

# An Embodied Approach to Cognitive Systems: A Dynamic Neural Field Theory of Spatial Working Memory

**John P. Spencer (john-spencer@uiowa.edu)**

Department of Psychology, E11 Seashore Hall  
Iowa City, IA 52242 USA

**Gregor Schöner (gregor.schoener@neuroinformatik.ruhr-uni-bochum.de)**

Institut für Neuroinformatik, Ruhr-University  
44780 Bochum, Germany

## Abstract

Spatial cognition has typically been studied as isolated parts—spatial perception, spatial memory, spatial attention, and so on. Although this approach has had many successes, it has failed to produce a detailed understanding of how the piece-meal processes that make-up spatial behavior occur together in time in a complex, behaving organism with a densely connected, highly interactive neural system. In the present report, we describe a neurally-plausible theory of spatial working memory, the Dynamic Neural Field Theory (DNFT). The DNFT specifies how information activated in spatial working memory (SWM) changes from second-to-second relative to perceived reference frames and long-term memories of previously responded-to locations. Moreover, this theory captures key aspects of the development of this cognitive system.

**Keywords:** spatial cognition, cognitive development, working memory, long-term memory, dynamic systems theory, modeling, neural networks.

## Introduction

It's late at night. A coffee-addicted professor sits at his computer, surrounded by the myriad of objects that have taken up residence on his cluttered desk. With little forethought, he reaches behind a stack of papers, picks up a coffee cup, takes a sip, and places it back down, all the while staring intently at the computer screen. Not convinced by his own arguments, he gets up to retrieve a manuscript from the filing cabinet across the room. He sits back down, thinks, writes, and deletes. He stares off into space. With eyes transfixed on nothing, he reaches behind a stack of papers, picks up the coffee, and takes another sip. Still pursuing his muse, he retreats to the lab where he locates several graphs of new data on the lab coordinator's immaculate desk. Ten minutes later, he returns to the office and begins typing furiously. The ideas are hot, the coffee is cold, and his reaches for the occluded cup are as precise and effortless as ever.

In this report, we describe a new theory of spatial working memory—the Dynamic Neural Field Theory—that captures key aspects of the dynamics of spatial cognition in situations such as these, that is, the time-dependent

processes that underlie coordinated spatial behavior. Such behavior requires that people remember the locations of important objects in the local surrounds with enough fidelity to coordinate a myriad of second-to-second decisions, actions, and attentional shifts. Moreover, the local “map” that is used in one workspace must be coordinated with other maps as people move from context to context—from the office desk, to the filing cabinet, to the desk in the lab. This requires the real-time, contextually-specific integration of past and present, of longer-term memory with short-term “working” memory.

The processes that underlie spatial cognition have been primarily studied by isolating different aspects of what it takes to be spatially skilled. Researchers have examined the visual challenges in such situations (Luck & Vecera, 2002; Wolfe, 1998), how people calibrate and update sensorimotor coordinate frames (Darling & Miller, 1993; John F. Soechting & Flanders, 1989), and the spatial reference frames people use over short and long time scales in small- and large-scale spaces (T. P. McNamara, Halpin, & Hardy, 1992; T. P. H. McNamara, John A.; Hardy, James K., 1992; Pick, Montello, & Somerville, 1988). And, more recently, neuroscientists have shed light on the neural processes involved in these different challenges, from cells in dorsal cortical areas like parietal cortex that are involved in coordinate transformations (Andersen, 1995), to cells in prefrontal cortex involved in the integration of “what” and “where” (Rao, Rainer, & Miller, 1997).

Although we have learned a lot about these pieces of the puzzle, these advances in understanding remain pieces. To date, there is no theory that effectively integrates them. Thus, the goal of our theoretical efforts was to develop a model of spatial working memory capable of integrating the diverse processes that underlie spatially-grounded behavior.

