Provided for non-commercial research and education use. Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright



New Ideas in Psychology 26 (2008) 227-251

NEW IDEAS IN PSYCHOLOGY

www.elsevier.com/locate/newideapsych

Moving to higher ground: The dynamic field theory and the dynamics of visual cognition

Jeffrey S. Johnson^{a,*}, John P. Spencer^{a,b}, Gregor Schöner^c

^aDepartment of Psychology, University of Iowa, USA ^bIowa Center for Developmental and Learning Sciences, University of Iowa, USA ^cRuhr-Universität, Bochum, Germany

Available online 29 August 2007

Abstract

In the present report, we describe a new dynamic field theory that captures the dynamics of visuospatial cognition. This theory grew out of the dynamic systems approach to motor control and development, and is grounded in neural principles. The initial application of dynamic field theory to issues in visuo-spatial cognition extended concepts of the motor approach to decision making in a sensori-motor context, and, more recently, to the dynamics of spatial cognition. Here we extend these concepts still further to address topics in visual cognition, including visual working memory for nonspatial object properties, the processes that underlie change detection, and the 'binding problem' in vision. In each case, we demonstrate that the general principles of the dynamic field approach can unify findings in the literature and generate novel predictions. We contend that the application of these concepts to visual cognition avoids the pitfalls of reductionist approaches in cognitive science, and points toward a formal integration of brains, bodies, and behavior. © 2007 Elsevier Ltd. All rights reserved.

Keywords: Dynamic systems theory; Dynamic field theory; Visual cognition; Change detection

1. Introduction

More than a decade ago, researchers and theorists from a range of disciplines called for a re-thinking of cognition and adaptive behavior using the tools and concepts of dynamical systems theory (Beer, 1995; Kelso, 1995; Thelen & Smith, 1994; van Gelder, 1995). In

*Corresponding author.

0732-118X/\$-see front matter © 2007 Elsevier Ltd. All rights reserved. doi:10.1016/j.newideapsych.2007.07.007

E-mail address: jeffrey-johnson-2@uiowa.edu (J.S. Johnson).

contrast to the information-processing approach in cognitive science, proponents of this new direction proposed that cognition could be viewed as arising from the adaptive realtime activity of a complex organism embedded in a physical environment. Since that time, a number of dynamic systems approaches to cognition have emerged (e.g., Tabor, in press; Townsend & Busemeyer, 1995; van Geert, 1998; van Gelder, 1998). The wealth of peerreviewed journal articles, books, chapters, and edited volumes that have appeared over the last decade attests to the power of the central concepts and rich tools available. In our view, the field en masse has shown that dynamical systems concepts can effectively be applied to brain and behavior and, furthermore, that the dynamic systems approach constitutes one of the most promising new directions in the study of behavior.

Although this new direction is promising, there is also a considerable amount of diversity among dynamical systems approaches. In our view, such diversity is good, but it also makes it challenging to compare and contrast different viewpoints. Because of this, we have opted to focus on one thread in the larger body of "dynamic systems approaches"— the motor approach that emerged from the work of Kelso, Thelen, Schöner, and others (Kelso, 1995; Kelso, Scholz, & Schöner, 1986; Schöner & Kelso, 1988a; Thelen & Smith, 1994). In particular, our goal for the present paper is to trace how we have begun to tackle a central challenge to this approach put forth decades ago: this approach is great for motor systems, but can it provide a useful framework for understanding cognition?

2. The dynamic systems approach to motor control and development

The approach to visual cognition described in the present article grew out of a program of research and theory that has come to be known as the 'motor approach' (see Spencer et al., 2006; Spencer & Schöner, 2003). This approach originated in the work of Scott Kelso, Gregor Schöner, and colleagues looking at coordinated motor activity (Kelso, 1995; Kelso et al., 1986; Schöner & Kelso, 1988a, 1988b), and the work of Esther Thelen and her colleagues on infant sensori-motor development (Thelen & Smith, 1994; Thelen & Ulrich, 1991).

According to the motor approach, human behavior reflects a dynamic balance among *stability*, *instability*, and *flexibility*. In dynamic systems terms, a system is said to be stable when it settles into one of its preferred 'attractor' states, wherein behavioral or neural states can persist in the face of systematic or random perturbations (Braun, 1994). The ability to maintain such stable states (e.g., stably reaching for a cup of coffee despite a bump on the arm) is a key component of ongoing adaptive behavior. However, of equal importance is the ability to flexibly change behavior when the circumstances require it. To achieve this, the current stable state must be destabilized so that a new stable state (i.e., a new behavior or mental state) can arise. In dynamic systems terms, a change of a system that leads a particular state to become unstable is referred to as an instability (see Braun, 1994, Chap. 4).

A second idea central to the motor approach is that behavior is *softly assembled* from multiple component processes, rather than being 'programmed' and micro-managed by a centralized controller (see discussion in, Clark, 1997; Thelen & Bates, 2003). That is, organized behavior reflects a temporarily stable assembly of various factors that happen to be available at a given point in time, and that can be created and dissolved as circumstances change. In the case of reaching for an object, these factors include neural plans to move the hand to a particular location in space, oscillatory processes that regulate

the timing of the movement, forces generated by the contraction of muscles to move the hand through space, as well as gravitational forces acting on the limbs. Third, such processes are often *non-linear* in nature, where small changes in one component (movement speed at the end of a reach) can lead to big changes in behavior (knocking the cup over). Finally, the dynamics of action change over *multiple time scales*, from the millisecond-to-millisecond unfolding of a reach, to a sequence of actions in context—to a coffee cup, then the computer, and then a stack of papers—to the many reaches to coffee cups made over weeks, months, and years.

Using these basic concepts, the motor approach has shown an impressive ability to provide a detailed 'collective' picture of behavior. This feature has made the approach attractive to experimentalists. As a result, conceptual theory (such as notions of attractors and their disappearance through instabilities, the ideas of emergence and self-organization) as well as formal theory (based on the mathematics of stochastic differential equations) have had close ties to experimental work, both qualitatively and quantitatively. Moreover, novel predictions have been possible in some cases (Schöner, 1989; Schöner & Kelso, 1988b), and several formal models of phenomena have been proposed that are consistent with all known facts (Schöner, Haken, & Kelso, 1986).

Although the motor approach has been shown to be a suitable framework for thinking about issues in motor control and development, in its original form it is not particularly well suited for addressing aspects of higher-level cognition. In particular, the motor approach fails to capture the nature of the representational states that underlie behavior (for in-depth discussion of our use of the 'R-word', see, Spencer & Schöner, 2003). To illustrate this limitation, consider three different situations where a person must search for a book on a table. In one case, the book is clearly visible and has a unique jacket clearly distinguishing it from the other books on the table; in another case, the book is clearly visible, but is surrounded by similarly colored 'distractor' books; in a third case, vision of the book is obstructed by a stack of other books. Further, let us assume that in each of these situations, the person makes an identical movement—an accurate, efficient reach that successfully makes contact with the desired book. Although the concepts of the motor approach can be used to characterize the resultant stable reach, this approach fails to capture differences in the representational states underlying these movements. In particular, in the first case, there was a high degree of certainty regarding where to move, whereas in the second case the decision to move was much less certain, and in the third case the movement was based on a memory of the book's location, rather than on-line visible information. Characterizing such differences in the representational states underlying behavior stands as a central challenge for any dynamic systems approach to higher-level cognition.

In the present paper, we highlight several recent extensions of the motor approach that address the challenge of developing a dynamic systems approach to cognition. We begin by describing the Dynamic Field Theory (DFT), which combines the concepts of activation and neural dynamics to implement basic cognitive processes, overcoming the limitations of the motor approach. Next, we trace the historical development of this framework, from decision making in a sensori-motor context to the dynamics of spatial cognition. Finally, we discuss a series of recent projects that seek to move the DFT to higher ground, addressing visual working memory (VWM) for non-spatial object properties (e.g., color or orientation) and the process of change detection, as well as the binding of multiple non-spatial features into object representations.

