Grounding Cognitive-Level Processes in Behavior:

The View from Dynamic Systems Theory

Larissa K. Samuelson, Gavin W. Jenkins & John P. Spencer

Department of Psychology and DeLTA Center

University of Iowa

Keywords: Marr, Levels of Analysis, Dynamic Systems, Representations, Cognitive Processes, Word Learning, Emergence

Abstract

Marr's seminal work laid out a program of research by specifying key questions for cognitive science at different levels of analysis. Because Dynamic Systems Theory focuses on time and interdependence of components DST research programs come to very different conclusions regarding the nature of cognitive change. We review a specific DST approach to cognitive-level processes: Dynamic Field Theory. We review research applying dynamic field theory to several cognitive-level processes: object permanence, naming hierarchical categories, and inferring intent, that demonstrate the difference in understanding of behavior and cognition that results from a DST perspective. These point to a central challenge for cognitive science research as defined by Marr—emergence. We argue that appreciating emergence raises questions about the utility of computational level analyses and opens the door to insights concerning the origin of novel forms of behavior and thought (e.g., a new chess strategy). We contend this is one of the most fundamental questions about cognition and behavior.

Introduction

Marr's landmark book *Vision* (1982) has had numerous lasting influences on the field of cognitive science. Most notably, Marr defined a research program for the field of vision but more generally, his tri-level approach created key questions for cognitive science at different levels of analysis: what abstract problem is the system designed to solve, what algorithm does it use on what representations, and how is this implemented in the brain. While these questions enable explanations at each level that are relatively independent, Marr emphasized that an understanding of any complex system at only one level is ultimately insufficient—"…one must be prepared to contemplate different kinds of explanation at different levels of description that are linked, at least in principle, into a cohesive whole…" (Marr, 1982 pp. 20). Similarly, he wrote, "To understand fully a particular machine carrying out a particular information-processing task, one has to [understand both hardware and information processing]. Neither alone will suffice," (Marr,1982 pp. 5). Nevertheless, Marr stressed the computational level and stressed the necessity of analyzing the task presented to the system as critical to defining the scope of the research question.

We argue that compared to other approaches, the focus on time and integration across system components that are central to Dynamic Systems Theory (DST) lead to different answers to Marr's questions and, critically, suggests that starting with the computational-level analysis can lead to serious misunderstandings of behavior (see also, McClelland et al., 2010; other papers in this topic). Below we present an overview of the DST approach, focusing on Dynamic Field Theory (Schöner et al., in press; Spencer, Perone & Johnson, 2009), which translates DST concepts into formal models that integrate neural, cognitive and behavioral levels. We then provide three examples that demonstrate the contrast in conclusions reached when the research question is defined by an abstract computational-level analysis versus an DST analysis focused on behavioral change. We highlight the role of formal dynamic systems approaches like DFT for grounding our understanding of cognitive phenomena, real-time processes and neural population dynamics. We close by considering the implications of these examples for the question for this topic: what is the role of process-oriented theories and cognitive constructs in understanding brain/behavior relationships?

Dynamic Systems Theory

Dynamic Systems Theory focuses on the processes of change over time in complex systems. It views behavior, including cognition, as emerging from the interaction of multiple softly assembled components that are mutually influential and evolve over multiple embedded timescales (Beer, 2000; Fischer & Bidell, 1998; Kelso, 1995, 2000; Lewis & Liu, 2011; Newell & Molenaar, 1998; Port & Van Gelder 1995; Thelen & Smith, 1994; ; Spivey 2007; van der Maas & Molenaar, 1992, van Gelder, T., 1998). Appreciating that behavior is emergent and softly assembled means understanding that behavior is the product of multiple components brought together in a moment of time based on the particular context, task, and history of the organism. The interaction of these components is not pre-specified or deterministic; thus, the particular assembly and resulting behavior are unique and variable from moment-to-moment and across specific contexts. For example, the particular muscles and joint angles used to pick up your coffee cup will change based on many factors including the starting point of the reach, the introduction of obstacles in the path, or the weight of a new watch on your wrist (see also Spencer et al., 2006, Thelen & Smith, 1994 for additional examples). The implication of this is that identical behavioral outputs can be the result of very different specific processes in different contexts or in organisms with different histories. Mutual interactivity means interactions proceed in both directions; not only does attentional selection influence the contents of visual working memory, but the contents of visual working memory reciprocally influence selective attention (Hollingworth et al. in press; Schneegans, et al. 2014; Schneegans et al., in press). Viewing timescales as embedded means appreciating that the different timescales cognitive science often considers are not independent and cannot be studied without recourse to each other.

