Title
The Emergence of Rule Use: A dynamic neural field model of the DCCS

Permalink
https://escholarship.org/uc/item/92s6h86p

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 30(30)

ISSN
1069-7977

Authors
Buss, Aaron
Spencer, John P.

Publication Date
2008

Peer reviewed

The Emergence of Rule-Use: A Dynamic Neural Field Model of the DCCS

Aaron Buss (aaron-buss@uiowa.edu) and **John P. Spencer (john-spencer@uiowa.edu)**
Department of Psychology, University of Iowa, 11 Seashore Hall East
Iowa City, IA 52242 USA

Abstract

The Dimensional Change Card Sort (DCCS) task has been extensively used to study the development of executive functioning in 3- to 5-year olds. Despite the interest in using the DCCS to study neural development and cognitive deficits in special populations, there has not been a formal neurally-based theory of children's behavior and development in this task. Here, we present a dynamic neural field (DNF) model that captures children's behavior, as well as development, in the standard DCCS task. We also show simulations of two other variants of the standard task—the Negative Priming version and the Absolute Negative Priming version—that highlight the extensive coverage of our new theory.

Keywords: neural network models; dynamic systems theory; DCCS; executive function; perseveration

The DCCS and Neural Development

The Dimensional Change Card Sort task (DCCS) and other rule-based card sorting tasks provide an index of executive functioning both in early development and disordered populations (Zelazo et al., 2003; Zelazo, 2006). In the DCCS task, children sort cards either by the dimension of shape or color and are then instructed to switch and sort by the other dimension. Target cards are displayed on the trays where children place their selections, and the test cards that are sorted match each target card along one dimension (see figure 1). Typically, 3-year-olds perseverate and continue to use the first set of rules after they have been instructed to switch. This task, then, measures aspects of cognitive flexibility, and children's behavior has been described in terms of rule-representation (Zelazo et al., 2003), object re-description (Kloo & Perner, 2005), inhibitory or dis-inhibitory control (Zelazo et al., 2003), or attentional control (Kirkham, Cruess, & Diamond, 2003).
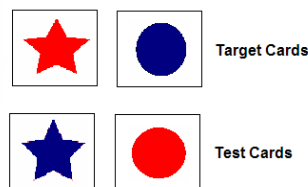

Figure 1: Example of cards used in the DCCS.

More recently there has been a growing interest in using such tasks as an index of neural functioning. Specifically, these tasks have been linked to the development of the prefrontal cortex (Bunge & Zelazo, 2006; Crone et al., 2006) and have been used to explore cognitive deficits in children with autism (Zelazo et al., 2002; Colver, Custance, & Swettenham, 2002) and ADHD (Mulas et al., 2006). Despite this interest in neural functioning and

developmental process, there is no neurally-based theory that captures developmental changes in performance in this task. The most comprehensive theory, the Cognitive Complexity and Control theory (CCC), describes children's performance in terms of hierarchical rule-representation (Zelazo et al., 2003; Bunge & Zelazo, 2006); however, describing children's behavior in terms of 'rules' raises some critical questions about the nature of the cognitive processes at work. It is unclear, for instance, how a hierarchical rule structure could be implemented in real-time in a nervous system. Similarly, ties to known changes in neural development have remained largely at the descriptive level. Morton and Munakata (2001) have made attempts to move explanations of performance in the DCCS in a more neurally plausible direction with a recent PDP model. This PDP model has generated novel predictions which have been empirically supported. However, this model has only been used to capture a subset of the data and, therefore, does not achieve the same degree of coverage as CCC theory.

In the present report, we describe a Dynamic Field Theory of children's behavior in the DCCS and present simulations of core findings in this task that previously have only been captured heuristically by CCC theory. In contrast to the PDP model, we use a more neurally-grounded view of the processes that underlie the representations of shapes, colors, and space in this task using a generic two-layered neural architecture. The proposed model provides an account of children's behavior and development that is couched in a general framework originally proposed to capture the integration of 'what' (i.e., ventral stream) and 'where' (i.e., dorsal stream) visual pathways.