3. Towards a dynamic systems approach to cognition

To extend the concepts of the motor approach into cognition, we have developed a new theoretical framework-the DFT-which is in a class of bi-stable neural networks first developed by Amari (1977; see also, Wilson & Cowan, 1972). This approach provides a basis for characterizing the representational states underlying behavior through the concept of activation, which is familiar from mathematical psychology, connectionism, and theoretical neuroscience (Churchland & Sejnowski, 1992; Williams, 1986). In our dynamic field approach, this concept takes the form of an activation field, defined over the metric dimension represented. In the example of reaching for a book described above, activation might be distributed across the dimension of reachable locations, a continuous metric dimension stretching from a far left location to a far right location. A localized peak of activation within this field indicates that a target object (a book) has been detected at a particular location. Such a peak might be built up via perceptual input that specifies where the cup is located within reachable space. Importantly, just as an unfolding action can be described using the dynamic systems concept of stability, patterns of activation can live in different attractor states: a resting state; an input-driven state where input can form stabilized "peaks" of activation in the field, but peaks go away when input is removed; and a self-sustaining state where activation peaks remain stable even in the absence of inputa form of working memory central to our work.

These attractor states are made possible through strong interactions among neurons within and between different fields. Specifically, stable patterns of activation are achieved through a locally excitatory/laterally inhibitory form of interaction among neurons. According to this type of interaction, neurons that respond selectively to similar values along the relevant dimension (e.g., similar reachable locations) excite one another whereas neurons that respond to very different values (e.g., different reachable locations) inhibit one another. This form of interaction allows self-sustained peaks of activation to be maintained in the absence of continuing input (i.e., when the target object is occluded or moves out of view). Critically, movement in and out of these attractor states is softly assembled in real-time depending on a variety of factors. For instance, activation patterns in dynamic fields can "rise above" the current input pattern, achieving a self-sustaining state where subsequent inputs are suppressed. That said, this state is still open to change: in the presence of continued input, the network might "update" its decision to focus on one item over another. This points toward flexibility—how activation patterns can smoothly and autonomously go from one stable state to another.

In addition to retaining many of the central concepts of the motor approach, the dynamic field approach has maintained a tight interface between theory and experiment. Although the DFT was originally developed as an account of simple motor decisions in cued saccadic eye-movement experiments, the approach has since been broadly applied to problems in movement preparation (Erlhagen & Schöner, 2002; Schutte & Spencer, 2007), including how pre-information is integrated with current input to produce reaching responses. Moreover, the DFT has been used to re-think several important issues in the development of reaching behavior (Thelen, Schöner, Scheier, & Smith, 2001) and visual perception (Schöner & Thelen, 2006), as well as to spatial discrimination (Simmering, Spencer, & Schöner, 2006, in press), and the perception of motion patterns (Hock, Schöner, & Giese, 2003).

Additionally, the dynamic field approach has extended the motor approach through the inclusion of representational states that have close ties to neurophysiology. For example,

recent approaches (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999) have demonstrated that the state of dynamic fields can be directly estimated from firing rates of populations of cortical neurons using population coding ideas pioneered by Georgopolous and colleagues (see review in Georgopoulos, 1995). These techniques have been used to directly observe, for instance, the patterns of neural activation underlying movement direction in motor and premotor cortex, providing evidence of preactivation when precues are given (Bastian, Riehle, Erlhagen, & Schöner, 1998). Note also that the general principles of dynamic fields can be usefully integrated with more biophysically realistic approaches that attempt to incorporate the details of neruotransmitter action, timing properties of neurons, and so forth (Compte, Brunel, Goldman-Rakic, & Wang, 2000).

Other applications have demonstrated that the DFT is committed to, and can answer the challenges inherent in, the development of an *embodied* account of cognition. By this view, cognition is always tightly linked to the sensory and motor surfaces and the structured environments in which behaviors are immersed. Embodiment is often painfully evident in infancy, but behavioral evidence from Barsalou (1999), Spivey (Spivey-Knowlton, Tanenhaus, Eberhard, & Sedivy, 1998), and others demonstrates that even the abstract concepts formed by adults for use in "off-line" reasoning are nonetheless intimately linked to sensori-motor systems. Neural evidence from Damasio and Damasio (1994) paints a similar picture, suggesting that the brain runs 'simulations' of past events, effectively reactivating perceptual and motor systems during many cognitive tasks.

Schöner and colleagues (Iossifidis & Schöner, 2006; Steinhage & Schöner, 1998) have demonstrated that the DF approach can realize an embodied view of cognition by using dynamic fields to organize the behaviors of autonomous robots. Such robots are able to autonomously orient towards and retrieve target objects in cluttered real-world environments while avoiding obstacles that can appear and/or change position unexpectedly. In addition, these robots are able to autonomously generate goal-directed sequences of behaviors. This impressive behavioral flexibility is achieved using multiple coupled dynamic fields together with dynamical systems typically used within the motor approach (e.g., dynamical models that govern the control of the robot's effectors). Such examples provide a concrete demonstration that the DFT can be seamlessly integrated with the motor approach to provide a unified view of cognition *and* action.

4. The DFT of spatial cognition

Although the DFT has been highly generative, most applications have addressed decision making in a sensori-motor context. In this section, we describe more recent developments that move beyond the sensori-motor case toward more typical memory paradigms. This phase of development started within the domain of spatial cognition (Schutte, Spencer, & Schöner, 2003; Spencer & Schöner, 2003; Spencer, Simmering, Schutte, & Schöner, 2007). This was an ideal next step because spatial systems are still clearly linked to sensori-motor systems, but their treatment required us to address several new issues, including working memory for locations, long-term spatial memory, perceptual reference frames including egocentric and allocentric encoding, and so forth.

Fig. 1 shows a simulation that implements our DFT of spatial cognition (Spencer et al., 2007). The full model of spatial cognition consists of seven layers, two of which are concerned with the transformation of visual inputs from retinal to object-centered coordinates, an issue that we ignore here for the sake of brevity. Therefore, the simulations



Fig. 1. A simulation of the dynamic field theory performing one spatial recall trial. In each panel, location is across the x-axis, activation on the y-axis, and time on the z-axis. The model consists of five layers: (A) a perceptual field; (B) a long-term memory field associated with this perceptual field; (C) a shared layer of (inhibitory) interneurons; (D) a spatial working memory field; and (E) a long-term memory field associated with the spatial working memory field. Solid arrows show excitatory connections between layers, and dashed arrows show inhibitory connections between layers. Brackets describe the behavioral functions of subsets of the layers. See text for additional details.

we focus on consist of five-layers: A perceptual field (PF, Fig. 1A), a long-term memory field associated with the perceptual field (LTM_{PF} , Fig. 1B), a layer of inhibitory interneurons (Inhib, Fig. 1C), a spatial working memory field (SWM, Fig. 1D), and a long-term memory field associated with the working memory field (LTM_{SWM} , Fig. 1E). In each case, the direction of the targets in the task space is shown along *x*, *y* shows the activation of each site in the field, and *z* captures the elapsed time from the start of the trial.

In this section, we discuss two central issues that lay the groundwork for the sections that follow. These concrete examples illustrate how the dynamical concepts discussed above (i.e., stability, instability, flexibility, etc.) are realized in our approach to cognitive processes. First, we discuss the integration of perception and working memory, focusing on the three layers that form the centrepiece of the model shown in Fig. 1 (PF, Inhib, and SWM). These three layers contribute to an important class of effects seen in the spatial

cognition literature—geometric biases in working memory for location. Second, we discuss the integration of shorter and longer time scales, focusing on the integration of working and long-term memory and an associated class of effects—experience-dependent biases in location memory.

