Published in final edited form as:
Monogr Soc Res Child Dev. 2014 June ; 79(2): vii–103. doi:10.1002/mono.12096.

The Emergent Executive: A Dynamic Field Theory of the Development of Executive Function

Aaron T. Buss and **John P. Spencer**
University of Iowa Department of Psychology Delta Center

Abstract

A dynamic neural field (DNF) model is presented which provides a process-based account of behavior and developmental change in a key task used to probe the early development of executive function—the Dimensional Change Card Sort (DCCS) task. In the DCCS, children must flexibly switch from sorting cards either by shape or color to sorting by the other dimension. Typically, 3-year-olds, but not 4-year-olds, lack the flexibility to do so and perseverate on the first set of rules when instructed to switch. In the DNF model, rule-use and behavioral flexibility come about through a form of dimensional attention which modulates activity within different cortical fields tuned to specific feature dimensions. In particular, we capture developmental change by increasing the strength of excitatory and inhibitory neural interactions in the dimensional attention system as well as refining the connectivity between this system and the feature-specific cortical fields. Note that although this enables the model to effectively switch tasks, the dimensional attention system does not 'know' the details of task-specific performance. Rather, correct performance emerges as a property of system-wide neural interactions. We show how this captures children's behavior in quantitative detail across 12 versions of the DCCS task. Moreover, we successfully test a set of novel predictions with 3-year-old children from a version of the task not explained by other theories.

Early childhood is a time of rapid change in the functional organization of cognition. The period between 2 and 4 years is particularly dramatic, highlighted by the emergence of executive function (EF)—the processes that underlie active behavioral control and cognitive flexibility. Executive function has pervasive influences on cognition as children learn to regulate their behavior in context-specific ways. For example, measures of executive function have been shown to be robustly predictive of language development (Im-Bolter, Johnson, & Pascual-Leone, 2006; McEvoy, Rogers, & Pennington, 1993), mathematical abilities (Bull & Scerif, 2001) and measures of general cognitive functioning (e.g., school performance, IQ, and psychopathology; Liss et al., 2001; Pennington & Ozonoff, 1996). Further, EF measures during childhood robustly predict physical health, substance dependence, personal finances, and criminal offending outcomes as measured at age 32 (Moffitt et al., 2011). Theories of EF, then, can potentially have far reaching influences on development.

Early theories of EF proposed that the emergence of cognitive control reflected developmental changes in a central executive system—a central resource that controls other

---

aspects of cognition (Baddeley, 1986; Duncan, Emslie, Williams, Johnson, & Freer, 1996; Duncan, Johnson, Swales, & Freer, 1997; Norman, & Shallice, 1986). This view was anchored, in part, to evidence that core executive functions could be localized to lateral pre-frontal cortex (PFC), a large region anterior to the precentral sulcus. Lateral PFC is one of the slowest developing brain regions (Giedd et al., 1999) and evidence from patient populations showed that impairments of PFC leads to dysfunctions of behavioral control (Baddeley, Della Sala, Papagno, & Spinnler, 1997; Milner, 1963) and behaviors that mimic the performance of young children (Dempster, 1992; Diamond, 2005). Recent neuroimaging studies show, however, that EF is not localized solely to PFC. Rather, cognitive flexibility and control emerge from an extensive network of regions within frontal and posterior cortical areas (for review, see Dosenbach et al., 2007; Fair et al., 2007; Morton, 2010). Moreover, there are extensive changes to this network over development, including changes in cortical volume, structural differentiation, functional connectivity, and the dynamics of neural activation (Barnea-Goraly et al., 2005; Crone, Donohue, Honomichl, Wendelken, & Bunge, 2006; Fair et al., 2007; Gogtay et al., 2004; Kelly et al., 2009; Lenroot & Giedd, 2006; Moriguchi & Hiraki, 2009; Sowell, Trauner, Gamst, & Jernigan, 2002; Stevens, Pearlson, & Calhoun, 2009; Stevens, Skudlarski, Pearlson, & Calhoun, 2009; Tsujimoto, 2008). Thus, the emergence of EF does not reflect changes in a single area of the brain, but instead reflects the development and organization of neural structures throughout the brain.

Although the emergence of EF is a major achievement, it is not clear how children develop the ability to flexibly control their own behavior. A primary challenge to achieving this understanding is the involvement of multiple processes in EF: *inhibition* which prevents acting on irrelevant exigencies in the environment or built up habits, *working memory* which maintains task or context relevant information, and *task switching* which allows for the updating of one's goals and behavior across different contexts (Collette et al., 2005; Davidson, Amso, Anderson, & Diamond, 2006; Garon, Bryson, & Smith, 2008; Lehto, Juujärvi, Kooistra, & Pulkkinen, 2003; Miyaki, Friedman, Emerson, Witzki, & Howerter, 2000).

These components emerge in task-specific contexts at different points during early childhood. Response inhibition develops in its most rudimentary form late in infancy as indexed by the A-not-B task (Marcovitch & Zelazo, 1999; Thelen, Schöner, Scheier, & Smith, 2001), but continues to develop during early childhood in more complex contexts that pit internalized rules against prepotent behaviors such as in the Simon task (Gerardi-Caulton, 2000), flanker task (Rueda et al., 2004), Stroop task (Carlson, 2005) and Go/No-Go task (Cragg & Nation, 2008; Dowsett & Livesey, 2000). Similarly, working memory increases in capacity between 3 and 5 years as measured in change detection tasks (Isaacs & Vargha-Khadem, 1989; Logie & Pearson, 1997; Simmering, 2008), and continues to increase into childhood and adolescence (Cowan et al., 2005; Isaacs & Vargha-Khadem, 1989; Logie & Pearson, 1997; Simmering, 2008; Vicari, Bellucci, & Carlesimo, 2003). Finally, task-switching shows severe impairments around 2 years of age: children are able to reliably engage in single rule tasks, but perseverately use only one rule if the task contains more (Zelazo, Reznick, & Piñon, 1995). Switching from one rule set to another is possible in some tasks at 3 years of age (Brace, Morton, & Munakata, 2006; Fisher, 2011; Müller, Dick, Gela, Overton, & Zelazo, 2006; Zealzo, Müller, Frye, & Marcovitch, 2003), but it is not until 4 years of age that children can reliably switch from one rule set to another using the same stimulus set (Müller et al., 2006; Zelazo et al., 2003). Importantly, switch costs, as indexed by reaction times, are still seen throughout childhood, adolescence, and adulthood (Diamond & Kirkham, 2005; Morton, Bosma, & Ansari, 2009).

These basic components can be used to guide and structure experimentation, making investigations of EF more tractable; however, the multi-component nature of EF makes it difficult to specify an integrated developmental theory. To formalize such a theory, one must understand not only the development of the components themselves, but also how their interaction changes over time. Moreover, recent neural evidence points to the need for theories to speak to both behavioral and neural levels of analysis.

In the present report, we propose a new theory of the development of executive function—Dynamic Field Theory (DFT)—that can span behavioral and neural levels to provide an integrative account of behavior over development. DFT uses real-time neural population dynamics to generate stable patterns of neural activity which correspond to active decisions in-the-moment. We anchor this theoretical account to data from a particular task that provides a common frame of reference for understanding developmental changes in EF—the Dimensional Change Card Sort (DCCS) task. This task requires aspects of all three components of EF, making it an ideal probe of EF in early development. Moreover, the vast literature using the DCCS task provides robust constraints for a developmental theory.

The rest of this paper is divided into six main sections. In the section below, we review the literature on the DCCS, including theories that account for aspects of children's performance in this task. Next, we describe a dynamic field theory of executive function and how this theory captures performance in the standard version of the DCCS task. In the third section, we probe the ability of the model to move beyond current theories by quantitatively simulating children's performance from multiple variants of the DCCS task. In section four, we generate and test novel predictions regarding the role of spatial locations in feature binding in the DCCS task. Data from this study are not consistent with any other theory of the early development of EF. In section five, we present further simulation results highlighting the ability of the model to capture asymmetrical switch costs associated with different attentional manipulations in the DCCS task. We conclude by contrasting the DFT with other theories in the literature as well as evaluating the strengths and limitations of this theoretical approach to the emergence of EF in early development.

The Development of Executive Function: A Case Study of the DCCS Task

The DCCS has been extensively used over the last two decades to probe the development of cognitive flexibility in early childhood. In this task, children must switch from sorting cards based on shape or color to sorting based on the other dimension using verbal rules provided by the experimenter (e.g., “Sort by shape/Sort by color”). Trays mark two sorting locations where target cards are affixed. These target cards provide cues as to which features go where (e.g., a blue-circle and a red-star). The test cards that children sort are typically constructed so that they match either target card along one dimension (e.g., a blue-star and a red-circle;

see Figure 1A). Thus, there is direct conflict when making a decision for a given card since it could go to either location depending on the dimension used for sorting.

The DCCS is an ideal task to study the development of EF for two reasons. First, flexible rule-use in this context emerges from the interaction of EF components. Specifically, this task requires inhibition to prevent responding based on the irrelevant dimension, working memory to maintain representations of the relevant task rules, and task switching in order to update these processes after the rule-switch (Garon et al., 2008). Second, this task reveals rapid and drastic changes in children's executive function and neural dynamics in early development. Although 4-year-olds have little trouble switching rules, 3-year-olds perseverate and continue using the first set of rules after they are instructed to switch. This perseveration is robust and persists despite constant reminders that the rules have changed. Research has also revealed changes in a network of brain areas associated with rule-switching which emerges after age 3 and becomes more refined into adulthood (Moriguchi & Hiraki, 2009; Morton, et al., 2009). Thus, recent data suggest that this simple task can provide insight to the wide spectrum of EF and associated neural changes in early childhood.

A long line of research has revealed many complexities in children's rule-use in the DCCS task. Below, we review the evidence organized around the different components of executive function as well as other themes highlighted across the different studies. This review shows the complex nature of children's rule-use, explaining why single-cause explanations (e.g., inhibitory control) fall short of capturing the full range of effects present in the literature as well as why this task is ideal for developing an integrative theory of EF.

Review of Behavioral Evidence with the DCCS

The role of inhibition—Several studies have probed the role that inhibitory processes play in children's perseveration in the DCCS task. This has been accomplished by altering different features of the target and test cards between the pre- and post-switch phases. These data show that children do not perseverate based on a fixed set of features or rules; rather, the dimensions and features on the cards participate in a more subtle pattern of interactions that implicates a more generalized form of inhibition.

In a Negative Priming version (NP), the features that were *relevant* for the pre-switch sort (e.g., color) are changed for the post-switch sort (e.g., shape; see Figure 1B). With the pre-switch features changed, the stimuli no longer afford being sorted by the color values used during the pre-switch phase; thus, it is not required that children inhibit those values during the post-switch phase. However, 3-year-olds fail to switch rules in this version suggesting they have difficulty overcoming the negative priming that occurs as the irrelevant features are ignored or suppressed during the pre-switch phase (Müller, et al., 2006; Zelazo et al., 2003). That is, 3-year-olds appear to have too much inhibition of the post-switch dimension that they cannot overcome. In a Partial-Change version (Zelazo et al., 2003) the features that were *irrelevant* during the pre-switch are changed in the post-switch phase (see Figure 1C). With these features changed, the rules for the post-switch feature values would not be negatively primed going into the post-switch. Nevertheless, 3-year-olds also have difficulty switching in this task presumably due to their inability to inhibit the pre-switch rules and attend to the pre-switch dimension (Zelazo et al., 2003). In this case, it seems that 3-year-olds have too little inhibition to suppress a prepotent response pattern.

The seemingly contradictory results of the NP and Partial-Change versions leaves one wondering whether *any* changes in the features of the cards improves performance in this task. Indeed, results from a Total-Change version where the features of both dimensions are changed show that 3-year-olds can reliably switch to sorting by the other dimension (Zelazo et al., 2003; see Figure 1F). Thus, if there is nothing to inhibit and the post-switch features have not been negatively primed, children can switch rules. These data indicate that single cause accounts such as inhibition or negative-priming at the level of specific features alone are not sufficient to fully capture the underlying processes giving rise to perseveration.

The role of working memory—The DCCS does not place high demands on working memory capacity since only two rules are relevant for each game and these rules are repeated throughout the pre and post-switch phases. Nevertheless, evidence shows that the strength of representations which are the basis for the different sets of rules critically influences performance. For example, 3-year-olds can switch rules if the representations utilized for the post-switch phase are stronger than those utilized for the pre-switch phase. Yerys and Munakata (2006) demonstrated this by manipulating the name used for the pre- and post-switch sorting games. Specifically, 3-year-olds' post-switch performance is improved if the pre-switch game is simply called a 'sorting' game and the post-switch game is given a standard informative name (such as the 'shape' or 'color' game). This effect is also seen if the manipulation occurs at the level of specific stimulus features. For example, 3-year-olds' post-switch performance is improved if novel shapes or colors with novel labels are used during the pre-switch phase while familiar features and labels are used for the post-switch game. In either case, an informative label can be maintained during the post-switch phase which can out-compete the less-familiar or non-informative representation used during the pre-switch phase.

The role of task switching—Various manipulations to the transition between the pre- and post-switch phases have been shown to facilitate correct rule-switching. Specifically, if the post-switch rules are given while the target cards are removed or if the concepts of shape-rules and color-rules are explained in detail before the child sorts, 3-year-olds have less difficulty switching rules (Mack, 2007). Further, if children are instructed to play a 'silly' version of the pre-switch game and to match the test cards to the opposite of the pre-switch features during the post-switch phase (e.g., matching red to blue and blue to red) 3-year-olds are able to switch rules (Kloo, Perner, Kerschhuber, Dabernig, & Aichhorn, 2008; however, see also Brooks, Hanauer, Padowska, & Rosman, 2003 who show children are worse with bi-dimensional stimuli than with uni-dimensional stimuli in this 'silly' version of the task). This highlights that children's representation of the task-switch is critical to perseveration or success in the DCCS.