Dynamic Field Theory

The Dynamic Field Theory was initially applied to issues in behavioral development within the context of infants' perseverative responses in the Piagetian A-not-B task (Thelen et al., 2001). Next, this theory was extended to explain decision-making and spatially-grounded behavior in spatial working memory tasks (see Schutte, Spencer, & Schöner, 2003; Simmering, Schutte, & Spencer, 2007). More recently, this theory has been used to capture the representation of objects in a way that links features to a spatial frame of reference (Johnson, Spencer, & Schöner, in press). Because the DCCS task requires children to make decisions about what features should go where in the task space, this theoretical framework provides a natural fit to this task. Moreover, because the Dynamic Field Theory has been used to capture a variety of phenomena in early development, it provides a fertile foundation for probing the neural processes that underlie the development of 'rule'-like behavior and perseveration in the DCCS task.

Model Architecture

The basic architecture of the model and its mapping to cortical pathways is displayed in Figure 2. The Dynamic Neural Field (DNF) model proposed here utilizes a generic two-layered architecture initially explored by Amari (1977) and Amari and Arbib (1977) to capture the dynamics of neural activity in visual cortex. This architecture consists of a layer of self-excitatory working memory neurons (WM) arranged by functional topography (e.g., neurons that code for nearby spatial locations are neighbors in the network). These neurons project activation (see green arrows in Fig 2) to a layer of inhibitory interneurons (Inhib) which project inhibition broadly back into WM (see red arrows). This two-layered architecture achieves a form of locally excitatory and laterally inhibitory interactions within WM that can sustain "peaks" of activation in the absence of input. In addition to these layers, we add a third long-term memory (LTM) layer that is reciprocally coupled to the WM layer. The WM and LTM layers implement a form of Hebbian learning: peaks in WM leave traces of activity at associated sites in LTM which, in turn, influence the strength of activation in WM.
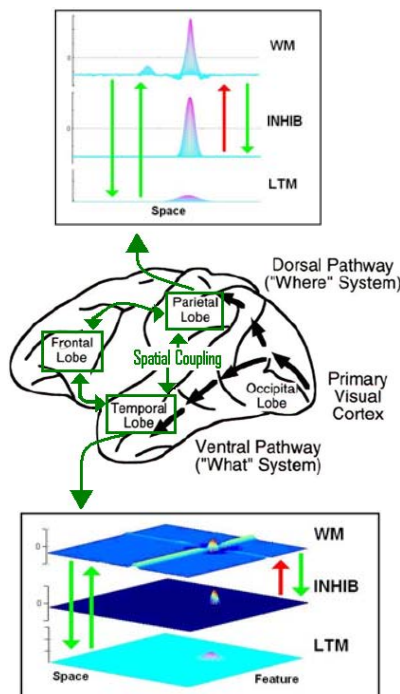

Figure 2: Model architecture and mapping to cortical pathways.

Dorsal stream processes, primarily devoted to encoding spatial information, are captured in a 1-dimensional spatial working memory (SWM) field (see top box in Fig 2). Ventral stream processes are captured by 2D feature-space working memory (FWM) fields with continuous feature dimensions (e.g., color) along one dimension and space along the other dimension. Activation peaks with the FWM fields, therefore, capture both the presence of, for example,

a red card in the task space as well as where that card is located. Note that feed-forward input to the FWM fields are typically precise along the feature dimension but imprecise (broad) along the spatial dimension. This reflects the tuning characteristics of neurons in many ventral stream areas which are precise in their representation of features but have broad spatial receptive fields.

Given the broad spatial tuning of ventral stream neurons, how does our model solve the "binding" problem of vision? That is, how would the model know that a "red" peak in the color-space field should be matched with a "star" peak in a shape-space field? The answer is through spatial coupling, that is, the SWM and FWM fields are coupled along the spatial dimension. This achieves distributed but coupled peaks of activation across multiple cortical areas—our implementation of an "integrated" object (for related ideas, see Treisman & Gelade, 1980).

