

In: Sun, R (ed.): The Cambridge Handbook of Computational Cognitive Science.
2nd Edition. Cambridge University Press (in press, 2023)

6 Dynamical Systems Approaches to Cognition

Gregor Schöner

6.1 Introduction

Think of a child playing in the playground, climbing up on ladders, jumping, running, catching other kids. Or think of the child painting a picture, dipping the brush into a paint pot, making a sequence of brush strokes to sketch a house. These behaviors certainly are not driven by reflexes, are not fixed action patterns elicited by key stimuli, nor are they strictly dictated by stimulus–response relationships. They exhibit hallmarks of cognition such as selection decisions, sequence generation, and working memory. What makes these daily life activities intriguing is, perhaps, how seamlessly the flow of activities moves forward. No artificial system has ever achieved even remotely comparable behavior. While computer programs may play chess at grand master level, their ability to generate smooth flows of actions in natural environments remains extremely limited.

Emphasizing how cognition links to sensory-motor activity is part of the embodiment perspective on cognition (Shapiro, 2019). Cognition that is directed at objects in the world may interact with motor activation (for example, Chrysikou, Casasanto, & Thompson-Schill, 2017). But motor activation is not mandatory for cognition and may be negligible for mental acts that are not directed at physical objects (M. Wilson, 2002). It is certainly possible to think without overt or even covert motor activation.

A more refined view of embodiment is, instead, that cognition inherits properties from the sensory-motor processes from which it emerged evolutionarily and developmentally. Lifting spatial relations and movement representations through metaphor from the sensory-motor domain to abstract thought is an example (Lakoff & Johnson, 1999). The use of spatial representations in creativity (Fauconnier & Turner, 2002) and the idea that concepts are embedded in feature spaces (Gärdenfors, 2000) are other examples.

The dynamical systems perspective on cognition is linked to the embodiment perspective for good reasons (Beer, 2000; Port & van Gelder, 1995). Dynamical systems are characterized by state variables, whose values at any given moment in time predict their future values (Perko, 2001). The laws of motion of physics take the form of dynamical systems, with the initial conditions of the physical state variables determining the future evolution of those state variables. The

dynamical systems perspective on cognition refers, however, not to just any dynamical system, but to those with particular properties, most prominently, those with attractor states, that is, invariant solutions to which the system converges from any initial condition nearby (Van Gelder, 1998). Such attractor states are critical to control, that is, to steering a physical system to a desired state (Ashby, 1956). In control, sensors pick up deviations from the desired state and the controller drives change of the state variables in a direction that reduces such deviations. Control works in closed loop, in which the controller's action leads to changes in sensory signals, which in turn lead to changes in the controller's action. Embodied cognition typically takes place as organisms act in closed loop with their environment. To direct an action at an object, for instance, you first shift gaze to the object's location. As a result of this action, the visual stimulus changes. As you handle an object, its visual appearance changes. To avoid run-away behavior, closing sensory-motor loops through the environment requires dynamic stability.

The dynamical systems perspective on cognition postulates that cognitive processes share properties with the sensory-motor domain, most centrally, stability properties that enable cognitive processes to link to the sensory-motor surfaces, continuously or intermittently. Dynamical systems ideas go beyond the notion of control, however. Cognition is characterized by the multiplicity of possible states, the complexity inherent in combining many different states into new entities, and the capacity to generate new sequences of states never before encountered. One idea is to attribute that complexity to the self-organizing capacity of nonlinear dynamical systems (Schöner, 2014; Schöner & Kelso, 1988; Thelen & Smith, 1994), in which new states emerge from dynamic instabilities, multiple stable states may coexist, and graded change during learning and development may give rise to qualitative change of behavior or competence.

Dynamical systems ideas also go beyond embodiment in that the closing of the loop that requires stability properties may take place within the nervous system. Recurrent neural networks (see Chapter 2 in this handbook) are dynamical systems: When the inputs to some neurons depend on the outputs of those neurons, activation must be looked at in time: the previous outputs determine the current inputs, leading to an iterative form of computation. Even though some models use discrete time, these iterative update rules for neural activation really are dynamical systems. Their properties are critical for sequence generation (Elman, 1990), for working (Compte, Brunel, Goldman-Rakic, & Wang, 2000; Durstewitz, Seamans, & Sejnowski, 2000) and episodic memory (Rolls, Stringer, & Trappenberg, 2002), and for the generation of actions (see Chapter 35 in this handbook). Couched in terms of the dynamics of neural populations, dynamical systems ideas are effectively a refinement of the more general connectionist ideas.

A related source of dynamical systems ideas comes from neurophysics, the dynamics of neural membranes and synapses (Gerstner, Kistler, Naud, & Paninski, 2014). These electro-chemical processes introduce continuous state

dependence even to individual neurons and thus also to feed-forward, not just to recurrent neural networks. Stephen Grossberg's pioneering work (Grossberg, 1970) established how simplified models of the dynamics of neurons provide the core mechanisms of perception, movement generation, and cognition, building a neural-dynamic theory of essentially everything that can be reached by the methods of experimental psychology (Grossberg, 2021). The neurally grounded dynamical systems ideas reviewed below could be viewed as a variant of that framework in which a small set of principles is used to organize this vast territory. The mathematics underlying much of this work has been elaborated in a large literature which this chapter only reviews selectively (Ermentrout, 1998; Coombes, beim Graben, Potthast, & Wright, 2014).

One particular dynamical systems approach, the neurally grounded Dynamic Field Theory (DFT, see Schöner, Spencer, & DFT Research Group, 2016 for a book-length tutorial), is presented as a case study in some mathematical detail below. Its relation to other dynamical systems approaches, to other neurally grounded approaches, and to cognitive modeling in general, is discussed in the final section of this chapter.

6.2 The Foundation of Neural Dynamics

To examine how cognition may emerge from sensory-motor processes, consider first the sensory and motor periphery. Sensory surfaces like the retina, the cochlea, the skin, or the proprioceptive system, respond to physical stimuli that originate from the world. Hypothetically, patterns of stimulation could be as high-dimensional as the number of sensor cells. In reality, stimuli driving individual sensor cells are not independent of each other when stimulation comes from the world. Such stimuli are much lower-dimensional, reflecting the continuity of surfaces in vision and touch, or the properties of sound sources in auditory perception (Gibson, 1966). Low-dimensional descriptions of stimuli may entail the two spatial dimensions of the visual and auditory arrays, visual feature dimensions such as local orientation, texture, or color, auditory feature dimensions such as pitch, haptic feature dimensions like the direction of local stress vectors, or proprioceptive feature dimensions like joint angles and their rate of change. The motor surface could analogously be construed as the ensemble of muscles and their mechanical linkages that span the space of possible motor states. Again, the covariation of muscle activation observed as synergies makes that the space of possible motor patterns is lower in dimension (Latash, 2008).

The firing rate of sensory neurons varies monotonically with the physical intensity of stimulation (e.g., luminance, loudness, or the displacement of a skin element). When the firing rate of motor neurons varies, the level of force generation in muscles co-varies. Figure 6.1 illustrates how these two links to the sensory-motor periphery bracket neural dynamic architectures.

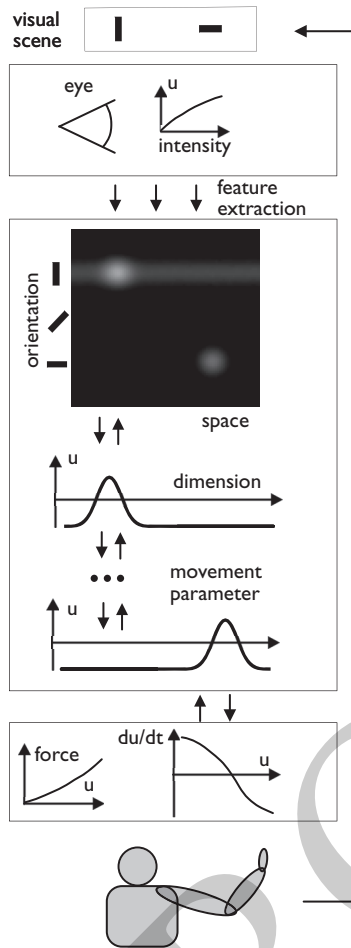


Figure 6.1 A schematic view of a neural dynamic architecture (center box) that is linked to sensory (top box) and motor systems (bottom box). Sensors transform physical intensity (e.g. luminosity impinging on the eye from the visual scene) into neural activation (here denoted by u). Forward neural networks extract feature dimensions that provide input to the neural dynamic architecture. Perceptual fields span such feature dimensions (here orientation and visual space) by virtue of that input connectivity. Coupled neural fields of varying dimensionality form the neural dynamic architecture. At the interface to the motor system, the pattern of connectivity sets fields up to span movement parameters. The neural dynamics of motor systems (often realized in the periphery by reflex loops) feeds into muscles that transform neural activation into force, driving the body's movement. Behavior unfolds in closed loop, in which actions impact on the visual scene.