The application of dynamic systems to psychological phenomena has its origins in motor control and perception-action (see for discussion Spencer & Schöner, 2003), but these approaches are often criticized for being overly metaphorical and not applicable to cognitive-level processes (Griffiths, Chater, Kemp, Perfors & Tenenbaum, 2010). Dynamic Field theory (DFT, Schöner, Spencer, & the DFT Research Group, in press) has emerged as a response to such critiques. DFT is an embodied, dynamic systems approach to cognitive-level processes based on an understanding of brain function at the level of neural population dynamics (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999; Jancke, et al., 1999). The basic building block of formal Dynamic Neural Field models is a field of metrically-organized neurons, that is neurons in this field are structured such that those close together have receptive fields that respond to similar feature values while those farther apart have receptive fields tuned to very different feature values (Schöner, Spencer, & the DFT Research Group, in press. These neurons interact according to a local excitation/lateral inhibition function (Spencer, Austin & Schutte, 2012), a common form of interaction in neural models of cortical function (Durstewitz, 2000). Neural fields, like local neural populations in the brain (Cohen & Newsome, 2009; Amari, 1977; Fuster, 2003), move into and out of attractor states, reliable patterns of activation that the neural population maintains in the context of particular inputs. For instance, when presented with visual input, neural populations in visual cortex create stable 'peaks' of activation representing a location and set of features.

Whereas DST provides a set of theoretical tools, mathematical formalizations, and empirical approaches for understanding complex systems (Lewis & Granic, 2000; Newell & Molenaar, 1998; Port & van Gelder, 1995; Thelen & Smith, 1994). DFT adds a formal approach for analyzing and understanding cognitive-level processes. Furthermore, because DFT links cognitive-level processes to both neural population dynamics and behavior, and focuses on how behavior evolves over time, it explicitly integrates brain and behavior and provides a formal method for understanding both behavior and behavior change. Critically, as we demonstrate in the following examples, explicitly linking cognitive-level processes to neural population dynamics and behavior, radically changes how phenomena are analyzed and our understanding of cognition.

The A-not-B error

The A-not-B error refers to the finding that infants (and adults under the right circumstances, Spencer & Hund, 2002) will continue to search for a hidden object in the place it has been previously found even after observing it being hidden in another location. Thus, the computational analysis suggested that the task was to remember that something existed even when it is out of sight (see for discussion McClelland et al, 2010; Thelen & Smith, 1994, Smith & Thelen, 2003), and the behavioral change was explained at the algorithmic and representational level in terms of a deficit in cognitive processing. Ten-month-old infants, who make the error reliably, were said to lack a specific representation called the “object concept” (Piaget, 1954) and did not understand that the object continued to exist after it was hidden. In contrast, 12-month-old infants, who are typically able to search correctly, were said to represent the object mentally even when it was not visually available.

In contrast, a DST analysis of this phenomenon revealed critical influences of the motor system and body by recognizing the role of infants’ ability to stabilize and repeat a reach to the same location (Thelen & Smith, 1994). Later research supported this analysis, demonstrating that the number of reaches to A, the posture of the infant, and the use of wrist-weights to change the specifics of the reaching trajectory all determined whether infants made the error (Smith, Thelen, Titzer & McLin, 1999). Likewise, a formal DNF model of developmental changes in the A-not-B error captured the relation between the body and neural populations that represent reaching direction (Smith et al., 1999; Thelen et al.; 2001) and made predictions about behavioral conditions that could both increase and decrease the likelihood of the behavior. Subsequent studies have supported the DNF models’ predictions for infants (e.g., Clearfield et al., 2009), toddlers (Spencer, Smith, & Thelen, 2001) and school-aged children (Hund and Spencer, 2003).

The final blow to the computational-level analysis that object permanence was the goal, and the algorithm-level explanation of a missing representation, however, was the demonstration that hidden toys were not even needed to produce the error—infants will make the error even when reaching to visible objects (Smith et al., 1999). Clearly, the A-not-B error cannot be about a failure to represent an object's existence if the object is in full view. Rather, DST behavioral analysis suggests that changes in motor components at two different timescales affected behavioral change: reaches made over the course of an experiment, and motor stability over the course of development. Thus, this work demonstrates that the behavior is not just about cognitive-level processes but about the body as well—the representation of things in the world is tightly coupled to the dynamics of the body (see also Eerland, Guadalupe & Zwaan, 2011; Zwaan & Kaschack, 2009 for related results in adults and Zwaan, van der Stoep, Guadalupe & Bouwmeester, 2012 for discussion). This has been recently demonstrated by instantiating the DNF model in an autonomous robot and quantitatively simulating experimental data from infants (Schöner, et al., in press).