The final aspect of the model architecture highlighted in Figure 2 is the contribution from frontal lobe areas. As we discuss below, we hypothesize that dynamic neural processes in the frontal lobe can selectively modulate the resting level of populations of neurons in SWM and FWM (e.g., boosting the excitability of color neurons). This implements a neurally-plausible mechanism that might underlie 'rule'-like behavior (Egner & Hirsch, 2005).

The DNF Model and the Standard DCCS

This framework provides useful machinery to capture the underlying dynamics of children's decision-making in the DCCS task. Any combination of shape and color can be represented in the model and the formation of peaks is a means by which the model makes decisions, that is, forms stable representations of 'what' should go 'where.'

Empirical results have shown that 3-year-olds have little difficulty using a particular feature dimension to sort the test cards during the pre-switch trials. As traces of decisions accumulate across multiple pre-switch trials, however, children have trouble switching to different rules. Conceptually, there is conflict on the switch trial between children's past behavior and the currently available perceptual cues. In the model, this conflict is reflected in competition between where LTM traces build up during the pre-switch trials and where the model is seeing the features on the target cards which mark the sorting bins. When faced with this competition, data suggest that LTM traces win the battle.

By 4 years of age, however, this is no longer the case. What explains this developmental difference? As we will see, our model suggests that a subtle boost in the excitability of neurons within a particular neural field (e.g., the shape-space field) driven by the task instructions is sufficient to tip the balance in favor of sorting based on perceptual cues during the switch trial rather than using LTM.

Inputs to the Model

The full model is shown in Figure 4 (without the inhibitory fields for simplicity). The top panel of this figure shows the

model at the moment when the test card is shown in the context of the two sorting bins and their associated target cards. The bottom panel of this figure shows the model generating a decision to place the card (the blue star) in the right bin (i.e., sorting by color).

There are six simulation figures within each panel: (1) activation in SWM with (2) its associated LTM layer (left column of the fields); (3) activation in a color-space FWM field and (4) its associated LTM layer (middle column); and, finally, (5) activation in a shape-space FWM field and (6) its associated LTM layer (right column). Note that the spatial dimension is shown along the x-axis in these figures. The y-axis for the spatial fields shows the strength of activation at each site in SWM and its LTM. Because the FWM fields are 2-dimensional, however, we indicate the degree of activation using the color scheme shown to the right, where red indicates strong activation.

There are two different kinds of input to the model. The *target input* (see Fig 3) goes into both the SWM and FWM fields. This sub-threshold input captures the spatial locations of the sorting trays (input to SWM), as well as the features located at each spatial position (input into the FWM fields). In Figure 3 (see also, Fig 1 and 4) the target inputs are a red star on the left and a blue circle on the right. *Test inputs* (see Fig 4 at Time-1) capture the feature information on the test cards presented to the model and only enter into the FWM fields. In Figure 4, the model is presented with a blue star test card. The test inputs are spread along the spatial dimension to reflect the nature of the task—the model (i.e., the child) must decide where the test card should be placed (i.e., where a peak of activation should be localized along the spatial dimension). Because the working memory fields are coupled along space and test inputs only come into FWM, the SWM field can be seen as an output layer, indicating where the model has decided to put the test input. For instance, in the lower panel of Figure 4, there is a strong peak of activation at a right location in the SWM field. Thus, the simulation has decided to place the test card in the right bin (i.e., the bin marked by the blue circle).
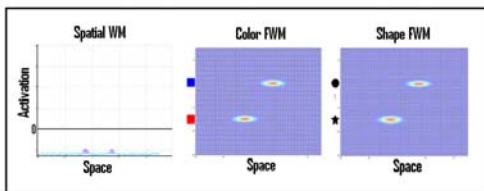

Figure 3: Target input to SWM and FWM for a red star on the left and blue circle on the right.

