates a perseverative reach.

Figure 4.12 shows how the behavioral history in an individual run of the experiment matters. In this instance, a fluctuation leads to the generation of a peak at the B location during an A trial. Such a reach to B on an A trial is called a spontaneous error. It leaves a memory trace at B and weakens the trace at A, predicting increased probability of a spontaneous error being observed again and, in particular, a reduced probability of perseverative reaching. Both are true in the experiment (Dineva & Schöner, 2007).

Why do older infants reach correctly? In the DFT account, this is due to a subtle shift in dynamic regime. Older infants are closer to the cooperative regime, that is, the regime

in which activation peaks are sustained in the absence of localized input (right column in Figure 4.9). This may arise because their overall level of activation is higher or because of characteristic changes in neuronal interaction (Schutte et al., 2003). They are thus capable of generating a sustained peak of activation when the specific cue is given and thus stabilize the reaching decision toward B during the delay against competing input from the memory trace at A.

In fact, whether or not people are capable of stabilizing decisions against competing influences depends on a variety of contextual factors. Toddlers and even 4-year-olds display perseverative errors, for instance, when toys are hidden in a sandbox (Spencer, Smith, & Thelen, 2001). After the toy is hidden, the sand is smoothed over




Figure 4.12. Another simulation run for the Dynamical Field Theory model of perseverative reaching. In this simulation, a fluctuation in activation leads to a spontaneous error on the fifth A trial: The field generates a peak at the B location rather than at the A location. This leads to a memory trace being laid down at B and the memory trace at A being weakened. As a result, a second spontaneous error occurs, and the model responds correctly on the B trials.

and no perceptual marker of the reaching location is left. On B trials after the delay, these children search for the toy in the sand at a location that is metrically attracted to the A location by as much as half the distance between A and B. In DFT, this metric bias comes from sustained peaks drifting during the delay, attracted by the memory trace input laid down during the A trials.

John Spencer and colleagues have extended this picture into the domain of spatial working memory (Spencer & Hund, 2002; Hund & Spencer, 2003; Spencer & Hund, 2003). They had children of various ages and adults point to a location on a featureless surface, at which variable amounts of time earlier a marker had been displayed. By varying the delay between presentation of the location and its probe, these researchers have been able to directly observe the drift of metric spatial memory. Drift occurs in the direction of locations previously held in spatial working memory and

away from any visual landmarks. Older children and adults drift less and more slowly. Here is a set of behavioral data that uncovers direct behavior signatures of the underlying neural picture of self-stabilized neural activation patterns storing metric information (Schutte et al., 2003)!

5. Discussion

5.1. *Is the Dynamical Systems Approach Embodied and Situated?*

Given the abstract mathematics behind the dynamical systems approach, it is legitimate to ask whether the approach does, in fact, embrace the theoretical stance of embodiment and situatedness as announced. Does the dynamical systems approach take seriously the link of cognitive processes to sensory and motor processes? Does it take into account the embedding of cognition within structured environments as well as within

the behavioral history of an organism? The more general issue is *autonomy*, that is, the continuous flow of experience under both the action of the individual and the time-varying sensory information about a changing environment. Cognition in naturalistic settings is based on such autonomy. We humans move seamlessly from decision to decision, generating action on our own time frame, only exceptionally reacting to stimuli, more typically actively controlling sensory flow. We update action plans as new sensory information arises. When we are interrupted, the last state of our cognitive processes continues to be available when we resume an action. Thus, cognitive processing is never really off-line.

Dynamical systems thinking is all about autonomy. The theoretical concept of stability, at the core of dynamical systems thinking, is key to understanding autonomy. Only if stability is warranted may cognitive processes be linked to on-line sensory input. Linking cognitive processes to motor systems that act in the real world requires control-theoretic stability. This includes the need to stabilize decisions against fluctuating sensory input and internal neuronal interaction long enough to physically act out what was decided. Conversely, in a system that is open to sensory input and that stabilizes decisions, the release from stability is required to bring about behavioral change. In fact, instabilities are crucial for understanding how the discrete behavioral events may emerge that structure the continuous flow of experience.

By applying dynamical systems concepts at multiple time scales, it is possible to understand how behavioral history as well as the history of stimulation matter. The accumulation of behavioral history is the basis of learning and adaptation. Individual differences may be preserved over time as differences in experience condition further differences in behavior.