Thus, this program of work demonstrates how a computational level analysis that defined the research question with respect to a cognitive-level process only—a missing representation—missed the crucial influence of the integrated body-brain system that is the basis of the A-not-B error. Some might argue that the dynamics apparent in this example reflect the motor system used in the behavioral response rather than cognitive-level processes. To counter this argument, we next turn to recent research on adults' and children's acquisition of names for hierarchical categories. This program of work has applied a DST perspective to a phenomenon—the suspicious coincidence effect (SCE)—previously thought to reflect abstract reasoning about category membership. A dynamic systems approach has integrated the SCE with the trajectory of vocabulary development and with lower-level processes such as perceptual comparison and working memory. In addition, this perspective grounds the SCE in neural population dynamics via dynamic field theory.

Generalization of names for hierarchical categories

The challenge in learning names for hierarchically-nested categories is in inferring the correct level of application of a novel word based on limited exposure. For example, when the novel word “Fep” is heard in association with a single Dalmatian, its meaning is ambiguous. Does “Fep” refer to only that particular dog, or to other Dalmatians, or even to other breeds of dog or other species of animal? Both children and adults must resolve these ambiguities when learning words for anything that might be part of a taxonomy.

Xu and Tenenbaum (2007) analyzed this generalization problem at an exclusively computational level, focusing only on Bayesian inductive reasoning as an overall explanation of learners’ strategy and logic, without respect to dynamics that might occur during the task or timescales beyond that of the task itself. Xu and Tenenbaum explicitly separated Marr’s levels in their approach, stating “Our analysis of word learning focuses on what Marr (1982) called the level of computational theory. We have tried to elucidate the logic behind word learners’ inductive inferences, without specifying how that logic is implemented algorithmically in the mind or physiologically in neural hardware,” (pp. 270).

Xu and Tenenbaum (2007) presented adults and children with one or more exemplars of a category and a novel label. The exemplars were either identical (multiple Dalmatians), members of the same basic level category (different breeds of dog), or members of a superordinate category (different species of animal). Participants were then asked whether the novel word generalized to test objects chosen from different taxonomic levels. Based on their computational-level analysis, Xu and Tenenbaum made a novel prediction—the “suspicious coincidence effect”— that both preschoolers and adults would generalize the novel words more narrowly when they were presented with three identical exemplars rather than just one. The idea is that it would be highly unlikely to hear the same word applied to three identical but separate instances of a category unless that word referred to that level of the category. This is what Xu and Tenenbaum found: “Fep” applied to one Dalmatian led to generalization to all dogs, whereas “Fep” applied to three identical Dalmatians led to generalization only to other Dalmatians. Xu and Tenenbaum suggested that rational Bayesian statistical reasoning was the basis for children’s and adults’ ability to quickly learn names. They presented a formal Bayesian model that fit the results from both adults and children. Furthermore, they convincingly argued that their model was the only one that would predict the suspicious coincidence effect and the only one to account for the data.

Subsequent research, however, suggests that details of the neural and behavioral process by which novel names are perceived, remembered, and generalized matter greatly in the suspicious coincidence effect, and change our understanding of the computational-level task. Spencer, Perone, Smith, and Samuelson (2011) examined the influence of stimulus presentation in the word generalization task. Their motivation was decades of research suggesting that generalization depends critically on variability in the presented instances and contexts of exposure, as well as whether exemplars are presented together (allowing direct comparison) or in sequence (forcing more global similarity analysis, see Spencer et al., 2011). In addition, prior work had shown a narrowing of representations in a visual working memory study when similar objects appeared near one another in time and space (Johnson, Spencer, Luck, & Schöner, 2009), a finding captured by a dynamic field model of visual working memory. In particular, these interactions produce narrowing of activated representations when multiple similar stimuli are seen close in time and space (see Fig. 1). Based on this prior work then, Spencer et al. (2011) presented adults with either three subordinate-level exemplars at once or sequentially and found that that the SCE was eliminated in the sequential condition. Furthermore, Spencer et al. (2011) were able to eliminate the SCE entirely by presenting six exemplars—an amount predicted to overwhelm working memory—sequentially. Thus, the DST analysis of the suspicious coincidence effect led to both a demonstration of when it does not occur and a critical contradiction of the rational analysis—more exemplars produced a reversal of the effect rather than the strengthening predicted by the Bayesian account. We have recently captured these data in a formal DNF model (Jenkins, Samuelson & Spencer, 2014; Samuelson, Spencer & Jenkins, 2013), thereby grounding the suspicious coincidence effect in neural population dynamics.