Pre-Switch Trials

The simulation shown in Figure 4 began with the presentation of target cards. This input pre-activated the two circular "hot-spots" in the color and shape FWM fields (see also Fig 3). In addition, we gave a slight boost to the resting level within the task-relevant field (e.g., neurons in the color field had a slightly higher resting level when the model played the "color" game). This implementation is

consistent with neurological evidence which suggest that the brain resolves perceptual conflict by boosting activity for the task relevant stimulus (Egner & Hirsch, 2005). One suggested hypothesis for how this is accomplished is by boosting the baseline level of activity in areas of the brain that processes task relevant information.
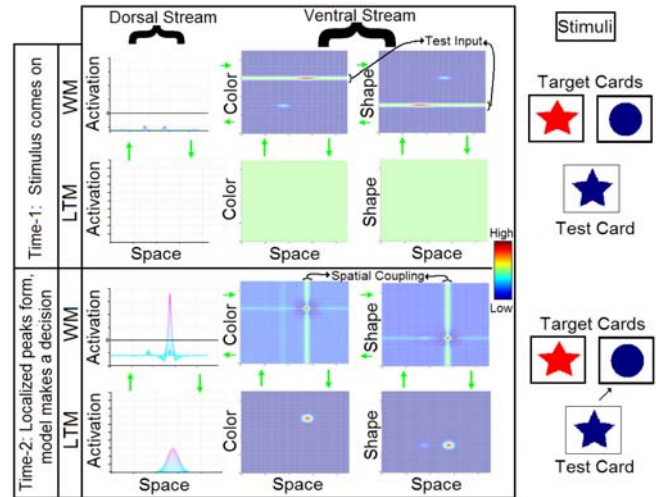

Figure 4: The full model (shown without the inhibitory fields) sorting a blue star by color.

After boosting the color field by 0.03 units, we presented a target card to the model to sort (see Time-1 in Fig 4). The presentation of the target card created "ridges" of activation at the specified feature values (blue, star) that spanned the spatial dimension. Next, activation peaks start to grow due to the locally excitatory/laterally inhibitory interactions in the FWM fields at the locations where the test inputs overlap with the target inputs. Because the color field received a slight boost, activation grows more quickly in this field. Once a self-stabilized peak in the color field becomes sufficiently strong, it begins to send activation to the SWM and shape FWM fields at the associated spatial location of the color peak (i.e., the right location, see Fig 4 at Time-2). This tips the balance in the shape FWM field such that activation begins to grow for that shape value at the right location (even though the target input for that feature is at the left location). The combination of the test input and spatial input from the color FWM field conspire with local interactions to build a peak in the shape field (see shape-space field in the lower panel of Fig 4).

At the same time, activation takes hold in the SWM field (which is now receiving robust spatial input at the right location from the peaks in the two FWM fields). A peak, then, grows at the right location in SWM and the model is said to have placed the card in the right sorting bin. This concludes the first trial. The same sequence of events plays out as the model is given the other test input (a red circle) which the model sorts in the left bin. Note that when the WM peaks form during each trial, they leave traces in their associated LTM layers (see Fig 4, lower panel). Such traces play a critical role during the switch trials.

Post-Switch Standard Trials

Figure 5a shows the state of the color and shape FWM fields going into the switch trials. Since the model consistently sorted the cards by color during the pre-switch trials, the LTM traces in color FWM are overlapping with the target input (i.e., the red card was always placed in the left bin; the blue card was always placed in the right bin). The story is quite different in the shape FWM field, however, because LTM traces in this field were always formed at the opposite locations of the target input for each feature. This sets up competition between LTM and the target input within the shape field at the start of the switch trials. It is this competition that ultimately will lead to perseveration when the model switches and plays the "shape" game.

The results of the full sequence of pre- and post-switch trials for the 3- and 4-year-old models are shown in Figure 5b and 5c (the test cards are displayed at the bottom of the figure at the time when they were presented to the model). The maximum values in the SWM field at the left and right target inputs are plotted over time. An above-threshold (>0) peak of the red line indicates a decision to sort to the right location, while a blue peak indicates a decision to sort to the left location.
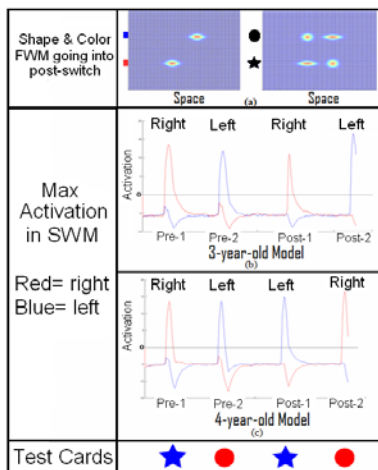