Even so, one may ask whether the link of cognitive processes to the sensory and motor surfaces is really always required. Are classical information-processing accounts and

many connectionist models not legitimately simplifying analysis by neglecting those links? Such simplification may, however, hide serious problems in the interface between the abstract information processing model and the sensory and motor processes through which cognition is realized. When, for instance, input units in a neural network are assumed to encode objects or symbols, this hides nontrivial processing, including segmentation, classification, and estimation of object parameters. Similarly, if a neuron encodes a motor output, this may hide the nontrivial processes, of control in realtime, including reactions to unforeseen perturbations of the effector system.

One qualitative form of demonstration and exploration of such issues hidden in the interfaces of cognitive models with the real world is to implement the models on physical robotic platforms. A robotic implementation requires complete specification of the entire path from the sensory surface to the cognitive model as well as on to the motor control system. Robotic function may require particular environmental conditions (e.g., uncluttered perceptual scenes), which expose hidden assumptions about simplifications and abstractions that may or may not limit the explanatory power of the cognitive model.

Dynamical systems thinking has been tested extensively with robotic implementations. In fact, an entire approach to robotic behavior has been developed based on dynamical systems thinking (Schöner, Dose, & Engels, 1995). Implementations have sometimes used very simple sensory systems and simple motor control strategies that did not include a detailed model of the plant (Bicho & Schöner, 1997). On the other hand, more complex systems with more highly developed cognitive processes have also been demonstrated (Bergener et al., 1999). The conceptual framework includes DFT employed to endow robots with representations (Engels & Schöner, 1995; Bicho et al., 2000; Erlhagen & Bicho, 2006; Faubel & Schöner, 2006). By generating complex behavioral sequences from attractor dynamics

that undergo instabilities, these implementations have demonstrated how autonomy emerges in dynamical systems architectures (Steinhage & Schöner, 1998; Althaus & Christensen, 2003).

Is dynamical systems thinking primarily limited to the sensorimotor domain? Historically, dynamical system thinking arose from work on movement coordination, and this review reflected that history. This history sets dynamical systems thinking apart from both the classical information processing approaches and connectionism and is reflected in the relationship between dynamical systems thinking and the embodied and situated conception of cognition. Much of the recent development of the dynamical systems approach, however, moves well beyond the sensorimotor domain. This newer work was only briefly touched, referring to accounts for working memory, category formation, and object recognition, for instance. DFT has been critical to the extension of dynamical systems thinking beyond the sensorimotor domain. In fact, through DFT, the concept of representation could be integrated into dynamical systems thinking (Spencer & Schöner, 2003). The DFT framework gives the concept of representation a very concrete, embodied sense. For instance, self-stabilized peaks induced by sensory information “re-present” that sensory information even when it is no longer available at the sensory surface. Similarly, self-stabilized peaks are “re-presentations” of motor decisions when they stabilize these decisions while they are not (yet) acted out. The strongly interacting neuronal dynamics of DFT are thus capable of achieving the two qualities that define representations: First, they enable the coupling to the sensory and motor surfaces through which representations establish and maintain links to the outside world. Second, they stabilize these representations, which thus continue to be effective even when no longer linked to sensory or motor systems. This is a sense of representation close to that invoked by the neuroscientist Joaquín Fuster (2005) as a universal feature of the cerebral cortex.

5.2. *Is the Dynamical Systems Approach Neurally Based?*

The second component of the embodied stance requires accounts to be consistent with principles of neural function. Do the abstract mathematical concepts of the dynamical systems approach live up to this requirement?

Biophysically, neurons really are, of course, little dynamical systems (Wilson, 1999). The historical origin of the concept of dynamic neural fields comes from biophysical models of cortical activation dynamics (Wilson & Cowan, 1972, 1973; Amari, 1977). The insight that cognitive function is best described in terms of neural dynamics is probably due to Grossberg (1980).