---------------------------- Insert Fig. 1 about here ----------------------------

In addition to these effects which emerge in real-time in the task, we have also found changes in the SCE over developmental timescales. In a series of studies, we have examined the relation between the category knowledge children bring to the SCE task and the strength of the effect they demonstrate. According to the Bayesian account, more knowledge should lead to a stronger suspicious coincidence effect. In contrast, Jenkins, Samuelson, Smith & Spencer (in press; see also Samuelson, Spencer & Jenkins, 2013) showed that children who had less knowledge of the familiar categories used in Xu and Tenenbaum's task showed a strong suspicious coincidence effect, whereas children entering with more knowledge of the English categories showed no effect. Thus, the suspicious coincidence follows a nonlinear, U-shaped curve over the developmental timescale of category learning.

Similar to Spencer et al. (2011), the discontinuity between the Bayesian model's predictions of category representations at the algorithmic level and the U-shaped curve of suspicious-coincidence behavior suggests that separating these levels of analysis leaves much of what is critical for understanding the behavior unexamined (see also Jenkins, Samuelson & Spencer, 2011). Rather, grounding the suspicious coincidence effect in neural population dynamics and time-dependent processes via the DNF model revealed critical influences that were not detectable when the abstract levels of analysis were explicitly separated. Thus, this work again highlights how a DST analysis results in a more integrated and grounded understanding of even a very high-level, abstract behavior.

Our final example focuses on another behavior that has been taken to represent one of the most abstract and potentially ungrounded abilities that humans demonstrate—inferring the thoughts of others. This work demonstrates how a DST analysis appreciating that behavior happens in the here and now of space and time creates an integrated neural, cognitive and perceptual-motor understanding of early referent selection.

Using space to infer reference

A central question in the field of early word learning is how is it that young learners are so good at inferring the meaning of novel words given that the possible referents in a naming situation can range from the objects present in the visual scene to properties of those objects, ongoing actions, and so on. Answers to this question range from the proposal that internal constraints limit the possibilities the child considers at the moment of naming (i.e., the whole object assumption, the taxonomic assumption and mutual exclusivity; see Markman, 1992 for review); to attentional biases such as the shape bias that direct children's focus to the most critical object features based on their prior word learning history (e.g., Landau, Smith & Jones, 1988). Also popular are social-pragmatic theories that suggest children use knowledge of speakers' intentions to determine the meaning of novel words.

Consider a seminal study by Baldwin (1993) examining young children's ability to read the referential intent of a speaker. A schematic of the task is presented in the far left column of Fig. 2. A novel object is presented to a 20-month-old child for exploration and manipulation on one side of a table. This object is then removed and a second novel object is presented on the other side of the table and the child is again allowed to reach for, grasp and explore the object. This is repeated for a set of familiarization trials. Both objects are then placed in separate opaque buckets on either side of the table. The experimenter looks into one bucket and says "Modi!" The object from the *other* bucket is then taken out and placed on its side of the table. It is removed after the child examines it and the other object is placed on the table. After examination, this item is also removed. Both objects are then placed on a tray on the center of the table. The tray is pushed toward the child, and the experimenter asks, "Can you get me the modi?" Children retrieve the object that was in the bucket the experimenter was looking in when she said the novel word .70 of the time. Baldwin interpreted this result as suggesting children understood the pragmatic use of eye gaze as an intentional cue (Baldwin, 1993).

---------------------------- Insert Fig. 2 about here ----------------------------

In contrast, we have used a DST analysis and a formal Dynamic Neural Field model to demonstrate that children's word-referent mappings in Baldwin's task are based instead on the neural encoding of the objects' spatial locations and the recall of those locations at the point of naming (Samuelson, et al., 2011). That is, during the familiarization trials, the children's behaviors—looking at the objects, reaching for them, manipulating them and attending as each is removed—create associations between each of the novel objects and their unique locations in the space of the task. Thus, when the experimenter looks into a bucket placed at one of those unique locations and says the name the child's memory of the object previously seen and acted on at that location is recalled and bound to the novel name. Thus, the child is able to link the novel name to the correct object via the space in which her body, her attention, her actions and the object itself occur.