5. Integration of perception and working memory: geometric biases in spatial recall

One of the most common tasks used in the spatial cognition literature is spatial recall. In this task, participants are briefly shown a single target item, the location of which must be remembered. This is followed by a delay interval of variable length (0-20 s), after which the participant indicates the remembered location of the target (e.g., by moving a computer mouse or a joystick). A number of studies of this type have demonstrated that metric memory for location shows delay-dependent biases away from perceived reference axes, such as the midline symmetry axis of a task space (Schutte & Spencer, 2003; Spencer & Hund, 2002).

The simulation shown in Fig. 1 illustrates how "drift" away from perceived reference axes is accounted for within the DFT. At the start of the trial, PF builds a small peak of activation at 180°, reflecting perception of a salient reference frame in the environment (e.g., two dots along the midline of the task space). Next, a target appears at 220°. This creates a peak of activation centered at that location. When the target disappears, a peak re-forms in PF at 180° as the system re-locks onto the reference cues in the task space. This is important because it allows the system to continually calibrate egocentric and object- or table-centered frames of reference during the delay (Spencer et al., 2007).

Panel D shows the effect of coupling PF to SWM. At the start of the trial, SWM receives relatively weak reference input from PF because the reference cue in the task space (i.e., the reference input to PF) is not very salient. Next, the target is turned on, passing strong target-related input into the working memory field. This event moves the working memory field into a strongly self-sustaining state where peaks of activation can be maintained in the absence of input. This maintenance arises due to the reciprocal coupling between the SWM and Inhib fields, which implements the locally excitatory and laterally inhibitory form of interaction described previously. Importantly, this occurs during the memory delay even though PF has re-acquired the reference frame. This sets the stage for memories to drift away from the reference frame because the system is effectively holding onto two stable states at once: it is locking onto the reference frame to stay calibrated with the world and it is actively maintaining a memory of the target location. The response drift evident in spatial recall experiments arises from the place where these two stable states meet-the shared inhibitory layer. In particular, reference-related input to the inhibitory layer near midline causes the peak of activation in SWM to "drift" away from the midline of the task space because there is stronger inhibition on the midline-side of the SWM peak than the "outer" side. In effect, the reference frame "pushes" the peak outward, exaggerating the "leftness" of the target location.

Importantly, the strength of the push depends on the proximity between the target location and the reference frame and the duration of the memory delay. For example, Spencer and Hund (Hund & Spencer, 2003; Spencer & Hund, 2002, 2003) have shown that the magnitude of the bias depends on the metric separation between the remembered location and the reference axis. Specifically, memory for targets presented at midline was found to be highly accurate, whereas memory for items presented to the left and right of

midline was biased away from this axis, with the greatest bias seen for targets $20-30^{\circ}$ from midline and smaller bias as the items were moved further and further away from the axis of symmetry. Additionally, we have shown that the magnitude of reference-related bias grows larger as the delay interval between presentation of the target and test is increased (Spencer & Hund, 2002).

6. Integration of short and long time scales: experience-dependent biases in spatial recall

The simulation shown in Fig. 1 also illustrates how shorter and longer-term time scales are integrated in the model. Panels B and E show LTM fields that are associated with the PF and SWM fields, respectively. Whenever there is above-threshold activation in PF or SWM, traces are left in their associated LTM fields. Reversely, activation traces in LTM feed back as excitatory input into each of these fields. Importantly, each of the LTM fields has a longer time scale. Thus, activation grows slowly in these fields and decays slowly in the presence of competing input. As can be seen in Fig. 1B, the dynamic interaction between PF and LTM_{PF} results in a robust trace of the midline reference frame at 180° . Importantly, there is virtually no trace of the target location. This reflects the fact that the target input was transient whereas the reference input was visible in the task space for most of the trial. Thus, LTM_{PF} retains a memory of the perceptual reference frame, which can be used to realign reference frames across trials (Spencer et al., 2007).

In the LTM_{SWM} field shown in Fig. 1E, however, the traces that are laid down reflect the distribution of targets seen during the session, rather than the perceptual reference frame. When the distribution of activation in LTM_{SWM} is centered around the target location, excitatory feedback to the SWM field can help to stabilize WM peaks against drift. However, when the model is presented with several close locations in a row and is then shown a target at a different (but not too distant) location, the WM of the new location is biased towards the previously responded-to targets. Attraction of working memory peaks toward the LTM traces of other targets arises as the LTM_{SWM} field accumulates traces of activation reflecting the distribution and frequency of targets responded to over the course of an experimental session. These traces then serve as an extra source of excitatory input back into SWM fields, producing attraction of working memory peaks towards the location of frequently responded to targets.

Results from several studies confirm the presence of systematic experience-dependent distortions in memory for spatial locations. For example, 6-year-olds, 11-year-olds, and adults make response errors in the direction of previously responded-to locations when recalling locations in the absence of salient location cues (Spencer & Hund, 2002, 2003). This effect varies as a function of target frequency, with more frequently responded-to locations exerting more of an attractive effect. Thus, the reciprocal interplay between SWM and LTM_{SWM} leads to the emergent formation of experience-dependent spatial categories that can influence subsequent performance in the task.

7. A dynamic field approach to non-spatial feature working memory

The DFT of spatial cognition provides a neurally plausible account of characteristic biases seen in spatial recall tasks. Although significant progress has been made with this model of spatial cognition, the type of cognition addressed is still strongly tied to action in space. Thus, an important goal is to extend this framework to address working memory for non-spatial metric feature dimensions such as color and orientation. To capture this, we have introduced the concept of a feature WM field (FWM), which has all of the characteristics of the SWM field described above, but the metric dimension along which activation is defined is featural rather than spatial in nature (e.g., hue, orientation, line length). Note that the FWM field captures more than just a re-labeling of an axis in our model. The claim here is that WM for metric features shares all of the properties captured by dynamic neural fields—WM as stabilized peaks of activation, coupling among perception, WM, and LTM, metric interactions leading to "drift", and so on. Do we see the same signatures of dynamic fields observed in spatial cognition when examining non-spatial feature working memory?

To investigate this possibility, we have conducted a recent experiment examining experience-dependent drift in working memory for color using the color estimation task depicted in Fig. 2A. For each trial, participants viewed a 2-s memory display containing a single colored triangle presented at a fixed location at the center of the screen. This was followed by a variable delay interval (0, 5, 10, or 15 s), and the appearance of a 110° colored arc containing the target color. To prevent spatial coding of the color, target colors were randomly positioned at one of eight possible angular positions within the arc.



Fig. 2. (A) Color estimation task used to probe WM for color. For each trial, participants viewed a memory display containing a single colored target, followed by a variable delay interval (0, 5, 10, or 15s), and the appearance of a colored arc containing the target color. Participants responded by moving a mouse cursor to the position on the arc that matched the color they were holding in memory. (B) Continuous color space used to generate each of the colors used in the experiment. Throughout the experiment, different groups of participants saw an identical 'B' target paired with an 'A' target that was either close or far in color space, and either clockwise (CW) or counterclockwise (CCW) from the 'B' target. See text for additional details.

Participants responded by moving a mouse cursor to the position on the colored arc that matched the color they were holding in memory.

To look for delay- and experience-dependent biases in WM for color, separate groups of participants were required to estimate either Close colors (30° separation in color space) or Far colors (80° separation). Participants only saw two different colors throughout the experiment: a 'B' color that remained the same across the Close and Far conditions, and an 'A' color that was either clockwise (CW) or counterclockwise (CCW) in color space from the 'B' target (see Fig. 2B). If WM for non-spatial features uses the same dynamic neural principles captured in our model of SWM, then participants in each condition should form LTM traces of the target colors across trials, but the nature of the LTM traces should differ across conditions. In the context of our model, the activation traces in LTM should then serve as a source of feedback excitation to WM, producing systematic distortions in WM for color. For example, in the Close condition, LTMs of the target colors should blend together, leading to strong activation around an average remembered color (i.e., somewhere midway between the 'A' and 'B' colors seen throughout the session). Consequently, WM for the 'B' target in the Close conditions should be attracted toward the 'A' target across delays, and thus, in opposite directions depending on the direction of 'A' (CW or CCW). By contrast, LTMs in the Far condition should be distinct, and there should be no systematic trends as a function of the direction of the 'A' target (CW, CCW) across conditions. Note that these predictions hold even though the B color is identical across Close and Far conditions.