6.2.1 Activation

Neural dynamic models abstract from some of the physiological details of neural activity. Real neurons in the brain carry a negative electric potential

inside their cellular membrane. Input from the synapses on a neuron's dendritic tree may induce increases (for excitatory synapses) or decreases (for inhibitory synapses) of the electric potential, which travel to the neuron's soma. If the electric potential near the soma exceeds a threshold, a spike or action potential is generated in which the electrical potential briefly becomes positive. Action potentials travel down the axon and activate synaptic connections on the output side, inducing post-synaptic potential changes on the dendritic trees of downstream neurons. In neural dynamics, the electrical potential is replaced by an activation state, u , that has abstract units. The mechanisms of spiking and synaptic transmission are simplified by modeling the output of a neuron as a sigmoid threshold function, $\sigma(u)$ (illustrated in Figure 6.3), which provides input to any down-stream neuron. This simplification is shared with most connectionist models and provides a good approximation for the activity in populations of neurons.

6.2.2 Activation Fields

Neurons in the brain receive input that ultimately comes from the sensory surfaces (Figure 6.1) and reflects patterns of stimulation from the world. The pattern of forward connectivity extracts feature information about such stimuli and creates cortical and subcortical maps, in which neural firing is characterized by tuning curves and receptive field (see Chapter 3 of Schöner, Spencer, & DFT Research Group, 2016 for tutorials on the core neurophysiological concepts). Modeling activity in such neural maps as neural fields amounts to neglecting the discrete sampling of the sensory surface and feature spaces by individual neurons. Because there are no known behavioral signatures of that discrete sampling, this is a useful approximation that helps keep track of the continuity of the underlying sensory and motor spaces. (There are also more specific neuro-anatomical arguments for that approximation based on the relative homogeneity of cortical layers and the strongly overlapping dendritic trees of neighboring neurons, see H. R. Wilson & Cowan, 1972 and Coombes et al., 2014). This leads to the notion of neural activation fields, $u(x)$, that are "defined" over spatial or feature dimensions, x (illustrated in Figure 6.2). They can be defined that way only because the forward connectivity from the sensory surface generates inputs to the fields that reflect the spatial and feature dimensions of possible stimuli.

Activation fields can be analogously defined for motor representations. Neurons in the motor areas of cortex and of subcortical structures have tuning curves that characterize how the firing rates of neurons vary when a voluntary movement is varied. For instance, neurons in motor and premotor cortex have broad tuning curves to the hand's movement direction in space (Schwartz, Kettner, & Georgopoulos, 1988). Similar tuning to movement parameters such as movement extent, or the direction of required force, can be observed. For any specific motor act, activation is localized along such motor dimensions. (This is true even though neighboring neurons do not always have similar tuning




Figure 6.2 *Activation fields span metric spaces whose dimensions are determined by the connectivity to and from each field. Activation patterns (fat line) represent particular values along the dimensions through peaks, stabilized by local excitatory and global inhibitory interaction. Peaks are induced, but not uniquely specified, by input (thin line), reflecting the capacity of fields to make decisions.*

curves in the motor domain. What matters is neighborhood in connectivity, not neighborhood on the cortical surface.)

In Dynamic Field Theory (DFT), localized peaks of activation are the units of representation. In the sensory domain, a localized peak of activation reflects the presence of an object on the sensory surface that can be described by a value along each of a set of feature dimensions. In the motor domain, a localized peak of activation reflects the preparation of a particular motor act. Fields further removed from the sensory and motor surfaces may come to represent more abstract mental states.

The level of activation of a peak may reflect sensory or motor variables. For instance, neural activation levels in visual feature fields may reflect local contrast (Grabska-Barwińska, Distler, Hoffmann, & Jancke, 2009). Neural activation levels in primary motor cortex may reflect the speed of the hand's movement in space (Moran & Schwartz, 1999). As discussed below, however, the activation levels of peaks are largely determined by neural interaction within fields, and are only in a secondary way modulated by feed-forward neural connectivity.

6.2.3 Field Dynamics

Activation fields are formalized mathematically as functions, $u(x, t)$ of the field dimension, x , and of time, t . (For now, consider one dimension only so that x is a scalar.) The evolution in time of activation fields is modeled in DFT by integro-differential equations of this general form:

$$\begin{aligned} \tau \dot{u}(x, t) = & -u(x, t) + \text{resting level} + \text{external input}(x, t) \\ & + \text{interaction}[x, x', \sigma(u(x', t))] \text{ for all } x' \text{ across the field}. \end{aligned} \tag{6.1}$$

The general form of this equation is inherited from models of the dynamics of neural membrane potentials (see Trappenberg, 2010 or Gerstner et al., 2014 for textbook treatment). Activation relaxes in exponential form to the equilibrium state, $u = \text{resting level} + \text{input}$, on the time scale of about 10 msec (so, $\tau = 10$ ms).

Inputs to a field that arise through forward connectivity from a sensory surface set up a field to represent a sensory feature dimension. In DFT architectures, input may also arise from the output of other activation fields. *Neural interaction* is input that arises from the output of the same field, a form of recurrent connectivity: the evolution of activation at a location, x , of the field depends on the output of activation at all other locations, x' , of the field. A core postulate of DFT is that neural interaction is organized to make localized activation peaks attractors of the neural dynamics. Local excitatory interaction stabilizes peaks against decay. Inhibitory interaction over larger distances stabilizes peaks against diffusive spread. Signatures of such a spatial pattern of neural interaction have been observed within populations of cortical neurons in a variety of cortical areas (Georgopoulos, Taira, & Lukashin, 1993; Jancke et al., 1999).

This pattern of connectivity within a field is mathematically modeled by an interaction kernel, $w(x - x')$, illustrated in Figure 6.3. In that description, neural interaction is homogeneous, that is, it has the same form and strength anywhere in the field. That enables neural activation fields to stabilize peaks anywhere along the dimension they represent. In DFT, neural interaction is postulated to be sufficiently strong to dominate the neural dynamics, so that activation may persist purely supported by interaction, without the need for input from outside the field. Strong interaction enables many of the core cognitive functions of DFT architectures, including detection and selection decisions, working memory, and sequence generation. Such strong, homogeneous neural interaction within populations of neurons characterizes DFT models as special cases of generic connectionist models (see also Section 6.6.3).

A concrete mathematical formulation of the field dynamics often used in DFT is:




Figure 6.3 (A) Sigmoidal threshold functions such as the one illustrated here, $\sigma(u) = 1/(1 + \exp(-\beta u))$, characterize the capacity of neural activation, u , to affect down-stream neural dynamics. Only sufficiently activated field locations contribute to output. (B) Homogeneous kernels, $w(x - x')$, depend only on the distance, $x - x'$, between field locations. The neural interaction kernel illustrated is positive over small distances (local excitation) and negative over larger distances (global inhibition). Inhibitory interaction may fall off with distance (not shown).

$$\tau \dot{u}(x, t) = -u(x, t) + h + s(x, t) + \int dx' w(x - x') \sigma(u(x', t)) \quad (6.2)$$

where the resting level is designated by $h < 0$, and external input is designated by $s(x, t)$. In this form, the neural dynamics of activation fields can be mathematically analyzed (Amari, 1977), characterizing the *qualitative dynamics*, that is, the attractor states and their instabilities. A variety of other mathematical formalizations are available (see Coombes et al., 2014 for a modern review, Gerstner et al., 2014 for textbook treatment), whose qualitative dynamics is overall consistent with that of Equation 6.2.

6.2.4 The Detection Instability and Its Reverse

The qualitative dynamics of neural fields comprise two categories of attractor solutions (Figure 6.4). *Input-driven* attractors are subthreshold patterns of activation shaped by input to which neuronal interaction contributes little. Neuronal interaction contributes massively to *self-stabilized* peaks, lifting activation above the input-driven level and suppressing activation outside the peak. That these are qualitatively different attractors can be seen from the fact that they coexist bistably under some conditions and are separated by a dynamical instability, the *detection instability* (see Bicho, Mallet, & Schönner, 2000 for an analysis; see Figure 6.4 for an explanation).




Figure 6.4 Detection descisions in dynamic fields. (A) For weak input (thin solid line: input plus resting level), only the subthreshold input-driven state (fat dashed line) is stable. (B) For stronger input, both the subthreshold input-driven state (fat dashed line) and the self-stabilized peak (fat solid line) are stable. In this bistable regime, which attractor activation converges to depends on the activation pattern present when the inputs first arise (initial condition). (C) For strong input, only the self-stabilized peak is stable. In the detection instability, the subthreshold input-driven state becomes unstable (transition from (B) to (C)). In the reverse detection instability, the self-stabilized peak becomes unstable (transition from (B) to (C)).

The detection instability is observed, for instance, when the amplitude of a single localized input is slowly increased. Below a critical level, the subthreshold input-driven solution, $u(x) \approx h + S(x) < 0$, is stable (for slowly varying $S(x, t)$ which can be approximated as $S(x)$). At appropriate settings of the parameters of the interaction kernel (Amari, 1977), a self-stabilized peak of activation centered on the localized input coexists as a stable stationary state. When the amplitude of localized input reaches a critical level, the subthreshold solution becomes unstable and disappears. This is caused by activation passing through the threshold of the sigmoidal function, so that neural interaction sets in, driving the growth of the peak beyond the level specified by input.

