

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Object Permanence as Relational Stability or How to Get Representation from the Dynamics of Embodiment

#### **Permalink**

<https://escholarship.org/uc/item/0vn2w030>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 29(29)

#### **ISSN**

1069-7977

#### **Author**

Luo, Jun

#### **Publication Date**

2007

Peer reviewed

# Object Permanence as Relational Stability or How to Get Representation from the Dynamics of Embodiment

Jun Luo (jun.luo@utoronto.ca)

Cognitive Science Program, University College, University of Toronto  
15 King's College Circle, Toronto, ON M5S 3H7 Canada

## Abstract

The dynamic systems approach is sometimes embraced in conjunction with causal coupling and in opposition to representation-based explanation. The dynamic systems model of the A-not-B phenomenon (Thelen et al, 2001) supposedly epitomizes this anti-representation alternative. It is argued here that this model fails to capture the crucial effects of body rotation in the A-not-B task and that we need a notion of relational stability to account for these effects. Relational stability outstrips causal coupling but may ground a renewed understanding of object permanence where representation is necessitated by and achieved through the dynamics of embodiment.

**Keywords:** embodiment; representation; dynamic systems approach; A-not-B; object permanence; relational stability.

## Introduction

Some advocates of the dynamic systems perspective in cognitive science take it to be an anti-representation alternative (Thelen & Smith, 1994; van Gelder, 1998; Thelen et al. 2001). This anti-representation stance, however, is not shared by many dynamic systems theorists of neural systems (e.g. Seung 1996, Zhang 1996, Eliasmith & Anderson 2002) or even by theorists explicitly embracing the “dynamic systems approach” (e.g. Elman, 1995; Beer, 1998). For discussion I will use the acronym “DSA”, with deference to Thelen and Smith (1994), to mean the dynamic systems perspective (i) restricted to psychological explanation solely in terms of *causally coupled* dynamic processes and (ii) taken to be *anti-representational*, but reserve the unabridged “dynamic systems approach” for a general and possibly inclusive emphasis on the value of dynamic systems analysis in cognitive science. The restriction to causal coupling is a central tenet of the DSA: it is to both provide a positive alternative to representation and steer it clear of the oft-laid accusation that dynamic systems modeling describes behavior without explicating the underlying mechanisms.

To deliver on the promise of the DSA, Thelen et al (2001) offered a *dynamic systems model of the A-not-B phenomenon* (henceforth the DSM) from child development. Since Piaget (1954) identified it, this phenomenon has attracted the attention of generations of psychologists. Contemporary A-not-B experiments typically use a hiding device with two wells. The child (8- to 12-month-old) sits on the floor or in the caregiver's lap *watching* the experimenter hide a toy in one of the wells. After hiding the toy, the wells are covered and the experimenter withholds the device to impose a delay of several seconds. The child is then allowed to reach and retrieve the toy. Such a trial is repeated a few times at the same well (called ‘A’ by convention) and the child typically

reaches correctly. The same procedure is then done at the other well (called ‘B’). The child makes an “A-not-B error” if, during the ‘B trial’, they go back to search at A. Because the child repeats their A-trial response in a B trial, the A-not-B error is also called the “perseverative response”. The child outgrows such perseveration at about one year of age.

For Piaget this phenomenon marks a transitional stage between the beginning of a sense of “object permanence” (that occluded objects persist) and a robust grasp of it: the child *can* already retrieve a toy from a single hiding location, but they *cannot* yet cope with its “visible displacement” among multiple hiding locations. Many contemporary researchers used this phenomenon to “showcase” their research programs: for Diamond (1990) it taps into the involvement of the prefrontal cortex, for Munakata (1998) it reveals the interplay of graded “active” and “latent” representation, for Marcovitch & Zelazo (1999) it is a window into conscious control, for Newcombe & Huttenlocher (2000) it illustrates the developmental interaction of locomotion and spatial coding. Thelen & Smith (1994) themselves already attempted an informal account of this phenomenon in their book that launched the DSA in developmental psychology. But, with the exception of Thelen & Smith, all such views are couched in explicitly representational terms: *spatial code*, *location representation*, *conscious control*, and so on.