The role of feedback and demonstration cues—Three-year-olds are able to correctly switch rules if they are given sufficient demonstration cues or feedback. Specifically, if children receive direct feedback on their performance (Bohlmann & Fenson, 2005), if children see the post-switch rules demonstrated (Towse, Redbond, Houston-Price, & Cook, 2000), or if children are told to wait and think about the rules before they sort a given card in the post-switch (Deák, Ray, & Pick, 2004), then their post-switch performance improves. Thus, even within the standard task structure, children are not completely rigid in their perseveration but can successfully switch given enough instruction.

The role of conflict—Children's ability to switch rules in the DCCS task is not simply based on the consistency of the features on the cards between the pre- and post-switch phases—w*hat* children sort also matters. For example, conflict between the dimensions during the pre-switch is necessary for perseveration. Zelazo et al. (2003) and Müller et al. (2006) showed that 3-year-olds no longer perseverate in the standard or NP versions if the test cards match the target cards along both dimensions during the initial sorting phase (e.g., sorting red stars to red stars and blue circles to blue circles; see the no-conflict versions in Figure 1D-E). Further, if conflict is decreased by only using one feature within the irrelevant dimension during both the pre- and post-switch phases (see Relational Complexity version in Figure 1G) children have less difficulty switching rules (Halford, Bunch, & McCredden, 2007).

Visual conflict can be further eliminated by using pictures of cartoon characters in lieu of target cards (the characters being characterized as ‘wanting’ one feature or another; Perner & Lang, 2002) or by completely removing target cards and having children sort to empty trays (Towse, et al., 2000). Under these circumstances, 3-year-olds have little trouble switching rules. The results with target cards absent are particularly interesting given the supposedly heightened demands on rule-representation in the absence of visual cues.

The role of feature binding—Switching rules can also be facilitated if the ‘objecthood’ of the images on the cards is eliminated by separating the features on the cards (e.g., an outline of a star next to a patch of blue; see Figure 1H; Diamond, Carlson, & Beck, 2005; Kloo & Perner, 2005; Zelazo et al., 2003). Further, 3-year-olds can switch rules if the dimensions are separated into four sorting locations with univalent target cards (e.g., a black outline of a star or circle and patches of red or blue) so that different pairs of trays are used for the color and shape rules (all four trays are displayed throughout the pre- and post-switch; Rennie, Bull, & Diamond, 2004). These versions highlight that children's difficulty is not just a function of the ‘rules’ in the task, but is also influenced by the nature of the objects to which children apply the rules. Specifically, children benefit from being able to apply the shape and color rules to different objects. This suggests that processes of object representation or selective attention are central aspects of children's rule-use.

Review of Neural Evidence with the DCCS

The DCCS is primarily studied in early childhood during an age when use of fMRI is very challenging. A more recent technology—functional Near-Infrared Spectroscopy (fNIRS)—is able to overcome the challenges of using fMRI with young children. fNIRS measures changes in cortical hemodynamics using the optical properties of oxy- and deoxy-hemoglobin (Cui, Bray, Bryant, Glover, & Reiss, 2011; Minati, Visani, Dowell, Medford, & Critchley, 2011). Although promising, there is only one study to date that has used fNIRS to probe the neural systems involved in the DCCS task. In contrast, more is known about the neural basis of the basic components of EF and how these neural systems change over development. Thus, in this section, we discuss what is known about the neural systems involved in the development of EF, as well as the neural systems implicated in the DCCS task in particular. As we review below, the components of EF are not uniquely associated with any single brain area; rather, the emergence of EF is associated with a broad network of frontal and posterior regions.

Inhibition—Baird, Kagan, Gaudette, Walz, Hershlag, and Boas (2002) recorded fNIRS from frontal cortex while infants performed the A-not-B task. These researchers reported an increase in frontal activation associated with correct reaching to the B location. Similarly, Schroeter, Zysset, Wahl, and von Cramon (2004) recorded fNIRS while 7- to 13-year-olds performed a Stroop task and found increased activity in right frontal cortex associated with greater inhibitory control.

Durston, Thomas, Yang, Ulug, Zimmerman, and Casey (2002) compared go-no-go performance and associated fMRI data between a group of 6- to 10-year-olds and adults. Adults were both faster and more accurate on ‘go’ trials. Both age groups showed stronger activation on ‘no-go’ trials in ventral PFC bilaterally, right dorsolateral PFC, and right parietal cortex. The difference in activation between trial types was larger for children than for adults suggesting that the task was less demanding for adults. Further, Hwang, Velanova, and Luna (2010) recorded fMRI while a group 8- to 12-year-olds, 13- to 17-year-olds, and 18- to 27-year-olds performed an anti-saccade task. While anti-saccade performance improved significantly between each age group, an increase in connectivity between frontal and parietal areas was observed along with a decrease in connectivity within parietal cortex. Thus, emerging inhibitory control is supported by strengthening interactions between frontal and parietal areas.

Working Memory—As with inhibition, a network of cortical regions has been implicated in the development of WM. Thomason, Race, Burrows, Whitfield-Gabrieli, Glover, and Gabrieli (2009) studied a group of 7- to 12-year-olds and adults with fMRI and found that activation in frontal and parietal regions scaled with the number of items to be remembered. Further, this study found left lateralization for verbal WM but right lateralization for spatial WM. Tsujii, Yamamoto, Masuda, and Watanabe (2009) used fNIRS to examine spatial WM with 5- and 7-year-olds and found that 5-year-olds did not show lateralized specialization of frontal cortex while 7-year-olds did show this pattern. Thus, the lateralization of spatial WM emerges early in childhood.

Klingberg, Forssberg, and Westerberg (2002) reported stronger superior frontal and intraparietal cortex activity with increasing WM capacity between 9 and 18 years. Edin, Macoveanu, Olesen, Tengér, & Klingberg (2007) further showed that stronger fronto-parietal connectivity is associated with improvements in WM (see also, Edin et al., 2009). Thus, over development the fronto-parietal network supporting visual WM becomes right lateralized and increasingly sensitive to the size of the WM load. Note, further, that the neural basis of WM and inhibitory control share many aspects. For instance, McNab, Leroux, Strand, Thorell, Bergman, and Klingberg (2008) compared neural activation associated with visual WM and inhibitory control and found common patterns of activation in right inferior frontal gyrus, middle frontal gyrus, and right parietal areas.

Task-Switching—Different lines of fMRI data have revealed dissociable components of rule-switching in tasks that only require shifting of responses instead of shifting of attention as required in the DCCS. With adults, the ventro-lateral prefrontal cortex (VLPFC) and rostro-lateral prefrontal cortex (RLPFC) are more active when demands on rule-representation are high (i.e., when sets of response mappings conflict with one another), while the supplementary- and pre-supplementary-motor areas (SMA/pre-SMA) together with the basal ganglia are more active when switching rules or reconfiguring a rule-set (Crone, Wendelken, Donohue, & Bunge, 2006; for a review of the neural circuitry underlying rule-use, see Bunge, 2004; Bunge et al., 2005).

Critically, research has revealed differences in these neural dynamics over development. Specifically, 8- to 12-year-olds display a less-differentiated pattern of neural activation than adults. These children engage the pre-SMA/SMA for rule-representation, not simply rule-switching as with adults. Further, while adults show heightened VLPFC activation for bivalent rules on both rule-repeat and rule-switch trials, children show greater activation in this area for both univalent rules and switch trials (Crone, Donohue, et al., 2006).

DCCS—Moriguchi and Hiraki (2009) used fNIRS to probe the early development of EF in the DCCS task. They found increases in frontal activity between 3 and 5 years associated with rule-shifting in the DCCS. Thus, frontal activity supports the early emergence of flexible rule-use. Further, Morton et al. (2009) used fMRI to explore differences in neural activation between 11- to 13-year-olds and adults when shifting dimensional attention in the DCCS task. This study showed that activation in superior parietal cortex, dorsolateral PFC, pre-supplementary motor area, inferior frontal junction, and fusiform gyrus was associated with shifting rules in the DCCS. Over development, there are increases in activation in superior parietal cortex, superior frontal sulcus, and fusiform gyrus. Thus, performance in the DCCS task involves a network of frontal and posterior areas which overlap with the regions implicated in WM and inhibition.

Summary of Behavioral and Neural Data

The rich behavioral literature on the DCCS provides robust constraints for the development of a theory of EF. Children fail to switch rules in very specific ways. First, 3-year-olds only have difficulty during the post-switch phase. Further, post-switch performance is typically all-or-nothing, that is, the majority of children either fails every trial or passes every trial during the post-switch phase. Perseveration is robust and persists despite changes to either dimension. Nevertheless, 3-year-olds can switch under particular circumstances when the demands on inhibition, WM, or task switching are decreased.

The neuroimaging literature has revealed that the components of EF are associated with a network of areas rather than any single area and developmental changes in EF can be attributed to increased interactivity between frontal and posterior areas (see also, Dosenbach et al., 2007; Fair et al., 2007). Further, the one study that has examined the neural dynamics of perseveration and switching in the DCCS task during early childhood has shown increases in frontal activation when 3- and 5-year-olds switch rules. Finally, fMRI data with adults shows a specialized network of cortical areas related to different aspects of rule-switching that emerges through childhood and adolescence. Although the picture regarding the neural basis of EF and rule-switching is incomplete, there are sufficient data to take seriously the bridge between behavioral and neural processes. In the next section, we discuss how current theories account for this pattern of behavioral and neural data.

Current Theories of the Development of Executive Function

Our review of the empirical literature on the DCCS task reveals a complicated pattern of results where *everything seemingly matters*. The DCCS, then, presents a formidable theoretical challenge. Are existing theories up to this challenge? Theories that address the development of executive function within the context of the DCCS task run the gamut from an information processing theory framed around hierarchical rule representation (Zelazo et al., 2003), to conceptual accounts of the task framed around attentional inertia (Kirkham, Cruess, & Diamond, 2003) or re-description (Kloo & Perner, 2005), to a formal connectionist model framed around active versus latent representations (Morton & Munakata, 2002). Next, we discuss the extant accounts in turn, building from abstract conceptual accounts at one end of the theoretical spectrum, to formally implemented models at the other end of the spectrum.

Cognitive Complexity and Control Theory—The most comprehensive account of children's performance in the DCCS task to date is the Cognitive Complexity and Control theory (CCC; Zelazo, 2004; Zelazo et al., 2003). CCC is an information-processing theory which conceptualizes children's behavior and development around hierarchical rule-representation. This theory contends that for children to be successful in the post-switch phase of the DCCS, they need to be able to reflect on the two sets of rules and construct a representation of a rule structure that can integrate the rules for the different featural dimensions. This enables them to select the appropriate rules given the game being played.

To construct a set of rules for the color or shape game, specific feature values (i.e., *antecedent conditions*) are connected to a location where the feature is to be sorted (i.e., *consequences*). This rule representation takes the form of an if-then production rule that reads, ‘if red, then place the card here, but if blue, then place the card there.’ Sets of antecedent conditions and consequences are linked at the next higher level in the hierarchy in a relationship between the dimensions of the rules (i.e., *setting conditions*) which specify when the different sets of rules should be used (see Figure 2). Integrating this representation with the representation of the individual rule sets also takes the form of an if-then statement reading, for example, ‘if color game and if red, then sort here, but if shape game and if circle, then sort there.’

To engage a set of rules during the pre-switch, the child must make an active decision to use a particular set of rules given the conflict created by the test card matching both of the target cards. One set of rules is used and gains a high level of activation, while the other rule set is inhibited and decreases in activation. Going into the post-switch phase, children must engage the second setting condition via a top-down process driven by reflection on the rule-structure. Without the representation of setting conditions in a second level of the rule hierarchy, children default to the more active set of rules that were initially used during the pre-switch. As the prefrontal cortex develops, this cortical area is able to represent increasingly complex rule-structures, allowing the child to construct a rule-hierarchy that allows for the top-down selection of the appropriate setting condition (see Bunge & Zelazo, 2006). Three-year-olds, then, are able to construct separate branches of antecedent conditions and consequences and are able to use these representations to sort cards during the pre-switch or answer questions about the post-switch rules. Nevertheless, 3-year-olds lack the ability to integrate both branches of rules under the setting conditions at a second level in the hierarchy.

Using these concepts of rule-representation and activation, the CCC theory can explain results of the NP and Partial-Change versions. In either situation, the rules that apply to the pre-switch dimension are more active during the post-switch phase either because the rules for the post-switch dimension were inhibited during the pre-switch (Negative Priming version; Figure 1B) or the rules for the pre-switch dimension have a high level of activation after being used in the pre-switch phase (Partial-Change version; Figure 1C). This imbalance of activation—even when one set of features changes—requires a rule-hierarchy for the top-down selection of the appropriate rules during the post-switch phase. Zelazo et al. (2003) further state that negative priming and the persistent activation of rule-sets depends on “the active selection of one pair of rules against a competing alternative” (p. 95). Thus, in the no-conflict Standard (Figure 1D) and no-conflict NP versions (Figure 1E), children are able to use two un-integrated branches of rules successively because the processes of inhibition and negative priming no longer have significant contributions to the activation of rules within the hierarchy. With the shape and color rules acquiring equal levels of activation, simply telling the child to use a different set of rules provides sufficient activation to use those rules during the post-switch phase.