At the detection instability, peaks are created. As peaks are the units of representation, this amounts to a decision that sufficient input is detected to create an instance of representation. If input increases continuously in time, the detection instability occurs at a particular, discrete moment in time when input reaches a critical level. The detection instability is thus instrumental in creating discrete events from time-continuous neural processing, a feature critical to understanding how sequences of neural processing steps arise in neural dynamics (Section 6.4).

Once a peak has been created, it is stable. If input falls below the critical level, the self-stabilized peak persists within a bistable range of input amplitudes. If localized input shifts along the field dimension, the peak tracks that input (Amari, 1977). So while self-stabilized peaks are separated from input-driven activation by the detection decision, they continue to be responsive to input.

Self-stabilized peaks become unstable in the *reverse detection instability* when activation falls below the critical level at which interaction is engaged. This may happen because input falls below a lower critical level, or because inhibitory input pushes activation levels down. At the reverse detection instability, activation is no longer supported by local excitatory interaction and begins to decay, converging to the subthreshold input-driven activation state. So the reverse detection instability causes the deletion of a peak, removing a unit of representation. Again, a time-continuous change may be transformed into an event.

6.2.5 Sustained Activation

There are conditions under which self-stabilized peaks of activation may remain stable even in the absence of any input beyond the resting level (Amari, 1977). Such a *sustained* peak of activation is illustrated in Figure 6.5. This dynamic regime comes about when excitatory interaction in the field, once engaged, is sufficiently strong to keep activation at positive levels, bridging the gap from the negative resting level. This may be because excitatory interaction simply is strong or because the resting level is closer to zero, so that the gap is easy to bridge. In fact, an increase of the resting level can shift the neural dynamics from a regime without to a regime with sustained activation peaks.

Sustained activation is the standard picture for how working memory is neurally realized (Fuster, 1995). Sustained peaks of activation may thus provide




Figure 6.5 *In a sustained peak of activation (fat line), a peak of positive activation persists in the absence of any localized input. Note that activation outside the peak is suppressed below the resting level (marked by the thin horizontal line) by inhibitory interaction. The positive activation level within the peak, induced by some earlier stimulation, is stabilized by local excitatory interaction.*

a neural mechanism for metric working memory. Localized input may induce a peak through the detection instability. The activation peak remains stable after the input is removed. The peak's location in the field retains the metric information about the earlier localized input. This metric information is preserved only to the extent to which no other localized inputs act on the field. Such inputs, even when they are small, may induce drift of the peak, both by attracting to locations with excitatory input and by repelling from locations with inhibitory inputs. Both effects have been observed behaviorally (Schutte & Spencer, 2009; Schutte, Spencer, & Schöner, 2003). Such metric distortions of working memory may be misread as evidence for underlying categorical representations (Spencer, Simmering, & Schutte, 2006).

Capacity limits are natural for DFT models of working memory (J. S. Johnson, Simmering, & Buss, 2014; Simmering, 2016): as the number of peaks increases, the total amount of inhibitory interaction increases, ultimately pushing peaks below the reverse detection instability. This emergent nature of the capacity limit is in contrast to the idea of a fixed number of slots and consistent with ability to modulate capacity by distributing resources (J. S. Johnson et al., 2014) and with other indices of a graded capacity of working memory (Schneegans & Bays, 2016).

6.2.6 Selection

When inhibitory interaction is sufficiently strong, only a single peak may be stable at any given time. This enables selection decisions as illustrated in Figure 6.6. In response to an input distribution that has multiple local maxima, the field generates a single peak positioned over one of those local maxima. That selection decision may be combined with a detection decision if the field is in a subthreshold pattern of activation when input first arises. The location that first reaches threshold wins the neural competition created by inhibitory interaction. Because the peak that emerges is a full self-stabilized peak whose shape and total activation does not reflect how close the selection decision was, this enacts a “winner takes all” mechanism. In some connectionist neural networks, such a normalization step is implemented by a separate mechanism (such as an algorithm reading out the location of the maximum, “argmax”). The decision




Figure 6.6 Selection decisions in dynamic fields. (A) When input on the left is sufficiently much stronger than input on the right, only the left-most peak remains stable. (B) In response to bimodal input (thin solid line), a dynamic activation field may be bistable, supporting a stable peak centered over either local maximum (fat solid and dashed lines). (C) When input on the right is sufficiently much stronger than than input on the left, only the right-most peak remains stable.

may be biased by earlier activation patterns, so that the selected location is not necessarily the location of maximal input. In fact, selection decisions are stable: When input at the selected location becomes weaker or input at another location becomes stronger, the selected peak persists. The limit to that stability occurs in the *selection instability*: When input at a new location becomes sufficiently strong, it lifts activation at that location above threshold in spite of inhibitory interaction, inducing a new peak that then suppresses the earlier peak. (Technically, the field may be bi- or multistable and one of those attractors loses stability).

A subtle, but important property of dynamic fields arises when selection occurs in response to broadly distributed input or to a homogeneous boost to the entire field. In the *boost-driven detection instability*, a field creates a single peak whose location represents a selection decision. Selection is sensitive to small inhomogeneities in the field from input or from a memory trace (Section 6.5): The peak arises at one of the locations with slightly higher initial activation level. In a sense, the boost-driven detection instability amplifies small differences into a full self-stabilized peak at one location, while other locations with very similar initial activation levels are suppressed.

Neural noise and noise originating in sensory inputs are important in DFT due their role at such instabilities. Noise may create a momentary selection advantage for one location which is then amplified into a macroscopic

decision. Only at instabilities does noise play such a role. While far from an instability, peaks are much too stable to be spontaneously suppressed or switched. Nondeterministic aspects of behavior are accounted for in DFT by the amplification of noise around instabilities. The generic mathematical formalization of neural noise in DFT is Gaussian white noise, added to the rate of change of activation (Equation 6.2). (Technically, this makes the neural dynamic model a stochastic differential equation. Because the Ito and Stratonovich calculus do not differ for additive noise, there is no need to specify either framework, see pages 35–37 in Oksendal, 2013). Typically, noise is assumed independent at each field location (spatial correlations can be modeled by a noise kernel).

6.2.7 Neural Dynamic Nodes

So far, all illustrations have been from one-dimensional fields, but the same solutions and instabilities are obtained in two-, three-, or four-dimensional fields (on limits to that later). What about zero-dimensional fields? Those could be thought of as small populations of neurons, mathematically described by a single activation variable, $u(t)$, subject to a neural dynamics of this general form

$$\tau \dot{u}(t) = -u(t) + h + s(t) + w_{\text{exc}} \sigma(u(t)), \quad (6.3)$$

where w_{exc} is the strength of self-excitation (really the net result of excitatory interaction within the small population). These dynamics have stable states analogous to those of neural dynamic fields: a subthreshold activation state ($u_0 \approx h + S < 0$, the “off” state) and a suprathreshold activation state ($u_1 \approx h + s + w_{\text{exc}} > 0$, the “on” state).

What the activation of such a *neural dynamic node* means is determined by the pattern of connectivity of its input and output. Concept nodes, for instance, may be linked to a variety of feature fields, so that particular ranges of feature values may activate such a node, and conversely, a node may provide input to those feature fields, supporting the form of cuing discussed next.

6.3 Neural Dynamic Architectures

6.3.1 Binding

When neural dynamic fields simultaneously represent dimensions that have different meanings, new functions emerge from the dynamic instabilities. Figure 6.7 shows a joint neural representation of visual space (only its horizontal dimension for ease of illustration) and of a visual feature, orientation. Such a joint representation could come about due to feed-forward connectivity from the visual array that extracts visual position and local orientation (e.g. making use of Gabor filters). Figure 6.7 also illustrates two fields that represent each dimension separately and are coupled reciprocally to the joint representation.




Figure 6.7 Core principle of a neural dynamic architecture for visual search. A visual scene (A) consisting of a vertical and a horizontal object provides input to a two-dimensional field (B) over space (horizontal spatial dimension) and orientation (local orientation feature dimension). That input (light gray blobs) is localized along both dimensions. A one-dimensional field defined over the orientation feature dimension (C) has a peak at the vertical orientation representing a search cue. That peak provides ridge input into the two-dimensional field, which induces a peak where the ridge overlaps with the blob input. Projecting suprathreshold activation, summed along the orientation feature dimension, onto a one-dimensional field over space (D) induces a peak at the spatial location of the vertical object.

A peak in the joint field binds the location of a visual object to its orientation. Summing activation along either dimension and projecting onto the separate fields induces peaks there, effectively extracting the individual feature values from the bound representation. Conversely, individual feature values represented by peaks in the separate fields can be bound together by projecting two ridges into the joint field, one along orientation, the other along space. Under appropriate conditions, the joint field reaches the detection threshold only at the intersection of the two ridges, generating a peak there that binds the two feature values together. Note that such binding requires that only one object is represented at a time. If a separate field had peaks at more than one feature value, the projections would intersect at more than one location, inducing “illusory conjunctions” of feature values that belong to different visual objects.