Against this background, we may understand the weight of the DSM for the DSA (Thelen et al 2001, p.2, e.a.):

Our message is: if we can understand this particular infant task and its myriad contextual variations in terms of *coupled dynamic processes*, then the same kind of analysis can be applied to *any task at any age*. If we can show that “knowing” cannot be separated from perceiving, acting, and remembering, then these processes are always linked. There is no time and no task when such dynamics cease and some *other mode of processing* kicks in. Body and world remain ceaselessly melded together.

The A-not-B task is representative of “any task at any age” because it had been *invariably* taken to mark the beginning of representational capacities for knowing that objects persist, remembering where they are, etc. Against this, the DSA proposes to treat it *solely* in terms of causally coupled dynamic processes. The hope is: if the DSA were to succeed on this task, it could deal with, solo, any task later in life, whether or not it is traditionally viewed as representational.

I shall note that the DSA's take on the A-not-B task itself is also revisionary. While accepting the canonical design,

the DSA theorists emphasize that a *hidden toy* is unnecessary for the rise of perseverative reaching (Smith et al, 1999). This, they reason, shows that the task is not representational even in traditional terms because the behavior persists when there is nothing to represent. The point here is not that occlusion makes no difference, but rather that it is only one among the “myriad contextual variations” and not essentially different from the others: in charting the complex “dynamics of embodiment” that gives rise to the final reach to A or B, we do not need to appeal to anything that is recognizably individuated representation of a hidden toy.

Below I will conduct a case study of the DSM. I will focus on whether it does justice to what the DSA wears on its sleeves: the dynamics of embodiment. Specifically, I will examine the DSM in light of the empirical evidence for the effects of body rotation. Body rotation, I suppose, is *bona fide* “dynamics of embodiment.” In doing so, I completely accept the view that the dynamics of embodiment pervades our cognitive life—and venture to raise it one, by following it through to arrive at representation.

### How the DSM Works

The heart of the DSM (Figure 1) is a motor planning (MP) neural field (Amari, 1977). Field locations are connected by shift-invariant, locally excitatory and globally inhibitory recurrent weights. The highest activation in the field after the delay codes the body-centered direction in which the child reaches. A parameter  $h$  controls interaction within the field: high  $h$  allows activation at more field locations to go above a threshold to strongly excite nearby ones and inhibit distant ones and low  $h$  suppresses such interaction. With the right  $h$ , localized activation peak can sustain itself in the absence of input and inhibit the rise of new peak from further input. Since a uniformly unexcited state is also stable, the *bistability* makes the Amari neural field good for modeling spatial memory: spatially-selective activation peak can be induced by strong transient input and then sustained stably by the field itself (e.g. Camperi & Wang, 1998).




Figure 1: The DSM of the A-not-B task. Dashed borders indicate components assumed.

The DSM exploits bistability to model the A-not-B phenomenon. During the A trials, a “preshape” is formed in the supra-trial bodily memory (BM) through integrating a thresholded version of MP activity. This BM preshape is highest for target A because the highest MP activity has been directed to A in the A trials. During the B trial, the BM preshape in turn biases MP activation in favor of A. When the MP field is not bistable (*low h*), after the initial strong “specific input” from B (i.e. cueing and hiding at B), it goes

through a transient to the resting level that is offset by the weak “task input” from A and B (i.e. visible covers at A and B). With enough delay, BM influence on MP will dominate and direct the system towards A to exhibit the perseverative response of *younger* children. When the MP field is bistable (*high h*), however, it can sustain a localized activation peak derived from specific input from B to counteract the BM bias for A. This way, the system can remain directed to B and exhibit *older* children’s correct reach.

### The Role of Body Rotation

In the A-not-B task, while the child sits on the floor or the caregiver’s lap, their head and upper body are typically allowed to—and do—rotate. This raises a question: *What role, if any, do such upper body rotations play in the task?*

When the upper body rotates, the upper-body-centered locations of potential targets shift (Figure 2). This does not pose too big a challenge, however, if there is only a *single* potential target (e.g. a covered well) that is constantly visible in the scene. In this case, so long as the child remains motivated to reach after the delay, the unique visible target can guide the actual reach. What makes the A-not-B task difficult, as Piaget (1954) emphasized, is the presence of *multiple* potential targets that have all become relevant due to the toy’s “visible displacement” among them. The (first) B trial is special precisely because it is the first time in the task when cueing is at a *second* target. Moreover, because the A and B targets typically look similar (e.g. with identical covers), task input cannot sufficiently discriminate them. One might thus expect that the choice of target must be based on memory derived from earlier specific input.