---------------------------- Insert Fig. 3 about here ----------------------------

We have tested several predications of this account (Figs. 2 and 3) and quantitatively simulated children's behavior with the same DNF model used to capture the suspicious coincidence effect. This work demonstrates, for instance, that children can bind the novel word to the correct object even in when the experimenter points to an empty location on the table. We have also demonstrated that space is special in facilitating these mappings: associating the potential referents with unique colors and providing the name in the presence of one of these colors does not support mapping.

Furthermore, children learn words better when their parents keep objects in consistent spatial locations when teaching them. Thus, a nonobvious factor—the history of where objects have been placed in a task—matters in young children's early word learning. This initially surprising finding fits with research showing that both adults and children will look back to the location in which a fact or sound was previously presented when trying to recall that information (Richardson & Spivey, 2000; Richardson & Kirkham, 2004). It is also consistent with the use of space for reference in sign languages and in gestural communication. A computational-level analysis that defines the task as inference from the intentions of another person does not predict the connection between space and word learning. Thus, a DST analysis again provides fundamentally different view on cognition and an integrated understanding of behavior grounded in terms of both neural dynamics and sensori-motor processes.

Conclusion

Each of the examples above show how a DST approach results in a radically different view of cognition compared to research that starts with a computational level analysis. We do not think these differences are a coincidence. Rather, they reflect a deep challenge in trying to infer a computational-level theory from an inherently non-linear, complex, and emergent system (see also McClelland et al., 2010). Emergence—the idea that behavior arises through the interaction of many components over time without recourse to explicit coding and without needing to be hardwired—played a key role in each of the examples we reviewed. For instance, according to Thelen et al. (2001), A-not-B errors are reduced when excitatory neural interactions increase sufficiently to support a qualitatively new attractor state—the self-sustaining state where neural activation patterns are actively maintained even in the absence of input. Recent work demonstrates how this new attractor state can arise through a variant of Hebbian learning as infants repeatedly reach to different locations in space (Schöner, et al., in press; Perone & Spencer, 2013). Thus, Hebbian learning gives rise to a new emergent ability—the ability to actively and flexibly remember a cued target location.

Of course, viewing cognition and behavior as the emergent product of a complex, time-extended system presents challenges for how we do our science. One of the lasting influences of Marr's seminal work was to lay out a framework for cognitive research and a method of analysis that enabled investigators to define the scope of the problem in a tractable way that was rigorous but not reductionist. Fortunately, DST also provides organized ways of analyzing a system and defining the scope of a research question. Specifically, if components are strongly coupled, we have to care about them and their interactions as a set because they all matter for the behavior or cognition in question. If, however, components are weakly coupled, then we can study each more independently (at least within context). A "subsystem" in DST, therefore, is defined as a collection of strongly coupled components that function as an integrated system, actively resisting perturbations from, for instance, external inputs, and showing only weak coupling to other components (Schöner, 1995). Thus, while scientific examination requires carving the system into analyzable sub-systems, the joints used to carve a dynamic system are defined relative to the specific behavior or phenomenon under examination; they are the places where a *behavioral* analysis suggests weak coupling among components. Such joints are not always readily apparent. For instance, one might think that learning names for hierarchically nested categories might be immune to perceptual-level processes, but this was not the case with the suspicious coincidence effect.

Viewing cognition and behavior as the product of a complex, time-extended system also means that the major unit of analysis and subject of study is defined in terms of a process of change and the trajectory of behavior over time (Beer, 1995, 2000; van Gelder & Port, 1995; Spivey, 2007). DST approaches strive to understand the next state of the system based on its current state. They recognize the history of the system as a critical influence on its current and future states. Thus, to understand cognition, we must examine thinking with respect to that history and over multiple timescales. This is evident in the referent selection example, where a child's history of past events enables the mapping of novel words to referents without having to read the minds of social partners. This DST focus on time and trajectories opens the door to both developmental approaches, defined not as the study of children but as the study of the system at multiple points in time (Smith & Thelen, 2003; Thelen & Smith, 2006), and to an appreciation of individual differences (see, e.g., Hollenstein & Lewis, 2006).