## Requirements for an “integrated” approach to spatial cognition

The example of reaching for an occluded coffee cup highlights several requirements that a theory of spatial cognition must meet. No previous models have handled all of these challenges. I

**Metric memory for locations.** To successfully reach for an occluded coffee cup, one must remember relatively precise, spatially continuous, metric information, not just

qualitative, categorical information. The neurophysiological substrate of SWM fits quite naturally with this picture. This substrate involves graded, metric representations that evolve continuously in time under the influence of current sensory information as well as the current activation state (Constantinidis & Steinmetz, 1996; Rao et al., 1997; Smyrnis, Taira, Ashe, & Georgopoulos, 1992). What is the nature of this metric information? Evidence suggests that many metric dimensions are used—head-centered, body-centered, hand-centered, allocentric, and object-based (e.g., Graziano, Hu, & Gross, 1997; J.F. Soechting & Flanders, 1991). Given this, we contend that it is important not just to specify which dimension is used in a given context, but to know the processes involved in calibrating, coordinating, and re-establishing metric, spatial information. *The dynamic neural field model described here provides a window on these processes.*

**Forming a stable working memory.** Given that spatial memory must be linked to multiple sensory and motor systems as well as to other internal processes (e.g. those dealing with reference frames), and given that the state of these multiple subsystems may vary in time due to the complex behavior of the organism in changing environments, the maintenance of spatial information in working memory requires processes that stabilize metric information against variable influences (Spencer & Schöner, 2003). Nervous systems can generate stability in a variety of ways, for instance, by monitoring and updating spatial information using sensory feedback. The resultant stable states can be usefully characterized using the concepts of dynamic systems theory (Braun, 1994). In this framework, the space of possible states of a system is spanned by state or behavioral variables (Schöner & Kelso, 1988). For every possible state (or value of the state variables), a vector predicts in which direction and at which rate the system's state will evolve. Stable states are then values of the state variables at which the rate of change is zero and to which the system converges from nearby values. *The DNFT uses this concept of stability.*

**Spatial categories.** As discussed above, reaching for an occluded object requires precise metric information. Does this indicate, however, that categorical or "coarse" spatial information is not needed? To the contrary, several studies have shown that metric information is supplemented by categorical information (Huttenlocher, Hedges, & Duncan, 1991; Kosslyn, Chabris, Marsolek, & Koenig, 1992). For instance, memory for a location is enhanced by remembering where the target is relative to a visible landmark or reference axis (Huttenlocher, Newcombe, & Sandberg, 1994; Tversky, 1981). Similarly, memory is enhanced by remembering where the target has been in the past (Hund & Spencer, 2003; Spencer & Hund, 2003). Although there are several models of how spatial category information is used (Huttenlocher et al., 1991; Kosslyn et al., 1992; T. P. McNamara & Diwadkar, 1997), two central issues have not been addressed in the literature: (1) how people form spatial categories, and (2) how people integrate categorical and metric information in real-time. *The DNFT*

*specifies a real-time, neurally-plausible process for spatial category use and category formation.*

**Updating and re-establishing reference frames.** Spatial working memory must be "grounded" relative to a reference frame. Typically, this is thought of in perceptual terms where spatial information is kept current relative to perceptual landmarks in the world. Spatial information can also be considered relative to action frames of reference, for instance, the reaching motion needed to acquire a coffee cup. Within this context, there is a need (a) to ground spatial information within a frame of reference, (b) to keep this information calibrated and updated relative to on-line changes in sensori-motor reference frames, and (c) to flexibly re-established the reference frame when ties to the sensori-motor context are cut (e.g., after intervening actions in a different workspace). Several neurally-plausible approaches to (a) and (b) have been proposed (e.g., Pouget, Deneve, & Duhamel, 2002); however, the challenges involved in (c) have not been addressed (but see Burgess, 2002; Burgess, Becker, King, & O'Keefe, 2002 for related work). *The DNFT proposes a mechanism by which reference frames can be re-established.*

## A dynamic neural field theory of spatial working memory

The DNFT takes a first step toward addressing each of the challenges a model of spatial cognition must overcome.