On this basis, the two branches of dynamics systems thinking reviewed in this chapter may be roughly mapped onto two forms of neuronal coding. In rate coding, different levels of firing rate are assumed to represent different states of a sensor or effector system. This form of coding is typical at the periphery of the nervous system. Motor neurons, for instance, bring about different levels of force production in muscles when active at different rates. Interestingly, even in these simplest cases, the actual physical state of the associated muscle–joint system depends on the resistance encountered and on local reflex loops requiring a proper dynamical understanding of the embodied and situated system (Latash, 1993). The uniquely instantiated dynamical systems approach generalizes the rate code principle to include forms of instability from which qualitatively new neural functions may emerge as the neuronal dynamics change gradually. Even invertebrate nervous systems, in which rate coding is the prevalent form of neural representation (Bullock, Orkand, & Grinnell, 1977), provide examples of such instabilities. In these systems, neurons may switch allegiance among different patterns of neuronal activity that are responsible for different elementary behaviors. When sufficiently large populations of neurons switch, a macroscopic change of behavior may result, for instance, a switch to another pattern

of coordinated activity in Central Pattern Generators (Nusbaum & Beenhakker, 2002).

The other form of neuronal representation is based on the space code principle, which states that what a neuron represents is determined by its position within the neuronal network. The firing rate then merely expresses how strongly the information represented by each neuron contributes. The space code principle is typically assumed to be prevalent in the central nervous system of vertebrates. In vertebrate nervous systems, space coding is postulated for cortex but also such subcortical structures as the thalamus, colliculus, or the hippocampus. Dynamic neuronal fields are direct abstractions of neuronal representations based on the space coding principle.

The concept of a dynamic neuronal field adds two assumptions to the space code principle. First, dynamic fields postulate a topology, in which neurons representing metrically similar contents interact excitatorily, whereas inhibitory interaction relates neurons across the represented metric dimension. This principle is consistent with the existence of topographic maps, in which neighborhood relationships on the sensory or motor surface are preserved in neuronal representation. Within such maps, local excitatory interaction is typically observed (Douglas & Martin, 2004). The overlapping patterns of input activation together with local excitatory interaction justify the continuous description of neuronal patterns of activation on which DFT is based.

A topographical layout of functional maps is not the only way in which this basic assumption of DFT can be realized. Population coding is a more general way for dynamical fields to be realized (Jancke et al., 1999; Erlhagen et al., 1999). The conception of population coding is based on the observation that cortical and subcortical neurons typically have broadly overlapping tuning functions, so that for any given perceptual or motor state, many neurons are active (Georgopoulos, Schwartz, & Kettner, 1986). This is true of most cortical maps, but also of brain structures without

apparent topographical order such as the motor and premotor cortices with respect to movement parameters such as the direction of end-effector motion or the spatial direction of end-effector force (Georgopoulos et al., 1992). In some instances, researchers were able to show that interaction is excitatory among neurons coding for metrically similar values of movement parameters (Georgopoulos, Taira, & Lukashin, 1993). A systematic mapping of neuronal population coding onto dynamic fields can be based on the concept of Distributions of Population Activation, in which not only the most active neurons and their preferred stimulus or motor state are taken into account, but the entire distribution of activation is interpreted (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999; Bastian, Schöner, & Riehle, 2003).

Dynamic field theory makes the second major assumption that under some circumstances neuronal interaction can be dominant in the sense that activation patterns are not necessarily dictated by afferent input but may be stabilized by interaction from "within" a neuronal representation. Neuroanatomically, the vast majority of neuronal connections are not part of a unidirectional feed-forward path from the sensory to the motor surfaces (Braitenberg & Schüz, 1991). This fact speaks in favor of the assumption that interaction may be dominant. There is still not much general understanding of the strength and effectiveness of neuronal interaction compared with neuronal input from the sensory surfaces. This may in large part be a consequence of the methodological bias toward input-output characterizations, for which we have a large ensemble of powerful techniques available. By contrast, the identification of strong forms of neuronal interaction is comparatively more difficult and methodologically not systematized.

Does any structure in the brain "read" population codes or cortical maps? In other words, could one ignore the structure of neuronal representations built on the space code principle and instead study cognition at a "higher" level, at which symbols and

their manipulation might be directly represented? Although there have been theoretical analyses of how population codes may be “read” (Seung & Sompolinsky, 1993; Denève, Latham, & Pouget, 1999), it is difficult logically to answer this question. What one can say is that a transformation into a different coding regime for higher brain function is unnecessary and not justified by any data. Furthermore, it is easy to conceive of dynamical principles that govern neuronal function all the way through to the effector systems. In this view, stable patterns of neuronal activation ultimately steer the periphery into dynamical states, from which behavior emerges, without any need to ever abstract from the space-time contiguous processes that embody cognition.