Figure 5: Simulation results for the standard version of the DCCS

Going into the post-switch phase, the model is 'told' to play the shape game. For the 3-year-old model, the resting level of the shape field is boosted by 0.03 (the same amount that the color FWM was boosted in the pre-switch). As inputs come in for the post-switch test cards, peaks try to form at two sites for the shape feature—at the target input and at the site activated by input from LTM. Because two peaks are trying to form in the shape field, they grow more slowly due to laterally inhibitory interactions.

Critically, at the same time two peaks are trying to form in the shape FWM field, a single peak is emerging in the color FWM field at sites that match the target input. Just as in the pre-switch phase, the local interactions in the color field build a peak which sends spatial information to the shape FWM field. This, in turn, helps activation grow at the location of LTM activation. Once the shape peak forms, peaks in both FWM fields drive up activation in SWM and the model makes its decision. The model was given the test cards in the same order for the post-switch as in the pre-switch. As Figure 5b shows, the 3-year-old model perseverates (as do 3-year-old children): it produces the same spatial responses for the pre- and post-switch phases.

Development in the Standard DCCS Task Correct rule-switching in the model is achieved through stronger modulation of the resting level of the shape FWM field during the post-switch. When the 4-year-old model is 'told' to sort by shape going into the post-switch, the shape FWM field is boosted by 0.6 (as opposed to 0.03). Now, as the inputs come in during the post-switch, the shape field has enough excitatory energy to resolve the conflict created by sorting decisions made during the pre-switch phase. Once a peak builds where the test input overlaps with the target input, spatial information is sent to the color FWM field as well as SWM. As figure 5c shows, the 4-year-old model produced opposite spatial responses in the post-switch, indicating that the model switched its rule-use.

The Negative Priming Version

Interestingly, 3-year-olds still perseverate in a Negative Priming (NP) version of this task (Müller et al., 2006; Zelazo et al., 2003). For this version, the features for the dimension that was relevant for the pre-switch are changed for the post-switch sort. For example, if children sorted red and blue stars and circles by color in the pre-switch, they would then be asked to sort *green* and *yellow* stars and circles by shape in the post-switch.

Pre- and Post-Switch NP Trials

For the NP version, the pre-switch portion of the task is exactly the same as the simulations of the standard task. Figure 6a shows the color and shape FWM fields at the start of the post-switch phase. The color field now has target inputs at new feature values: thus the target input and LTM activation no longer overlap in color FWM. Importantly, though, there is still competition in the shape FWM field for each feature value. With competition still present in the shape field and only a 0.03 boost during the post-switch (for the 3-year-old model), this field continues to be delayed in forming a peak and still allows the color field to build a peak more quickly at the site where the test input overlaps with the target input. Figure 6b shows the spatial responses for the 3-year-old model. Again, the model perseverates and the same pattern of spatial responses is produced for both the pre- and post-switch phases. Even though the localized activation in the color field is weaker now because target input and LTM no longer overlap, the competition in the shape field slows down peak formation sufficiently to allow the color FWM field to form a stable peak first and drive a perseverative response as in the simulations of the standard task.

Development in the NP Version Correct rule-switching, again, comes from stronger modulation of the resting level for the post-switch sort. For this version, though, the resting level of the shape FWM field only needed to be boosted by 0.3 (as opposed to 0.6 for the standard task). There is less of a boost required because of the lack of overlap in the color FWM field. This means there is weaker spatially localized activation in this field and, thus, it takes longer to form a peak. In this sense, the shape FWM needs less of a head-start in order to localize a decision based on its target inputs.
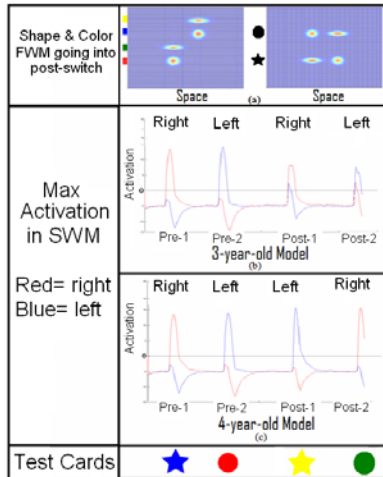

Figure 6: Simulation results for the NP version of the DCCS.