Figure 2: Relative shift of potential target locations

There is, however, a simpler “embodied strategy”: remain fixated on the cued target during the delay. Since the child mostly reaches to where they look (see below), if, throughout the delay, they fixate on B, they will be able to reach correctly. With fixation doing the heavy-lifting of “target selection”, memory derived from specific input is only needed to facilitate the child’s motivation to reach after delay at the target in front of oneself, would be sufficient for success in the A-not-B task (Piagetian Stage 5 behavior). The prediction, put bluntly, is: *without delay-period body rotation there will be no A-not-B error.*

Empirical evidence supports this analysis. To minimize straining to the correct well during delay, Diamond (1985) instructed parents to hold tight the child’s upper body in her A-not-B study; she also called the child to look at the center so as to break their fixation on the correct well during delay.

Since neither measure completely eliminated straining to or fixation on the target, she analyzed their effects and found that “when the infant’s strain or gaze was *uninterrupted and maintained throughout the delay*, success rate was significantly higher than on comparable trials where strain or gaze was not thus maintained” (p.873, e.a.). Similarly, Horobin & Acredolo (1985) found that “infants who maintained visual orientation *only* toward B after the toy was hidden in that location were most likely to direct correct search there” (p.125) but those who looked to A during B trial delay were more likely to perseverate. Gratch & Landers (1971) found that (i) if the child looked at A before the delay ended they would reach to A and (ii) older children’s orientation to B predicted their reaching to B. These results agrees with Harris’s (1973) finding that lowering a curtain at A at the beginning of delay led to A-not-B response and Smith et al’s (1999) finding that tapping a rod at either A or B at the end of delay increased the child’s reaching in that direction. These last two studies did not report the child’s orientation, but the perturbation presumably got them to turn to it.

In a related study of 9- and 16-month-old infants’ performance in a free-choice two-location hide-and-retrieve task, Cornell (1979) found that when children maintained gaze on the container with the hidden object, they were correct in 1254 of 1260 trials, but if they gave a glance to the empty container before opening one of the containers, they were correct in only 1176 among a total of 2580 trials.

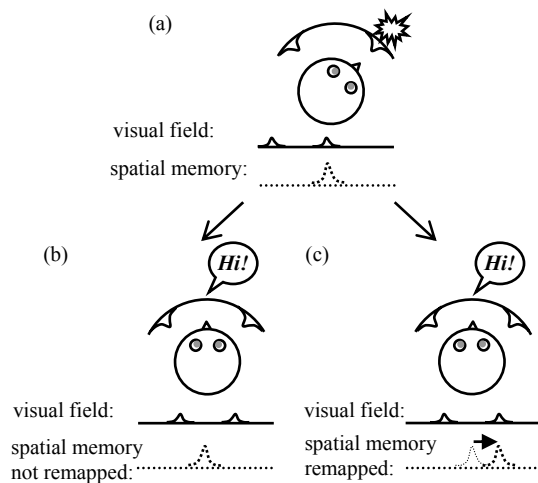


Figure 3: Child (a) turns to cued target at beginning of trial and memory activation peak for target is formed at center of memory field. Child is called to center during delay: (b) without remapping the memory peak now corresponds to center of scene; (c) with appropriate remapping it remains corresponded to target on the right.