The core mechanism of visual search combines these two directions of coupling. Localized input into the joint field from the visual array is boosted by a ridge of input from a peak in the orientation field that represents the search cue (Figure 6.7). This induces a peak in the joint field only at those locations that overlap with the ridge (a form of biased competition (Desimone, 1998)). A visual object is thus selected, whose orientation matches the search cue represented by the peak in the orientation field. Based on this core mechanism, a comprehensive DFT model of visual search (Grieben et al., 2020) addresses conjunctive search and the autonomous sequential selection of candidate objects.

Binding dimensions by a joint neural field is neurally costly, however, as every possible combination of feature values across dimensions requires dedicated activation variables. Such binding scales poorly with the number of dimensions. Using only 100 neurons per feature dimension, the binding of orientation, color, texture, movement direction, and visual space, for instance, would take 10^{12} neurons, as much as in the entire brain (see Eliasmith & Trujillo, 2014 for a discussion of such scaling issues). The form of conjunctive feature binding relevant for visual search and many other tasks must be more flexible and efficient. Feature Integration Theory (Treisman, 1980) provides a cue. Feature dimensions may each be individually bound to visual space by joint neural representation, consistent with the fact that neurons tuned to different feature dimensions all have spatial receptive fields. But there is no need for all combinations of feature dimensions to be represented by particular neurons. Instead, a stack of neural fields, each spanning visual space and one or a small number of other feature dimensions may together represent the ensemble of features. Binding the different feature dimensions of a particular visual object now occurs through the shared spatial dimension. Bidirectional excitatory interaction along the shared spatial dimension (a cylinder-shaped input pattern to each feature/space field) enables search for conjunctions of features (Griegen et al., 2020). The same mechanism can be used to explain how change detection for feature conjunctions may be achieved (Schneegans, Spencer, & Schöner, 2016).

6.3.2 Coordinate Transforms

Binding different dimensions through joint neural representations enables active coordinate transforms, which are relevant to many sensory-motor and cognitive tasks. To direct action at an object, for instance, visual information in retinal coordinates must be transformed into coordinates anchored in the body (to which the arm is attached). Such a transform depends on (is steered by) an estimate of gaze direction (Schneegans, 2016; Schneegans & Schöner, 2012). The body-centered object location must be further transformed into a frame centered on the initial position of the hand to extract movement parameters such as direction and extent (Schöner, Tekülve, & Zibner, 2019).

The bottom half of Figure 6.8 illustrates an active coordinate transform in a much more cognitive context, perceptually grounding a spatial relation like “the vertical bar to the left of the horizontal bar.” In a spatial representation of the visual array that is centered on the reference object, the “horizontal bar” (bottom of the figure), it is easy to conceive of a pattern of connectivity that would define the relational concept “to the left of.” The connectivity would activate a neural node representing that concept only when activation falls into an appropriate spatial region to the left of the field’s center (Lipinski, Schneegans, Sandamirskaya, Spencer, & Schöner, 2012). An active coordinate transform of the original visual array into a frame centered on the reference




Figure 6.8 A neural dynamic architecture for the grounding of spatial relations. The visual scene on top provides input to a two-dimensional field over orientation and space. Nodes for “vertical” and “horizontal” orientation (circles on top left, filled for activated node) are reciprocally connected to matching regions in a one-dimensional orientation field. The orientation-space field projects onto two spatial fields, “target” and “reference,” by summing along the orientation dimension. These are reciprocally coupled to the diagonal two-dimensional transformation field, which is, in turn, reciprocally coupled to a spatial field that represents the target centered on the reference. Nodes for “to the left of” and “to the right of” are reciprocally coupled to corresponding spatial regions of that spatial field.

object would enable generalizing this pattern of connectivity to reference objects anywhere in the visual array. That transformation would be steered by the reference object’s location in the original frame of reference.

Neural implementations of active coordinate transforms can be based on a joint representation of the original space and a space representing the steering dimension (Pouget & Snyder, 2000). Such representations are observed as gain fields in area LIP of parietal cortex (Andersen, Essick, & Siegel, 1985) and

elsewhere. In the example, the joint representation binds the visual array containing potential target objects of the relation to a spatial representation of the reference object. The projections from the target and reference spaces into the joint representation takes the form of two ridges. Where these ridges meet, a peak is induced that binds the spatial locations of target objects to those of reference objects. Projection from the joint representation onto the transformed space sums outputs along an appropriate subspace. In this example, summing along the diagonal yields a spatial representation centered on the reference object.

6.3.3 Architectures

The neural dynamics in architectures such as the one illustrated in Figure 6.8 can be characterized in terms of dynamic concepts for the individual fields like the detection instability and the capacity for selection. This is not trivial, and only true because of the stability postulate for meaningful activation states. The dynamic stability of such states implies structural stability under change of dynamics. When the dynamics (the equation) change in a continuous way, attractors remain stable (Perko, 2001). Coupling among fields can be viewed as a continuous change of the dynamics by thinking of the coupling strength as being increased from zero. So in tying function to attractor states, DFT models avoid the classical problem of analog computing in which solutions may be completely changed when a new component is added.

Fields retain their dynamic properties within limits that are reached exactly when the coupling within neural architectures induces instabilities. That makes DFT architectures intrinsically flexible. The architecture shown in Figure 6.8 illustrates this point. To perceptually ground spatial relations such as “the vertical bar to the left of the horizontal bar,” this architecture performs visual search first for the reference (“the horizontal bar”), then for the target object (“the vertical bar”). The top half of Figure 6.8 is simply the mechanism for visual search from Figure 6.7. The search cue is provided by concept nodes that may activate either the feature representation of “vertical bar” or of “horizontal bar” by virtue of their connectivity with the feature field defined over orientation. The output of visual search in the orientation-space field projects both to a field representing the spatial location of the reference object and to a field representing the spatial location of possible target objects. By boosting the reference spatial field when the reference object is searched, only that field can reach the detection instability based on the search output. By boosting the target spatial field when, in the next step, target objects are searched, only that field can build peaks. This way, the outcome of the visual search can be directed into either field by boosting the destination field. In connectionist models, such steering of projection is achieved by multiplicative “gating” connections to the projections among neural populations (O’Reilly, 2006).

6.4 Autonomous Sequence Generation

The visual search for target and reference must be performed sequentially. How may such sequences of processing steps arise in neural dynamic systems? And how do the transitions among such steps arise at discrete moments in time from the time-continuous neural dynamics? Figure 6.9 illustrates how the detection instability can be harnessed to bring about such transition events (Sandamirskaya, 2016; Sandamirskaya & Schöner, 2010). A neural field, labelled here the *intention* field, represents an ongoing mental or motor act by a suprathreshold peak of activation. The peak's location specifies the intended act, for instance, the feature value of the object that must be visually searched. That intentional state predicts a sensory or internal outcome that counts as its *condition of satisfaction* (a term borrowed from Searle, 1983). The prediction is realized through neural connectivity, which may have to be learned, to a neural field that represents the condition of satisfaction. The intention to visually search the target predicts an internal outcome, a peak in the joint feature/space field at the cued feature value. The predictive input alone is not sufficient, however, to push the condition of satisfaction field through the detection instability. A peak is formed in that field only when the predicted input arises from a sensory surface (for real motor acts) or from another neural representation (for mental acts).

The condition of satisfaction field inhibits the intentional field globally by providing a negative boost. So once it builds a peak, that inhibition pushes the intentional field through the reverse detection instability, leading to the decay of

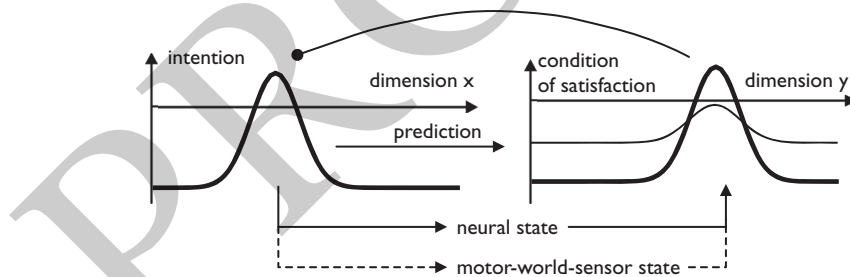


Figure 6.9 The neural dynamic mechanism for sequence generation is based on a pair of neural fields, the *intention* and the *condition of satisfaction* fields, which may be defined over different dimensions. A peak in the *intention* field (fat line on the left) drives the mental or motor act by projecting onto the rest of the neural dynamic architecture. It also provides input (thin line on the right) to the *condition of satisfaction* field that predicts the outcome of a successful completion of the intended mental or motor act. When signals from inside the neural dynamic architecture or from sensory systems provide input that overlaps with that prediction, the *condition of satisfaction* field generates a peak. Through inhibitory projection onto the *intention* field (top line with a filled circle at its end), the peak in the *condition of satisfaction* field may then suppress the peak in the *intention* field and subsequently become unstable itself.

the peak there. This removes the predictive input from the condition of satisfaction field, pushing that field below the reverse detection instability and leading to the decay of that peak as well. The end result of this cascade of instabilities is that both intention and condition of satisfaction fields are returned to a sub-threshold state of activation. The intended act has successfully terminated.