All this having been said, the dynamical systems approaches sketched here do remain at a certain level of abstraction. More detailed neuronal realism may entail taking into account the spiking nature of neuronal interaction, cellular processes both at the level of activation dynamics and their modification by processes of learning and adaptation. What is the right level of abstraction for an understanding of neuronal function? This question may be difficult to answer in general. It is useful to keep in mind, however, that all answers to this question are subject to critical assessment. If one assumes primacy of the micro level, then the flight toward the microscopic would find no end. Why wouldn't molecular or even atomic levels of description be privileged over cellular descriptions, for instance?

So what are the arguments in favor of the particular neural level of description at which our dynamical systems approach is so effective? The answer lies within the embodied stance: It is mass action in the nervous system that is correlated with those motor and sensory parameters to which cognition is sensitive. Neural activity of populations of neurons in various parts of the central nervous system modulate their temporal evolution with the demands of cognitive tasks. The time courses of population activation are predictive of behavioral

events (Shadlen & Newsome, 2001; Schall, 2001; Bastian et al., 2003), and the metrics of distributions of population activation are predictive of the metrics of behavioral responses (Cisek & Kalaska, 2005). Similar arguments can be made for the instantiated dynamical systems approach (Schöner & Kelso, 1988a).

The level of description of DFT makes explicitly the assumption that the temporal discreteness of neuronal spiking is unrelated to cognitive and behavioral events. Such events must therefore be understood as emerging from an underlying temporally continuous process. Analogously, the assumption is made that the discrete nature of neurons is unrelated to any cognitive or behavioral discreteness. In particular, the formation of discrete cognitive categories is understood as emerging from an underlying continuum of neuronal representation.

5.3. *What Kind of Account Does Dynamical Systems Thinking Generate?*

If all practical difficulties were removed, what would an ultimate dynamical systems account of cognition look like? It is easier to say what it would not look like. It would not look like the ultimate information-processing model of cognition, with all cognitive processing units listed and their pathways of information exchange identified. Nor would it be like the ultimate connectionist network model, the layers of which would encode all existing neural representations and the network topology of which would reflect the neuronal architecture.

In fact, in dynamical systems thinking, the conceptual interaction with experiment, proposing new questions and new measures, has been more important than the models that resulted. In that sense, dynamical systems thinking is primarily aimed at developing a generative theoretical language that facilitates the uncovering of regularities in nervous systems. Models are tools to test concepts both for internal consistency and, through quantitative theory-experiment relationships, for consistency with nature. Models are not by themselves the main goal

of a dynamical systems approach to cognition.

This emphasis on concepts over models is fostered by a central property of dynamical systems, sometimes designated as *emergence*. Attractors in a dynamical system emerge when the conditions are right (when parameters have particular values). The dynamical system relaxes to a new attractor when the history is right, so that the initial condition puts the system in a position from which the attractor is reached. Both may occur in response to changes that are not specific to the newly realized attractor. For instance, a change of resting level in a dynamic field may lead to the attractor state, in which a localized peak is sustained in the absence of localized input. The resting level does not specify any particular location for a peak, nor that a peak must be generated. But everything else being equal, a peak may persist stably only when the resting level is sufficiently large. That attractor is the basis of such cognitive functions as working memory or sensorimotor decision making. These capacities may thus emerge from a neural field dynamics in the here and now, in response to inputs or global changes. These capacities are not enclosed in a particular module that sits somewhere in the brain, waiting to be invoked. Instead, the same neuronal dynamics may under some circumstances have the cognitive functions of working memory and decision making and, under other circumstances, lose these functions. Any individual contribution to the neuronal dynamics is thus *multifunctional*.

Conversely, there are multiple ways a new stable state may emerge as well as disappear again (*multicausality*). There is quite possibly no single parameter that is strictly necessary or is always sufficient for a given cognitive function to emerge. Even a complete understanding of the dynamics of the neural cognitive system is not by itself sufficient to predict all possible ways cognition may unfold when an organism is immersed in a new and rich environment. When processes of adaptation and learning incorporate parts of the environment and of experience into the system, the resultant com-

plexity may become inextricable (Rosen 2005).