**Metric working memory.** There is general agreement that some form of sustained activation is the most plausible neuronal substrate for short-term spatial memory (Constantinidis & Steinmetz, 1996; Fuster, 1995; Miller, Erickson, & Desimone, 1996). Exactly how sustained activation is neurally realized, however, is not clear. One class of models achieves a stable memory state using bi-stable networks in which a stable state of sustained activation coexists with an "off-state" (Amari, 1989; Amari & Arbib, 1977; Compte, Brunel, Goldman-Rakic, & Wang, 2000). Within the "on" state, locally excitatory and laterally inhibitory interactions among neurons create sustained activation patterns.

The DNFT is in this class of neural networks (see also, Erlhagen & Schöner, 2002; Thelen, Schöner, Scheier, & Smith, 2001). To describe the theory, consider an activation field defined over a metric spatial dimension,  $x$ , the exact nature of which we shall examine below. The continuous evolution of the activation field is described by an activation dynamics, that is, a differential equation which generates the temporal evolution of the field by specifying a rate of change,  $du(x,t)/dt$ , for every activation level,  $u(x,t)$ , at every field location,  $x$ , and any moment in time,  $t$ . The basic stabilization mechanism of the field is modeled by an inverse relationship between the rate of change and the current level of activation. This means that at high levels of activation, negative rates of change drive activation down, while at low levels, positive rates of change drive activation up. The activation level that emerges is a function of the balance of different inputs and interactions in the field. For example, when a negative resting level,  $h < 0$ , coexists with

a source of excitatory input,  $S > 0$ , then the resulting stable state of the activation dynamics

$$\tau du(x,t)/dt = -u(x,t) + h + S(x)$$

is  $u(x) = h + S(x)$ , the level at which positive and negative rates of change balance so that  $du/dt = 0$ . Note that  $\tau$  is a parameter that fixes the time scale of the activation field.

When the rate of change of activation at a field site,  $x$ , depends not only on the activation level,  $u(x,t)$ , and current inputs,  $S(x)$ , but also on the activation levels,  $u(x', t)$ , at other field sites,  $x'$ , then the activation dynamics are interactive. Locally excitatory interaction is described by a kernel,  $w(x-x')$ , such that

$$\tau du(x,t)/dt = -u(x,t) + h + S(x,t) + \int dx' w(x-x') \sigma(u(x',t))$$

Only sufficiently activated sites,  $x'$ , contribute to interaction. This is expressed by passing activation level through a sigmoidal function:

$$\sigma(u) = 1/(1 + \exp(-\beta u))$$

Such threshold functions are necessarily non-linear and are the basis for the bi-stability that structures the activation dynamics. Because cortical neurons never project both excitatorily and inhibitorily onto targets, the inhibitory lateral interaction must be mediated through an ensemble of interneurons. A generic formulation (Amari & Arbib, 1977) is to introduce a second, inhibitory activation field,  $v(x,t)$ , which receives input from the excitatory activation field,  $u(x,t)$ , and in turn inhibits that field:

$$\tau_u du(x,t)/dt = -u(x,t) + h_u + S(x,t) + \int dx' w(x-x') \sigma(u(x',t)) - c \int dx' w_i(x-x') \sigma(v(x',t))$$

$$\tau_v dv(x,t)/dt = -v(x,t) + h_v + \int dx' w(x-x') \sigma(u(x',t))$$

**Stabilizing the contents of working memory via spatial categories.** The set of equations above describes a neurally-plausible bi-stable network for SWM. Although sustained activation peaks in this network are stably in the

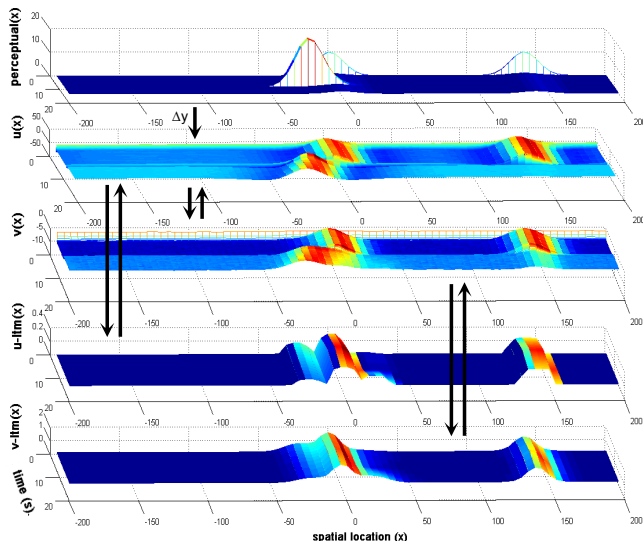


Figure 1. The DNFT.