If uninterrupted gaze naturally leads to correct reach, *why is turning away, spontaneous or induced, so detrimental to the child’s performance?* One explanation is that it invalidates the embodied strategy of fixation: with body rotation, the correct target is no longer always right in front of the

child. Spatially-selective memory spanning the delay is thus necessary for guiding the final reach. But if spatial memory is needed, a further challenge arises: *How is the correspondence maintained amid body rotations between target-selective spatial memory and relatively shifting targets?* If sustained spatial memory were not correctly updated to compensate for body rotation, it would miss the target after body rotation (Figure 3). With delay-period body rotation, not only spatial memory is needed, so is a *dynamic remapping* of such memory that ensures its correspondence with the target (Pouget & Sejnowski, 2002). Therefore, *the presence or absence of memory remapping may in part explain younger children’s failure and older children’s success.*

Such a perspective on the A-not-B task was actually proposed a long time ago (e.g. Bremner, 1978). Recently, Newcombe (2001) specifically named *dead reckoning* as a possible mechanism behind older children’s success. Dead reckoning, if we may recall, uses feedback about one’s own movement to dynamically remap spatial memory so that one remains correctly oriented to a target. This, according to our analysis, is precisely something useful in the A-not-B task.

To sum up, body rotation is a central part of the A-not-B phenomenon—without body rotation there won’t be A-not-B error, but with body rotation correspondence between memory and target needs to be dynamically maintained.

### Limitations of the DSM

The DSM, however, cannot account for the role of body rotation because, first, its assumed *static coupling* between BM activity (which in real baby is presumably in *body-base* coordinates) and MP activity (which in real baby is presumably in *upper-body* coordinates) forces on it an implicit conflation of upper-body and body-base coordinates, which in turn bars it from even recognizing upper body rotation.

Moreover, if the child makes the A-not-B error, under typical experimental design, there are at least two body rotations in that trial: first to B when cued and then to A; if the child is called to the center during the delay as in Diamond (1985), there will be three: first to B when cued, then to the center when called, and then either to A (error) or to B (correct reach). Given what Thelen et al (2001; see also Smith et al., 1999) convincingly argued—that motor planning for the final reach happens continuously throughout the whole trial, it follows that planning for the final reach *overlaps in time* with planning for delay-period body rotations. This is supported by the empirical findings surveyed above that where the child faces strongly predicts where the child reaches. The DSM, however, ignores planning for the intervening body rotations. Moreover, due to its bistability, the MP field *cannot* subserve planning for these rotations if it is to also, during the delay, sustain an activation peak that statically corresponds to the target. Thus, the DSM’s MP field embodies a second conflation of two closely related yet crucially differentiable processes: *motor planning* and *working memory*. Spencer & Schöner (2003, p.400)’s renaming the MP field “working memory” seems to be a terminological shift that betrays its ambiguous theoretical role.

The problem is really one of *time scales*. Body rotation prompted by specific input (cueing) or perturbation (e.g. calling) is towards what is *now* attention-grabbing but may be of no persistent interest. Motor planning for it accordingly takes place on the fast time scale of *orienting response*. Planning for the final reach, in contrast, involves tracking an earlier-cued target of lasting interest and takes place on the slow time scale of *working memory*. The DSM conflates these time scales and thus cannot accommodate memory remapping, which must work on the *fast* time scale of orienting response to protect the memory with its *slow* dynamics from being invalidated by the orienting response.

Due to these limitations, the DSM can explain neither why the child needs to turn away from B to make the A-not-B error, nor how they may reliably succeed in spite of having, during the delay, turned away from the target.

## Towards Relational Stability

### To Remember or Not To Remember

Recall that the hidden toy is not necessary for perseverative response to rise and the DSA theorists take this to show that memory of a *hidden toy* is not needed to explain the data. That much may be conceded. It might be further suggested that the child's sustained motivation to reach is not specific to the toy either, but is instead directed at the visible target. I.e., their interest *in the toy* after they remove the cover may be a complete renewal upon its reappearance. This, too, might be conceded. Either way, however, there are good reasons for us to believe that the involvement of memory with specific content is inevitable—so long as, to succeed, the child has to sustain, during the delay, differential interest or motivation towards the two visible potential targets.

What makes the A-not-B task challenging, again, is that the toy's "visible displacement" implicates *multiple* potential targets. Since differential selection of target *after the delay* of a trial is here expected to be based on the experimenter's differential manipulation of target (specific input) *before the delay* of the same trial, trial-specific and target-selective memory must be involved for there to be systematic success. Memory of the toy may indeed be unnecessary, but *target-differentiating* memory based on earlier specific input is indispensable. Spatial gaps due to occlusion are not the only factor necessitating memory. Temporal gaps—whether or not they are also entangled with occlusion—are after all what memory is in general supposed to span. If the DSM has collapsed hiding a toy and waving the well cover into "specific input", it has also incorporated *sustained* target-selective MP activity to explain Stage 5 behavior.