Results have confirmed the central predictions of the DFT: estimates of the 'B' color drifted toward the 'A' colors over delays in the Close conditions (Johnson, Spencer, & Schöner, 2007). This was not the case in the Far conditions. These data suggest that properties central to the dynamic field approach to spatial cognition generalize to non-spatial working memory systems.

8. A DFT approach to multi-item VWM and change detection

The findings discussed above suggest that working memory for color does indeed share many of the same properties as SWM, including the delay- and experience-dependent drift of WM peaks towards LTM traces. However, the color memory task only required participants to remember a single item across the delay interval, and to estimate that color at test. This task is quite different from the standard task used to probe VWM: change detection. A typical change detection task is depicted in Fig. 3. In such tasks, observers are shown a *sample array*, which contains one or more items (e.g., colored squares) that the observer must try to remember. After a brief *delay interval* (e.g., 1s), a *test array* is presented, and observers compare the test array with the sample array. In most experiments, the test array is identical to the sample array on 50% of trials and differs in the value of a single item (e.g., one square changes color) on the remaining trials.

Research using this paradigm has revealed a number of important properties of VWM (for in-depth discussion of the use of change detection tasks in the study of VWM, see Luck, in press; Rensink, 2002). First, information is encoded very rapidly in VWM, with estimated encoding rates of approximately 20–50 ms/item (Gegenfurtner & Sperling, 1993; Shibuya & Bundeson, 1988; Vogel, Woodman, & Luck, in press). Second, VWM has a very limited storage capacity of approximately three to four items (Cowan, 2001; Luck & Vogel, 1997; Phillips, 1974), and, in many cases, appears to store items in the form of



Fig. 3. Change-detection paradigm used to investigate working memory for spatial and non-spatial object properties. Participants are shown a sample array containing a variable number of colored squares, which they must hold in memory across a brief delay interval (different textures on each of the squares represent different colors). This is followed by the appearance of a test array that is either identical to the sample array or differs from it in the color of a single item. When the test array appears, participants make an unspeeded two-alternative-forced-choice response, indicating whether the colors in the test array are the same as or different than the colors that were present in the sample array. Adapted from Luck and Vogel (1997).

integrated object representations rather than as individual features (Luck & Vogel, 1997; although see Olson & Jiang, 2002; Vogel, Woodman, & Luck, 2001; Wheeler & Treisman, 2002; Xu, 2002, 2004 for important qualifications of these results). Additionally, electrophysiological and functional Magnetic Resonance Imaging (fMRI) studies of change detection have begun to isolate the neural substrates that underlie these characteristics of VWM (e.g., Pessoa & Ungerleider, 2004; Todd & Marois, 2004; Vogel & Machizawa, 2005; Xu & Chun, 2006).

Although studies of change detection have begun to make significant contributions to our understanding of VWM at both the behavioral and neural levels, few theoretical models have been formulated within a neurally plausible framework that could effectively address both lines of research. In particular, no current theory addresses both the maintenance of information in VWM *and* the process by which the test array is compared with the contents of WM, generating the same and different responses required by the task. In this section, we describe a new dynamic field model of multi-item working memory and change detection that provides a neurally plausible framework for addressing both of these issues. In generalizing the DFT to handle these issues, we needed to address two concrete challenges: First, we needed to scale up from single-item memory to allow the simultaneous maintenance of multiple items in WM; second, we needed to develop a process-based account of comparison/change detection. These are treated in turn in the following sections.

9. Multi-item VWM and capacity limits

As discussed above, VWM has been found to have a limited capacity of around three to four items. How many items can be maintained concurrently in the DFT? Recall that the DFT is in a class of models that use a locally excitatory/laterally inhibitory form of interaction to sustain activation in WM. When lateral inhibition is broad (i.e., global) and relatively strong, such models implement a winner-take-all process, whereby only a single localized peak of activation can be sustained at a given time. That is, they have a capacity

of a single item. Thus, as Trappenberg (2003) has pointed out, such networks often show too severe a limit in VWM capacity. From this view, the question that must be answered is not "why is the capacity of WM so small?", but how can capacity be increased to match the behavioral findings?

In the dynamic field model described here, we achieve multi-item memory by using a "Mexican Hat" form of interaction among neurons, where inhibition is much stronger near the focus of excitation than at more peripheral sites. This form of interaction allows the locally excitatory interactions associated with each peak to be isolated by lateral inhibition, while keeping the total amount of inhibition in the field low enough that multiple items can be maintained simultaneously. However, as more items are added to working memory (i.e., more peaks are added to the field), the overall amount of inhibition is also increased, which, together with metric interactions among peaks, provides a natural basis for capacity limits.

This functionality is depicted in the simulations shown in Fig. 4A–C. The simulation in Fig. 4A shows the formation of three stable peaks of activation in FWM in response to the appearance of a visual display containing multiple colored objects. These peaks are retained in FWM throughout a subsequent delay interval even though the display has been removed. In Fig. 4B, a fourth item is added to the display and, once again, all items in FWM are stably maintained throughout the delay in the absence of input. However, when a fifth item is added to the display (Fig. 4C), inhibition in FWM begins to outweigh excitation, and two of the peaks are suppressed, leaving only three peaks that survive the delay interval. Together, these simulations demonstrate multi-item working memory in the



Fig. 4. Multi-item working memory and capacity limits in the DFT using a "Mexican Hat" form of interaction, where inhibition is much stronger near the focus of excitation than at more peripheral sites: (A) stable maintenance of three items (e.g., three colors) in WM in the absence of input. WM remains stable when a fourth item is added (B), but the addition of a fifth item (C) raises the amount of inhibition in the field beyond a critical point, which suppresses all but three of the peaks in WM, producing the capacity limits found in behavioral studies.

DFT, and provide a natural, neurally plausible explanation for the existence of capacity limits found in behavioral tasks.

10. The comparison process in change detection

The next issue we need to address is the process by which the contents of WM are compared to available perceptual information (i.e., the test array). The starting point for our investigations of this issue is the observation that comparing an item in working memory to another item present in the environment involves the integration of perception and WM. Recall that the PF, Inhib, and SWM fields at the heart of the dynamic field model of spatial cognition described previously serve exactly this role, allowing the contents of WM to be encoded within a perceptually available reference frame. In the present context, this raises an interesting question: Could the same architecture that deals with the integration of perception and WM in space generalize to handle the integration of perception and WM during comparison in change detection tasks?

To explore this possibility, we have extended a recent dynamic field model of position discrimination (Simmering et al., 2006). This model has the same structure as the DFT of spatial cognition described above, with the exception that the metric dimension along which activation is distributed is a non-spatial feature such as color. The model consists of a feature-selective perceptual field that provides afferent input to a layer of inhibitory neurons and to an excitatory FWM. Excitatory and inhibitory interactions among these three fields allow the network as a whole to function as a "difference" detector. Specifically, peaks in WM activate similarly tuned neurons in the Inhib field via excitatory feedback, and these neurons in turn send inhibitory feedback to PF. This inhibitory input to PF suppresses the firing of neurons in PF that are tuned to the values currently activated in FWM. Because of this, PF is only able to build a new peak of activation when a change occurs at test, that is, when a new item appears in the task space. This leads to a natural basis for "same" and "different" responses: when a peak is present in PF at test, the model responds "different", and, conversely, when no peaks are present in PF but peaks remain in FWM, the model responds "same".