What happens next depends on the neural dynamic architecture. The three classical conceptions for serial order can all be realized in neural dynamic architectures (Henson & Burgess, 1997). First, in the *gradient* conception, intentional states are competing for activation and the most activated one wins. This happens in many neural dynamic architectures. An example is the DFT account of visual search referenced above in which object locations are selected for attention based on the amplitude of summed inputs (Grieben et al., 2020). Second, in the *chaining* conception, an intentional state has a successor that is becoming activated once the intentional state is terminated. In neural dynamic terms, such successor relationships may be expressed by specific coupling structure. For instance, among sets of intentional states, asymmetrical inhibitory coupling may prevent certain states from becoming activated while others are active. Termination of one intentional state may then release other intentional states from inhibition and allow them to become activated. This is how the sequential search for target and reference objects is organized in the DFT architecture of grounding relations (Figure 6.8) (Richter, Lins, & Schöner, 2017, 2021).

Third, the *positional* conception combines chaining with the idea that a neural representation of ordinal position in a sequence points to its contents by neural projection. A neural dynamic architecture realizing positional serial order (Sandamirskaya & Schöner, 2010) is illustrated in Figure 6.10. A set of neural dynamic nodes is coupled to enable their sequential activation along an implied ordinal dimension. Two nodes, an intention and a working memory node, represent each ordinal position. All intention nodes are coupled inhibitory, so that only one of them can be active at any time. Each intentional node activates its memory node which remains activated (sustains activation by self-excitation) after the intention node has been deactivated. Each memory node provides excitatory input to the intention node of its successor within the ordinal set. This leads to the successive activation of intentional nodes along the ordinal dimension each time a condition of satisfaction is reached (Sandamirskaya, 2016). Content is associated with each ordinal position by synaptic connectivity from each intention node to relevant feature fields (which may be learned, see below). So when an intentional node at a particular ordinal position becomes activated, it induces peaks in the feature fields it projects to, which then drive further processes or actions in the architecture. These peaks also provide input to the condition of satisfaction field that predicts the outcome of the intention (connectivity which may again be learned).

In effect, this system will go through the neural processes associated with each ordinal position in serial order. The processing steps may entail actual




Figure 6.10 A neural dynamic mechanism for serial order in (A) is added to the intention/condition of satisfaction system of Figure 6.9. Circles denote neural dynamic nodes, above threshold when filled, below threshold when open. Gray shading indicates subthreshold activation above resting level. The lower row depicts ordinal intention nodes whose projection onto regions of the intentional dimension (irregular arrows) gives contents to each ordinal step. The upper row is matching memory nodes. Each ordinal intention node activates its memory node (vertical arrow), which preactivates the successor ordinal intention node (diagonal arrows). All ordinal intention nodes are inhibited by the condition of satisfaction field (line with a filled circle at its end). Inhibitory coupling among ordinal intention nodes is not shown. Illustrated is an activation state while the system is in the first step of a serial order task.

motor behavior that may take variable amounts of time. For instance, the agent modelled by Sandamirskaya & Schöner, 2010 was taught a serial order of colors which it then searched for in a new environment. Finding an appropriately colored object at any given step would then take variable amounts of time. During that time, the intention to search for the current color would remain stable against distractors (e.g. objects with colors that are to be searched at other steps in the sequence). A similar demonstration for a robot arm is reviewed in (Tekülve, Fois, Sandamirskaya, & Schöner, 2019). In other cases, the processing steps may be entirely neural, but their duration may still vary depending on activation levels and their distance from instabilities. An example is the building of a mental map by processing spatial relations (Kounatidou, Richter, & Schöner, 2018), in which the time needed to induce an entry into the map depends on how many items are already present (due to inhibition from those). This robustness of sequential processing is critical to scaling such neural dynamic architectures beyond a limited set of demonstrations. Connectionist architectures for serial order do not address this problem of stabilization against variable timing of events. In the classical architectures, time is either discretized so that one item is activated on each step (Elman, 1990) or is based on transient activation patterns that generate a regular pattern of serial recall (Botvinick & Plaut, 2006).

6.4.1 Multi-Layer Fields and More Complex Neural Dynamics

In the brain, neurons make only one type of synapse on their targets, either excitatory or inhibitory. This principle, sometimes referred to as Dale's law, gives the notions of "excitatory" and "inhibitory" neuron their meaning. From the interplay of excitatory and inhibitory populations, more complex neural dynamics emerge that may deliver further cognitive and motor function. Only some basic ideas are reviewed here (see Buonomano & Laje, 2010; Schöner et al., 2019; Sussillo, Churchland, Kaufman, & Shenoy, 2015; Tripp & Eliasmith, 2016 for further reading).

The neural dynamics reviewed up to this point violate, in part, Dale's principle. For instance, the interaction kernel of Equation 6.2 (Figure 6.3) postulates that activation at one field location has excitatory connections to nearby locations and inhibitory connections to locations further removed in the field. In the brain, the inhibitory influence must be mediated by inhibitory interneurons that are excited by the activation field and that, conversely, project inhibitorily onto the activation field, a pairing of excitatory and inhibitory populations. In fact, the model of Equation 6.2 is an approximation of such a more realistic two-layer model (Amari, 1977). The approximation is valid when inhibition is sufficiently fast dynamically, but fails when the time needed to build up inhibition matters. This is relevant to understanding the time course of decision making (Wilimzig, Schneider, & Schöner, 2006), for instance, in which early decisions are influenced more strongly by excitatory input and interaction that promote averaging among inputs, while late decisions are more strongly influenced by inhibitory interaction that promotes selection. Excitatory and inhibitory neural populations also play different roles during learning (see Section 6.5).

More complex arrangements of layers of excitatory and inhibitory neural populations lead to new functions. Inspired by the so-called canonical microcircuit of neocortex (Douglas, Martin, & Whitteridge, 1989), a model with two excitatory and one inhibitory layer has been proposed that accounts for change detection in visual working memory tasks (J. Johnson, Spencer, Luck, & Schöner, 2009; Schneegans et al., 2016). Multilayer structures also account for match and mis-match detection such as those occurring for each examined item in visual search (Grieben et al., 2020). Pairs of excitatory–inhibitory populations may generate time courses, either as active transients or as periodic oscillations. These may be used to model the generation and coordination of movement (see, for instance, Knips, Zibner, Reimann, & Schöner, 2017; Schöner et al., 2019).

6.5 Memory Formation and Learning in Neural Dynamics

Learning is the change of behavior or thought that is driven by experience. In DFT terms, learning is the change of the neural dynamics of a system that is driven by the activation patterns themselves and their sensory-motor

consequences. The simplest forms of such learning from experience are probably sensitization and habituation (Thompson & Spencer, 1966). Sensitization is the lowering of the threshold for a motor behavior or percept over its repeated experience. Habituation is the increase of the threshold across experience. In DFT, these two simple forms of learning can be modeled by the laying down of a memory trace of activation fields. Sensitization is modeled by a memory trace for excitatory fields that locally lifts the resting level making it easier to induce a peak at locations that had previously been activated. Habituation is modeled by a memory trace for inhibitory fields that locally makes it easier to build inhibition and thus more difficult to build peaks in the associated excitatory field.

The mathematical formalization of the memory trace in DFT has taken a variety of forms which are all largely equivalent. The evolution of the memory trace, $u_{\text{mem}}(x, t)$, of an activation field, $u(x, t)$, is described as a dynamical system on the somewhat slower time scale, τ_{mem} :

$$\tau_{\text{mem}} \dot{u}_{\text{mem}}(x, t) = -u_{\text{mem}}(x, t) + \sigma(u(x, t)). \quad (6.4)$$

The memory trace is thus a local low-pass filter of the activation field. The equation must be modified to express the understanding that $\dot{u}_{\text{mem}}(x, t) = 0$ if activation in the field, $u(x, t)$, is nowhere above threshold (see Erlhagen & Schöner, 2002 for a formalization). That means that there is no spontaneous decay of the memory trace, which decays only by interference, that is, decays at locations without activation when at the same time the memory trace builds at other activated locations. More refined models postulate a slightly faster time scale for building the memory trace than for the decay of the memory trace (see Sandamirskaya, 2014, for review). The coupling from the activation field, $u(x, t)$, into the memory trace may be described by a kernel, spreading activation to neighboring sites.

The memory trace couples back into the neural dynamics of the field by providing excitatory input, for example, in this form:

$$\tau \dot{u}(x, t) = -u(x, t) + h + S(x, t) + \int dx' w(x-x') \sigma(u(x', t)) + c_{\text{mem}} u_{\text{mem}}(x, t) \quad (6.5)$$

with coupling strength, c_{mem} (which can be expanded to include a kernel). Typically, the strength of input from the memory trace is small compared to other inputs and to neural interaction, so that the memory trace amounts to a small local adjustment of the resting level. One may thus think of the memory trace as *reshaping* the activation field.