The Absolute Negative Priming Version

Zelazo et al. (2003) and Müller et al. (2006) further explored the NP variant of the task by testing an Absolute Negative Priming version (ANP). They showed that perceptual competition during the pre-switch was required in order for 3-year-olds to perseverate in the NP version. In the ANP version, there is no perceptual competition during the pre-switch phase because children sort test cards that match the target cards along *both* dimensions. Now when the features of the previously relevant dimension are changed in the post-switch, 3-year-olds have little trouble switching their rule-use.

Pre- and Post-Switch ANP Trials

The pre-switch phase for the ANP version is slightly different from that of previous simulations. Now, there is no competition established in any of the fields. Since the test cards match the target cards along both dimensions, the LTM activation overlaps with the pre-switch target input for both FWM fields. Figure 7a shows the FWM fields at the start of the post-switch phase. As in the NP version, the feature values for the color field are changed. Going into the post-switch phase, the shape field is only boosted by 0.03. With only this small boost and no competition for spatial locations within the feature values, the shape FWM field is now able to form peaks where the test input overlaps with target input. It shares its spatial information with the

color FWM field, drives a peak where the color feature value overlaps with the spatial information, and these FWM peaks, in turn, drive a spatial response based on shape. Figure 7b shows results of the simulation of a 3-year-old model, which switches rules, of the ANP version.
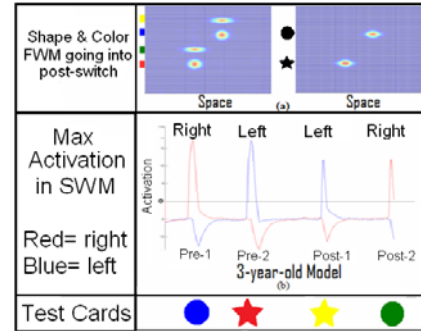

Figure 7: Simulation results for the ANP version.

Conclusions and Future Directions

The model provided here demonstrates that the neural mechanisms that underlie dorsal and ventral stream visual pathways—when combined with a simple mechanism that modulates the resting level of neural populations—can give rise to emergent rule use. This provides the first formally grounded theory of perseveration and development in the DCCS task that is closely tied to known neural properties. Ultimately, perseveration is captured as the competition between the target input and the LTM memory traces that are formed during the pre-switch. With two sites of activation for the features that are relevant for the post-switch, extra inhibition is built-up, it takes longer for that field to form a peak, and the other feature field forms a peak sooner and drives a perseverative response.

Correct switching occurs through stronger boosting of the post-switch feature field. This modulation of the resting levels of the feature fields is thought to occur through frontal lobe processes. Gaining an understanding of these processes is the next step to providing a full account of perseveration and task-switching. Ultimately, the boosting of the relevant dimension should be mediated by the level of conflict detected in the task environment and the trial-to-trial and developmental history of the system. Over development, what exactly is changing to allow correct switching? It is possible that the conflict detection areas become better tuned to detect perceptual conflict and thus improve on how it processes its inputs. Perhaps the area of the frontal lobe responsible for boosting the relevant feature field becomes stronger or more efficient, thus allowing for differential modulation in the face of varying levels of conflict. Further modeling will need to determine how this modulation can be formally implemented by capturing neural activity in areas of the prefrontal cortex, such as the ventrolateral or rostrolateral prefrontal cortex, and other frontal lobe regions that have been implicated in task-switching.