By structuring the rule-hierarchy around the different dimensions of the visual features (i.e., the setting conditions), CCC theory is able to generate various predictions that have been empirically supported. For example, Zelazo et al. (2003) showed that 3-year-olds are unable to switch between single rules for features within different visual dimensions or setting conditions. For example, 3-year-olds were asked to sort a single test card containing an image of a green car within the standard task-space structure containing two sorting locations with target cards. When playing the green game, the card was to be sorted by color, but if playing the car game the card was to be sorted by its shape. Even with these simpler branches containing a single rule under each setting condition, 3-year-olds perseverated on the rule they used first. Zelazo et al. (2003) further showed that 3-year-olds are able to use four rules at once (e.g., using four color rules) and switch between sets of rules if they are under the same setting condition (e.g., switching between two sets of color rules using red and blue, or green and yellow). Thus, the number of rules under different setting conditions does not necessarily matter. This indicates that children's difficulty does not stem from a limitation in memory capacity or from a general difficulty in switching rules. Rather, 3-year-olds only have trouble when the pairs of rules require attention to different visual dimensions and, thus, span separate branches of the rule-hierarchy.

Although CCC theory generalizes to a broad range of effects, it has several critical limitations. Most critically, CCC theory does not specify how children learn to construct complex rule-hierarchies over development. Moreover, this theory does not explain the real-time process of forming and engaging rules or how the dynamics of activation or inhibition work within the hierarchical system. In this sense, although CCC uses neural concepts, it is not well-positioned to integrate brain and behavior. Rather, ties to neural dynamics and development have remained largely at the descriptive level (Bunge & Zelazo, 2004). Further, CCC is not able to shed light on the multi-faceted nature of EF. CCC theory explains performance in the DCCS using a specialized rule-switching system which limits its applicability to other components of EF. Finally, there are empirical data that highlight a fundamental limitation of an abstract rule representation framework. Specifically, CCC theory is unable to explain asymmetries in switching ability based on the attentional salience of different features used in the task (Fisher, 2011; Honomichl & Chen, 2010).

Other Conceptual Accounts—An alternative account, the Attentional Inertia hypothesis (Kirkham et al., 2003), moves away from the concept of rule-representation and instead attributes children's perseveration to inflexible attention. In particular, when using a particular set of rules, children must selectively attend to a visual dimension of the stimulus. Children perseverate because attention becomes stuck on the featural dimension to which it is initially applied. Shifting attention, then, requires the active inhibition or suppression of the current deployment of attention. By inhibiting and disengaging, attention becomes free to shift to a different dimension of the stimuli. With poorly developed inhibitory control, the child is unable to disengage, and is thus unable to refocus her attention on the new dimension to use the post-switch rules.

To support their account, Kirkham et al. (2003) showed that boosting attention to the post-switch features by prompting children to re-label the test cards during the post-switch by the relevant dimension significantly improved children's ability to switch rules (however, see also Müller, Zelazo, Lurye & Liebermann, 2008, for a failure to replicate these results in a series of experiments). Kirkham et al. (2003) further showed that sorting cards face-up impaired 4-year-olds’ ability to switch rules. They suggest that this manipulation provided an even stronger pull on attention to the pre-switch dimension, making attention even more inflexible. Diamond and Kirkham (2005) also found evidence for attentional inertia with adults in a timed version of the DCCS. Adults showed significantly slower reaction times on the trials immediately following a switch in rules, suggesting that it even takes adults extra time to shift their attention to another featural dimension.

Another alternative account put forth by Kloo et al. (2008) and Kloo and Perner (2005)—the Re-description hypothesis—frames children's difficulty in the DCCS around an inability to re-describe objects. They suggest that children use a more general matching strategy along dimensions (e.g., match by color) than a hierarchy of rules. For this strategy to work, however, children need to engage descriptions of the objects along the correct dimension. Children fail to switch in this case, because they have difficulty re-describing the test cards by the new dimension. If children can apply a new label to a different object, they should have little trouble switching rules. This appears to be the case: 3-year-olds succeed with separated dimensions where they can describe a patch of color as red and the outline of the shape as a star (Perner & Lang, 2002; Zelazo et al.,2003). Further, when the post-switch is introduced as a ‘silly’ version of the first game and children are instructed to match the pre-switch features to the opposite target cards, children do not need to re-describe the test cards by the new dimension and they typically succeed. What develops to allow correct rule-switching, according to this view, is “a conceptual understanding that things can be described differently under different perspectives” (p. 132, Kloo et al., 2008).

The attentional inertia and re-description hypotheses offer different ways of describing children's behavior in the DCCS task. The attentional inertia hypothesis lacks a clear definition attention, which is problematic given the complexity of this cognitive construct (for a review, see Luck & Vecera, 2002). Specifically, it is unclear what type of attention is involved, what mechanisms are involved in changes in attention, or how the influence of inhibitory control on attention changes over development. Similarly, the re-description hypothesis does not offer an account of what processes underlie children's ability to apply flexible descriptions or what mechanisms produce changes in children concepts which support flexible descriptions. Thus, it is difficult to determine exactly how well these accounts capture behavior, how they generalize beyond the DCCS task, and what develops to enable flexible behavior.

Connectionist Models—Morton and Munakata (2002) took an important step toward formalizing the neural and developmental mechanisms underlying the DCCS by implementing this task in a connectionist model (see Figure 3). In the model, relevant dimensions are represented within a set of pre-frontal cortex (PFC) nodes. These nodes reflect the dimensionality of the stimulus, representing either shape or color as the relevant dimension. When the model *actively* represents either the shape or color ‘rules’ in the PFC nodes, the feed-forward connections between the hidden units and output units for the relevant features are strengthened. As inputs are processed and ‘sorted’ according to the pre-switch dimension, *latent* connections are established through a Hebbian process between the hidden layer and output nodes. Importantly, only the features which are used for sorting undergo Hebbian learning. This creates a bias in the system toward making decisions based on that dimension. For the model to overcome these latent connections and correctly sort by the post-switch rules, the PFC nodes need to have a relatively strong active representation of the current rules to shift the balance of activation between hidden layer and output layer toward making decisions based on the post-switch dimension. As the PFC develops (i.e., as the strength of recurrent activation in the model's PFC nodes is increased), the model is able to actively represent the relevant rules which exert a stronger influence on the decision-making process. This facilitates flexible rule-use.

This model has generated various predictions about children's rule-use in the DCCS task. For example, Morton and Munakata (2002) gave their model practice with unidimensional post-switch inputs before the start of the post-switch phase. This increased the strength of latent traces for the post-switch features and improved the performance of the model. Brace, Morton, and Munakata (2006) tested this prediction by administering a training phase between the pre-switch and post-switch phases. The training phase consisted of a set of cards that started with univalent stimuli for the post-switch features. Throughout the training phase, the pre-switch features were gradually ‘morphed’ into the image on the cards until they contained fully bivalent stimuli when the child started the post-switch phase. This training phase significantly improved three-year-olds' post-switch performance over that of children who had a training phase with irrelevant features or dimensions.

In a second test of the model, Jordan and Morton (2008) explored the role that environmental support can have in promoting the active representations of new rules. They showed that children are better able to switch rules with the use of congruent flankers. For instance, if color were the relevant post-switch dimension, patches of color would flank the test-card image. With this extra environment support, children can activate the relevant task rules more strongly and correctly switch rules.

The account offered by Morton and Munakata (2002) has many strengths. It uses a neurally plausible implementation of rule-use and has generated multiple empirical predictions that have been successfully tested. Nonetheless, this theory does not achieve as much coverage of the DCCS literature as CCC theory. In particular, the model has not been generalized to account for other versions of the task to explain, for example, why children still perseverate in the NP version or are able to switch in versions where there is no conflict during the pre-switch phase. Indeed, there are no dimensions per se in the model; there are only associations between the PFC and hidden units representing the particular features of the task. Thus, it is unclear how the model would capture the introduction of new features at different phases of sorting. Interestingly, the model has not been used to quantitatively capture children's performance or changes over development, although Morton and Munakata (2002) have demonstrated that correct switching increases as the recurrent connections for the PFC nodes are strengthened.

Summary—The theoretical perspectives reviewed here offer differing views of what a ‘rule’ is and how the ability to use rules changes over development. CCC theory has been used to conceptually integrate an extensive portion of the literature but lacks specification of several key processes and an ability to interface with neural processes. The connectionist model put forth by Morton and Munakata (2002), at the other extreme, uses formal neural concepts but has not been used to explain a broad array of findings in the literature.

Our primary aim in developing the DNF model described below is to achieve extensive theoretical coverage while adhering to, and anchoring our concepts in, formalized and general neural principles. This integrative focus is inspired by recent data showing not only the complexity of explaining cognitive flexibility at the neural level, but also the promise that this level of explanation might offer for new ways to test and constrain developmental theories of EF. In the present paper, we take a first step toward this brain-behavior link by grounding our theory of executive function and the DCCS task within Dynamic Field Theory. Critically, DFT sheds light on the integration of EF components in the DCCS task, and provides insight into the neural processes associated with different aspects of rule-use and EF. Moreover, the model offers useful clarity about the link between cognitive concepts and their neural instantiation.

Dynamic Field Theory

Dynamic Field Theory (DFT) grew out of the principles and concepts of dynamical systems theory initially explored in the ‘motor approach’ pioneered by Gregor Schöner, Esther Thelen, Scott Kelso, and Michael Turvey (Kelso, Scholz & Schöner, 1988; Schöner & Kelso, 1988; Thelen & Smith, 1994; Turvey & Shaw, 1995). The goal was to develop a formal, neurally-grounded theory that could bring the concepts of dynamical systems theory to bear on issues in cognition and cognitive development. DFT was initially applied to issues closely aligned with the cognitive aspects of motor systems such as motor planning for arm and eye movements (Erlhagen & Schöner, 2002; Kopecz & Schöner, 1995). Subsequent work extended DFT, capturing a wide array of phenomena in the area of spatial cognition, from categorical biases to changes in the metric precision of spatial working memory from childhood to adulthood (Schutte, Spencer, & Schöner, 2003; Simmering, Schutte, & Spencer, 2008). More recently, these models have been used to capture how objects are neurally represented in a way that links features to a spatial frame of reference (Johnson, Spencer, & Schöner, 2008), how object recognition can emerge from associating features with labels (Faubel & Schöner, 2008), and how young children learn words in a social context using a common spatial frame to bind words and objects together (Samuelson, Perry, Smith & Spencer, in press).

The model we present here is not solely a rule-use model; rather, it is an extension of an object WM model proposed to capture the integration of ‘what’ (i.e., ventral stream) and ‘where’ (i.e., dorsal stream) visual pathways (Johnson et al., 2008; see also Samuelson et al., in press). This model represents ‘integrated’ objects by binding together different neural populations that encode and remember feature values (e.g., colors, orientations) using a shared spatial dimension. Here, we couple a dimensional attention system to this model that can autonomously modulate the baseline activation level of different cortical populations. By ‘boosting’ the resting level of neural populations tuned to particular dimensions (e.g., color), the model can selectively bind the features of objects together in a rule-like fashion and sort these objects in space. We then use this integrated DNF model to probe whether specific developmental changes in the dimensional attention system are sufficient to capture the emergence of flexible rule-use in the DCCS task between 3 and 4 years.

In the following sections, we describe the DNF model. We begin with an overview of several central concepts in DFT, including a discussion of neural population dynamics within multi-layered cortical fields, as well as the variant of Hebbian learning we use to capture changes in neural dynamics over a trial-to-trial timescale (see, Faubel & Schöner, 2008; Lipinski, Spencer, & Samuelson, 2010; Samuelson et al., in press; Simmering et al., 2008). Next, we describe the object WM model (Johnson et al., 2008; see also, Samuelson et al., in press). This is followed by an overview of the dimensional attention system that we couple to this model to enable flexible rule-use, as well as a discussion of the developmental hypotheses explored in the present study. Third, we step through how the DNF model sorts cards in a rule-like fashion in the DCCS task, perseverating early in development and switching rules later in development. Finally, this section concludes with a discussion of how the DNF model captures the multi-component nature of EF.

Basic Concepts of DFT

DFT simulates thinking in the form of neural population dynamics, that is, patterns of neural activity within cortical fields consisting of neurons ‘tuned’ to continuous metric dimensions (e.g., space, color, orientation). These fields are organized so that neighboring locations in a field have similar receptive fields—they respond maximally to similar feature values. Moreover, neighboring neurons share lateral excitatory interactions, such that the activity of one neuron can boost the activity of its neighbors. By contrast, neurons with very different receptive fields share inhibitory interactions, such that the activity of a local group of neurons can inhibit the activity of neurons tuned to different feature values. When inputs to such a field are strong enough to reach an activation threshold (an activation level of 0), these inputs combine with lateral interactions among neurons (local excitation, surround inhibition) to form “peaks” of activation that stably represent a particular feature value. These peaks represent a type of neural decision that a particular feature value is, for instance, present in the world. Moreover, if neural interactions are strong enough, peaks can remain stably activated through time—even in the absence of sensory stimulation. In this sense, peaks can capture key properties of ‘working’ or ‘active’ memory (Compte, Brunel, Goldman-Rakic, & Wang, 2000; Edin et al., 2007; Edin et al., 2009).

In the present report, we use two-layered neural fields composed of an excitatory layer of neurons (which we call a working memory (WM) field) coupled to a layer of inhibitory interneurons (Inhib field). The WM field is self-excitatory. When neural activity in this layer exceeds an activation threshold (>0), these neurons pass excitation to the Inhib field. Inhibitory interneurons, in turn, pass broad inhibition back into WM. This dynamic back-and-forth among layers creates the local excitatory and laterally inhibitory interactions needed to form localized peaks within the field (see the Appendix for a more detailed discussion and the two-layered neural field equations).

Our DNF models also use a variant of Hebbian learning which allows neural populations to learn as they acquire a history in the task. Specifically, as neurons in a WM field are activated, they leave a trace in a Hebbian layer. These traces—which vary in strength from 0 to 1, much like a synaptic connection in a connectionist model—are projected back onto the neural population, effectively boosting previously-activated neurons in the field. As a consequence, Hebbian traces (HT) enable specific neural sites in the field to become activated more quickly on subsequent trials. This results in a priming effect, facilitating the response of a local population of neurons to a familiar stimulus.