Dynamical systems thinking is in that sense open ended. It is not aimed, even in principle, at an ultimate model, which would include process models of all cognitive, sensory, and motor capacities of the human. Instead, it is aimed at understanding constraints for learning and development, for how individual differences may manifest themselves in different contexts, how individual learning and developmental histories may lead to the same function. So although we may never be able to predict how a child moves about in the playground, which swing or slide she will select, we may very well understand how progress in her motor skills may improve her spatial orientation or how perceptual experience with a set of objects will impact on what she pays attention to when naming a new object.

Appendix: Dynamical Field Theory of Perseverative Reaching

The dynamic field theory of perseverative reaching has its roots in a metaphorical dynamical systems account of Piaget's A not B error (Thelen & Smith, 1994). A first formalization into a mathematical model was reported in (Thelen et al., 2001). Conceptual errors in that earlier account were corrected by Evelina Dineva, and it is her model that I review here (Schöner & Dineva, 2006; Dineva & Schöner, 2007).

A dynamical activation field is defined over the space of movement directions, ϕ , of the infant's hand. This is the "reaching" field, $u(\phi, t)$. Its dynamics has the form of an interactive neuronal field (Amari, 1977) receiving a set of inputs:

$$\begin{aligned} \tau \dot{u}(\phi, t) = & -u(\phi, t) + h \\ & + \int d\phi' w(\phi - \phi') \sigma(u(\phi', t)) \\ & + S_{\text{task}}(\phi) + S_{\text{spec}}(\phi, t) \\ & + S_{\text{boost}}(t) + u_{\text{habit}}(\phi, t). \end{aligned} \tag{4.3}$$

Here, $\tau \dot{u}(\phi, t) = -u(\phi, t) + h$ sets the activation field up as a dynamical system with resting state $u(\phi) = h < 0$, a homogenous stationary stable state in the absence of input and interaction, to which the field relaxes on the time scale, τ . Interaction consists of input from other field locations, ϕ' , which is excitatory or inhibitory, depending on the interaction kernel

$$w(\phi - \phi') = w_{\text{excitatory}} \exp\left[-(\phi - \phi')^2 / 2\Delta\right] - w_{\text{inhibitory}}. \quad (4.4)$$

For sufficiently close field locations ($|\phi - \phi'| < \Delta$), the intra-field connectivity is predominantly excitatory ($w_{\text{excitatory}} > w_{\text{inhibitory}}$), for larger distances it is inhibitory. Only sites, ϕ' , with sufficiently positive levels of activation contribute to interaction, as controlled by the nonlinear sigmoidal function

$$\sigma(u) = \frac{1}{1 + \exp(-\beta u)} \quad (4.5)$$

whose parameter, β , controls how nonlinear the interaction term is.

The input functions, $S_{\text{task}}(\phi)$, $S_{\text{spec}}(\phi, t)$, and $S_{\text{boost}}(t)$ model the experimental scenario. The task input, S_{task} , captures the visual layout of the workspace and is modeled as a sum of two gaussians centered over the two movement directions, in which the two locations, A and B, lie. The specific input, $S_{\text{spec}}(\phi, t)$, captures the experimenter's effort to attract attention to the A location on A trials and to the B location on B trials. It is modeled as a gaussian centered on the corresponding location that is nonzero only during the time interval during which the experimenter performs this stimulation. The boost, $S_{\text{boost}}(t)$, captures the effect of pushing the box with the two lids into the reaching space of the infant. It is modeled as a positive constant present only during the time interval when the box is in reaching space.

Finally, the formation of a habit of reaching is modeled by a second dynamical activation field, $u_{\text{habit}}(\phi, t)$, which evolves over a longer time scale, τ_{habit} , and forms a memory trace of locations in the reaching field, $u(\phi, t)$, at which sufficient levels of activation have been generated. The dynamics of this memory trace is modeled as follows:

$$\begin{aligned} \tau_{\text{habit}} \dot{u}_{\text{habit}}(\phi, t) &= [-u_{\text{habit}}(\phi, t) + c_{\text{habit}}\sigma(u(\phi, t))] \\ &\times \Theta\left(\int d\phi' \Theta(u(\phi', t))\right). \end{aligned} \quad (4.6)$$

The last term turns the memory trace mechanism off if there is no positive activity anywhere in the reaching field. This makes use of the step function, $\Theta(u) = 1$ if $u > 0$ and $\Theta(u) = 0$ while $u \leq 0$. Thus, during epochs in which there is no activation in the reaching field, the memory trace remains unchanged. This captures the observation that inter-trial intervals do not seem to matter much in the A not B paradigm. In fact, perseverative tendencies persist through considerable delays.