The DNF model reported here is the first account of the DCCS task that has implicated a role of the spatial layout of

the target cards in children's perseveration. In this task, children are not just sorting shapes and colors, but are tying those features to space. Since the formal model uses space as a critical binding dimension, it makes specific predictions about the role space should play in children's performance. For example, switching the location of the target cards in the post-switch of the NP version should elicit correct switching in 3-year-olds (e.g., instead of the star on the left and circle on the right, the star would be moved to the right with the circle on the left). This would create overlap between the target inputs and the LTM traces in the posts-switch, task-relevant feature field. Thus, peaks would form in the shape FWM field without competition. Similarly, 3-year-olds should perseverate if the target cards switched locations in the post-switch of the ANP version. The post-switch feature field would then have competition between the target inputs and LTM traces as is the case for the normal NP version. We are currently testing these predictions.

Rule-use and executive function has long been thought to require an intelligent director, a 'central-executive' that coordinates behaviors or schemas, sets goals, and assesses success/failure; however, with little further implementations, the basic model developed to handle a wide array of spatial and cognitive tasks can display rule-like behavior, perseveration, and development without a central executive. In this framework, 'rule'-like behavior is emergent from dynamic interactions of fields of neurons that are based in known neural properties.

Acknowledgements

References

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

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

Bunge, S. A., & Zelazo, P. D. (2006). A brain-based account of the development of rule use in childhood. *Current Directions in Psychological Science, 15(3),* 118-121.

Colvert, E., Custance, D., & Swettenham, J. (2002). Rule-based reasoning and theory of mind in autism: a commentary on the work of Zelzo, Jacques, Burack, and Frye. *Infant and Child Development, 11,* 197-200.

Crone, E. A., Donohue, S. E., Honomichl, R., Wendelken, C., and Bunge, S. A. (2006). Brain regions mediating flexible rule use during development. *The Journal of Neuroscience, 26(43),* 11239-11247.

Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience, 8(12),* 1784-1790.

Johnson, J.S., Spencer J.P., and Schöner, G. (in press). Moving to higher ground: the Dynamic Field Theory and the dynamics of visual cognition. In F. Garzòn, A. Laakso, T. Gomila (Eds.). Dynamics and Psychology [special issue]. *New Ideas in Psychology.*

Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science, 6(5),* 449-476.

Kloo, D., & Perner, J. (2005). Disentangling dimensions in the dimensional change card sort. *Developmental Science, 8(1),* 44-56.

Morton, J. B., & Munakata, Y. (2001). Active versus latent representations: a neural network model of perseveration, dissociation, and decalage. *Developmental Psychobiology, 40,* 255-265.

Mulas, F., Capilla, A., Fernández, S., Etchepareborda, M. C., Campo, P., Maestú, F., Fernández, A., Castellanos, F. X., & Ortiz, T. (2006). Shifting-related brain magnetic activity in attention-deficit/hyperactivity disorder. *Biological Psychiatry, 59,* 373-379.

Müller, U., Dick, A. S., Gela, K., Overton, W. F., & Zelazo, P. D. (2006). The role of negative priming in preschoolers' flexible rule use on the Dimensional Change Card Sort. *Child Development, 77(2),* 395-412.

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.

Simmering, V. A., Schutte, A. R., & Spencer, J. P. (2007). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. In S. Becker (Ed.), *Computational Cognitive Neuroscience [special issue]. Brain Research.* doi:1031016/j.brainres.2007.06.081.

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

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97-136.

Zelazo, P. D. (2006). The Dimensional Change Card Sort (DCCS): a method of assessing executive function in children. *Nature Protocols, 1(1),* 297-301.

Zelazo, P. D. (2004). The development of conscious control in childhood. *Trends in Cognitive Sciences, 8(1),* 12-17.

Zelazo, P. D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development, 11,* 37-63.

Zelazo, P. D., Jacques, S., Burack, J. A., & Frye, D. (2002). The relation between theory of mind and rule use: evidence from persons with autism-spectrum disorders. *Infant and Child Development, 11,* 171-195.

Zelazo, P. D., Müller, U., Frye, D., & Marcovitch, S. (2003). The development of executive function. *Monographs of the Society for Research in Child Development, 68(3).*