A central question with any neural network framework is how the concepts of the theory are anchored to the neural reality of the brain. On this front, DFT is well-grounded. The layered architecture we use was initially developed to capture neural activation patterns within visual cortex (Amari, 1977, 1980; Amari & Arbib, 1977; see also, Compte et al., 2000; Wilson & Cowan, 1972). More recently, our colleagues have demonstrated that the population dynamics within neural fields can be directly derived from multi-unit neurophysiology using the Distribution of Population Activation approach (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schöner, & Riehle, 2003; Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999; Jancke et al., 1999). This approach enables researchers to directly test DNF models using neurophysiological measures in non-human animals. More recent efforts suggest that it might be possible to capture aspects of human neural function by mapping the neural population dynamics in DNF model to the hemodynamic responses measured using fMRI (Edin et al., 2007).

A DNF Model of Object WM

The model we propose here is based, in part, on a recent model of object working memory shown in Figure 4 (see Appendix for equations). This model is composed of a one-dimensional spatial field, a two-dimensional color field, and a two-dimensional shape field. For the spatial field, the spatial dimension is plotted along the x-axis while activation is plotted along the y-axis. For the feature fields, the spatial dimension is plotted along the x-axis, the feature dimension (i.e., shape or color) is plotted along the y-axis, and the strength of activation within these fields is depicted by the color (see the inset scale). Note that neural population dynamics within two-dimensional neural fields have the same properties discussed above, except now lateral neural interactions occur along both dimensions. This creates a ‘Mexican-hat’ profile with a peak of excitation surrounded by a circular trough of inhibition. Note also that each WM field in Figure 4 consists of 3 layers: an excitatory layer, a layer of inhibitory interneurons, and a Hebbian layer. For simplicity, we only show activation in the excitatory layer in all figures. This makes sense given that the inhibitory and Hebbian layers have their impact by modulating neural activity in the excitatory layer. We highlight this in some cases, showing, for instance, the impact of the Hebbian traces on neural activity in the excitatory layer.

The architecture of the object WM model was inspired by the properties of the primate visual system which has distinct processing pathways for visual information (Haxby et al., 1991; Ungerleider & Mishkin, 1982): a dorsal (‘where’) pathway that encodes the spatial locations of objects (Anderson, 1995) and a ventral (‘what’) pathway that encodes object features such as color or orientation within different populations of neurons tuned to different visual dimensions (Desimone & Gross, 1979; Xiao, Wang, & Fellman, 2003). This creates a ‘binding’ problem in vision (Treisman, 1996; Treisman & Gelade, 1980): given that different populations of neurons are tuned to different feature dimensions in the ventral pathway, how does the brain know which features go with which other features to quickly form a novel, integrated object representation?

Figure 4 shows a concrete example of this challenge: when shown the red rectangle and the pink diamond in Figure 4A, how does the brain know that the red hue—actively represented by neural populations in the fusiform area (see green circle in the brain inset; Simmons, Ramjee, Beauchamp, McRae, Martin, & Barsalou, 2007)—should be coupled to neural populations actively representing aspects of the rectangular shape in the lateral occipital complex (see yellow circle in the brain inset; Drucker & Aguirre, 2009; Kourtzi, Erb, Grodd, & Bülthoff, 2003). One candidate solution to this problem is to capitalize on the bimodal nature of neural populations in the ventral stream. In particular, ventral stream neurons are sensitive to visual features but are also coarsely receptive to spatial information (see Aggelopoulos & Rolls, 2005; DiCarlo & Maunsell, 2003; Op De Beek & Vogels, 2000; for a review, see Kravitz, Vinson, & Baker, 2008). Thus, in the DNF model, neural populations are bound together by a common spatial dimension anchored to the precise spatial representations found in the parietal cortex (see purple circle in the brain inset; for a review, see Silver & Kastner, 2009). Note that this solution to the binding problem shares elements with Feature Integration Theory proposed by Treisman and Gelade (1980).

The sequence of simulations in Figure 4 illustrates how the DNF model ‘binds’ visual features together. In Figure 4A, two objects are presented to the model simultaneously—a red square on the left and a pink diamond on the right. This produces feature inputs that overlap along the spatial dimension in the color and shape fields due to the coarse spatial receptive fields of neural populations in the ventral pathway. Left alone, these bimodal cortical fields would have difficulty binding the correct colors to the correct shapes (for a demonstration, see Johnson et al., 2008). The model resolves this ambiguity through spatial coupling: all of the WM fields share a common spatial dimension and are coupled to the spatial field which can precisely localize spatial inputs (see small bumps in the spatial field in Figure 4A which captures the initial detection of an object to the left and right). When activation goes above threshold in one field, spatial information projects to the other WM fields at the associated spatial location. This is shown in Figure 4B which shows the fields after peaks have emerged. As can be seen, the model correctly binds ‘red’ and ‘square’ on the left and ‘pink’ and ‘diamond’ on the right. Note the vertical ridges of sub-threshold activation at the left and right locations in the 2D fields. This reflects the coupling across space which gives rise to an integrated object: a distributed pattern of peaks in multiple feature fields, all coupled to a shared spatial dimension. Figure 4C shows the fields after the inputs have been turned off. This reveals the contribution from the HTs: the model has formed a memory trace that the red square was on the left and the pink diamond was on the right. Consequently, when these objects are presented again, peaks will build more readily at these particular feature-space conjunctions.

Autonomous Dimensional Attention and the Development of Rule-Use

The model shown in Figure 4 is able to actively form working memories and Hebbian associations for object features in the task space, but how does it sort cards and behave in a rule-like fashion in the DCCS task? For this, we need to add an additional concept—dimensional attention. Figure 5 shows the complete model architecture which adds a frontal dimensional attention system (see blue circle in brain inset; Morton et al., 2009) to the object WM model. This frontal system is reciprocally coupled to cortical fields in the ventral and dorsal pathways (see Crone, Wendelken, Donohue, van Leijenhorst, & Bunge, 2006; Dosenbach et al., 2007)[i].

The dimensional attention system consists of two attentional nodes labeled ‘color’ and ‘shape’. The nodes have the same dynamics as neural fields: they are self-excitatory, mutually-inhibitory, and can show robust, above-threshold activation when sufficiently stimulated. In this sense, then, the nodes can be viewed as localized neural populations in frontal cortex that can enter a robust ‘peak’ state when the system is actively attending to a particular type of information. These nodes also learn as they are repeatedly activated in a

---

[i] Note that the fronto-parietal connections did not play a role in the simulations reported here. Consequently, we do not discuss them further, although recent studies show developmental changes in fronto-parietal connectivity during childhood (Crone, Donohue, et al., 2006).

task. In particular, the frontal system uses the same type of Hebbian learning described previously. Hebbian traces at the level of the frontal system boost the baseline level of activation in the relevant node which can effectively prime ‘color’ or ‘shape’ responding.

The information that each node actively represents is reflected in the pattern of connectivity between the frontal nodes and the posterior fields in the dorsal and ventral pathways. As is shown in Figure 5, the ‘color’ node has strong excitatory connections to the color field, and weaker excitatory connections to the shape field. Consequently, when the ‘color’ node goes into an active or ‘on’ state, the color field receives a global activation boost. This makes it easier for peaks to form in the color field which can bias the model to initially form peaks in the posterior neural system based on color information. Reversely, the ‘shape’ node has strong excitatory connections to the shape field, and weaker excitatory connections to the color field. Consequently, when the ‘shape’ node goes into an active or ‘on’ state, the shape field receives a global activation boost. This creates a bias in the model to initially form peaks in the posterior neural system based on shape information.

The activation dynamics within our model are consistent with neurophysiological evidence (e.g., Egner & Hirsch, 2005; Lepsien & Nobre, 2007; Zanto, Rubens, Bollinger, & Gazzaley, 2010; Zanto, Rubens, Thangavel, & Gazzaley, 2011) which has demonstrated boosts in baseline activity for cortical areas processing task-relevant information under situations with high executive demands. For example, in a face/name Stroop task, boosts of baseline activity were observed in the fusiform face area on trials that required participants to categorize faces when conflicting names were presented. Note that the patterns of connectivity in the model are fully reciprocal; thus, the growth of neural activity in the posterior neural fields can also impact the activation of the frontal system.

In the present work, we established the frontal-posterior connectivity patterns by hand. This allowed us to probe whether specific changes in the connectivity pattern over development were sufficient to capture the emergence of rule-use in the DCCS. In particular, we created a 3-year-old model that had poorly organized connectivity: activation of the ‘color’ node, for instance, boosted the baseline level of activation in the color field, but also boosted the shape field. By contrast, the 4-year-old model had well-organized connectivity: activation of the ‘color’ node selectively boosted the color field and had little impact on the shape field (see Appendix for details). As we discuss below, these changes in connectivity were critical in capturing key aspects of young children's performance in the DCCS task.

In addition to refining the connectivity between the frontal and posterior neural systems in the model, we also implemented a specific developmental hypothesis central to other work using DFT—the spatial precision hypothesis (Schutte & Spencer, 2009; Spencer, Perone, & Johnson, 2009). According to this hypothesis, excitatory and inhibitory neural interactions become stronger over development. Here, we asked whether it was sufficient to boost the strength of excitatory and inhibitory neural interactions in the dimensional attention system to capture developmental change.

Note that, although both developmental changes in the model—changes in the pattern of connectivity and changes in the strength of neural interactions in the dimensional attention system—were implemented ‘by hand’, recent work suggests that these changes could emerge via an autonomous learning mechanism. We highlight this direction in the General Discussion by describing recent work using DFNs which shows how patterns of connectivity between nodes and fields can be learned by standard forms of Hebbian learning as the system learns the correlation between, for instance, the use of the word ‘color’ and the importance of task-relevant hue values in the task space (see Faubel & Schöner, 2008; Sandamirskaya & Schöner, 2010).

The DNF Model and the DCCS Task

Figure 6 shows the young 3-year-old model as it sorts cards in the DCCS task. The top panel shows the activation of the dimensional nodes across the pre-switch and post-switch phases. The bottom panels show the posterior neural system at particular points during the task indicated by the vertical lines in the top panel.

Figure 6A shows the *target inputs* presented to the model that capture the presence of the target cards and trays used in the DCCS. For this example, the sorting trays contain a blue-circle on the left and a red-star on the right. In the spatial field, there are subthreshold ‘bumps’ of activation at the left and right locations. These serve to “pre-shape” or prime activation in the fields to build peaks at the locations of the trays, that is, to sort cards to these locations (for related ideas, see Erlhagen & Schöner, 2002). In the color field, there are two subthreshold hotspots of activation. These correspond to the blue item to the left and the red item to the right. There are also two subthreshold hotspots in the shape field. These correspond to the circle at the left location and the star at the right location. Thus, between the shape and color fields, there are inputs for a blue circle at the left tray and a red star at the right tray.

Test inputs are presented to the feature fields as ridges of activation for a feature across all spatial locations. This captures the nature of the task: the model, like children, must spatially localize the features presented on the test cards; that is, the model must decide to place the card to the left or right. Figure 6B shows the test inputs when a red circle is presented to the model. Notice that there is spatial conflict between the two feature fields for the features presented on the test card. The test input for the red feature overlaps with the target input at the right location. The test input for the circle feature, on the other hand, overlaps with the target input at the left location in the shape field. With these conflicting cues, the model has difficulty spatially localizing the inputs in a systematic fashion.

Critically, however, the model is ‘told’ to play the color game at the beginning of the pre-switch phase. This is implemented as a sub-threshold input to the ‘color’ node (note the slight boost in activation for the ‘color’ node in the top panel at timestep 300). As a test card is presented in panel B, the nodes become activated based on the inputs from the feature fields (see boost in activity starting at timestep 600). The self-excitatory and inhibitory interactions between these nodes are weak. Thus, neither node is boosted above threshold (i.e., above 0). Nevertheless, the color input is sufficient to tip the scales in favor of sorting the test card based on color. Figure 6C shows this: the model has built peaks of activation at the right location, effectively binding the red and circle features to the right location in the task space. Figure 6D shows the consequence of this sorting decision: there is a HT (see white circles) in the color field that associates the red feature and the right sorting location; similarly, there is a HT in the shape field that associates the circle feature and the right sorting location.

The same sequence of events play out as the model sorts the second type of test card highlighted in panels E-G. In Figure 6E, the model is shown a blue star to sort and is told to play the color game. This boosts the activity of the ‘color’ node (see top panel), biasing the model to form a peak of activation in the color field. Consequently, the model forms a peak of activation to the left side of the color field, binds the star feature to the left as well, and effectively sorts the blue star to the left bin (see Figure 6F). Figure 6G shows the result of repeatedly sorting red circles to the right and blue stars to the left during the pre-switch phase: the model has robust HTs of where the test cards were sorted on the previous trials (see white circles). Critically, this results in *cooperation* in the color field—the HTs overlap with the target inputs—and *competition* in the shape field—the HTs associate the star with the left and the circle with the right, while the target inputs indicate a circle on the left and a star on the right.

Figure 6H shows the consequence of this pattern of cooperation and competition for the 3-year-old model on the first post-switch trial—the model perseveratively sorts the red circle to the right. Why does this happen? As can be seen in the top panel, the post-switch phase begins with a boost to the ‘shape’ node as the model is told to play the shape game. The ‘shape’ node is given the same strength of input as the ‘color’ node during the pre-switch phase. Nevertheless, this slight boost is not sufficient to beat out the cooperation in the color field (i.e., the overlapping HTs and target inputs), especially since there is competition in the shape field which slows down the rise of activation.