“on” state, they are inherently unstable with respect to the metric information they represent. One manifestation of this metric instability is the “drift” of sustained peaks under the influence of noisy inputs that are common in the nervous system (Compte et al., 2000). Peak drift can also be induced by small, localized input gradients into the excitatory layer of the field which attract sustained peaks if they are positioned sufficiently close to the gradient (Amari & Arbib, 1977). Conversely, small localized inputs into the inhibitory layer cause peaks to drift away from the input gradient.

How might such gradients arise? A specific mechanism is through long-term memory traces of activation patterns. Whenever and wherever above threshold activation is present in WM, traces of activation can be slowly built up. This can be modeled through a simple linear activation dynamics of an additional set of fields—the LTM fields—which receive inputs from the corresponding layers of WM. Conversely, LTM traces feed back as excitatory inputs into the corresponding layers of WM:

$$\tau_{trace} du_{trace}/dt = -u_{trace} + \sigma(u);$$

$$\tau_{trace} dv_{trace}/dt = -v_{trace} + \sigma(v);$$

$$\tau_u du/dt = \dots + c_{u,trace} u_{trace} + noise$$

$$\tau_v dv/dt = \dots + c_{v,trace} v_{trace} + noise$$

A LTM trace of the excitatory layer will generate a small source of input that stabilizes WM peaks near the locations at which peaks have been activated earlier. Such excitatory memory traces form the neural substrate of spatial categories. Conversely, LTM traces of the inhibitory layer will generate a source of input that repels memory items from field sites that have been activated earlier. Such traces provide long-term discriminative information, amplifying activation differences based on past experiences. If excitatory memory traces are the substrate from which spatial categories are built, then inhibitory memory traces maximize the differences between categories.

**Updating and re-establishing reference frames.** To this point, we have described a neural mechanism for SWM and spatial categories but have remained vague on the

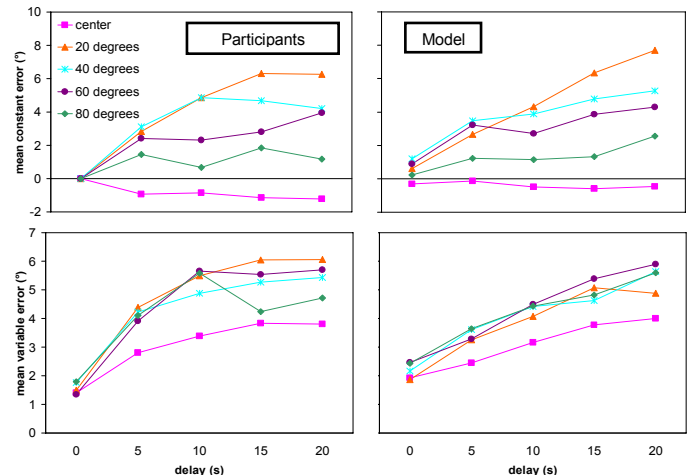


Figure 2. Simulations of data from Spencer & Hund (2003)

dimension,  $x$ , over which activation peaks are defined and how that dimension is linked to a sensori-motor reference frame. In some cases, this issue is relatively straightforward. For instance, the moment when the location of a coffee cup is first detected, the rich perceptual scene creates a context in which spatial relationships are defined and can be kept in register by relatively well-understood neural mechanisms (Devene & Pouget, 2003; Graziano et al., 1997).