When an activation peak has been induced in the reaching field, then the memory trace mechanism leads to increase of the memory trace in locations on which the peak is centered, whereas activation elsewhere in the memory trace decays toward zero. Thus, a dynamical balance emerges between different locations at which peaks are induced in different trials. The constant, c_{habit} , determines the amplitude of the memory trace. In Dineva's implementation of this dynamics, a memory trace is laid down only during the interval when the box is in the reaching space (that is, while the boosting input is present). At the end of a trial, the peak in the reaching field is deleted, and the field starts the next trial from its resting state (Figures 4.11 and 4.12).

Neuronal activity in the nervous system has a stochastic component. To account for fluctuations in activation, which make the outcome of reaching decisions nondeterministic, the model contains stochastic forces. The generic model for such forces is additive

gaussian white noise. This noise may be spatially uncorrelated at different field sites or also be correlated, modeling stochastic input distributed by input kernels.

To simulate DFT models, the equations must be numerically solved on a computer using standard numerical procedures (Kloeden & Platen, 1999). Simulating the experimental paradigm typically requires programming the time courses of sensory inputs that describe the experience in such paradigms. Under some circumstances, this may include the need for sensor and motor models, in which the sensory consequences of a motor act driven from the DFT model is also taken into account (Steinhage & Schöner, 1998).

References

- Althaus, P., & Christensen, H. I. (2003). Smooth task switching through behaviour competition. *Robotics and Autonomous Systems*, *44*, 241–249.
- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*, 77–87.
- Amit, D. J. (1994). The hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, *18*(4), 617–626.
- 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.
- Bergener, T., Bruckhoff, C., Dahm, P., Janssen, H., Joubin, F., Menzner, R., et al. (1999). Complex behavior by means of dynamical systems for an anthropomorphic robot. *Neural Networks*, *12*(7), 1087–1099.
- Bicho, E., Mallet, P., & Schöner, G. (2000). Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research*, *19*, 424–447.
- Bicho, E., & Schöner, G. (1997). The dynamic approach to autonomous robotics demonstrated on a low-level vehicle platform. *Robotics and Autonomous Systems*, *21*, 23–35.
- Braitenberg, V., & Schüz, A. (1991). *Anatomy of the cortex*. Berlin: Springer Verlag.
- Brooks, R. A. (1991). New approaches to robotics. *Science*, *253*, 1227–1232.
- Bullock, T. H., Orkand, R., & Grinnell, A. (1977). *Introduction to nervous systems*. San Francisco: W.H. Freeman.
- Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, *3*(45), 801–814.
- Denève, S., Latham, P., & Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nature Neuroscience*, *2*(8), 740–745.
- Dineva, E., & Schöner, G. (2007). Behavioral history matters: A dynamic field account of spontaneous and perseverative errors in infant reaching. Manuscript submitted for publication.
- Douglas, R. J., & Martin, K. A. C. (2004). Neural circuits of the neocortex. *Annual Review of Neuroscience*, *27*, 419–451.
- Engels, C., & Schöner, G. (1995). Dynamic fields endow behavior-based robots with representations. *Robotics and Autonomous Systems*, *14*, 55–77.
- 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., & Bicho, E. (2006). The dynamic neural field approach to cognitive robotics part of the 3rd neuro-it and neuroengineering summer school tutorial series. *Journal of Neural Engineering*, *3*(3), R36–R54.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545–572.
- Faubel, C., & Schöner, G. (2006). Fast learning to recognize objects: Dynamic fields in label-feature space. In *Proceedings of the Fifth International Conference on Development and Learning ICDL 2006*. Bloomington, IN: Indiana University.
- Fuster, J. M. (2005). *Cortex and mind—Unifying cognition*. New York, NY: Oxford University Press.
- Gelder, T. van, & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.