How is the model, like 4-year-olds, able to overcome these biases and switch rules? Figure 7 shows a simulation of the 4-year-old model across the pre- and post-switch phases. As with the 3-year-old model, the simulation begins by telling the model to play the color game. Even though the input to the ‘color’ node is the same strength we used with the 3-year-old model, the ‘color’ node for the 4-year-old model becomes robustly active due to the stronger neural interactions present in the dimensional attention system. That is, the inputs to these nodes are sufficient to lift neural interactions to a point where self-excitation can rise above-threshold and actively suppress the other dimensional node. This selectively boosts the color field due to the refined pattern of connectivity between the frontal and posterior systems, and the model sorts the red circle to the right (see Figure 7C). Similarly, when the model is shown a blue star, it sorts this card to the left bin due to the robust boost to the color field from the dimensional attention system (see Figure 7F). Note that this sorting behavior results in the same pattern of HTs and target inputs going into the post-switch phase: as with the 3-year-old model, there is cooperation in the color field and competition in the shape field.

Critically, when the 4-year-old model is asked to play the shape game, the developmental changes to the dimensional attention system enable robust task-switching. In this case, input to the ‘shape’ node lifts this neural population sufficiently to achieve an above-threshold activation state. This sends a strong boost of activity to the shape field and suppresses the activity of the ‘color’ node. Consequently, the 4-year-old model is able to overcome the competition in the shape field and successfully sort the red circle to the left (see Figure 7H).

Summary

The DNF model consists of multiple, reciprocally coupled neural fields with featural and spatial properties common to neural populations in the dorsal and ventral visual pathways. Through neural interactions within these populations, stable peaks of activation can emerge that capture decisions to sort objects (i.e., cards) to the left or right in the DCCS task. Hebbian processes operating within these fields serve to accumulate traces over the course of the task which can influence subsequent decisions.

To sort cards in a rule-like fashion, we introduced a frontal dimensional attention system that effectively boosts the baseline level of activity in the posterior neural fields (e.g., Egner & Hirsch, 2005; Lepsien & Nobre, 2007; Zanto et al., 2010; Zanto et al., 2011). In addition, we implemented a specific type of developmental change in the dimensional attention system to capture the emergence of rule-use between 3 and 4 years: the older model had stronger excitatory and inhibitory neural interactions in the frontal system and a more precise pattern of connectivity with the posterior system. As a consequence, the 3-year-old model perseverated: this model was not able to achieve above-threshold activation of the ‘shape’ node; consequently, it perseveratively sorted by color due to cooperation in the color field and competition in the shape field. By contrast, the 4-year-old model switched rules: this model robustly activated the ‘shape’ node, boosted the baseline activity in the shape field, and correctly sorted the red circle to the circle target card on the left.

In the next section, we demonstrate that the DNF model can explain a host of findings in the literature and capture these findings in quantitative detail. Before pursuing that goal, however, it is useful to consider two questions: how does the DNF model differ from other theories in the literature, and how does the DNF model relate to the broader literature on the multi-component nature of executive function? We consider each question in turn.

The DNF model differs from existing theories of the development of EF in several critical respects. Relative to CCC theory and other conceptual accounts of children's performance in the DCCS task (e.g., the attentional inertia account), the DNF model is neurally grounded. In particular, the model is based on initial work probing the neural population dynamics that underlie elementary forms of decision-making in simple sensory and motor situations (Amari, 1977, 1980; Amari & Arbib, 1977; Bastian et al., 2003). Moreover, the specific architecture we proposed was based on an earlier model of object WM that captures how people quickly form integrated object representations in working memory across distributed neural populations (Johnson et al., 2008; Samuelson et al., in press). Finally, we added a frontal dimensional attention system to the object WM model that can boost the baseline levels of activation in posterior cortical fields, consistent with recent neural evidence (e.g., Egner & Hirsch, 2005; Lepsien & Nobre, 2007). Thus, we contend that this theory is well-positioned to integrate neural and behavioral dynamics. This is an exciting advance in that it opens the door to probing how this theoretical account maps onto the emerging literature using functional neuroimaging, an issue we return to in the General Discussion.

Another key difference between the DNF model and CCC theory is that our theory is not strictly hierarchical. For instance, the dimensional attention system does not ‘know’ all the details of the rules; rather, this system creates a bias, but then ‘off-loads’ the detailed decision-making to the feature WM fields. Similarly, the feature fields ‘know’ the details of the specific stimulus-response mappings, but these fields do not ‘know’ the higher-order goals of the task (e.g., sort by shape or color). This type of distributed neural organization can enable flexible behavior and the ‘soft assembly’ of rule-like behavior, that is, success or failure can come from multiple sources. We highlight this in the next section.

The other primary account of children's performance in the DCCS task is the connectionist model proposed by Morton and Munakata (2002). The central difference between the DNF model and the connectionist model is that the latter model does not learn about the unattended feature values. Thus, when the connectionist model sorts a red circle to a red star during the pre-switch phase, it has no memory for the shapes of the objects. Our model, by contrast, binds features together during every sort. In a later section, we use this feature of the DNF model to make a novel prediction that is not consistent with the connectionist model.

Although there are substantive differences between the DNF and connectionist models, they share an important commonalty: these models do not have separate components that correspond to the different aspects of EF. Rather, EF emerges from a continuous, distributed, and fully interactive neural system. In this context, what insights can the DNF model provide about the components of EF? We consider this question first in a much simpler context: previous work examining inhibitory control in early development in the Piagetian A-not-B task. In this task, infants are cued to reach to one location (the A location) over a series of trials. They are then cued to reach to a different location (B). Infants younger than 10 months typically make the A-not-B error and perseveratively reach back to A when cued to reach to B. Thus, young infants lack the inhibitory control to suppress the prepotent response to A.

Thelen et al., (2001) simulated infants’ performance in this task using a DNF model. As the model made a series of reaches to A, a Hebbian trace accumulated and boosted the overall amount of activation associated with the A location. When the young infant model was cued to B and then forced to wait during a few second delay, the peak associated with the cuing event at B fell to baseline levels. Consequently, the model perseveratively responded to A based on the strong HT of the A location. To overcome this bias, the older infant model had stronger neural interactions. This enabled the model to maintain a robust peak associated with the B location, that is, the model actively maintained a working memory of B.

This example is informative in that it highlights the challenge of relating cognitive concepts such as ‘inhibitory control’ to mechanisms in a neural model. The cognitive concept of inhibition reflects the ability of an individual to suppress a prepotent representation or behavior. Simulations of the DNF model certainly capture this: the older infant model effectively suppresses the prepotent response to A. Critically, however, this is achieved by having a robust working memory. In particular, the active maintenance of a peak associated with the B location suppresses the strong trace of A via lateral inhibition. Thus, does the

DNF model capture a change in inhibitory control over development or a change in working memory? The answer is ‘both’.

The story becomes more complex in the DNF model presented here in that the components of EF are distributed throughout the neural system. Task-switching is certainly enabled by the selective activation of the dimensional nodes; however, the 4-year-old model is able to switch tasks due to an increase in the strength of neural interactions that enables the ‘color’ or ‘shape’ node to enter a robust working memory state. Furthermore, the robust activation of a dimensional node in the post-switch phase actively suppresses the prepotent bias to response based on the previous rule—a clear form of inhibitory control. Thus, all three aspects of EF are evident in the neural interactions within the dimensional attention system.

There is a similar convolution of cognitive concepts in the posterior neural system. Here, a boost of input from the frontal system can bias the system to respond based on a particular featural dimension. This boost helps overcome the prepotent response accumulated during the pre-switch trials, but the actual inhibition of these HTs arises from the lateral inhibition generated by building a working memory peak in the feature fields.

In summary, then, the DNF model captures the multi-component nature of EF, but in a neurally grounded way that can help clarify the link between the cognitive concepts used in the behavioral literature and the neural mechanisms that operate in a complex neural system. For this link to have meaning, however, we must first demonstrate the utility of the theory. We do this in the following sections by first quantitatively simulating a host of behavioral effects in the literature. Next, we generate a set of novel behavioral predictions and test these predictions with 3-year-old children. Critically, these predictions are not consistent with any other theory of children's performance in the DCCS task. We then demonstrate that the DNF model can capture several additional findings from a recent study that present a challenge to other theories.

Quantitative Fits of Core Findings in the DCCS Literature

In this section, we present—for the first time—quantitative simulations of a diverse array of findings from the literature on children's performance in the DCCS task. Specifically, we show that the DNF model reproduces the pattern of data across versions of the task that manipulate the features between the pre-switch and post-switch phases, as well as versions that manipulate the degree of conflict during the pre-switch phase. These variants of the task have played a central role in evaluating existing theories of children's performance. Thus, the simulation results in this section establish that the DFT achieves extensive coverage of the extant literature at a level of quantitative detail that has not yet been achieved by any other theory.

In the previous section, we described how the model can autonomously sort cards by modulating the resting level of populations of neurons tuned to specific features. This was accomplished through the neural dynamics in a dimensional attention system that was coupled to the feature fields. Here, we simplified the neural dynamics within the dimensional attention system to make the task of quantitatively simulating results from many studies in the literature tractable. In particular, we replaced the autonomous dimensional attention system with two nodes that passed static activation values to the posterior neural system during the pre- and post-switch phases. This eliminated the rise and fall dynamics of the dimensional system evident in, for instance, the top panel of Figure 7, but retained the key consequence of these neural dynamics—a boost in the baseline level of activation in a particular feature field.

We used two parameters to specify the activation passed from the dimensional nodes to the feature fields. First, we used a *boost* parameter which specified the strength of the global boost. The 3-year-old model used small boost values to reflect the weak neural interactions within the dimensional attention system and the imprecise connectivity between the dimensional nodes and the feature fields. Second, we used a *shift* parameter which specified the fidelity of the boost during the post-switch phase. The 3-year-old model used weak shift values, reflecting the difficulty this model had selectively activating a single dimensional node during the post-switch phase given the Hebbian learning that boosts the dimensional node active during the pre-switch phase.

Figure 8 shows an example of the simplified dimensional attention system. In this example, the boost value is 0.4 and the shift value is .75. Thus, on the pre-switch trials (trials 1-6 along the x-axis), the resting level of the color field was boosted by .4 units, while the shape field remained at baseline. On the post-switch trials (trials 7-12 along the x-axis), we applied the shift value. Thus, on trial 7, 75% of the boost value (0.4*.75 = 0.3) was applied to the post-switch field—the shape field—while the remaining 25% remained with the color field.

To capture developmental differences between 3 and 4 years, we used the two *boost* distributions shown in Figure 9. As this figure shows, the 3-year-old model had a mean boost value of 0.35 and the 4-year-old model had a mean boost of 0.5. This reflects the increase in the strength of neural interactions in the dimensional attention system, as well as an increase in the precision of the connectivity between the frontal and posterior systems. We also used two *shift* distributions (see Figure 10). The 3-year-old distribution encompassed a broad range of values centered on 0.5 while the 4-year-old distribution was skewed more toward values between 0.6 and 1. This reflects an increase in the efficiency with which children are able to switch attention. Specifically, as the strength of excitatory and inhibitory interactions within the dimensional attention system increases, the relevant dimensional node can be selectively activated during the pre- and post-switch phases (see Figure 7). Note that the boost and shift values were selected independently for each simulation. This provides a conservative starting point for probing the development of flexible rule-use in this system.

Methods

Simulations were conducted in Matlab 7.5.0 (Mathworks, Inc.) on a PC with an Intel® i7™ 3.33 GHz quad-core processor (the Matlab code is available from the authors on request). The same model parameters were used for all conditions; the only differences across simulations were the inputs presented to the model to reflect differences in the DCCS tasks that were simulated.

For all simulations, the model was given six test cards (three of each possible combination of features) for both the pre- and post-switch phases. Throughout each simulation, target inputs were presented at specific feature and spatial values to capture the relevant details of the targets cards for the pre-switch and post-switch phases. At the start of each trial, the model was presented with ridges of input for the features present on the test cards. The model's response on each trial was determined based on the location (left or right tray location) of a peak in the spatial field. For example, if the model made a decision to sort a card to the left tray, then it would bind the features on the test card to the left location in the feature fields and build a corresponding peak of activation at the left location in the spatial field. Thus, the model generates a ‘left’ or ‘right’ decision on every trial. The model is scored as correct based on the spatial location of the target inputs to the shape and color feature fields and the ‘game’ being played (color game or shape game).

Each trial was simulated for 1500 timesteps, with the test card stimulus presented for 1000 timesteps. The model always generated a response by the end of the 1500 timestep interval. Next, we decreased the resting-level of the WM fields to delete any peaks and prepare the field for the subsequent presentation of a test card. For all of the figures and simulation examples in the Results section, color is the pre-switch dimension and shape is the post-switch dimension.

Three batches of 100 simulations (i.e., 100 children) were conducted for each condition for each age group to get rates of perseveration and switching for quantitative fits. As in the literature, the model was required to sort at least 5 out of 6 pre-switch test cards correctly. Further, a model was characterized as passing if 5 or more post-switch test cards were correctly sorted and as failing if 1 or fewer post-switch cards were sorted correctly. The model parameters can be found in the Appendix. As noted previously, these parameters were held constant for all simulations in both age groups.

Results

Grounding rule-use in a process-based model that generates active responses on every trial imposes robust constraints when trying to capture the details of 3- and 4-year-olds’ performance. First, because we simulated both the pre- and post-switch phases, the model must show high levels of correct sorting during the initial sorting phase. More importantly, the model must also capture the all-or-none aspect of children's performance during the post-switch. That is, the model must either completely perseverate and sort all of the cards incorrectly, or completely switch and successfully sort all of the cards correctly. The literature reveals a very specific pattern of results depending on the presence or absence of conflict and different changes to the features on the cards. In the following section, we show that the model is able to quantitatively capture these details.