What happens, however, when the sensory link to this context is interrupted, such as when we get up and leave the room and later return? In this case, a connection must be made between LTM traces originally represented in some frame of reference,  $x$ , and the current sensory layout. Such a connection can be made if, during the original activities, the perceptual structure is allowed to induce peaks of activation in SWM that leave LTM traces of the context. Such traces can later be re-activated, and a pattern-matching process can be used to detect a match and estimate the amount of shift,  $\Delta y$ , needed to bring the remembered and current sensory frames into register. In this way, context-based memories can be re-established to help organize and coordinate ongoing spatial activities. Although this aspect of the model is currently under development, we contend that the LTM mechanisms described previously provide an entry point into this problem. That said, this is, at present, an underdeveloped aspect of our approach (though see Steinhage & Schöner, 1998 for evidence that the dynamic field framework we adopt here can address these challenges within an autonomous robotics setting).

**The DNFT and strong ties to behavior.** To illustrate how the DNFT captures real-time behavior, consider our spatial estimation task (Spencer & Hund, 2002; Spencer & Hund, 2003). In this task, participants are seated at a large empty table and a target object is displayed for 2s. After delays ranging from 0 – 20s, participants are asked to reproduce the location of the target by, for instance, pointing to the location. As the memory delay increases, adults' responses are systematically biased away from the midline symmetry axis of the table and toward an average or "prototypical" target location.

Figure 1 shows a simulation of the DNFT during a single trial in this task. The top layer—the perceptual field—captures the perception of the spatial context and the target presentation. The next two layers are the excitatory,  $u$ , and inhibitory,  $v$ , layers of SWM. The bottom two layers show activation in the excitatory,  $u_{ltm}$ , and inhibitory,  $v_{ltm}$ , memory traces. Within each layer, spatial location is along  $x$ , where  $0^\circ$  is the midline of the space and positive locations are rightward;  $y$  captures time from the start (back of figure) to the end of a trial; and  $z$  shows activation.

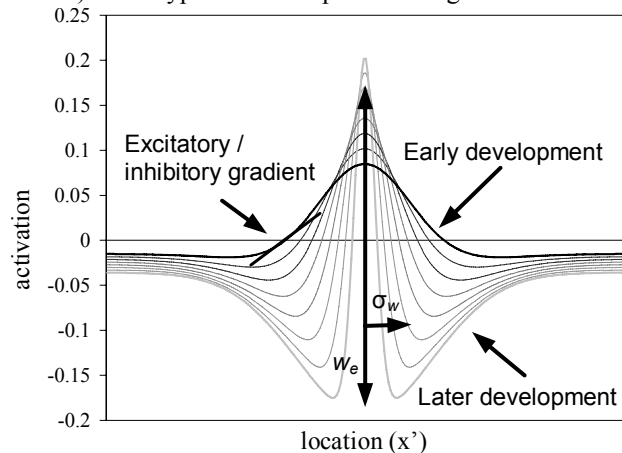
The simulation begins with SWM in "reference" mode. In this mode, the resting level,  $h$ , of SWM is raised, allowing multiple reference peaks to form, driven by activation in the perceptual field. In the simulation, two reference peaks have formed reflecting participants' perception of the midline axis and the right edge of the table. Next, the target is turned on at  $-20^\circ$ . This event triggers a lowering of the resting level in SWM to move this

field into "memory" mode. In this mode, the field selects the dominant input—the target—and a self-sustaining peak forms at the location in SWM associated with the position of the target. This "on" or "peak" state is stably maintained during the memory delay; however, the peak "drifts" systematically away from  $0^\circ$ , that is, away from midline. Consequently, when the model responds at the "go" signal by moving to the location associated with maximal activation, the model makes a leftward error. This is caused by the inhibitory input from the inhibitory memory trace ( $v_{ltm}$ ). Note that this bias is slightly counteracted by the excitatory memory trace ( $u_{ltm}$ ).

Figure 2 shows quantitative fits of the model to results from Spencer and Hund (2003). The model provides an excellent fit to both constant (top panel) and variable (bottom panels) errors. Importantly, these simulation results were generated with a single parameter setting.