- Georgopoulos, A. P., Ashe, J., Smyrnis, N., & Taira, M. (1992). The motor cortex and the coding of force. *Science*, 245, 1692–1695.
- Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental rotation of the neural population vector. *Science*, 243, 1627–1630. (Reprinted in Kelner, K. L., Koshland, D. E., Jr. (Eds.), *Molecules to models*, AAAS, 1989.)
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neural population coding of movement direction. *Science*, 233, 1416–1419.
- Georgopoulos, A. P., Taira, M., & Lukashin, A. (1993). Cognitive neurophysiology of the motor cortex. *Science*, 260(5104), 47–52. (Comment in *Science* 1994 Mar 4; 263:1295–1297)
- Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge, UK: Cambridge University Press.
- Ghez, C., Favilla, M., Ghilardi, M. F., Gordon, J., Bermejo, R., & Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research*, 115, 217–233.
- Goodale, M. A., Pélisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320, 748–750.
- Grossberg, S. (1980). Biological competition: Decision rules, pattern formation, and oscillations. *Proceedings of the National Academy of Sciences (USA)*, 77, 2338–2342.
- Grossberg, S., Pribe, C., & Cohen, M. A. (1997). Neural control of interlimb oscillations. I. Human bimanual coordination. *Biological Cybernetics*, 77, 131–140.
- Heisenberg, M., & Wolf, R. (1988). Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *Journal of Comparative Physiology*, 163, 373–388.
- Hock, H. S., Bukowski, L., Nichols, D. F., Huisman, A., & Rivera, M. (2005). Dynamical vs. judgemental comparison: Hysteresis effects in motion perception. *Spatial Vision*, 18(3), 317–335.
- Hock, H. S., Gilroy, L., & Harnett, G. (2002). Counter-changing luminance: A nonfourier, nonattentional basis for the perception of single-element apparent motion. *Journal of Experimental Psychology: Human Perception and Performance*, 28(1), 93–112.
- Hock, H. S., Kelso, J. A. S., & Schöner, G. (1993). Perceptual stability in the perceptual organization of apparent motion patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 63–80.
- Hock, H. S., Kogan, K., & Espinoza, J. K. (1997). Dynamic, state-dependent thresholds for the perception of single-element apparent motion: Bistability from local cooperativity. *Perception & Psychophysics*, 59, 1077–1088.
- Hock, H. S., Nichols, D. F., & Espinoza, J. (2004). When motion is not perceived: Evidence from adaptation and dynamical stability. *Spatial Vision*, 17(3), 235–248.
- Hock, H. S., & Ploeger, A. (2006). Linking dynamical perceptual decisions at different levels of description in motion pattern formation: Psychophysics. *Perception & Psychophysics*, 68(3), 505–514.
- Hock, H. S., Schöner, G., & Giese, M. A. (2003). The dynamical foundations of motion pattern formation: Stability, selective adaptation, and perceptual continuity. *Perception & Psychophysics*, 65, 429–457.
- Hund, A. M., & Spencer, J. P. (2003). Developmental changes in the relative weighting of geometric and experience-dependent location cues. *Journal of Cognition and Development*, 4(1), 3–38.
- Jancke, D., Erlhagen, W., Dinse, H. R., Akhavan, A. C., Giese, M., Steinhage, A. (1999). Parametric population representation of retinal location: Neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience*, 19, 9016–9028.
- Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system*. Cambridge, MA: MIT Press.
- Keijzer, F. A. (2001). *Representation and behavior*. Cambridge, MA: MIT Press.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. MIT Press.
- Kloeden, P. E., & Platen, E. (1999). *The numerical solution of stochastic differential equations* (2nd ed.). Berlin: Springer-Verlag.
- Köhler, W. (1920/1939). *Physical gestalten*. London: Routledge & Kegan. (Reprinted in W.D. Ellis (Ed.), *A source book of gestalt psychology*.)
- Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural, dynamic fields. *Biological Cybernetics*, 73, 49–60.
- Latash, M. (1993). *Control of human movement*. Champaign, IL: Human Kinetics.