Standard—The first issue we probed was whether the model quantitatively captures children's performance in the Standard DCCS task. Figure 11 shows the rates of perseveration and switching for the model and literature sampled for 3- and 4-year-olds (Halford et al., 2007; Müller et al., 2006; Zelazo, Frye, & Rapus, 1996; Zelazo et al.,2003). As can be seen, the model exhibits a high rate of perseveration similar to 3-year-olds and also closely matches the level of correct switching for 4-year-olds. As discussed above, the only difference between these simulations across age is the strength of the *boost* for each age group and the *shift* of the resting-level modulation between the pre- and post-switch phases.

The left panel in Figure 11 explains why the 3-year-old model perseverates. Going into the post-switch phase, there is cooperation of HTs and target input in the pre-switch field (see color field in Figure 11). That is, the model has always sorted the red test cards to the left and the blue test cards to the right. By contrast, there are competing HTs and target inputs in the post-switch field: the model is ‘seeing’ a star on the left and a circle on the right, but has HTs at the opposite locations (see shape field in Figure 11). Thus, the model must resolve competitive memories for the post-switch features in the context of cooperation in the pre-switch field. In these circumstances, the model requires a strong boost to the shape field to sort by the post-switch dimension. Thus, the weak boost of the relevant field for the 3-year-old model is unable to inhibit the perseverative biases accumulated during the pre-switch. The 4-year-old model, on the other hand, has a stronger boost which serves to strongly activate the post-switch field. This, in turn, inhibits the pre-switch dimension allowing the model to correctly sort by the post-switch rules.

No-conflict Standard—If conflict during the pre-switch phase plays a critical role in 3-year-olds’ perseverative behavior, then eliminating this conflict should have a significant impact on their performance. Zelazo et al. (2003) investigated this issue using a no-conflict standard version where the pre-switch test cards match the target cards along both color and shape dimensions. Three-year-olds were able to switch rules in this case. Figure 12 shows simulation results and data with 3-year-olds (Zelazo et al., 2003). The model quantitatively reproduces the relatively high switching rate displayed by 3-year-olds—both models and children sort correctly over 60% of the time.

As can be seen in the left panel, no-conflict cards during the pre-switch establish overlapping HTs and target inputs in both the shape and color fields. That is, the inputs for the target cards overlap with the HTs established during the pre-switch phase. Although there is still overlap of HTs and target inputs in the pre-switch field, there is no longer spatial competition between HTs and target inputs for the post-switch features. Thus, the 3-year-old model with only a weak boost is able to correctly bind the test cards to locations based on the post-switch features. This shows how memories for the spatial layout of features can eliminate the need for inhibition which critically influences the amount of boost required to successfully sort during the post-switch phase.

Negative Priming—Our simulations of the standard condition indicate that there are two key factors that play a role in 3-year-olds’ perseverative behavior—(1) overlap between HTs and target inputs in the pre-switch field and (2) conflict between HTs and target inputs in the post-switch field. That is, what the model remembers, and what the model ‘sees’. One question is whether each source plays a central role. This is addressed in the Negative Priming version (NP) where the features that were relevant during the pre-switch phase are changed in the post-switch phase. This eliminates the overlap in the pre-switch field, but retains the conflict in the post-switch field.

Figure 13 shows quantitative fits of the model with this version of the task where the colors were changed in the post-switch phase. Like the empirical data (Müller et al., 2006; Zelazo et al., 2003), the 3-year-old model shows a high rate of perseveration and the 4-year-old model shows a high rate of switching. The left panel in Figure 13 shows the model just before the start of the post-switch phase. Now, new target inputs are present in the pre-switch color field at the values for yellow and green (see y-axis), while the HTs for red and blue are positioned at the location where those features were sorted during the pre-switch phase. In the pre-switch field, then, there is no longer overlap of HTs and target inputs. There is, however, still a need to inhibit the HTs in the post-switch field since this activation conflicts with the target inputs in the shape field. Again, this conflict slows peak building in this field during the post-switch trials. Consequently, peaks first emerge in the color field and the 3-year-old model perseveratively sorts by this dimension. Thus, perseveration at 3 years can emerge solely due to conflict in the post-switch field. The 4-year-old model overcomes this conflict with a stronger resting-level boost to the shape field; consequently, this model correctly executes the post-switch rules.

No-Conflict Negative Priming—The no-conflict Negative Priming (NNP) version verifies that conflict along the post-switch dimension is the critical source of error in the NP version. This version eliminates the post-switch conflict by asking children to sort test cards that match the target cards along both color and shape dimensions during the pre-switch phase. Results from this version are shown in Figure 14. As in the NP version, there are new target inputs (green, yellow) for the pre-switch dimension at the start of the post-switch trial. Now, however, there are overlapping HTs and target inputs in the shape field because the model sorted no-conflict cards during the pre-switch phase. This eliminates conflict in the post-switch field, eliminates the need for inhibition of the HTs, and the 3-year-old model, like children, is able to switch rules (Müller et al., 2006; Zelazo et al., 2003).

Partial-Change Version—The Partial-Change version is the opposite of the Negative Priming version in that it probes whether the other source of perseveration—overlap in the pre-switch field—is sufficient to drive children's errors. In this version, the features for the post-switch dimension are changed while the features for the pre-switch dimension remain constant throughout the procedure. As can be seen in Figure 15, changing the shape features going into the first post-switch trial eliminates conflict in the shape field, that is, there is no longer conflict between the HTs and target inputs. There is, however, still overlap between HTs and target inputs in the color field. As the simulation results show, this overlap is sufficient to drive perseveration with only a weak boost provided to the post-switch feature field—activation builds more quickly in the color field and the model, like children, tends to perseverate (Zelazo et al., 2003).

Total Change Version—In the total change version, all of the features that were present during the pre-switch are changed for the post-switch phase. Thus, as can be seen in the left panel of Figure 16, this version eliminates both the overlap in the pre-switch (color) field and the competition in the post-switch (shape) field. New target inputs are present in both the shape and color fields while the HTs from the pre-switch phase can still be seen at the location where those features were sorted. Thus, the model does not need to inhibit these

HTs. Figure 16 shows that both 3-year-olds and the model do not perseverate in the total change version (Zelazo et al., 2003). Thus, even though the 3-year-old model has just a small boost in the resting level of the shape field on the post-switch trials, this is sufficient—under this condition—to correctly switch rules.

Relational Complexity Version—The relational complexity version examines the role of the relational propositions that must be represented in a hierarchical rule-structure like that used in CCC theory. This was accomplished by reducing the number of irrelevant features during the pre- and post-switch phases (see Figure 1G). Here, children received 2 conflict and 2 no-conflict test cards and were scored as passing if they sorted all four post-switch cards correctly. Three-year-olds—and the model—show better switching performance in this version relative to the standard task (Halford et al., 2007). Figure 17 shows the fit of the model with data from Halford et al. (2007), as well as the WM fields at the beginning of the post-switch phase. Now there is only conflict for one of the post-switch features, but there is also a stronger HT for this feature. This reduces the spatial competition for this feature and allows a weak representation of the rules to accurately guide rule-use.

Brace, Morton, and Munakata (2006) —Brace et al., (2006) showed that 3-year-olds' post-switch performance was improved when given a training phases between the pre- and post-switch phases. Specifically, they gave children a series of cards that started with images that contained only the post-switch features. Over a series of 6 cards the pre-switch features were gradually 'morphed' into the image until they contained fully bivalent stimuli at which point the post-switch phase began.

We administered this to the model in the same fashion. The model was given 6 pre-switch trials. At the start of the training phase, the model was given only a post-switch feature. Over the 6 training cards, the strength of the pre-switch inputs were increased to reflect the increasing salience of the pre-switch features as they were 'morphed' into the stimuli. As can be seen in Figure 18, the model benefited from this training phase similar to 3-year-olds. Specifically, 83% of the models were able to correctly switch rules compared to 88% of 3-year-olds averaged across the Training and Training Plus Instructions conditions reported by Brace et al., (2006). In the model, similar to the model reported by Morton and Munakata (2002), this benefit resulted from the formation of HTs that overlapped with the post-switch target features. This reduced the spatial competition between the target inputs and HTs established during the pre-switch phase and allowed the young model to switch rules.

Summary of Quantitative Fits

A central innovation of the DNF model is that it integrates the dynamics of object WM with a simple form of dimensional attention—boosting the resting level of different neural populations when cued by the label 'shape' or 'color'. In this sense, the dimensional attention plays a rather general role; the rest of the work is done by the real-time neural dynamics of the WM fields and the trial-to-trial process that builds HTs in the context of the task. In this sense, the DNF model formalizes how the multiple factors that contribute to executive function interact: the 'rule-like' processes involved in task-switching, the neural dynamics of working memory, and the neural dynamics of inhibition during selection.

The simulations reported here showed that the model was able to quantitatively simulate children's performance from eight different variants of the DCCS task all with the same parameters, including two variants reporting data from both 3- and 4-year-olds. Note that our ability to simulate 4-year-olds' performance was limited by the fact that the majority of studies in the literature do not report data for this age group. We note, however, that the 4-year-old model does, in fact, switch correctly in all variants of the task reported here. Note also that the four developmental parameters we changed to create the 4-year-old model were all necessary to achieve accurate switching performance. Changing only the boost or shift distributions, for instance, produced high levels of intermediate switching (i.e., sorting between 2 or 4 out of 6 cards correctly during the post-switch). Thus, the number of parameters changed over development was necessary to capture the specific pattern of rule-switching present in the literature.

In summary, three features of our theory thus far make it stand out from previous explanations of children's performance in the DCCS task: (1) the DNF model can capture children's performance in quantitative detail across a host of studies; (2) the model specifies how the multiple components of EF interact during a period of rapid developmental change in rule use; and (3) the model simulates behavior using neural dynamics, making it possible to integrate behavioral and neural dynamics. In the next section, we pursue another stand-out aspect of the DFT—its ability to generate a set of novel predictions that are not consistent with other theories in the literature.

Empirical Test of the DNF Model: The Role of Space

The model we have presented is able to achieve extensive coverage of the existing literature through the interaction of spatially-specific memories for features and a simple form of dimensional attention. Critically, all of the effects simulated above emerged from spatially-specific competition or cooperation between target inputs and HTs in different feature fields. This leads to the prediction that the DCCS task is not just about 'rules'—*space should also matter*. This is a radical departure from previous accounts: spatial information is not an aspect of any theory of EF, rule-use, or the DCCS task. And, as we highlight next, the spatial properties of the DNF model predict that we can reverse two known effects in the literature simply by moving where the target cards are located in the task space. Note that we are not claiming the DCCS task is *only* about space—an issue we will return to later. Rather, in the present section, we highlight a key way in which our theory diverges from all others.

To test whether space plays a role in children's rule-use in the DCCS task, we constructed two new experimental conditions based on the Negative Priming (NP) and the no-conflict NP (NNP) versions simulated above. In the NP version, there is conflict in the post-switch feature field and 3-year-olds perseverate, while there is cooperation in the post-switch feature field in the NNP version and 3-year-olds switch rules. Critically, if the layout of HTs and target inputs in these two versions is indeed the source of 3-year-olds' success and failure, then swapping the spatial locations of the target cards before the post-switch trials should reverse this pattern of results.

Recall that in the NP version (Müller et al., 2006; Zelazo et al., 2003), the features that were relevant for the pre-switch game are changed for the post-switch game and 3-year-olds still perseverate. The model perseverated in this version because of the conflict in the post-switch (shape) feature field. Figure 19 illustrates the effect of swapping the spatial locations of the target cards in the NP version. Panel A shows the WM fields of the model in the NP version at the start of the post-switch phase reproduced here for comparison. Note the conflict between the HTs and the target inputs in the post-switch (shape) field which leads to the 3-year-old model's failure to switch rules. Figure 19B shows the WM fields of the model going into the post-switch phase in the NP “space swap” version (NPS). Here, the locations of the target cards have been swapped at the start of the post-switch phase (see display in panel B). Consequently, in Figure 19B the locations of the target inputs now overlap with the HTs in the post-switch field which should facilitate rule-switching in 3-year-olds. That is, 3-year-olds in the NPS version should switch rules, while 3-year-olds in the NP version should perseverate—even though everything, the features and the rules, is identical in these versions except the location of the target cards during the post-switch phase.

In the NNP version (Müller et al., 2006; Zelazo et al., 2003), on the other hand, the test cards match the target cards along both dimensions during the pre-switch phase and 3-year-olds are able to switch rules. Figure 20A shows the model of this version reproduced here for comparison. The 3-year-old model is able to switch in this version because the conflict in the post-switch feature field is removed. The NNP-SpaceSwap (NNPS) version is shown in Figure 20B. This version swaps the locations of the target cards as before, but this leads to the opposite prediction. In particular, the target inputs in the post-switch field now conflict with the HTs formed during the pre-switch phase. Thus, 3-year-olds in the NNPS version should perseverate, while 3-year-olds in the NNP version should switch correctly—even though, once again, everything is identical in these versions except the location of the target cards during the post-switch phase.

Figure 21 shows quantitative simulation results of the predicted cross-over interaction using the 3-year-old model from before. The NP version shows high rates of perseveration while the NNP version shows high rates of switching. In the SpaceSwap conditions, however, this relationship is reversed. Here, the NP-SpaceSwap condition shows a high rate of switching while the NNP-SpaceSwap condition shows a high rate of perseveration. Thus, the model predicts that simply swapping the locations of the target cards will push performance in opposite directions for these two conditions. Note the because the features and rules remain the same in the space-swap conditions, all other existing theories predict that this manipulation should have no impact on children's performance. We tested the DNF predictions in the following experiment.