These ideas are illustrated in the simulations shown in Fig. 5A and B. Both simulations show three peaks of activation that are built following the presentation of a sample array (e.g., three colored squares). Note that these peaks are only transiently sustained in PF (see A1 and B1 in Fig. 5A and B, respectively), but are maintained throughout the delay interval in the FWM field (see A2 and B2). Each of the peaks in WM activates similarly tuned neurons in the Inhib field (see A3 and B3), which then projects localized inhibition back to PF, inhibiting similarly tuned neurons in that field (see A4 and B4). At the end of the delay, a test array is presented to probe WM for color. In panel A, the test items are the same as the items being held in working memory. As a result, each of the peaks in PF remain below threshold at test due to strong localized inhibition from the Inhib field at those locations (A5). Thus, the three peaks remain stable in WM at test and a "same" response is produced (A6). In panel B, however, the test display contains a new color that was not present in the memory array; thus, this novel input comes in at a relatively uninhibited region of PF. In this case, a self-stabilized peak builds in PF (B5), which suppresses the peak associated with the changed item in WM (B6). The presence of a peak in PF at the end of the response interval (B5) leads to a "different" response.



Fig. 5. The generation of "same" (A) and "different" (B) responses in a dynamic field model of VWM and change detection. The presentation of a stimulus array (e.g. three colored squares) leads to the formation of multiple peaks of activation in each layer. However, the peaks in PF decay once the stimulus is removed, whereas they are maintained throughout the delay interval in FWM. Each of the peaks in WM activate similarly tuned neurons in Inhib, which sends inhibitory feedback to PF, suppressing the firing of similarly tuned neurons in that field. As a result, PF is only able to build a peak when a change occurs at test (B5), which serves as the basis for a "different" response. See text for additional details.

11. Novel predictions in change detection

The proposed model represents the first neurally plausible theory of VWM that addresses both the maintenance of information in VWM and the process of change detection. When we combine these new ideas with our understanding of the metric- and delay-dependent interactions underlying VWM for colors (see above), we can go a step further, generating novel and counterintuitive predictions in the context of change detection. Recall that the local excitation/lateral inhibition function underlying sustained activation in the DFT can lead to interactions between peaks when more than one item is being held in WM. The specific form of the interaction depends critically on how similar the items are along a given dimension (e.g., color). One consequence of such interaction is depicted in Fig. 6. Panels A and C show a time slice through the FWM field during the delay interval of a change detection task and the associated perceptual fields (panels B and D). As can be seen in panel A, when targets are far apart in color, the peaks of activation associated with each color in FWM do not interact. Consequently, the resulting peaks are fairly wide and relatively high energy (strong excitation and strong inhibition). This produces a region of inhibition in the perceptual field that is also fairly broad and deep. In contrast, when the peaks in FWM are very near to each other (i.e., very similar colors), as they are in panel B, the inhibition associated with one peak extends to the "far side" of the other peak, making each peak sharper, and producing a much narrower region

240



Fig. 6. Metric interactions in WM leading to enhanced change-detection for close features: (A) time-slice through the FWM field during the delay interval of a change detection task showing WM peaks representing two far color targets separated by 160 units. Relatively broad and high-energy peaks in WM produce correspondingly broad and deep inhibition in PF (B) via inhibitory feedback. (C) With close colors, peaks are narrower and somewhat lower energy, which produces narrower and shallower inhibition in PF (D), making it easier for peaks to build in PF when a new item is presented at test. See text for details.

of inhibition in the perceptual field. Because inhibition is narrower and weaker in PF when the items in WM are close in feature space, it is easier for a "different" peak to be built when a change occurs at test. Thus, our model makes the counterintuitive prediction that *change detection will be enhanced when items are highly similar!*

We have confirmed this counterintuitive prediction in a recent study comparing color change detection accuracy for close vs. far colors (see also, Luck, Lin, & Hollingworth, 2005). In this study, participants viewed three target items (colored squares), two that were close in color space and one that was far away. Individual target items were presented at different spatial locations one at a time for 200 ms each, and were separated by a 500 ms inter-stimulus interval (ISI). We used a sequential-presentation paradigm to ensure that any differences in change-detection performance were due to interactions between items in WM, rather than arising as a result of color-contrast effects that can occur when multiple uniquely colored items are presented simultaneously. The presentation of the third target item was followed by a 1000-ms delay interval and the appearance of a single test item at one of the locations previously occupied by a target item. The test item was either identical to the item that was at that position originally (50% of trials), or was changed to a new color that was 30° away in color space.

Using this paradigm, we found that change detection performance was significantly better for close versus far colors at each probe position. More recently, we have shown that this effect generalizes to the feature dimension of orientation, and we are currently exploring whether these results also generalize to WM for spatial locations. These findings provide strong evidence for a general prediction made by our model: items in WM interact, giving rise to time- and metric-dependent distortions of WM representations. Note that these predictions run counter to several recent proposals regarding the nature and neural mechanisms that underlie VWM representations (Raffone & Wolters, 2001; Vogel et al., 2001).

12. A DFT approach to the "Binding" problem

The extension of the DFT of spatial cognition to address non-spatial object properties and the processes underlying change detection represents an important step towards the development of a comprehensive theory of the dynamics of visual cognition. In the present section, we describe our initial efforts as we take yet another step towards higher level cognition. The general goal of this work is to move toward the integration of "where" and "what" visual systems (Ungerleider & Mishkin, 1982), building on our theory of spatial cognition and our initial forays into FWM. We contend that such an integrated framework can offer novel insights into visual cognition by linking this area more strongly to the rich, embodied, spatial system we have already developed to handle reference frames, SWM, LTM, and so forth. Here, we focus on one piece of that effort: the question of how the visual system links, or "binds", individual features of objects to spatial locations, enabling goal-directed action towards objects in space. Moreover, we sketch our ideas regarding how multiple features are bound to each other, providing a basis for the creation of multifeature object representations.

As one progresses through the ventral object recognition pathway of the primate visual system, from primary visual area V1, through extrastriate areas V2-V4, and on to areas TEO and TE of the inferior temporal lobe, several obvious changes in neural response properties can be observed (see discussion in Luck, Girelli, McDermott, & Ford, 1997). First, there is an increase in the complexity of the features coded by individual neurons. For example, whereas neurons in V1 respond preferentially to rather simple stimuli such as oriented line segments, cells in TE may respond to complex stimuli such as faces (Desimone & Gross, 1979; Desimone, Albright, Gross, & Bruce, 1984; Tanaka, 1996). Second, there is a dramatic increase in receptive field sizes and an accompanying decrease in the spatial resolution of receptive fields for individual neurons (Desimone & Gross, 1979; Gross, Rocha-Miranda, & Bender, 1972). Additionally, although feature selectivity becomes more complex at higher levels of the visual system, the individual features of complex objects (e.g., color, form, size, and direction of motion) are coded in a distributed manner through the parallel activation of large numbers of neurons (Fujita, Tanaka, Ito, & Cheng, 1992; Komatsu & Ideura, 1993; Llinás & Paré, 1996). Although this type of encoding can be quite efficient, it can also lead to substantial problems when multiple objects are presented simultaneously, as in visual search experiments and most real-world visual tasks. For example, when two or more objects are presented simultaneously, the individual features making up the objects are coded in partially independent neural populations with overlapping spatial receptive fields. As a result, it can be difficult to determine which features belong together as attributes of a single object, an example of the binding problem in vision (Damasio, 1989; von der Malsburg, 1981, 1995).