**The development of spatial working memory.** SWM must be conceptualized in a way that can interface with critical developments in spatial cognitive abilities. For instance, early in development, human infants do not succeed in stabilizing remembered spatial information. This is most dramatically illustrated in the Piagetian A-not-B task (Piaget, 1954; Smith, Thelen, Titzer, & McLin, 1999). Beyond infancy, we still see the challenge of stabilizing metric working memory. In a version of the A-not-B task where children search for toys buried in the sand, memory for a B location is distorted in a graded manner toward A (Spencer, Smith, & Thelen, 2001). Importantly, this bias increases systematically as the memory delay increases (Schutte & Spencer, 2002), and, with increasing age, the amount of deviation toward A decreases.

The DNFT offers a unified account of these developmental effects. According to our *spatial precision hypothesis*, the spatial precision of neural interactions becomes more precise and more stable over development (Schutte, Spencer, & Schöner, 2003; Spencer & Hund, 2003). This hypothesis is captured in Figure 3. Each curve



**Figure 3.** The spatial precision hypothesis in this figure shows an example of a Gaussian interaction function,  $w(x-x')$ , at some point in development where the spatial precision of local excitatory interactions is given by  $\sigma_w$ , the strength of lateral inhibition is given by  $w_i$ , and the

overall strength of interaction is given by the excitatory scaling parameter,  $w_e$ . We capture developmental changes in interaction by coupling variations in the spatial precision parameter,  $\sigma_w$ , to changes in the scaling parameter,  $w_e$ :

$$w_e = S_e / \exp(\sigma_w / \alpha_e),$$

where  $S_e$  specifies the overall strength of the developmental modulation of interaction and  $\alpha_e$  specifies the steepness of the exponential modulation in interaction over development.

Using this equation, we can capture the proposed developmental changes in neural interaction by only changing a single parameter,  $\sigma_w$ . This is illustrated in Figure 3 where we varied  $\sigma_w$  quantitatively to produce the different curves. Two changes in interaction are apparent. As the interaction functions move from early development (darker lines) to later development (lighter lines), the spatial precision of interaction narrows, and the excitatory / inhibitory gradient becomes steeper. This results in relatively unstable self-sustaining peaks early in development that are sensitive to input across a broad spatial range, as well as stable self-sustaining peaks later in development that are only sensitive to input at narrow separations. This can explain, for instance, the reduction in A-not-B-type effects in the sandbox task over development.

## Conclusions

The DNFT provides the first formal theory of spatial working memory that integrates sensori-motor, working memory, and long-term memory processes in a neurally-plausible framework that is grounded by a close interplay between theoretical and experimental work. We contend that this theory takes an important first step toward an understanding of the processes that govern human activity in space—how people think about space, how people organize spatial activities, and the local “maps” of the world people bring with them from context to context. And, critically, this theory offers novel insights into the development of the spatial working memory system.

## Acknowledgments

NIMH RO1 MH62480, NSF BCS 00-91757, and NSF HSD 0527698 awarded to John P. Spencer

## References Cited

Amari, S. (1989). Dynamical stability of formation of cortical maps. In M. A. Arbib & S. Amari (Eds.), *Dynamic Interactions in Neural Networks: Models and Data* (pp. 15-34). New York, NY: Springer.

Amari, S., & Arbib, M. A. (1977). Competition and cooperation in neural nets. In J. Metzler (Ed.), *Systems Neuroscience* (pp. 119-165). New York: Academic Press.

Andersen, R. A. (1995). Encoding of intention and spatial location in the posterior parietal cortex. *Cerebral Cortex*, *5*, 457-469.

Braun, M. (1994). *Differential equations and their applications* (4 ed.). New York, NY: Springer Verlag.

Burgess, N. (2002). The hippocampus, space, and viewpoints in episodic memory. *Quarterly Journal of Experimental Psychology A*, *55A*(4), 1057-1080.

Burgess, N., Becker, S., King, J. A., & O'Keefe, J. (2002). Memory for events and their spatial context: Models and experiments. In e. al (Ed.), *Episodic Memory: New Directions in Research*.

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.