Method

Participants—Seventy-six 3-year-olds between 38 and 46 months of age (*M*= 41.67 mos., *SD*= 2.68 mos.) were included in the final analysis (*N*= 19 for all four groups). An additional 9 children were dropped due to experimenter error (*N*= 4), failing to pass the pre-switch phase (i.e., sort at least 4 cards correctly, see Zelazo et al., 2003; *N*= 3), or parental interference (*N*= 2).

Materials—The pre-switch target cards for all versions were a red star at the right location and a blue circle at the left location. The pre-switch test cards for the NP versions were red circles and blue stars. The pre-switch test cards for the NNP versions were red stars and blue circles. Depending on the task-order, the post-switch cards were either red and blue triangles and squares (for the shape-color task order), or green and yellow stars and circles (for the color-shape task order). In the Space-Swap versions, the locations of the target cards were swapped for the post-switch phase (see Figures 19-20). For all versions, 5 pre-switch and 5 post-switch test cards were used with no more than three of any particular card appearing in any sorting phase.

Procedure—Sessions were run individually with the child and experimenter in a quiet room. The experimenter began by showing the cards to the child and saying that they were going to play a couple of games. The experimenter stated the name of the first game, how the cards were to be sorted, and the specific rules for that game. For example, the experimenter said, “First, we are going to play the color game. In the color game, we sort the cards by color. So all of the red ones go here and all of the blue ones go there.” The experimenter then demonstrated a card for each rule saying, for example, “See, this one is red so it goes here.” The experimenter presented the test cards one at a time to the child. When presenting a card, the experimenter did not provide a label but asked, “Where does this one go in the color game?”

After the pre-switch cards had been sorted, the experimenter stated that they were all done with that game and were going to play a new game. The experimenter removed the target cards and replaced them with the post-switch target cards, stated the name of the new game, how the cards were now to be sorted, and the specific rules for the new game just as in the pre-switch. No demonstration of the post-switch rules was supplied. The post-switch test cards were presented just as in the pre-switch. Throughout both phases, no direct feedback was provided. If the child incorrectly sorted a card, the experimenter simply reminded the child of the rules by saying, for example, “Remember, we are playing the color game, so we are sorting the cards by color. All of the red ones go here and all of the blue ones go there.” Note that the space swap conditions were identical to the standard conditions except for a single change—after the target cards were removed at the end of the pre-switch phase, they were re-positioned at the opposite spatial locations.

Results

Children were categorized as passing the post-switch phase if they sorted at least 4 out of 5 cards correctly and as failing if they sorted 1 or fewer cards correctly (Zelazo et al., 2003). This resulted in dropping 3 children in the NP condition, 1 child in the NP-SpaceSwap condition, 2 children in the NNP version, and 0 children in the NNP-SpaceSwap condition from further analysis for intermediate responding. The percent of children who correctly switched rules in each condition is shown in Figure 22.

The first question we examined was whether data from the No Swap conditions replicated previous findings. This was indeed the case: significantly more children in the NP group perseverated than in the NNP group ($\chi^2(1)= 10.186$, $p= 0.001$). Next, we tested the predictions of the DNF model by comparing the NP and NP-SpaceSwap conditions. As predicted, significantly more children switched rules in the NP-SpaceSwap version than in the NP version (($\chi^2(1)= 4.48$, $p= 0.03$). We then compared the NNP and the NNP-SpaceSwap conditions. Again as predicted, significantly more children perseverated in the NNP-SpaceSwap condition than in the NNP condition ($\chi^2(1)= 3.995$, $p= 0.04$). In a final analysis, we examined whether performance in the SpaceSwap conditions differed from one another. This was not the case ($\chi^2(1)= 0.696$, *ns*).

Discussion

Results from the present study showed that a simple manipulation of the spatial structure of the DCCS task can have surprising effects on the behavior of 3-year-olds. Specifically, by swapping the location of the target cards between the pre- and post-switch phases, we reversed the pattern of results from the NP and NNP versions. Three-year-olds typically perseverate in the NP version; however, swapping the locations of the target cards significantly improved their performance. Three-year-olds typically switch rules in the NNP version; however, swapping the locations of the target cards significantly impaired their performance. The DNF model predicted this cross-over effect *a priori*; no other theory provides an account of such effects because the features and rules are identical in the standard and space-swap variants of the task. The present study, therefore, provides a compelling demonstration of the utility of our theory.

Beyond Space: The Role of Feature-Saliency and Attentional-Weights

The work we have presented thus far has highlighted the interaction of spatially-specific feature inputs from the target cards and spatially-specific HTs for features that accumulate as decisions are made. We turn now to new versions of the DCCS reported by Fisher (2011). These versions attempted to isolate the processes of voluntary and automatic shifts of attention by manipulating the saliency of dimensional features or the number of features per dimension. These versions provide a particularly useful probe of the DNF model because in this design the spatial locations of the target cards were randomized from trial to trial. Thus, these versions allow us to examine whether the DNF model can capture children's performance in tasks that are explicitly removed from spatial influences.

Feature-Saliency

In one set of conditions, Fisher (2011) manipulated the saliency of the features for each dimension. This manipulation was motivated by two factors. First, previous research has shown that salient features automatically capture attention (Smith et al., 1996; Treisman & Gelade, 1980). Thus, salient features facilitate an automatic shift of attention, while less salient features require a voluntary shift of attention. Second, data suggest that automatic attention is robust by 3 years of age, but voluntary attention is still developing (Fisher & Sloutsky, 2005; Smith, Jones, & Landau, 1996). Based on these observations, Fisher (2011) proposed that changes in voluntary attention might play a critical role in the emergence of rule-use in the DCCS task. If this is the case, then post-switch trials that require a shift to a more salient dimension should be easier for 3-year-olds because such trials engage automatic attention which is robust at age 3. By contrast, a switch to a less salient dimension should be harder because this requires voluntary attention which is less robust by age 3.

To test this, Fisher (2011) used a version of the DCCS task with similar colors (red and pink) but distinct shapes. If voluntary attention is particularly hard for 3-year-olds to recruit, they should have particular difficulty switching to color. Switching to shape, by contrast, should be automatic. Figure 23 shows the stimuli Fisher (2011) used. In addition to varying the salience of the stimuli, Fisher (2011) eliminated potential spatial biases by randomizing the location of the target cards on each trial. She also included no-conflict test cards on half of the trials. These cards matched a target along both dimensions, allowing for an assessment of whether children were actually staying on task.

Results of this study are shown in Figure 24. Children performed quite well on the pre-switch trials and on all no conflict trials. Critically, 3-year-olds had significantly more difficulty switching to color rules than to shape rules on conflict trials. That is, when they had to switch to the salient shape dimension and play the shape game, 3-year-olds tended to sort cards correctly; when they had to switch to the less salient color dimension and play the ‘color’ game, children perseverated.

Can the DFT capture this pattern of results? To probe this, we gave the model inputs and the task structure taken directly from the methods used by Fisher (2011). We used the same parameters used previously for the 3-year-old model. Figure 24 shows that the simulation results closely follow the pattern of behavior displayed by 3-year-olds. In particular, the model showed a high rate of correct sorting for all trials during the pre-switch phase. During the post-switch phase, the model again had a high rate of correct sorting for no-conflict trials with both shape and color rules. Like 3-year-olds, however, the model had significant difficulty sorting conflict cards by color rules, but not by shape rules.

What is the source of this asymmetry in the model? The asymmetry arises from competition in the color field due to the metrically close target inputs. As seen in the left panel of Figure 24, the target inputs are very close along the color dimension, but are far apart along the shape dimension. Interestingly, this competition is evident even during the pre-switch trials as reflected in the peak-build latency in the shape and color fields. Specifically, the peak-build latency for the pre-switch phase when the model sorted by color was 547 time-steps, while the latency for the pre-switch phase when the model sorted by shape was 425 time-steps. Thus, the model required slightly more time to build a peak in the color field. This results from the overlap along the feature dimension which leads to competition as the field tries to build peaks at two neighboring color values. That is, extra inhibition is needed to suppress the activation of a nearby task-relevant color. This effect is exacerbated in the post-switch phase when the model tries to sort by color. In this case, the accumulation of HTs at both spatial locations creates even more conflict. Consequently, the model perseverates and sorts by shape.

Interestingly, we can ask the reverse question: if the 3-year-old model typically perseverates, why does it succeed here when asked to sort by shape during the post-switch trials? Although the behavioral result is the opposite, the explanation is the same: because the colors are metrically close, it is hard for the model to build peaks in the color field due to competition. Consequently, the shape field wins out. Note that the randomization of space actually helps in this case. In contrast to the Standard DCCS task where there is competition in the post-switch field and overlap in the pre-switch field, here there is competition in both fields going into the post-switch phase.

In summary, the same model and parameters as in the simulations above quantitatively captures data from Fisher's first experiment. This is non-trivial for the model because spatial locations were randomized. Thus, all effects must come from the metrics of the feature dimensions themselves. In this sense, then, these simulations demonstrate that our account of the development of EF generalizes beyond spatially-grounded effects.

Attentional Weights

Another set of conditions presented in Fisher (2011) manipulated the number of features per dimension. This was inspired by work suggesting that attention is a fixed resource (Nosofsky, 1984). Thus, the amount of attention given to a particular object or feature is a function of the total number of features within the scope of attention. As more items are added, less attention is given to each item. Consequently, attention becomes less attached to any particular feature value and can be moved more easily to new features during the post-switch phase. Again, assuming automatic, but not voluntary, attention is robust by 3 years, Fisher (2011) predicted that children should show better rule-switching behavior with more items.

Fisher (2011) tested this idea by comparing 3-year-olds’ ability to switch rules with 2 versus 4 features per dimension. Sorting 2 features should place a heavy demand on voluntary attention because each item receives more of the attentional resource; thus, 3-year-olds should have difficulty switching rules. Sorting 4 features, however, should create ‘lighter’ allotments of attention. This should facilitate rule switching.

These predictions were tested in a version of the task comparable to the feature saliency condition. Children were assigned to a condition with either 2 or 4 features per dimension (see Figure 25). Eight trials were administered during the pre- and post-switch phases, half of which were no-conflict trials. Again, the spatial locations of the target cards were randomized on each trial. Under these conditions, 3-year-olds had significantly more difficulty switching with two features per dimension than with four features per dimension on conflict trials (see Figure 26).

Once again, we examined whether the DNF model could capture these results. Inputs to the model and the task structure were taken directly from the methods used by Fisher (2011). As above, we used the same parameters for the 3-year-old model. Figure 26 shows the fit of the model. Once again, the model follows the pattern of behavior displayed by 3-year-olds, although the model did slightly better than children in these conditions. In particular, the model—like children—showed correct sorting for all trials during the pre-switch phase. During the post-switch phase, the model again had a high rate of correct sorting for no-conflict trials with both 2 or 4 features per dimension. Finally, like 3-year-olds, the model had difficulty sorting conflict cards with 2 features, but not with 4 features.

Again, we can probe the source of this asymmetry. Two factors are important. The first is the fact that individual features are presented less frequently. The second is that the decisions about these features are more spatially distributed since the model is making decisions at four locations instead of two. These factors conspire to produce weaker HTs in the four-feature condition at the end of the pre-switch phase. In particular, the average strength of HTs with four features was 0.25 units, while this strength with two features was 0.48 units. Thus, using four features creates weaker memory traces for where the features were sorted in space. Consequently, there is less conflict heading into the post-switch phase and weaker HTs overall. This creates less need for inhibition and allows the resting-level modulation to have a stronger influence on the decision-making process. With only two features, HTs are more concentrated at fewer locations. This creates more conflict and a greater need for inhibition.

Summary of Quantitative Fits across All Conditions

The simulations in this section demonstrate the versatility of the DNF model, highlighting the role of both feature metrics and the distribution of experience through space and time. These simulations, in conjunction with the previous results, reveal the many factors that influence rule-switching in early childhood. Rule-use between 3 and 4 years is softly-assembled ‘in-the-moment’ based on a confluence of factors, including the target and test inputs, memories for features at specific spatial locations, the trial-to-trial dynamics of Hebbian learning, the boosting of attention to particular feature dimensions, and the metric details or number of the stimuli. All of these factors were captured by the DNF model with a single set of parameters at each age.

In total, the current report presented simulation results for 14 different conditions that reflect the model's performance across 57,000 real-time responses. Table 1 shows a summary of all simulated versions. The first column indicates the number of data points simulated for each version. For example, the *Standard 3-yo* line has 2 data points corresponding to the pre- and post-switch phases. (Recall that we verified the model showed the correct pattern of pre-switch performance in each case, even though we reported only post-switch performance in the simulation figures). The *2-Feature* version, on the other hand, has 4 data points corresponding to the conflict and no-conflict trials in both the pre- and post-switch phases. The second column lists the number of trials that were included in a single simulation of each version. Finally, the third column indicates the total number of trials across all simulations for each version.

To evaluate the model's performance across the 37 data points we quantitatively fit, we computed the root mean squared error (RMSE) between the simulated and empirical data. The overall RMSE was very good at 0.06. Thus, the model results were typically within 6% of the empirical values—an impressive result, particularly given that we used the same parameters across all variants of the DCCS task.

General Discussion

The goal of the present work was to take a first step toward a new theory of the development of EF using the DCCS task as a case study. The DNF model presented here moves beyond previous theories of the DCCS by integrating the processes of object WM and dimensional attention to construct a neurally grounded model of rule use. Critically, we were able to quantitatively capture a wide range of effects in the literature, focusing on a key transition in EF between 3 and 4 years. Further, we generated a novel set of predictions from the model that no other theory readily explains. In the sections that follow, we situate DFT in the context of current theories of the DCCS task as well as the broader literature on EF. We then discuss limitations of the theory and the important challenges that lie ahead as we pursue a theory of autonomous development and ties between DFT and the growing literature on the neural bases of EF.