The functional constraints for the dynamics of the memory trace come from accounts of behavioral experiments. The memory trace of excitatory fields was used to account for perseverative reaching in infants (Thelen, Schöner, Scheier, & Smith, 2001) and that work pointed to the absence (or very slow rate) of spontaneous decay. That work also suggested decay of the memory trace by interference (Clearfield, Dineva, Smith, Diedrich, & Thelen, 2009; Dineva & Schöner, 2018). The memory trace of inhibitory fields has been used to account

for infant habituation (Perone & Spencer, 2013, 2014; Schöner & Thelen, 2006). Earlier work on choice reaction times has shown how the memory trace may build estimates of the probability of choices from the frequencies of particular decisions (Erlhagen & Schöner, 2002), consistent with similar signatures in infant motor decision making (Dineva & Schöner, 2018).

From a connectionist perspective, the memory trace is an elaboration of the bias term, an offset to the sum over inputs that each model neuron performs. The bias term plays a limited role in neural network learning because it is just one input in addition to many synaptic inputs to the neuron. In DFT, in contrast, this term plays a much stronger role because the detection instability may amplify small differences in activation into macroscopic suprathreshold peaks. The dynamics of the memory trace does not model associative learning as it strengthens active neural representations irrespective of how they were activated. Associative learning through Hebbian strengthening of connections reflects coactivation of pre- and postsynaptic neural populations. Such a mechanism can also be used within the framework of DFT. The appropriate mathematical formalization makes use of time-continuous learning rules modeled as a dynamical system (Sandamirskaya, 2014), an approach that goes back at least to Grossberg, 1970. For examples of using this form of learning in DFT see Klaes, Schneegans, Schöner, & Gail, 2012; Sandamirskaya & Schöner, 2010; Sandamirskaya & Storck, 2015; Tekülve & Schöner, 2020.

6.6 Relation to Other Approaches

6.6.1 Relation of Dynamic Field Theory to Other Dynamical Systems Approaches

Neural dynamics as formalized in DFT was reviewed in this chapter as a concrete, mathematically specific case study of dynamical systems thinking in cognition. In DFT, meaningful thoughts and actions are generated by attractor states of neural populations whose stability enables linking cognitive processes to sensory-motor systems. Stability is generated by spatially organized neural interaction that erect localist neural representations. Multiple local neural activation patterns can be flexibly bound by such neural interaction within neural dynamic architectures. The time- and state-continuous neural dynamics gives rise to events at discrete moments in time through dynamic instabilities, that can be harnessed to generate sequences of mental or motor acts.

How is DFT positioned relative to other strands of dynamical systems thinking in cognition? The introduction to this chapter provided the embedding of dynamical systems ideas in embodiment. A body equipped with sensors, effectors, linked by a nervous system, and situated in an appropriately structured environment may give rise to meaningful and complex behavior (Braitenberg, 1984). Because behavior is ultimately critical to evolutionary success, one may think of physically embodied cognition as a form of “minimal

cognition,” from which all other forms of cognition may have emerged (Beer, 2000). DFT is consistent with this line of thinking (Schöner, Faubel, Dineva, & Bicho, 2016). DFT makes a distinction, however, between “behavioral” dynamics, in which the physical state of an agent or organism is critical, and “neural” dynamics, to which the physical state may, but need not, contribute. Through neural dynamics, DFT makes use of the notion of representation of thought as simply inner neural dynamic states that shape the evolution of further thought and action (Spencer & Schöner, 2003). (In the philosophy of mind, debates about the sense in which dynamical systems views are compatible with the notion of representation are based on a more nuanced view of representation reviewed, for instance, in Ramsey, 2007).

More radically, neural dynamic thinking as formalized in DFT is based on the hypothesis that embodiment, the evolutionary and developmental link of cognition to behavior, and the properties of cognitive processes that derive from that link, pervade all forms of cognition. The research program is to understand how abstraction from sensory-motor states and invariance against change of the sensory-motor rendering of experience are effortfully achieved by neural processes (for example, by coordinate transforms). This is in contrast to the research program of other approaches to cognition that postulate abstract, invariant representations from the beginning.

Emergence is a related notion used to characterize how specific competences arise once an embodied agent is situated in an appropriate environment. Over development, the demands on the environment may be relaxed as competences arise in ever broader and less specific contexts (Thelen & Smith, 1994). No single component process may be sufficient nor necessary to bring about a competence so that behavioral and developmental transitions may occur in multiple different ways, not following a unique causal path. On the one hand, DFT embraces this notion and provides concrete mechanistic accounts for how emergence in this sense may happen (Schöner, 2014). Near instabilities, for instance, a variety of small contributions to a neural or behavioral dynamics may push the system through a bifurcation and bring about change, which may then be consolidated by learning from experience. The inducing factors need not be causal for the competence in any broader sense. On the other hand, the notion of emergence is sometimes invoked to suggest that cause and effect cannot be identified. As a mechanistic theory, DFT is not aligned with such a view.

Two potential tensions between DFT and other approaches are worth examining. The alignment of DFT with the general role of models of cognition as informed by mathematical psychology is first addressed. The relationship of DFT to other neurally mechanistic approaches to cognition is discussed second.

6.6.2 Does Dynamic Field Theory Deliver Models or Neural Process Accounts?

Conceptually, dynamical systems accounts formalized in DFT are presented as neural process models of cognition. In many cases, including some of the

best known DFT models, the interface to sensory and motor systems is limited to a simple mapping of states of the model to events in the world. For instance, in the DFT account of perseverative reaching (Dineva & Schöner, 2018; Thelen et al., 2001), an intended movement was modeled by a peak in a neural field defined over movement direction. That peak's position at given moments in time was mapped onto the observed movement of the infant's reach toward a matching location. Inputs to the field were modeled by Gaussian functions centered on movement directions specified by putative sources of sensory information. How sense data provide these inputs and how a peak of activation actually drives the hand's movement was not part of the model (although an implementation of the model on a robot vehicle demonstrated that the link to sensory-motor systems can be established, in principle (Schöner, Faubel, et al., 2016)).

Mappings between model and experiment of this form are common in mathematical psychology and connectionist modeling. For DFT models, accounts for psychophysical data based on such mappings are strong when the captured experimental signature is linked to the model's deeper conceptual structure rather than being merely a reflection of judiciously chosen parameter values. The dependence of performance on the metrics of a task was structural in this sense in a number of models as it is directly linked to the interaction kernel. Examples are metric effects in reaction times (Erlhagen & Schöner, 2002), in change detection (J. Johnson et al., 2009), or in visual habituation (Schöner & Thelen, 2006). The dependence of performance on time is also often structural in this sense. Examples are the time courses of decision making (Wilimzig et al., 2006), of perceptual preference (Perone & Spencer, 2013), or of motor biases (Schutte & Spencer, 2009; Schutte et al., 2003). Because DFT models are strongly constrained by the imposed principles of stability, homogeneity (reducing the number of parameters strongly over connectionist models), achieving quantitative fit is not trivial (see Buss & Spencer, 2014; Samuelson, Smith, Perry, & Spencer, 2011 for two insightful case studies and Chapter 15 of Schöner, Spencer, & DFT Research Group, 2016 for a discussion).

Dynamic Field Theory models may be linked more directly to sensory and motor processes. A recent model of visual search (Grieben et al., 2020), for instance, takes visual input from a camera based on feed-forward feature extraction that is consistent with known neural projections. A neural dynamic model for the perceptual grounding of relations is similarly driven by real camera input (Richter et al., 2017, 2021). Both the sensory and the motor interface was physical and real in neural architectures for reaching movements (Bicho, Louro, & Erlhagen, 2010; Knips et al., 2017; Strauss, Woodgate, Sami, & Heinke, 2015). Such models come close to a neural process account in that they can "act out" the modeled behavior and thus prove that the interfaces to sensory-motor systems do not hide unsolved problems (such as when the input to a model neuron is assumed to reflect the detection, segmentation, a shape estimation of a visual object, a rather nontrivial task). Closest to true neural

process models come neuromorphic implementations of DFT architectures on robots with neuromorphic sensors (Kreiser, Aathmani, Quio, Indiveri, & Sandamirskaya, 2018; Milde et al., 2017).

Mapping neural dynamic models onto neural data is another way to constrain the interface between model and experiment. The distribution of population activation is a formalized method to estimate the activation state of neural fields from multiple single unit recordings (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999). The method uses the tuning curves of individual neurons to establish their contributions to a field defined over the probed sensory or motor dimension. This is how a neural dynamic model of population activity in primary visual cortex (Jancke et al., 1999) provided evidence for the neural interaction kernel (see Section 6.2). A neural dynamic model of population activity in motor and premotor cortex (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schöner, & Riehle, 2003) provided evidence for the integration of prior information. Through a neural dynamic model of saccadic selection mapped onto neural activity in superior colliculus, Trappenberg and colleagues have been able to link different components of that model to different subpopulations of neurons (Marino, Trappenberg, Dorris, & Munoz, 2012; Trappenberg, Dorris, Munoz, & Klein, 2001). Voltage-sensitive dye imaging provides neural data sets ideally suited to constrain DFT models this way (Markounikau, Igel, Grinvald, & Jancke, 2010).