One of the strongest sources of evidence suggesting that this is a real problem for the visual system is the finding that, under certain circumstances, the features of objects can be miscombined in normal perception. For example, in a series of experiments reported by Treisman and Schmidt (1982), participants were briefly shown multi-element displays of

colored letters. In some cases, the colors and letters were perceived correctly but in the wrong combination. For example, a participant could be presented with a display containing a red horizontal line and a green vertical line, and incorrectly report seeing a red vertical and a green horizontal line (see also, Ashby, Prinzmetal, Ivry, & Maddox, 1996; Cohen & Ivry, 1989; Prinzmetal, 1981).

To address this issue in a working memory context, we have developed the dynamic field model of multi-feature binding depicted in Fig. 7. The core of the new model consists of



Fig. 7. A dynamic field model of feature-binding. The model consists of multiple feature-space working memory fields (FSWM; Panels A and B) coupled to a single SWM field (C). The FSWM fields receive broad spatial and narrow featural inputs from the sensory surface, in keeping with cortical physiology. The presence of overlapping spatial receptive fields can make it difficult to correctly assign individual features to the correct objects when two or more objects are present simultaneously. We solve this problem through coupling with SWM, which maintains precise information about spatial location. See text for additional details.

multiple feature-space working memory fields (FSWM), which, when taken individually, capture WM for specific what-where conjunctions (e.g., a target color or orientation at a certain location), coupled to a single SWM field. Activation in feature-space fields evolves over time according to the same principles as in the fields described previously. For instance, neurons interact according to a local excitation/lateral inhibition function such that strong input (e.g., the presentation of a target color at a specific location) can create a self-sustaining WM peak. Such a peak of activation in a single FSWM field represents WM for a particular target feature at a given location, whereas the presence of peaks at the same spatial location across different fields implements a distributed representation of a simple multi-feature object. Additionally, to capture some of the constraints evident in higherlevel visual areas like IT, we tuned the feed-forward projections into these fields such that they are precise along the feature dimension but broad along the spatial dimension. As discussed previously, distributed encoding of object properties together with broad, overlapping spatial receptive fields can make it difficult to correctly assign individual features to the correct objects when two or more objects are present simultaneously. We propose to solve this problem by coupling each of the FSWM fields to a single SWM field that maintains precise information about spatial location. The presence of additional spatial input should "pull" activation into the right position along the spatial dimension, keeping peaks aligned in each of the FSWM fields and clarifying which features belong to which objects.

These ideas are illustrated in the simulations shown in Fig. 7A–F. For these simulations, the model was presented with three colored, oriented objects at three different locations (see visual display panel) for 200 time steps followed by 700 time steps with no additional inputs. To explore the impact of coupling to SWM on the positioning of peaks in the FSWM fields, we ran one simulation where the color-space field was uncoupled to SWM and the orientation-space field was strongly coupled, and a second simulation where both fields were strongly coupled to SWM.

The first simulation is depicted in Fig. 7A-C, which shows activation in two FSWM fields (a color-space field and an orientation-space field) and a SWM field in response to the visual display shown in the center of the figure. At this point, activation is relatively precise and well positioned along the feature dimensions, correctly indicating the presence of the colors red, orange, and green, and the orientations vertical, diagonal, and horizontal. However, in the FSWM fields, activation along the spatial dimension is relatively spread out and substantially overlapping, particularly in the color-space field, which is decoupled from the SWM field. The simulations in Fig. 7D–G show the formation of WM peaks in each of the three fields over the course of the next 700 time steps. As can be seen, peaks in the SWM field (panel D) and in the orientation-space field (panel E) are aligned in space, and each orientation is at the correct location. In contrast, the color-space field (panel F), which is uncoupled to space, has peaks positioned at the correct feature values, but the red peak is located near 5°, whereas the green peak is closer to 0° —the opposite of their positioning in the visual display. Additionally, the color-space field has incorrectly estimated the number of colors present in the display, forming a peak of activation at approximately 10° in space. However, when the color-space field is also coupled to SWM (panel G), peaks in all three fields are in alignment, correctly reflecting the three items present in the original visual display.

The dynamic field approach to multi-feature binding proposed here represents an explicit move to pursue a spatial solution to the binding problem, and is consistent in many

respects with the approach adopted by Treisman and colleagues (Treisman, 1996; Treisman & Gelade, 1980). However, unlike in *Feature Integration Theory*, visual features are not free-floating in our model, but are tied, albeit coarsely, to particular spatial locations during the initial presentation of a stimulus. In this respect, our approach is more similar to the *location uncertainty theory* of feature binding proposed by Ashby and colleagues (Ashby et al., 1996). According to this theory, incorrect combinations of features occur as a result of uncertainty about the precise spatial location of individual object features, rather than through random binding of features in the absence of attention, as proposed by Treisman. In the context of our model, inputs from SWM effectively provide converging evidence that helps to resolve the ambiguous spatial coding in the FSWM fields, similar to the role of spatial attention in other models (Desimone & Duncan, 1995; Luck et al., 1997). However, our model moves these ideas to a neurally plausible, process-based, and formal level.

The DFT is not the only neurally plausible approach that has addressed the integration of spatial and non-spatial information in some way. For example, Deco and Lee (2004) have recently proposed a neuro-dynamical model in which recurrent connections between higher-order visual areas in the dorsal ("where") and ventral ("what") visual streams and V1 serve to integrate location and identity information (see also, de Kamps & van der Velde, 2001; Deco, Rolls, & Horwitz, 2004). In most cases, such models are quite sophisticated on the 'what' side, providing detailed accounts of ventral stream processes that, for instance, are involved in the integration of multiple features into object representations (Deco & Rolls, 2004; Olshausen, Anderson, & Van Essen, 1993). However, these models often provide a limited view of dorsal stream processes. For example, several neurally based models (Itti & Koch, 2001; Mozer & Sitton, 1998) have employed the concept of a "salience map" that tags specific locations as important for attention or WM (see also, Treisman & Gelade, 1980; Wolfe, 1994). Critically, the proposed salience map is not linked to a particular frame of reference, it is not calibrated and adjusted relative to movement in the surrounds, it does not form spatial categories, and so on. In short, these models do not address many of the challenges we have overcome with the DFT (Spencer et al., 2007). Additionally, many of these approaches are targeted at a biophysical level of explanation, which involves the creation of neurally realistic models of neurons and neural networks that capture key elements of neuronal firing rates, neurotransmitter action, and so on. As a result, such models often do a good job of accounting for the behavior of neurons, but have relatively limited ties to the real-time behavior of subjects performing complex tasks. In contrast, the dynamic field approach to visual cognition introduced here has established a tight interface between neurally plausible modeling on one hand, and empirical investigations with behaving subjects on the other.

13. Conclusions

In the present paper, we traced a research trajectory demonstrating that the motor approach to dynamic systems theory can be effectively scaled up from sensori-motor systems into cognition. In particular, we traced movement from the motor approach into cognition via the DFT and then through three phases of exploration of cognitive phenomena. The first phase retained a strong sensori-motor flavor, focusing on motor control and sensori-motor development. From there we moved into spatial cognition, focusing on how interactions among perception, working memory, and long-term memory

can be used to explain delay-, metric- and experience-dependent drift of WM peaks away from perceived axes of reference and towards LTM traces of targets seeing during an experimental session. Finally, we used the DFT as a basis for addressing important issues in visual cognition, including WM for metric feature dimensions and the process of change detection, as well as the binding of spatial and non-spatial information into simple object representations.

Critically, as we traced each step of this research trajectory, we found that general principles of the DFT apply across the board. For example, the concepts of stability, instability, and flexibility were found to apply equally to both spatial and non-spatial working memory. In both cases, stability is realized through intrafield interactions among neurons, allowing the maintenance of self-sustained peaks of activation in the absence of input. The concept of instability can be used both to understand transitions from one stable state of the system to another, and metric-instabilities reflected in, for instance, the delay- and experience-dependent drift of WM peaks. As an example of the former, the formation of peaks in WM involves a transition from a stable resting state through an instability (bifurcation) into a new attractor state-the self-sustaining state. Finally, the flexibility of this approach is exemplified by the fact that the same three-layer model is capable of integrating perceived reference frames and WM in one case and comparing the contents of WM with perception in another. Additionally, we have shown that the same model using identical parameters is sufficient to capture behavioral data across four separate tasks in addition to developmental data using the spatial recall paradigm with children (Simmering, Schutte, & Spencer, 2007).