It is, of course, possible to do a computational level analysis that incorporates some of the critical components of a dynamic systems perspective—for example agent/environment interactions or the influence of the body. What we have argued, however, is that the concept of emergence is fundamentally at odds with a computational-level analysis because such analyses start by defining the agent's behavior in terms of goals that are independent of the neural, physical, and historical processes that produce behavior. This does not mean that computational-level thinking cannot still be impactful. It is clear that Marr's *Vision* has provided great benefit to cognitive science, and that computational analyses have led to advancements in many areas of cognitive science (see papers in *Perception* v41, 2012 for review and commentary). Nevertheless, the specific examples reviewed here demonstrate that a dynamic systems perspective provides a valuable alternative framework for cognitive science, because it appreciates that the neural-behavioral system is complex, non-linear, and emergent (see also McClelland et al., 2010).

Furthermore, taking emergence seriously opens the door to answers to one of the most critical questions concerning cognitive-level processes: where do new forms of behavior and thought such as riding a bicycle or a new chess strategy, come from (see also Poggio, 2012)? This is a question that other approaches to cognition typically ignore, sometimes going so far as to build in miniature pre-formed versions of new behaviors (Smith & Pereira, 2009; Smith, 2001; McClelland, 2010). In contrast, a DST perspective suggests that new behaviors can emerge organically as subtle changes in the components of the system—the strength of a muscle that stabilizes a reach, the presentation of objects sequentially in time, the association of an object with a specific location—softly assemble and produce changes in cognition and action. But does emergence require that we throw up our hands and say everything matters? No, we can (and should) be analytical about our research question. DST offers a way to determine a unit of study, that is, what a ‘subsystem’ is and how to separate one subsystem from another. And DST identifies the focus of study—trajectories of behavior and behavioral change through time. The examples here illustrate the comparative utility of this perspective, and how changing our focus can lead to a deeper understanding about the organization of behavior and how behavior changes over learning and development.

Acknowledgments

The authors wish to thank Cathleen Moore for helpful discussions of Marr's vision. Writing of this article was supported by Award Number R01HD045713 from the Eunice Kennedy Shriver National Institute of Child Health & Human Development to LKS, a National Defense Science & Engineering Graduate Fellowship to GWJ, and Award Number R01MH62480 from the National Institute of Mental Health to JPS. The content is solely the responsibility of the authors and does not necessarily represent the official views of the awarding agencies.

References

Amari, S. (1997). Dynamics of pattern formation in lateral inhibition type neural fields. *Biological Cybernetics, 27,* 77-87.

Baldwin, D.A. (1993). Early referential understanding: Infants' ability to recognize referential acts for what they are. *Developmental Psychology, 29*. 832–843.

Beer, R.D. (2000) Dynamical approaches to cognitive science. T*rends Cognitive Science*, *4*, 91-99

Beer, R.D. (1995). Computational and Dynamical Languages for Autonomous Agents. In Port, R.F. & van Gelder, T. (Eds.) *Mind as Motion: Explorations in the Dynamics of Cognition.* MIT Press, Boston, MA, pp 1-43.

Clearfield, M. W., Dineva, E., Smith, L. B., Diedrich, F. J., & Thelen, E. (2009). Cue salience and infant perseverative reaching: Tests of the dynamic field theory. *Developmental Science, 12*(1), 26-40.

Cohen, M. R. & Newsome, W. T. (2009). Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *The Journal of Neuroscience, 29*(20), 6635-6648.

Durstewitz, D. Neurocomputational models of working memory. *Nature Neuroscience.* 3, 1184 (2000).

Eerland, A., Guadalupe, T. M., & Zwaan, R. A. (2011). Leaning to the left makes the Eiffel tower seem smaller: Posture-modulated estimation. *Psychological Science, 22*(12), 1511-1514.

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.

Fischer, K.W. and Bidell, T.R. (1998) Dynamic development of psychological structures in action and thought. In Handbook of Child Psychology In Theoretical Models of Human Development (Vol. 1), 5th edn, (Lerner, R.M. and Damon, W., Eds.), pp. 117–121, Plenum.

Fuster, J. M. (2003). *Cortex and mind: Unifying cognition*, Oxford University Press.

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*(8), 357–64. doi:10.1016/j.tics.2010.05.004

Hollingworth, A., Matsukura, M. & Luck, S. J., (in press). Visual working memory modulates low-level saccade target selection: Evidence from rapid generalized saccades in the global effect paradigm. *Journal of Vision*.

Hollenstein, T., & Lewis, M. D. (2006). A state space analysis of emotion and flexibility in parent-child interactions. *Emotion, 6*(4), 656-662.