### To Remap or Not To Remap

If target-selective memory is needed, correspondence between memory—whether or not of a hidden toy—and target must also be maintained amid body rotations. The question is how. One possible alternative to the dead-reckoning sort of remapping is that the child uses landmarks. But there are good reasons for us to think that this is not the case in such a task at the age concerned.

Bushnell et al (1995) found that 1-year-olds' performance in a locomotor search task was facilitated by "direct landmark" (a distinctively colored cushion directly covering the toy) but *not* by "near landmark" (a distinctively colored cushion near the cushion above the toy). Since well covers in the A-not-B task are typically identical, we shall be primarily concerned with the use of near landmarks, which in turn seems to be ruled out by Bushnell and colleagues' finding. In general, as Gibson and Pick (2000, p.132) observed: infants start using direct landmarks as their locomotor ability emerges and begin to use near landmarks only with considerable locomotor experience.

More importantly, even if near landmarks were used, that does not seem to take away the need for dynamic remapping. Because cueing at A *or* B is specific to the current trial, the validity of landmark-target association is also limited to it. Were near landmarks available, the child would still need to, during a trial, specifically relate such landmarks to the *intra-trial* target cueing, which means selection and tracking of the landmark must in turn be trial-specific. If visual features do not uniquely distinguish the landmark, we are back in square one with the child having to sustain *landmark selectivity* during the delay and amid body rotations. If visual features do uniquely single out the landmark, the child still needs to, (i) sustain a trial-specific preference to these features, (ii) use these features to track the landmark, and (iii) relate in a trial-specific way the landmark thus tracked to the target that is spatially separate from it. These are very exacting demands for an infant and it is far from clear that (ii) and (iii) will not be vulnerable to body rotation if memory remapping is not already in operation.

Finally, neuroscience research has robustly identified in animals remapping of *remembered* target or direction in compensation for self-movement (Mays & Sparks, 1980; Duhamel et al, 1992; Graziano et al, 1998; Baker et al, 2003). Several neural computation models, based on these findings, have also been proposed (see Pouget & Sejnowski 2002 for a survey). These suggest that the dynamic remapping in question in the A-not-B task is not an isolated phenomenon, but has deeper roots and broader implications and is neurologically plausible. An embodied movement model of the A-not-B phenomenon that incorporates kinesthesia-based remapping (Zhang 1996) has been constructed (Luo 2004). Luo also showed that such remapping mechanism can be learned with a variant of the algorithm proposed by Stringer et al (2002) in a way that implicates locomotor development without presupposing landmark use.

### To Represent or Not To Represent

The success condition on memory remapping is that the sustained activation peak *dynamically corresponds* to the correct target amid body rotations. I.e., the target must be *relationally stabilized* even though both the activation peak and the target move—the activation peak relative to the field and the target relative to the rotating child. Such "relational stability" contrasts with the attractor stability—strictly speaking, *line* or *continuous* attractor stability (Seung 1996; Trappenberg 2004)—exploited by the DSM and related

work (Schutte et al, 2003), where the primary concern was to keep the activation peak *at a field location* in the presence of distracting inputs. In the face of body rotation, however, pinning down the peak is precisely to miss the relatively shifting target (Figure 3)! Amid the dynamics of embodiment, target tracking must be dynamically achieved, through shifting the activation peak!

Crucially, memory remapping is here *not* causally driven by the target “over there”. It is instead occasioned by and coordinated with the self-movement “around here”: body rotation both changes the spatial relation and generates the feedback needed for the remapping. Task input may still help sustain the memory. In this sense memory and its target are still “coupled”. Correct dynamic correspondence, however, is *not* a result of this weak coupling because task input does not differentiate the correct target from the incorrect one. It is rather the other way around: task input from the correct target can support the memory activation peak only if their correspondence is dynamically maintained.