Dynamic fields also embody the concepts of soft-assembly and self-organization. For example, color memory was found to reflect the interplay between LTM traces built up over the course of an experiment and basic maintenance processes underlying the formation of self-sustained peaks. Additionally, our new framework continues to provide for tight theory-experiment links, leading to novel predictions in both color estimation and change detection experiments. Moreover, the theoretical approach remains committed both to neural principles, as exemplified in our work on neurophysiology and the binding problem, and the concept of embodiment, exemplified by our ongoing work using dynamic fields to control autonomous robots.

The ideas presented here constitute a demonstration proof showing that the motor approach can scale up to cognition. At this point, one might be tempted to ask: So what? The field of visual cognition is already crowded with theoretical concepts ranging from verbal information processing concepts to biophysical models of cortical function. Why do we need dynamical systems concepts in this domain?

In our view, there are at least two central contributions that dynamic systems concepts can make to the study of visual cognition. First, a general trend in visual cognition, as in other areas of cognitive science, is to pursue greater and greater fractionation. The challenges of cognition are broken down into component problems that are addressed separately. Thus, researchers investigate working memory, long-term memory, attention, various perceptual processes, and so forth, in isolation, but rarely do people address the issue of "what then". That is, once a multitude of "parts" have been identified, what do we do? In general, it is assumed that when the critical subcomponents have been understood, we will be able to fit the parts together to make a coherent whole. Although this seems like a reasonable strategy, we suspect that the integration of parts will prove to be much trickier than anticipated. A good litmus test for whether the separate components postulated to underlie behavior can be integrated to form a single behaving organism comes from the field of autonomous robotics. As Brooks (1991) points out, the separable components approach has not faired particularly well in this area. In most cases, the components remain separate and are used to solve specialized problems in carefully circumscribed areas, for instance, in medical diagnosis. However, in the few cases where integration has been attempted, performance has been found to be much too slow to deal in a reasonable way with dynamic real-world environments. Indeed, such systems have only been found to work in highly simplified environments, such as more-or-less empty rooms with uniformly colored walls and dark baseboards, and uniquely colored target objects. Moreover, solutions that are found to apply in one environment have not easily transferred to performance in other environments. As a result, specific solutions are often tailor-made to fit each new situation, limiting behavioral flexibility.

The DFT addresses these challenges by thinking about cognition and action from a dynamic systems perspective, with an emphasis on the integrated, collective behavior of the organism (Schöner & Kelso, 1988a, 1988b). For example, although our model does acknowledge distinctions between, for instance, perception, working memory, and long-term memory, these processes are neither conceived of nor implemented as separate encapsulated subsystems in our model. Rather, the dynamic field model within which these processes are realized is a fully integrated dynamical system with specific classes of attractor states. Thus, we do not make assumptions about separability and, consequently, integration is not a problem. The challenge posed by an integrated systems perspective, however, is the challenge of analysis: If behavior results from one large integrated system, how can we make sense of the processes that underlie behavior? Our approach to this issue is to build cognitive/neural systems that are reentrant, reciprocally coupled, fully integrated dynamical systems, but that have sub-systems whose dynamic properties leave behavioral "signatures" that can be observed in particular situations (e.g. delay-dependent drift in location and color memory). In this way, we can engage in rigorous hypothesis testing—a core strength of the separable systems view—while staying committed to integrated systems.

A second, closely related contribution of dynamic systems concepts is that they can help us understand the integration of brains, bodies, and behavior. Specifically, in developing this approach, we have taken inspiration from the densely interconnected nervous system to re-think cognition. As stated in the last section, a centerpiece of this approach is to embrace the use of complex, dynamic neural networks to capture brain–behavior relations. Although neural networks have architectures that can be depicted as separate systems, they are—at their core—complex, reentrant, densely interconnected, complex systems that violate the assumptions of encapsulation and separability. Critically, dynamic systems concepts of attractors, bifurcations, instabilities, and so on, give us the conceptual tools to understand the behavior of complex neural networks and to harness their real-time potential in the service of goal-directed action. The examples discussed herein provide illustrations of this marriage between the mathematics of complex systems and the emerging science of brain–body–behavior relations.

References

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

- Ashby, F. G., Prinzmetal, W., Ivry, R., & Maddox, W. T. (1996). A formal theory of feature binding in object perception. *Psychological Review*, 103, 165–192.
- Barsalou, L. W. (1999). Perceptual symbol systems. Behavioral and Brain Sciences, 22, 577-660.
- Bastian, A., Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport*, 9, 315–319.
- Beer, R. (1995). A dynamical systems perspective on agent–environment interactions. *Artificial Intelligence*, 72, 173–215.
- Braun, M. (1994). Differential equations and their applications (4th ed.). New York: Springer.
- Brooks, R. (1991). New approaches to robotics. Science, 253, 1227-1232.

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

- Clark, A. (1997). Being there: Putting brain, body, and world together again. Cambridge, MA: MIT Press.
- Cohen, A., & Ivry, R. B. (1989). Illusory conjunctions inside and outside the focus of attention. Journal of Experimental Psychology: Human Perception & Performance, 15, 650–662.
- 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.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87–185.
- Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, 1, 123–132.
- Damasio, A. R., & Damasio, H. (1994). Cortical systems for retrieval of concrete knowlfedge: The convergence zone framework. In C. Koch, & J. L. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 61–74). Boston: MIT Press.
- de Kamps, M., & van der Velde, F. (2001). Using a recurrent network to bind form, color, and position into a unified percept. *Neurocomputing*, *38*, 523–528.
- Deco, G., & Lee, T. S. (2004). The role of early visual cortex in visual integration: A neural model of recurrent interaction. *European Journal of Neuroscience*, 20, 1089–1100.
- Deco, G., & Rolls, E. T. (2004). A neurodynamical model of visual attention and invariant object recognition. *Vision Research*, 44, 621–642.
- Deco, G., Rolls, E. T., & Horwitz, B. (2004). 'What' and 'where' in visual working memory: A computational neurodynamical perspective for integrating fMRI and single-neuron data. *Journal of Cognitive Neuroscience*, 16, 683–701.
- Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience*, 8, 2051–2062.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. Annual Review of Neuroscience, 18, 193–222.
- Desimone, R., & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. *Brain Research*, 178, 363–380.
- Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schöner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, 94, 53–66.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109, 545–572.
- Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature*, 360, 343–346.
- Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. Journal of Experimental Psychology: Human Perception & Performance, 19, 845–866.
- Georgopoulos, A. P. (1995). Motor cortex and cognitive processing. In M. Gazzaniga (Ed.), *The cognitve neurosciences*. Cambridge: MIT Press.
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of of neurons in inferotetmporal cortex of the macaque. *Journal of Neurophysiology*, 35, 96–111.
- Hock, H. S., Schöner, G., & Giese, M. A. (2003). The dynamical foundations of motion pattern formation: Stability, selective adaptation, and perceptual continuity. *Perception and 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, 3–38.