- Markman, A. B., & Dietrich, E. (2000). In defense of representation. *Cognitive Psychology*, 40, 138–171.
- Nichols, D. F., Hock, H. S., & Schöner, G. (2006). Linking dynamical perceptual decisions at different levels of description in motion pattern formation: Computational simulations. *Perception & Psychophysics*, 68(3), 515–533.
- Nusbaum, M. P., & Beenhakker, M. P. (2002). A small-systems approach to motor pattern generation. *Nature*, 417, 434–450.
- Ottes, F. P., van, Gisbergen, J. A. M., & Egermont, J. J. (1984). Metrics of saccade responses to visual double stimuli: Two different modes. *Vision Research*, 24, 1169–1179.
- Perko, L. (1991). *Differential equations and dynamical systems*. Berlin: Springer Verlag.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Poggio, T., & Reichardt, W. (1976). Visual control of orientation behaviour in the fly. *Quarterly Reviews of Biophysics*, 9, 377–438.
- Reichardt, W., & Poggio, T. (1976). Visual control of orientation behaviour in the fly: I. A quantitative analysis. *Quarterly Reviews in Biophysics*, 9, 311–375.
- Rosen, R. (2005). *Life itself: A comprehensive inquiry into the nature, origin, and fabrication of life*. New York: Columbia University Press.
- Schall, J. D. (2001). Neural basis of deciding, choosing and action. *Nature Reviews Neuroscience*, 2, 33–42.
- Scholz, J. P., Kelso, J. A. S., & Schöner, G. (1988). Dynamics governs switching among patterns of coordinations in biological movement. *Physics Letters*, A134, 8–12.
- Schöner, G. (2002). Timing, clocks, and dynamical systems. *Brain and Cognition*, 48, 31–51.
- Schöner, G., & Dineva, E. (2006). Dynamic instabilities as mechanisms for emergence. *Developmental Science*, 10, 69–74.
- Schöner, G., Dose, M., & Engels, C. (1995). Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16, 213–245.
- Schöner, G., Haken, H., & Kelso, J. A. S. (1986). A stochastic theory of phase transitions in human hand movement. *Biological Cybernetics*, 53, 247–257.
- Schöner, G., & Kelso, J. A. S. (1988a). Dynamic pattern generation in behavioral and neural systems. *Science*, 239, 1513–1520.
- Schöner, G., & Kelso, J. A. S. (1988b). A dynamic theory of behavioral change. *Journal of Theoretical Biology*, 135, 501–524.
- Schöner, G., & Kelso, J. A. S. (1988c). Synergetic theory of environmentally-specified and learned patterns of movement coordination: I. Relative phase dynamics. *Biological Cybernetics*, 58, 71–80.
- Schöner, G., Zanone, P. G., & Kelso, J. A. S. (1992). Learning as change of coordination dynamics: Theory and experiment. *Journal of Motor Behavior*, 24, 29–48.
- Schutte, A. R., & Spencer, J. P. (2002). Generalizing the Dynamic Field Theory of the A-not-B error beyond infancy: Three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, 73, 377–404.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74, 1393–1417.
- Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5, 534–546.
- Seung, H. S., & Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the National Academy of Science (U.S.A.)*, 90, 10749–10753.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916–1936.
- Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: The task dynamics of the a-not-b error. *Psychological Review*, 106(2), 235–260.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, XX(1), 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., & Schöner, G. (2003). Bridging the representational gap in the dynamical systems approach to development. *Developmental Science*, 6, 392–412.
- Spencer, J. P., Simmering, V. R., & Schutte, A. R. (2006). Toward a formal theory of flexible spatial behavior: Geometric category

- biases generalize across pointing and verbal response types. *Journal of Experimental Psychology: Human Perception and Performance*, 32(2), 473–490.
- Spencer, J. P., Simmering, V. R., Schutte, A. R., & Schöner, G. (2007). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In J. 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(5), 1327–1346.
- Steinhage, A., & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. In M. G. T. Schenker P. S. (Ed.), *Sensor fusion and decentralized control in robotic systems: Proceedings of SPIE* (Vol. 3523, p. 169–180). Boston, MA: SPIE Publishing.
- 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, MA: MIT Press.
- Trappenberg, T. P., Dorris, M. C., Munoz, D. P., & Klein, R. M. (2001). A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *Journal of Cognitive Neuroscience*, 13(2), 256–271.
- Turvey, M. T. (1990). Coordination. *American Psychologist*, 45(8), 938–953.
- Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Brain and Behavioral Sciences*, 21, 615–665.
- Vorberg, D., & Wing, A. (1996). Modeling variability and dependence in timing. In H. Heuer & S. W. Keele (Eds.), *Handbook of perception and action* (Vol. 2, pp. 181–262). London: Academic Press.
- Wilimzig, C., Schneider, S., & Schöner, G. (2006). The time course of saccadic decision making: Dynamic field theory. *Neural Networks*, 19, 1059–1074.
- Wilson, H. R. (1999). *Spikes, decisions, and actions: Dynamical foundations of neurosciences*. New York: Oxford University Press.
- Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12, 1–24.
- Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13, 55–80.