Neither is the remapping a passive, continuous evolution of system dynamics from the early strong coupling when the target was being cued. Rather the remapping is *endogenous* and *active*. It requires the compensation mechanism to kick in upon body rotation. But nothing in the environment or in the child-environment coupling necessitates the kicking in of this remapping. It is rather a matter of how the various capacities of the child get dynamically and successfully coordinated. And this coordination is *an architectural feature of the child*, not that of the coupling.

The point about the endogenous contribution to relational stability really concerns a dynamic division of labor: when the scene becomes quiescent, responsibility falls on the child to remain correctly directed to the real target. Accordingly, the child must be equipped with appropriate memory sustaining and remapping capacity. If the cueing were still on-going, or the target were uniquely singled out by its visual feature, or the embodied strategy of fixation were not upset, the environment could be shouldering much more of the responsibility for successful target tacking.

I take it that the appropriately remapped activation peak, diachronically individuated, *represents* the target, or more accurately, something like “this over there now is where I want to go.” This is because (i) the peak *stands in for* the target’s being the right one *in the absence of* sufficient visual distinction of the target and (ii) it is also *manipulated*—dynamically remapped—*in lieu of* specific input marking out the target as the earlier specific input is no longer available for direct causal coupling with it. These dual features of “standing in for” and “being manipulated in lieu of” make the memory peak representational. Crucially, however, (ii) is here prior to (i): remapping establishes a *dynamic correspondence* that allows the shifting peak to, during the delay, *stand in for* the (relatively shifting) target. Such a representation, rather than presupposing a static or even atomic correspondence with its target as under the traditional, logicist conception of representation, is subsequent to and constituted by dynamically achieved relational stability.

## To Reduce or Not To Reduce

Parallel things can already be said about the sustenance of memory peak in the DSM, where attractor stability allows temporal gaps to be spanned so that the *currently* quiescent scene may be treated in a way that is appropriate given *previous* cueing. The achievement of this diachronic relation with the past, which we may even call “temporally relational stability”, is again *endogenous*: It is based on the appropriate *architecture of the neural field*. The memory is here also manipulated—sustained—in lieu of causal coupling with the now bygone cueing and stands in for something like “this over there is where I want to go.”

The problem with the DSM is not that it is not representational, which it already is. It is rather that it is not dynamic enough: the static correspondence assumed by the DSM is ironically reminiscent of logicist representation! What the DSM lacks is adequate recognition of both the challenge from the dynamics of embodiment—that it can disrupt static correspondence based on attractor stability, and the potential in the dynamics of embodiment—that correspondence can be established dynamically.

How about the DSA, then, if the DSM already does not follow its script? With its exclusive commitment to explanation in terms of causal-coupling, the DSA (recall our convention with this acronym!) cannot even recognize relational stability as such, because neither the memory remapping nor the memory sustaining in question can be explained in terms of causal coupling with what the memory is *of*. The point is not that the dynamic systems theory, which the DSA uses, is irrelevant—after all the remapping mechanism in question has been fruitfully studied using dynamic systems theory (e.g., Zhang 1996). The point is rather that the relational stability needed for spanning intertwined temporal and spatial *detachment* cannot be reduced to dynamic patterns of causal coupling because no causal coupling is even possible across such detachment. The cognitive agent may still be “ceaselessly melded” with the environment, but just not the whole world, just not all at once.

## Concluding Remarks

The notion of relational stability is akin to Brian Cantwell Smith’s metaphor of “intentional acrobatics” for the dynamic achievement of reference (Smith, 1996). It also suggests an initial answer to Nora Newcombe’s question “What IS location at A?” (Newcombe, 2001), namely *location* representation presupposes embodied dynamic tracking or *localization*. But maybe such tracking dynamics is also a “dynamics of permanence” (Luo, 2004)? Maybe we can also understand the development of object permanence as partly a developmental achievement of relational stability? In the end, uncovering the limitations of the DSM and the DSA is just a step towards what seems to be an exciting alternative understanding of *object permanence and representation as achieved relational stability that is necessitated by, accountable to, and made possible through the dynamics of embodiment*.