Constantinidis, C., & Steinmetz, M. A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology*, *76*, 1352-1355.

Darling, W. G., & Miller, G. F. (1993). Transformations between visual and kinesthetic coordinate systems in reaches to remembered object locations and orientations. *Experimental Brain Research*, *93*, 534-547.

Devene, S., & Pouget, A. (2003). Basis function for object-centered representations. *Neuron*, *37*, 347-359.

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

Fuster, J. M. (1995). *Memory in the Cerebral Cortex: An Empirical Approach to Neural Networks in the Human and Nonhuman Primate*. Cambridge: MIT Press.

Graziano, M. S. A., Hu, X. T., & Gross, C. G. (1997). Coding the locations of objects in the dark. *Science*, *277*, 239-241.

Hund, A. M., & Spencer, J. P. (2003). Developmental changes in the relative weighting of geometric and experience-dependent location cues. *Journal of Cognition and Development*, *4*(1), 3-38.

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

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

Kosslyn, S. M., Chabris, C. F., Marsolek, C. J., & Koenig, O. (1992). Categorical versus coordinate spatial relations: Computational analyses and computer simulations. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(2), 562-577.

Luck, S. J., & Vecera, S. P. (2002). Attention. In S. Yantis (Ed.), *Stevens' Handbook of Experimental Psychology: Vol. 1. Sensation and Perception* (Vol. 1, pp. 235-286). New York: Wiley.

McNamara, T. P., & Diwadkar, V. A. (1997). Symmetry and asymmetry of human spatial memory. *Cognitive Psychology*, *34*, 160-190.

McNamara, T. P., Halpin, J. A., & Hardy, J. K. (1992). Spatial and temporal contributions to the structure of

- spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 555-564.
- McNamara, T. P. H., John A.; Hardy, James K. (1992). Spatial and temporal contributions to the structure of spatial memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 18(3), 555-564.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16(16), 5154-5167.
- Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic Books.
- Pick, H. L., Montello, D. R., & Somerville, S. C. (1988). Landmarks and the coordination and integration of spatial information. *British Journal of Developmental Psychology*, 6(4), 372-375.
- Pouget, A., Deneve, S., & Duhamel, J. (2002). A computational perspective on the neural basis of multisensory spatial representations. *Nature Reviews Neuroscience*, 3(9), 741-747.
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science*, 276, 821-824.
- Schöner, G., & Kelso, J. A. S. (1988). Dynamic pattern generation in behavioral and neural systems. *Science*, 239, 1513-1520.
- Schutte, A. R., & Spencer, J. P. (2002). Generalizing the dynamic field theory of the A-not-B error beyond infancy: Three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, 73, 377-404.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74(5), 1393-1417.
- Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: the task dynamics of the A-not-B error. *Psychological Review*, 106, 235-260.
- Smyrnis, N., Taira, M., Ashe, J., & Georgopoulos, A. P. (1992). Motor cortical activity in a memorized delay task. *Experimental Brain Research*, 92, 139-151.
- Soechting, J. F., & Flanders, M. (1989). Errors in pointing are due to approximations in sensorimotor transformations. *Journal of Neurophysiology*, 62(2), 595-608.
- Soechting, J. F., & Flanders, M. (1991). Arm movements in three dimensional space: computation, theory and observation. *Exercise Sports Science Review*, 19, 389-418.
- Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*, 131, 16-37.
- Spencer, J. P., & Hund, A. M. (2003). Developmental continuity in the processes that underlie spatial recall. *Cognitive Psychology*, 47(4), 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., Smith, L. B., & Thelen, E. (2001). Tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of 2-year-olds. *Child Development*, 72, 1327-1346.
- Steinhage, A. & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. Sensor Fusion and Decentralized Control in Robot Systems. *Proceedings of SPIE*, 3523, 160-180.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A dynamic field theory of infant perseverative reaching errors. *Behavioral and Brain Sciences*, 24, 1-86.
- Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology*, 13, 407-433.
- Wolfe, J. M. (1998). Visual Search. In H. Pashler (Ed.), *Attention*. London, UK: University College London Press.