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

Jancke, D., Erlhagen, W., Dinse, H. R., Akhavan, A. C., Giese, M., Steinhage, A., & Schöner, G. (1999). Parametric population representation of retinal location: Neuronal interaction dynamics in cat primary visual cortex. *The Journal of Neuroscience, 19(20)*, 9016–9028.

Jenkins, G. W., Samuelson, L. K., & Spencer, J. P. (2011). Come down from the clouds:

Grounding Bayesian insights in developmental and behavioral processes. *Behavioral and Brain Sciences, 34*(4), 204-206.

Jenkins, G.W., Samuelson, L.K., & Spencer, J.P. (2013). *Learning words in space and time: Contrasting Bayesian and Dynamic Neural Field Accounts of the Suspicious Coincidence Effect.* Manuscript submitted for publication.

Jenkins, G., Samuelson, L. K., Smith, J. R., & Spencer, J. P. (in press). Non-Bayesian noun generalization in 3-5-year-old children: Probing the role of prior knowledge in the suspicious coincidence effect. *Cognitive Science*.

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science, 20*(5), 568-577.

Kelso, J.A.S. (1995) *Dynamic Patterns*, MIT Press

Kelso, J.A.S. (2000) Principles of dynamic pattern formation and change for a science of human behavior. In Developmental Science and the Holistic Approach (Bergman, L.R. and Cairns, R.B., Eds.), pp. 63–83, Erlbaum.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3*(3), 299-321.

Lewis, M.D. and Granic, I. (2000) Emotion, Development, and Self-Organization: Dynamic Systems Approaches to Emotional Development, Cambridge University Press.

Lewis, M. D., & Liu, Z. (2011). Three time scales of neural self-organization underlying basic and nonbasic emotions. *Emotion Review, 3*(4), 416-423.

Markman, E. M. (1992). Constraints on word learning: Speculations about their nature, origins, and domain specificity. In Modularity and constraints in language and cognition (Gunnar, M. R. & Maratsos, M., Eds.), pp. 59-101, Erlbaum.

Marr, D. C. (1982). *Vision.* San Francisco: Freeman.

McClelland, J. L. (2010). Emergence in Cognitive Science. *Topics in Cognitive Science*, *2*(4), 751–770. doi:10.1111/j.1756-8765.2010.01116.x

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, *14*(8), 348–56. doi:10.1016/j.tics.2010.06.002

Newell, K.M. & Molenaar, P.C.M. (1998) Applications of Nonlinear Dynamics to Developmental Process Modeling, Erlbaum.

Perone, S. & Spencer, J. P. (2013). Autonomous visual exploration creates developmental change in familiarity and novelty seeking behaviors. *Frontiers in Psychology.* doi: 10.3389/fpsyg.2013.00648#sthash.q4WQcYVD.dpuf

Piaget, J. (1954). *The construction of reality in the child*. Abingdon, Oxon: Routledge Press.

Poggio, T. (2012). The levels of understanding framework, revised. *Perception*, 41(9) 1017-1023.

Port, R. F., & van Gelder, T. (1995). *Mind as motion: Explorations in the dynamics of cognition* Cambridge, MA, US: The MIT Press.

Richardson, D. C., & Spivey, M.J., (2000). Representation, space and hollywood squares: Looking at things that aren't there anymore. *Cognition, 76*(3), 269.

Richardson, D. C., & Kirkham, N.Z., (2004). Multimodal events and moving locations:

Eye movements of adults and 6-month-olds reveal dynamic spatial indexing. *Journal of Experimental Psychology: General, 133*(1), 46.

Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. *PLoS One, 6*(12) .

Schneegans, S., Spencer, J. P., Schöner, G., Hwang, S. & Hollingworth, A. (2014). *Dynamic interactions between visual working memory and saccade target selection.* Manuscript submitted for publication.

Schneegans, S., Lins, J. & Spencer, J.P. (in press). Integration and selection in dynamic fields: Moving beyond a single dimension. In G. Schöner, J.P. Spencer & the DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory. New York: Oxford University Press*.

Schöner, G. (1995). Recent developments and problems in human movement science and their conceptual implications. *Ecological Psychology*, *7*(4), 291-314.

Schöner, G., Faubel, C., Dineva, E. & Bicho, E (in press). Embodied neural dynamics. In G. Schöner, J.P. Spencer & the DFT Research Group (Eds.), *Dynamic Thinking: A Primer on Dynamic Field Theory. New York: Oxford University Press*.

Schöner, G., Spencer, J.P., & the DFT Research Group (in press). Dynamic Thinking: A Primer on Dynamic Field Theory. Oxford University Press.