6.6.3 Relation of Dynamic Field Theory to Other Neurally Grounded Theories of Cognition

Mathematically speaking, the neural dynamic models of DFT are special cases of general neural network models, characterized by dominant, recurrent connectivity that is organized homogeneously over low-dimensional spaces. The conceptual commitment to attractors as the functionally significant activation states is shared by a line of neural models of spatial orientation that are more strongly neurally mechanistic (reviewed in Knierim & Zhang, 2012). The emphasis on instabilities as the basis for detection and selection decisions, for how the capacity of working memory is limited, and how sequences are generated, is a defining feature of DFT.

The neural fields of DFT can represent continuously many different stable states as localized peaks thanks to their invariant pattern of interaction connectivity. With this localist form of representation, DFT foregoes the higher representational capacity and the associative function of distributed representation (Bowers, 2017). Attractor states in distributed representations arise in Hopfield networks whose neural dynamics have the same form as used in DFT, but whose interaction connectivity is not constrained to low-dimensional kernels (Hopfield & Tank, 1986). That interactive connectivity specifies particular vectors of neural activation as attractors. Hopfield networks may thus represent as attractors specific learned (or memorized) states rather than a range of states that may arise as a stable state for the first time. Hopfield networks also

do not enable targeted instabilities that may drive autonomous cognitive operations of the type reviewed in Section 6.4. The commitment of DFT to localist representations derives from that hypothesized limitation of distributed representations.

Most feed-forward neural networks, including the currently very successful deep neural networks, exploit the power of distributed representations, but use some form of localist representation at read-out, for instance, in the form of a winner-takes-all mechanism. One vision could be that neural dynamics of the DFT type happens at and beyond the classification decisions made in the final layers of feedforward networks. Most cognition does not depend on the continued presence of high-dimensional sensory stimulation. So it is thinkable that autonomous cognitive processing may take place primarily once the high-dimensional sensory information has been left behind. In fact, a possible view is that the generation of sequences of neural attractor states in DFT provides a, perhaps limited, form of symbolic processing that remains consistent with neural principles and with the need to link to sensory and motor systems (for a first step in this direction, see Sabinasz, Richter, Lins, Richter, & Schöner, 2020). In that view, the frameworks of logic-based cognitive processing and information processing would provide descriptions of what the neural processes unfolding in DFT architectures achieve. Probabilistic approaches to cognition could be similarly viewed as descriptions of the integrative function that the strong interaction within neural fields provides. At this time, this vision remains largely speculative.

An alternative to this vision is the framework of vector symbolic architectures (VSA) (Smolensky, 1990). VSAs exploit the property of random, high-dimensional neural activation vectors to be approximately orthogonal to each other. This makes it possible to combine vectors in various ways without losing access to the original component vectors (Gayler, 2003). VSAs thus enable a form of information processing using distributed neural representations. The difficulty of creating and sustaining such neural activation vectors in physiologically plausible neural networks has been viewed as a problem. The neural engineering framework (Eliasmith, 2005) represents such vectors by small populations of integrate and fire spiking neurons (Stewart, Tang, & Eliasmith, 2011), suggesting that VSAs could be implemented in the brain (Eliasmith et al., 2012). To continue to represent the high-dimensional vectors as they are passed from population to population in a neural architecture, the connectivity has to be chosen in a specific way that is informed by the original encoding function. That may raise doubts as to the neural viability of this framework.

6.7 Conclusion

In conclusion, dynamical systems thinking has evolved from its origins in the sensory-motor domain toward capturing increasingly abstract and invariant forms of cognition while retaining the principle of sensory-motor grounding

of cognitive processes. Stable states of neural activation, realized by neural populations localized in low-dimensional neural fields are the units of representation. Their dynamic instabilities lead to the emergence events at discrete moments in time from continuous-time dynamics. These enable sequences of neural processing steps and flexible binding of multiple localist representations within neural dynamic architectures. Research challenges remain to establish (or refute) the capacity of neural dynamic thinking to account for the extraordinary flexibility and productivity of higher cognition.

References

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Andersen, R. A., Essick, G. K., & Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. *Science*, 230(4724), 456–458.
- Ashby, R. W. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall Ltd.
- Bastian, A., Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information pre-shapes the population representation of movement direction in motor cortex. *Neuroreports*, 9, 315–319.
- Bastian, A., Schöner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, 18, 2047–2058.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3), 91–99.
- Bicho, E., Louro, L., & Erlhagen, W. (2010). Integrating verbal and nonverbal communication in a dynamic neural field architecture for human-robot interaction. *Frontiers in Neurorobotics*, 4(5), 1–13.
- 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.
- Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological Review*, 113(2), 201–233.
- Bowers, J. S. (2017). Grandmother cells and localist representations: a review of current thinking. *Language, Cognition and Neuroscience*, 32(3), 257–273.
- Braitenberg, V. (1984). *Vehicles. Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
- Buonomano, D. V., & Laje, R. (2010). Population clocks: motor timing with neural dynamics. *Trends in Cognitive Sciences*, 14(12), 520–527.
- Buss, A. T., & Spencer, J. P. (2014). The emergent executive: a dynamic field theory of the development of executive function. *Monographs of the Society for Research in Child Development*, 79(2), 1–103.
- Chrysikou, E. G., Casasanto, D., & Thompson-Schill, S. L. (2017). Motor experience influences object knowledge. *Journal of Experimental Psychology: General*, 146(3), 395–408.
- Clearfield, M. W., Dineva, E., Smith, L. B., Diedrich, F. J., & Thelen, E. (2009). Cue salience and infant perseverative reaching: tests of the dynamic field theory. *Developmental Science*, 12(1), 26–40.

- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*, 910–923.
- Coomes, S., beim Graben, P., Potthast, R., & Wright, J. (Eds.). (2014). *Neural Fields: Theory and Applications*. New York, NY: Springer Verlag.
- Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *353(1373)*, 1245–1255.
- Dineva, E., & Schöner, G. (2018). How infants' reaches reveal principles of sensorimotor decision making. *Connection Science*, *30(1)*, 53–80.
- Douglas, R. J., Martin, K. A. C., & Whitteridge, D. (1989). Microcircuit for neocortex. *Neural Computation*, *1*, 480–488.
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience Supplement*, *3*, 1184–1191.
- Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor networks. *Neural Computation*, *17*, 1276–1314.
- Eliasmith, C., Stewart, T. C., Choo, X., et al. (2012). A large-scale model of the functioning brain. *Science*, *338(6111)*, 1202–1205.
- Eliasmith, C., & Trujillo, O. (2014). The use and abuse of large-scale brain models. *Current Opinion in Neurobiology*, *25*, 1–6.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179–211.
- 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(1)*, 53–66.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109(3)*, 545–572.
- Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, *61*, 353–430.
- Fauconnier, G., & Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York, NY: Basic Books.
- Fuster, J. M. (1995). *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*. Cambridge, MA: MIT Press.
- Gardenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. Boston, MA: MIT Press.
- Gayler, R. (2003). Vector symbolic architectures answer Jackendoff's challenges for cognitive neuroscience. In P. Slezak (Ed.), *ICCS/ASCS International Conference on Cognitive Science* (pp. 133–138). CogPrints, Sydney, Australia, University of New South Wales.
- Georgopoulos, A. P., Taira, M., & Lukashin, A. (1993). Cognitive neurophysiology of the motor cortex. *Science*, *260(5104)*, 47–52.
- Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge: Cambridge University Press.
- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton Mifflin Co.
- Grabska-Barwińska, A., Distler, C., Hoffmann, K. P., & Jancke, D. (2009). Contrast independence of cardinal preference: stable oblique effect in orientation maps of ferret visual cortex. *European Journal of Neuroscience*, *29(6)*, 1258–1270.