## References

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Baker, J. T., Harper, T. M., & Snyder, L. H. (2003). Spatial memory following shifts of gaze. I. Saccades to memorized world-fixed and gaze-fixed targets. *Journal of neurophysiology*, 89, 2564–2576.
- Beer, R.D. (1998). Framing the debate between computational and dynamical approaches to cognitive science. Commentary on van Gelder (1998). *Behavioral and Brain Sciences*, 21, 630.
- Bremner, J. G. (1978). Egocentric versus allocentric spatial coding in nine-month-old infants: Factors influencing the choice of code. *Developmental Psychology*, 14, 346–355.
- Bushnell, E. W., McKenzie, B. E., Lawrence, D. A., & Connell, S. (1995). The spatial coding strategies of one-year-old infants in a locomotor search task. *Child Development*, 66, 937–958.
- Camperi, M., & Wang, X.-J. (1998). A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of Computational Neuroscience*, 5, 383–405.
- Cornell, E. H. (1979). The effects of cue reliability on infants' manual search. *Journal of Experimental Child Psychology*, 28, 81–91.
- Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on A-not-B. *Child Development*, 56, 868–83.
- Diamond, A. (1990). The Development and Neural Bases of Memory Functions as Indexed by the AB and Delayed Response Tasks in Human Infants and Infant Monkeys. *Ann. of the N.Y. Acad. of Sci.*, 608, 267–317.
- Duhamel, J. R., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255, 90–92.
- Eliasmith, C., & Anderson, C.H. (2002). *Neural Engineering: Computation, Representation and Dynamics. in Neurobiological Systems*. Cambridge: MIT Press.
- Elman, J. L. (1995) Language as a dynamical system. In Port, R.F. & van Gelder, T. (Eds.) *Mind as Motion*, pp. 195–225, MIT Press.
- Gibson, E.J. & Pick, A.D. (2000). *An Ecological Approach to Perceptual Learning and Development*. New York: Oxford University Press.
- Gratch, G., & Landers, W. F. (1971). Stage IV of Piaget's theory of infant's object concepts: A longitudinal study. *Child Development*, 42, 359–72.
- Graziano, M., Hu, X., & Gross, C. (1997). Coding the locations of objects in the dark. *Science*, 277, 239–241.
- Harris, P. L. (1973). Perseverative errors in search by young infants. *Child Development*, 44, 29–33.
- Luo, J. (2004). *The Dynamics of Permanence*. PhD dissertation, Indiana University.
- Marcovitch, S., & Zelazo, P. D. (1999). The A-not-B error: Results from a logistic meta-analysis. *Child Development*, 70, 1297–1313.
- Mays, L. E., & Sparks, D. L. (1980). Dissociation of visual and saccade-related responses in superior colliculus neurons. *Journal of Neurophysiology*, 43, 207–232.
- Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP Model of the A-not-B task. *Developmental Science*, 1(2), 161–184.
- Newcombe, N. (2001). A spatial coding analysis of the A-not-B error: What IS "Location at A"? Commentary on Thelen et al (2001). *Behavioral and Brain Sciences*, 24, 57–58.
- Newcombe, N., & Huttenlocher, J. (2000). *Making space: The development of spatial representation and reasoning*. Cambridge, MA: Bradford Books/MIT Press.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Pouget, A. & Sejnowski, T. J. (2002). Dynamical remapping. In M. A. Arbib (ed). *The Handbook of Brain Theory and Neural Networks*, 2nd edition (pp.335-338). Boston: MIT Press.
- Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74, 1393–1417.
- Seung, H. S. (1996). How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences USA*, 93, 13339–44.
- Smith, B. C. (1996). *On the origin of objects*. Cambridge, MA: Bradford Books/MIT Press.
- 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.
- Stringer, S. M., Trappenberg T. P., Rolls, E. T. & Araujo, I. E. T. (2002). Self-organising continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network: Computation in Neural Systems*, 13, 217–242.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: Bradford Books/MIT Press.
- 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.
- Trappenberg, T. P. (2004). Continuous attractor neural network. In L. N. de Castro & F. J. von Zuben (Eds.) *Recent Developments in Biologically Inspired Computing*. Idea Group Publishing.
- van Gelder, T. J. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 1-14.
- Zhang, K. (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16, 2112–2126.