- Iossifidis, I., & Schöner, G. (2006). Dynamical systems approach for the autonomous avoidance of obstacles and joint-limits for an redundant robot arm. Paper presented at *the IEEE 2006 international conference on intelligent robots and systems (IROS)*, Beijing, China.
- Itti, L., & Koch, C. (2001). Computational modeling of visual attention. Nature Reviews Neuroscience, 2, 194-203.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (2007). Delay- and experience-dependent biases in working memory for color, Manuscript in preparation.
- Kelso, J. A. S. (1995). Dynamic patterns: The self-organization of brain and behavior. Cambridge: MIT Press.
- Kelso, J. A. S., Scholz, J. P., & Schöner, G. (1986). Non-equilibrium phase transitions in coordinated biological motion: Critical fluctuations. *Physics Letters A*, 118, 279–284.
- Komatsu, H., & Ideura, Y. (1993). Relationship between color, shape, and pattern selectivities in the inferior cortex of the monkey. *Journal of Neurophysiology*, 73, 218–226.
- Llinás, R., & Paré, D. (1996). The brain as a closed system modulated by the senses. In R. Llinás, & P. S. Churchland (Eds.), *The mind-brain continuum*. Cambridge: MIT Press.
- Luck, S. J. (in press). Visual short-term memory. In S. J. Luck, A. Hollingworth (Eds.), *Visual memory*. New York: Oxford University Press.
- Luck, S. J., Girelli, M., McDermott, M. T., & Ford, M. A. (1997). Bridging the gap between monkey neurophysiology and human perception: An ambiguity resolution theory of visual selective attention. *Cognitive Psychology*, 33, 64–87.
- Luck, S. J., Lin, P. -H., & Hollingworth, A. (2005). Similarity and interference in visual working memory. Paper presented at *the conference name*. Retrieved Access Date from URL.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Mozer, M. C., & Sitton, M. (1998). Computational modeling of spatial attention. In *Attention* (pp. 341–393). London: UCL Press.
- Olshausen, B. A., Anderson, C. H., & Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience*, 13, 4700–4719.
- Olson, I. R., & Jiang, Y. (2002). Is visual short-term memory object based? Rejection of the "strong-object" hypothesis. *Perception and Psychophysics*, 64, 1055–1067.
- Pessoa, L., & Ungerleider, L. G. (2004). Neural correlates of change detection and change blindness in a working memory task. *Cerebral Cortex*, 14, 511–520.
- Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics*, 16, 283–290.
- Prinzmetal, W. (1981). Principles of feature integration in visual perception. *Perception & Psychophysics*, 30, 330–340.
- Raffone, A., & Wolters, G. (2001). A cortical mechanism for binding in visual working memory. *Journal of Cognitive Neuroscience*, 13, 766–785.
- Rensink, R. A. (2002). Change detection. Annual Review of Psychology, 53, 245-277.
- Schöner, G. (1989). Learning and recall in a dynamic theory of coordination patterns. *Biological Cybernetics*, 62, 39–54.
- 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). A dynamic theory of behavioral change. Journal of Theoretical Biology, 135, 501–524.
- Schöner, G., & Kelso, J. A. S. (1988b). Dynamic pattern generation in behavioral and neural systems. *Science*, 239, 1513–1520.
- Schutte, A. R., & Spencer, J. P. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74(5), 1393–1417.
- Schutte, A. R., & Spencer, J. P. (2007). Planning discrete movements using a continuous system: Insights from a dynamic field theory of movement preparation. *Motor Control*, 11, 166–208.
- Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychological Review*, *113*, 273–299.
- 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.
- Shibuya, H., & Bundeson, C. (1988). Visual selection from multielement displays: Measure and modeling effects of exposure duration. *Journal of Experimental Psychology: Human Perception & Performance*, 14, 591–600.

- Simmering, V. R., Spencer, J. P., & Schöner, G. (2006). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Perception & Psychophysics*, 68, 1027–1046.
- Simmering, V. R., Spencer, J. P., & Schöner, G. (in press). Generality with specificity: The dynamic field theory generalizes across tasks and time scales. Developmental Science.
- Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2007). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. In S. Becker (Ed.), *Computational Cognitive Neuroscience* [special issue]. Brain Research, doi:10.31016/j.brainres.2007.06.081.
- Spencer, J. P., Clearfield, M., Corbetta, D., Ulrich, B. D., Buchanan, P., & Schöner, G. (2006). Moving toward a grand theory of development: In memory of Esther Thelen. *Child Development*, 77, 1521–1538.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Spatial categories are formed using geometric and experience-dependent information. *Journal of Experimental Psychology: General*, 131, 16–37.
- Spencer, J. P., & Hund, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. Cognitive Psychology, 47, 432–480.
- Spencer, J. P., & 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., & Schöner, G. (2007). What does theoretical neuroscience have to offer the study of behavioral development? Insights from a dynamic field theory of spatial cognition. In J. M. Plumert, & J. P. Spencer (Eds.), *Emerging landscapes of mind: Mapping the nature of change in spatial cognitive development*. New York, NY: Oxford University Press.
- Spivey-Knowlton, M. J., Tanenhaus, M. K., Eberhard, K. M., & Sedivy, J. C. (1998). Integration of visualspatial and linguistic information: Language comprehension in real time and real space. In P. Oliver, & K. P. Gapp (Eds.), *Representation and processing of spatial expressions*. Mahwah, NJ: Erlbaum.
- Steinhage, A., & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. In P. S. Schenker, & G. T. McKee (Eds.), Sensor fusion and decentralized control in robotic systems: Proceedings of SPIE, Vol. 3523 (pp. 160–180). SPIE Publishing.
- Tabor, W. (in press). Dynamical insights into structure in connectionist models. In J. P. Spencer, M. Thomas & J. L. McClelland (Eds.), *Toward a new grand theory of development? Connectionism and dynamic systems theory re-considered*. New York: Oxford University Press.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. Annual Review of Neuroscience, 19, 109-139.
- Thelen, E., & Bates, E. (2003). Connectionism and dynamic systems: Are they really different? *Developmental Science*, 6, 378–391.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, 24, 1–86.
- Thelen, E., & Smith, L. B. (1994). A dynamic systems approach to the development of cognition and action. Cambridge: MIT/Bradford.
- Thelen, E., & Ulrich, B. D. (1991). Hidden skills: A dynamic systems analysis of treadmill stepping during the first year. *Monographs of the Society for Research in Child Development*, 56, 1–104.
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, 428, 751–754.
- Townsend, J. T., & Busemeyer, J. R. (1995). Dynamic representation of decision making. In R. F. Port, & T. v. Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition*. Cambridge: MIT Press.
- Trappenberg, T. P. (2003). Why is our capacity of working memory so large? Neural Information Processing-Letters and Reviews, 1, 97–101.
- Treisman, A. M. (1996). The binding problem. Current Opinion in Neurobiology, 6, 171-178.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Treisman, A. M., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107–141.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.
- van Geert, P. L. C. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, vygotsky and beyond. *Psychological Review*, 5, 634–677.
- van Gelder, T. (1995). What might cognition be, if not computation? The Journal of Philosophy, 91, 345-381.
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. Behavioral and Brain Sciences, 21, 1-14.
- Vogel, E. K., & Machizawa, M. G. (2005). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428, 748–751.

- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 92–114.
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (in press). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception & Performance*.
- von der Malsburg, C. (1981). *The correlation theory of brain function*. MPI Biophysical Chemistry, Internal Report 81-2.
- von der Malsburg, C. (1995). Binding in models of perception and brain function. Current Opinion in Neurobiology, 5, 520-526.
- Wheeler, M., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131, 48–64.
- Williams, R. J. (1986). The logic of activation functions. In D. E. Rumelhart, J. L. McClelland, & t. P. R. Group (Eds.), *Parallel distributed processing*, Vol. 1 (pp. 423–443). Cambridge: MIT Press.
- Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12, 1–24.
- Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, 1(2), 202–238.
- Xu, Y. (2002). Limitations of object-based feature encoding in visual short-term memory. Journal of Experimental Psychology: Human Perception & Performance, 28, 458–468.
- Xu, Y. (2004). An object benefit for encoding two color features of an object in visual short term memory. Paper presented at *the 45th annual meeting of the Psychonomic Society*, Minneapolis, MN.
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, 440, 91–95.