- Griegen, R., Tekülve, J., Zibner, S. K. U., Lins, J., Schneegans, S., & Schöner, G. (2020). Scene memory and spatial inhibition in visual search. *Attention, Perception, and Psychophysics*, 82, 775–798.
- Grossberg, S. (1970). Some networks than can learn, remember, and reproduce any number of complicated space-time patterns, II. *Studies in Applied Mathematics*, XLIX,(2), 135–166.
- Grossberg, S. (2021). *Conscious Mind, Resonant Brain: How Each Brain Makes a Mind*. Oxford: Oxford University Press.
- Henson, R. N. A., & Burgess, N. (1997). Representations of serial order. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Connectionist Representations* (pp. 283–300). New York, NY: Springer Verlag.
- Hopfield, J. J., & Tank, D. W. (1986). Computing with neural circuits: a model. *Science*, 233, 625–633.
- Jancke, D., Erlhagen, W., Dinse, H. R., et al. (1999). Parametric population representation of retinal location: neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience*, 19, 9016–9028.
- Johnson, J., Spencer, J., Luck, S., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science*, 20(5) 568–577.
- Johnson, J. S., Simmering, V. R., & Buss, A. T. (2014). Beyond slots and resources: grounding cognitive concepts in neural dynamics. *Attention, Perception, and Psychophysics*, 76(6), 1630–1654.
- Klaes, C., Schneegans, S., Schöner, G., & Gail, A. (2012). Sensorimotor learning biases choice behavior: a learning neural field model for decision making. *PLoS Computational Biology*, 8(11), e1002774.
- Knierim, J. J., & Zhang, K. (2012). Attractor dynamics of spatially correlated neural activity in the limbic system. *Annual Review of Neuroscience*, 35(1), 267–285.
- Knips, G., Zibner, S. K. U., Reimann, H., & Schöner, G. (2017). A neural dynamic architecture for reaching and grasping integrates perception and movement generation and enables on-line updating. *Frontiers in Neurobotics*, 11(9), 1–14.
- Kounatidou, P., Richter, M., & Schöner, G. (2018). A neural dynamic architecture that autonomously builds mental models. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 1–6).
- Kreiser, R., Aathmani, D., Quio, N., Indiveri, G., & Sandamirskaya, Y. (2018). Organizing sequential memory in a neuromorphic device using dynamic neural fields. *Frontiers in Neuroscience*, 12(717), 1–17.
- Lakoff, G. J., & Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York, NY: Basic Books.
- Latash, M. L. (2008). *Synergy*. New York, NY: Oxford University Press.
- Lipinski, J., Schneegans, S., Sandamirskaya, Y., Spencer, J. P., & Schöner, G. (2012). A neuro-behavioral model of flexible spatial language behaviors. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 38(6), 1490–1511.
- Marino, R. A., Trappenberg, T. P., Dorris, M., & Munoz, D. P. (2012). Spatial interactions in the superior colliculus predict saccade behavior in a neural field model. *Journal of Cognitive Neuroscience*, 24(2), 315–336.
- Markounikau, V., Igel, C., Grinvald, A., & Jancke, D. (2010). A dynamic neural field model of mesoscopic cortical activity captured with voltage-sensitive dye imaging. *PLoS Computational Biology*, 6(9), e1000919.

- Milde, M. B., Blum, H., Dietmüller, A., et al. (2017). Obstacle avoidance and target acquisition for robot navigation using a mixed signal analog/digital neuro-morphic processing system. *Frontiers in Neurorobotics*, *11*(28), 1–17.
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching movement. *Journal of Neurophysiology*, *82*, 2676–2692.
- Oksendal, B. (2013). *Stochastic Differential Equations: An Introduction with Applications* (6th ed.). Berlin and Heidelberg: Springer.
- O'Reilly, R. C. (2006). Biologically based computational models of high-level cognition. *Science*, *314*, 91–94.
- Perko, L. (2001). *Differential Equations and Dynamical Systems* (3rd ed.). Berlin: Springer Verlag.
- Perone, S., & Spencer, J. P. (2013). Autonomy in action: linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive Science*, *37*(1), 1–60.
- Perone, S., & Spencer, J. P. (2014). The co-development of looking dynamics and discrimination performance. *Developmental Psychology*, *50*(3), 837–852.
- Port, R., & van Gelder, R. (Eds.). (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.
- Pouget, A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience Supplement*, *3*, 1192–1198.
- Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge: Cambridge University Press.
- Richter, M., Lins, J., & Schöner, G. (2017). A neural dynamic model generates descriptions of object-oriented actions. *Topics in Cognitive Science*, *9*, 35–47.
- Richter, M., Lins, J., & Schöner, G. (2021). A neural dynamic model for the perceptual grounding of spatial and movement relations. *Cognitive Science*, *45*, e13405.
- Rolls, E. T., Stringer, S. M., & Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proceedings of the Royal Society B: Biological Sciences*, *269* (1496), 1087–1093. <https://doi.org/10.1098/rspb.2002.2009>
- Sabinasz, D., Richter, M., Lins, J., Richter, M., & Schöner, G. (2020). Grounding spatial language in perception by combining concepts in a neural dynamic architecture. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.
- Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. *PloS One*, *6*(12), e28095.
- Sandamirskaya, Y. (2014). Dynamic neural fields as a step toward cognitive neuro-morphic architectures. *Frontiers in Neuroscience*, *7*(276), 1–13.
- Sandamirskaya, Y. (2016). Autonomous sequence generation in dynamic field theory. In G. Schöner, J. P. Spencer, & T. DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory* (pp. 353–368). New York, NY: Oxford University Press.
- Sandamirskaya, Y., & Schöner, G. (2010). An embodied account of serial order: how instabilities drive sequence generation. *Neural Networks*, *10*, 1164–1179.
- Sandamirskaya, Y., & Storck, T. (2015). Learning to look and looking to remember: a neural-dynamic embodied model for generation of saccadic gaze shifts and memory formation. In P. Koprinkova-Hristova, V. Mladenov, & N. K. Kasabov (Eds.), *Artificial Neural Networks*, vol. 4 (pp. 175–200). New York, NY: Springer International Publishing.

- Schneegans, S. (2016). Sensori-motor and cognitive transformation. In G. Schöner, J. P. Spencer, & T. DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory* (pp. 169–196). New York, NY: Oxford University Press.
- Schneegans, S., & Bays, P. M. (2016). No fixed item limit in visuospatial working memory. *Cortex*, 83, 181–193.
- Schneegans, S., & Schöner, G. (2012). A neural mechanism for coordinate transformation predicts pre-saccadic remapping. *Biological Cybernetics*, 106(2), 89–109.
- Schneegans, S., Spencer, J. P., & Schöner, G. (2016). Integrating ‘what’ and ‘where’: visual working memory for objects in a scene. In G. Schöner, J. P. Spencer, & T. DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory* (chap. 8). New York, NY: Oxford University Press.
- Schöner, G. (2014). Dynamical systems thinking: from metaphor to neural theory. In P. C. M. Molenaar, R. M. Lerner, & K. M. Newell (Eds.), *Handbook of Developmental Systems Theory and Methodology* (pp. 188–219). New York, NY: Guilford Publications.
- Schöner, G., Faubel, C., Dineva, E., & Bicho, E. (2016). Embodied neural dynamics. In G. Schöner, J. Spencer, & T. DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory* (pp. 95–118). New York, NY: Oxford University Press.
- Schöner, G., & Kelso, J. A. (1988). Dynamic pattern generation in behavioral and neural systems. *Science*, 239(4847), 1513–1520.
- Schöner, G., Spencer, J. P., & DFT Research Group, T. (2016). *Dynamic Thinking: A Primer on Dynamic Field Theory*. New York, NY: Oxford University Press.
- Schöner, G., Tekülve, J., & Zibner, S. (2019). Reaching for objects : a neural process account in a developmental perspective. In D. Corbetta & M. Santello (Eds.), *Reach-to-Grasp Behavior: Brain, Behavior and Modelling Across the Life Span* (pp. 281–318). Abingdon: Taylor & Francis.
- Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychological Review*, 113(2), 273–299.
- Schutte, A. R., & Spencer, J. P. (2009). Tests of the dynamic field theory and the spatial precision hypothesis: capturing a qualitative developmental transition in spatial working memory. *Journal of Experimental Psychology. Human Perception and Performance*, 35(6), 1698–1725.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory : working memory for locations becomes more spatially precise over development. *Child Development*, 74(5), 1393–1417.
- Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience*, 8(8), 2913–2927.
- Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Shapiro, L. (Ed.). (2019). *Embodied Cognition* (2nd ed.). London: Routledge.
- Simmering, V. (2016). Working memory capacity in context: modeling dynamic processes of behavior, memory and development. *Monographs of the Society for Research in Child Development*, 81(3), 1–158.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1–2), 159–216.

- Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamic 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.
- Stewart, T. C., Tang, Y., & Eliasmith, C. (2011). A biologically realistic cleanup memory: autoassociation in spiking neurons. *Cognitive Systems Research*, 12(2), 84–92.
- Strauss, S., Woodgate, P. J., Sami, S. A., & Heinke, D. (2015). Choice reaching with a LEGO arm robot (CoRLEGO): the motor system guides visual attention to movement-relevant information. *Neural Networks*, 72, 3–12.
- Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7), 1025–1033.
- Tekülve, J., Fois, A., Sandamirskaya, Y., & Schöner, G. (2019). Autonomous sequence generation for a neural dynamic robot: scene perception, serial order, and object-oriented movement. *Frontiers in Neurorobotics*, 13, 208014669.
- Tekülve, J., & Schöner, G. (2020). A neural dynamic network drives an intentional agent that autonomously learns beliefs in continuous time. *IEEE Transactions on Cognitive and Developmental Systems*, 99, 1–12.
- 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.
- Thompson, R. F., & Spencer, W. A. (1966). Habituation: a model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, 73(1), 16–43.
- Trappenberg, T. P. (2010). *Fundamentals of Computational Neuroscience* (2nd ed.). Oxford: Oxford University 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.
- Treisman, A. M. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Tripp, B., & Eliasmith, C. (2016). Function approximation in inhibitory networks. *Neural Networks*, 77, 95–106.
- Van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Brain and Behavioral Sciences*, 21, 615–665.
- Wilimzig, C., Schneider, S., & Schöner, G. (2006). The time course of saccadic decision making: dynamic field theory. *Neural Networks*, 19(8), 1059–1074.
- Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12, 1–24.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636.