Smith, L. B. (2001). How domain-general processes may create domain-specific biases. In M. Bowerman & S. Levinson (Eds.), *Language Acquistion and Conceptual Development*. Cambridge University Press.

Smith, L., & Pereira, A. (2009). Shape, action, symbolic play and words: Overlapping loops of cause and consequence in developmental process. In S. P. Johnson (Ed.), *Neo-constructivism: The new science of cognitive development*. Oxford University Press.

Smith, L. B., & Samuelson, L. K. (2003). Different is good: Connectionism and dynamic systems theory are complementary emergentist approaches to development. *Developmental Science, 6*(4), 434-439.

Smith, L. B., & Thelen, E. (2003). Development as a dynamic system. *Trends in Cognitive Sciences, 7*(8), 343-348.

Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review, 106*(2), 235-260.

Spencer, J. P., Austin, A., & Schutte, A. R. (2012). Contributions of dynamic systems theory to cognitive development. *Cognitive Development, 27*(4), 401-418.

Spencer, J. P., Clearfield, M., Corbetta, D., Ulrich, B., Buchanan, P., & Schöner, G. (2006). Moving toward a grand theory of development: in memory of Esther Thelen. *Child Development*, *77*(6), 1521–38. doi:10.1111/j.1467-8624.2006.00955.x

Spencer, J. P., & Hund, A. M. (2002). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General, 131*(1), 16-37.

Spencer, J.P., Perone, S., & Johnson, J.S. (2009). The dynamic field theory and embodied cognitive dynamics. In J.P. Spencer, M.S. Thomas, & J.L. McClelland (Eds.) Toward Unified Theory of Development: Connectionism and Dynamic Systems Theory Re-Considered (p. 86-118). New York: Oxford University Press.

Spencer, J. P., Perone, S., Smith, L. B., & Samuelson, L. K. (2011). Learning words in space and time: Probing the mechanisms behind the suspicious-coincidence effect.

Psychological Science, 22(8), 1049-1057.

Spencer, J.P. & Schöner, G. (2003). Bridging the representational gap in the dynamic systems approach to development. *Developmental Science*, *6*, 292-412.

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 two-year-olds. *Child Development, 72*(5), 1327-1346.

Spivey, M. (2007). *The continuity of mind* New York, NY, US: Oxford University Press.

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Brain and Behavioral Sciences, 24*(1), 1-86.

Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.

Thelen, E., & Smith, L. B. (2006). Dynamic systems theories. In R. M. Lerner W. Damon (Ed.), (pp. 258-312). Hoboken, NJ, US: John Wiley & Sons Inc.

van der Maas, H.L.J. and Molenaar, P.C.M. (1992) Stage-wise cognitive development: an application of catastrophe theory. Psychol. Rev. 99, 395 – 417

van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences, 21*(5), 615-665.

van Gelder, T. & Port, R.F. (1995). It's about time: An overview of the dynamical approach to Cognition. In Port, R.F. & van Gelder, T. (Eds.) *Mind as Motion: Explorations in the Dynamics of Cognition.* MIT Press, Boston, MA, pp 1-43.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review, 114*(2), 245-272.

Zwaan, R. A., & Kaschak, M. P. (2009). Language in the brain, body, and world. In P. Robbins M. Aydede (Ed.), (pp. 368-381). New York, NY, US: Cambridge University Press.

Zwaan, R. A., van, d. S., Guadalupe, T., & Bouwmeester, S. (2012). Language comprehension in the balance: The robustness of the action-compatibility effect (ACE). *PLoS One, 7*(2)

Fig. Captions

Fig. 1. The narrowing effect revealed by Johnson et al., 2009. The two peaks on the right that are in close proximity in the dynamic neural field are narrower than the peak on then left even though the original inputs (lighter curves) were the same widths. This narrowing is caused by the inhibitory interactions of the field.

Fig. 2 Schematic of experimental tasks in Samuelson et al., 2011; including a replication of Baldwin, 1993 (far left column).

Fig. 3. Performance of children and model in Samuelson et al. 2011 Experiments 1-5. Children's percent of correct choices for each experiment (black bars) with standard deviations (range of error bars). *s indicate performance significantly above chance (.50 in a two item forced-choice task). The mean performance of the Dynamic Neural Field model (across 12 batches of simulations) for all experiments is also shown (white bars). Error bars show the standard deviation of the model's performance (across 12 batches of simulations) per condition, relative to the target means.