Contrasting DFT with Other Theories of the DCCS

DFT presents a novel approach to the development of EF that can be distinguished in various ways from current theories of the DCCS. The account we offered proposes that rule-use emerges as dimensional attention modulates the binding of cortical feature fields coupled along a shared spatial dimension. This confers several distinct properties to DFT. For instance, because the DNF model is based on neural population dynamics within cortical fields tuned to continuous, metric feature dimensions, the model was able to explain effects that arise from manipulations to the attentional saliency of the features across dimensions (Fisher, 2011; Honomichl & Chen, 2011). CCC theory and the connectionist model of Morton and Munakata (2002) are unable to account for these effects due to their abstract representational format of features and rules.

The DFT is also unique in the special role space plays in the binding of feature representations, the spatial nature of the target inputs, and the spatially-grounded learning that happens across trials. This feature of the model was central to the novel predictions we successfully tested in the present report. Moreover, our simulation results suggest that spatial conflict or cooperation within feature fields plays a central role in many different variants of the DCCS task.

We also proposed a specific developmental account that quantitatively captured the shift in children's performance in the DCCS task between 3 and 4 years. This account moves beyond the descriptive developmental accounts offered by CCC theory, the attentional inertia account, and the re-description account. Interestingly, simulations of the model showed that the *qualitative* shift in children's rule-use in the DCCS task can be captured by a *quantitative* shift in model parameters between 3 and 4 years. The qualitative shift in the operation of the dimensional attention system was evident in Figures 6 and 7: neural interactions were too weak for the dimensional nodes to rise above threshold in the 3-year-old model; by contrast, quantitatively stronger neural interactions at 4 years robustly engaged the dimensional nodes, enabling the model to shift from the color rule to the shape rule.

Our developmental account is similar to the account proposed by Morton and Munakata (2002). These researchers increased the strength of recurrent interactions in a connectionist model to boost the active representation of rules in a bank of PFC nodes. Stronger neural interactions in our DNF model achieve a similar effect. Nevertheless, our account differs in that we proposed a refinement of the pattern of connectivity between the frontal nodes and posterior cortical fields over development.

The DNF model also differs from the connectionist model in that it makes an active decision to sort each card to the left or right on each trial. By contrast, the connectionist model simulates a probability distribution of responses on a given trial. For instance, if a ‘left’ output node was active at a level of .6, Morton and Munakata (2002) interpreted this as a 60% chance that the model would respond ‘left’ on that trial. This is a conceptually different approach to capturing real-time human behavior. For instance, individual children do not respond correctly on 60% of the pre- or post-switch trials. Rather, they generally respond either correctly or incorrectly across all trials within each phase. The 60% value arises when children are grouped together within an experimental condition. In this sense, the DNF model provides a better conceptual match to the empirical data in that the behavior of each simulation better reflects the behavior of each individual child.

Most critically, the DNF and connectionist models differ in terms of what they learn from trial-to-trial. In particular, the connectionist model does not learn about the irrelevant feature on the test card as it sorts. Rather, latent traces are strengthened only along the dimension used for sorting on each trial. By contrast, our quantitative simulations show that learning along the irrelevant dimension played a critical role in several variants of the DCCS task. For example, in the No-Conflict version, the overlap of HTs and target inputs during the pre-switch phase for the irrelevant dimension provided a sufficient boost when these features became relevant during the post-switch phase. This allowed the model to correctly switch rules (for related results, see the No-Conflict Negative Priming version). It is unclear how the connectionist model would explain children's successful performance in this condition. Further, in the Negative Priming version, the conflict between HTs and target inputs along the irrelevant feature dimension was critical for simulating children's failure in this version of the task. Here, it is unclear how the connectionist model would capture these findings because latent traces during the pre-switch phase would only be strengthened for the specific pre-switch feature values. Once these values are changed during the post-switch phase, the model would not have a bias to continue sorting by the pre-switch dimension.

Finally, it is worth noting that although our model differs from previous accounts in many respects, the DFT was still able to account for key findings central to other theories of the DCCS. For example, the key prediction of the connectionist model was that scaffolding the post-switch dimension with a training phase would improve post-switch performance (Brace et al., 2006). The DNF model was able to quantitatively simulate this finding.

DFT and Executive Function

Although we focused on the DCCS task in the present report, the work presented here provides a critical step toward a general theory of the development of EF. In particular, unlike previous conceptual theories of the DCCS, the DFT can address the multi-component nature of EF. Critically, however, these components are not isolated parts of the model. Rather, working memory, inhibition, and task-switching emerge from system-wide interactions in the neural architecture. For instance, building a WM peak in a feature field can actively inhibit the prepotent response associated with a previously relevant feature value. Moreover, the formation of a WM peak can be biased by the boost provided by a dimensional attention node. Thus, all three processes conspire together in the DCCS task to enable flexible behavior.

Recent work highlights how these processes in DFT offer a robust framework for thinking about executive function more generally. Sandamirskaya and Schöner (2010) used DFT to explore how an autonomous robotic system might organize its behaviors in a real-world context in real time. Autonomous robotics provides an exciting platform for probing the details of EF because behavioral organization in the real-world provides novel challenges that are typically not present in laboratory tasks. For instance, in the DCCS task, the rule-switch is specified by the experimenter. But oftentimes, rule-switches must be internally and autonomously generated. This requires that one determine whether the current goal has been achieved by, for instance, perceiving that the correct state of affairs has been achieved in the world. Next, the current goal must be de-activated, and the new goal state must be brought into an active neural state.

Sandamirskaya and Schöner (2010) showed that an autonomous robot using DFT could actively organize its behavior in time. In one demonstration, for instance, a robot was taught a sequence of colors. The robot was then allowed to freely navigate in a playpen with an array of colored blocks. The task was to find the blocks in the correct sequence. This is a challenging robotics task because the robot must keep the current goal state actively in mind (e.g., find the green object) as it navigates around the playpen. Once it detects green, it must then navigate successfully to the object, avoiding obstacles along the way. When it arrives at the green block, the color green dominates the camera image. This was the ‘condition of satisfaction’ used by the robot to de-stabilize the ‘green’ goal and bring up the next goal in the sequence. The robotic architecture used by Sandamirskaya and Schöner effectively learned a color sequence and autonomously organized its behavior in the playpen to find the blocks in the correct sequence. Critically, the details of the robot's behavior were not pre-programmed; rather, the robot behaved in real time based on its own, autonomously generated neural dynamics. Note, further, that details of the robot's neural architecture had several similarities with the DNF model proposed here. Most critically, the robotic model modulated its own behavior through time by boosting cortical fields tuned to specific features—the same neural mechanism implemented here.

Neural Basis of EF and Rule-Use

A critical strength of DFT is the ability to explicitly tie the dynamics of the model to neural function. DFT operates based on general properties of neural activation distributed across cortical layers (Douglas & Martin, 1998), and the local-excitatory / laterally-inhibitory interactions that arise from such interactions. These interactions give rise to stable neural patterns (Durstewitz, Seamans, & Sejnowski, 2000) that form the basis of thinking and decision-making in DFT. Critically, these neural dynamics have been shown to be robustly linked to single- and multi-unit recordings in both visual and motor cortex (Bastian et al., 1998; Bastian et al., 2003; Erlhagen et al., 1999; Jancke et al., 1999).

An important next step for DFT is to link to neural measures observed in human cognitive tasks, such as the BOLD signal recorded with fMRI. Recent biophysical work suggests that the mapping of DFT to neural activity is possible (Deco & Rolls, 2004; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001). This can achieve multiple goals. First, we can ask whether the dynamics of different fields in the model can be explicitly localized to different brain areas. Localization in the present study was based on the extant fMRI literature, but it will be important for future work to probe these localization assumptions directly. Second, developing the mapping from field activity to BOLD signals in fMRI can constrain modeling efforts to simultaneously capture both behavior and neural dynamics (Ashby & Waldschmidt, 2008). Third, if the mapping between the DNF model and fMRI is robust, we can test the DNF model by generating not only behavioral predictions, but also *neural predictions* based on how neural activation evolves over time as the model sorts cards.

This could be particularly useful in the case of EF given the complexity of the frontal systems involved in cognitive flexibility. For instance, model-based approaches to fMRI might help unpack the frontal system in our current model, isolating different functional components that are involved in EF. In this way, we can use fMRI to go beyond simply asking where a cognitive function is localized to address how neural activity is functionally related to cognition (Ashby & Waldschmidt, 2008). This is a critical step in probing the neural mechanisms underlying EF.

Limitations of Model and Future Directions

In the present report, we showed how an autonomous dimensional attention system coupled to an object WM model could capture children's performance in the standard DCCS task. Nevertheless, when we simulated children's performance in quantitative detail, we simplified the dimensional attention system by using distributions of *boost* and *shift* parameters for each age group. This allowed us to probe whether specific developmental changes to the dimensional attention system were sufficient to capture children's performance. A critical next step in theory development, therefore, is to probe whether we can capture developmental changes in EF within a fully autonomous neural system.

In principle, each developmental change implemented by hand in our model could emerge from a Hebbian learning process. The first developmental change we implemented increased the strength of excitatory and inhibitory neural interactions within the dimensional attention system. It is likely that this modification can emerge from the same Hebbian process already incorporated in the DNF model if the model's behavior were simulated across many, many trials. In particular, as the model repeatedly activates the ‘color’ or ‘shape’ node, this will increase the strength of the Hebbian trace for a particular node. Stronger excitation can, in turn, lead to an effective strengthening of inhibitory interactions as the nodes competitively interact.

The second developmental change we implemented by hand refined the pattern of connectivity between the frontal and posterior systems. Again, in principle, this can be tackled with a Hebbian process that learns the patterns of covariance between, for instance, activation of a ‘color’ node and activation patterns in the color WM field (see, e.g., Sandamirskaya & Schöner, 2010).

But this points toward another critical piece of the developmental puzzle. Conceptually in the model, the dimensional nodes represent a mapping between a neural population representation in frontal cortex and cortical fields in posterior brain regions tuned to different feature dimensions. Left out of our developmental account, at present, is that fact that children are, in reality, learning something specific about dimensional attention as they learn dimensional labels in early word learning (e.g., Sandhofer & Smith, 1999; 2001). This highlights the need to link our theory to the literature on early word learning which will require consideration of the learning opportunities that children have and the developmental time course of dimensional label learning. Although challenging, linking to this literature will undoubtedly provide robust constraints for theory development as we attempt to integrate findings from multiple paradigms across different domains of cognitive development.

Note that related work with DFT has already moved in this direction. Specifically, Faubel and Schöner (2008) implemented fast word learning using dynamic neural fields in an autonomous, interactive robot. In their work, the perceptual features of objects were bound together along a label dimension similar to how features are bound through space in the model presented here. Using relatively few training trials, the robot was able to form categories of objects using the graded associations between features and labels which facilitated accurate recognition of an assortment of real-world items. This work demonstrates that an associative mechanism within dynamic neural fields can serve as a basis for functionally linking labels and feature fields. And, critically, labels in this context need not be associated with specific features; rather, labels such as ‘shape’ and ‘color’ can be associated with entire feature dimensions (for related work in early word learning with children, see Samuelson, Schutte, & Horst, 2009; Samuelson et al., in press). In summary, then, although our theory, at present, lacks an autonomous account of development, we contend that this key goal is within reach.

Finally, we note that there are variations of the DCCS task that are beyond what the current model architecture is capable of simulating in any principled fashion. For example, Yerys and Muankata (2006) showed that different manipulations to the labels for the games or the features used in a task can make switching easier for 3-year-olds. This effect may be within reach as we pursue links between the emergence of EF in development and how children use words for the different ‘games’ they play in the DCCS task. Similarly, Kloo and Perner (2005) showed that separating the dimensions on the *test* cards improved rule switching with 3-year-olds. In the current model architecture, the test card is given as ridges of input to the feature fields, effectively asking the model to localize these features in the task space given by the target cards. In a more complete model, however, these ridge inputs must be supplied by an early visual process that attends to the test card and identifies the features present at the fixated location. It is possible that the results of the Kloo and Perner study could be captured by sequentially or selectively attending to the features on the test card. Indeed, this version of the task might reveal how dimensional attention not only biases feature binding in working memory, but also the initial selection of where to attend in the task space.

Conclusion

EF is a pervasive aspect of cognition and any theory of EF is likely to have widespread ramifications. The theory we presented here shows promise in its ability to integrate the multiple components shown to underlie EF—response inhibition, working memory, and task switching. DFT was able to quantitatively capture empirical results from a canonical task used to examine the early development of EF—the DCCS task. Moreover, this theory generated a set of novel predictions that we successfully tested here. In this context, DFT offers the most comprehensive account of children's performance in the DCCS task to date and provides a robust starting point for a theory of the development of executive function.

Acknowledgements

The authors thank the parents and children who participated in this research. We also thank Gregor Schöner, Sammy Perone, Jeff Johnson, and Vanessa Simmering for critical discussions as we developed the dynamic neural field model, as well as undergraduate members of the SPAM Lab for their assistance with data collection. This work was supported by National Institutes of Health R01-MH062480 and National Science Foudation HSD0527698 awarded to JPS.

Appendix

Model Equations

The basic formulation for a 1-dimensional neural field tuned to, for instance, spatial information is given in Equation 1. The rate of change of activation in a cortical field, *w*, evolves over time, *t*, at each location in the field, *x*.

$$\begin{aligned}\tau_w \, \dot{w} \, (x,t) = - \, w \, (x,t) + h_w = S \, (x,t) + \int c_{ww} \left(x - x'\right) \Lambda \left(w \left(x', t\right)\right) dx' \\ + \int c_{wv} \left(x - x'\right) \Lambda \left(v \left(x', t\right)\right) dx'\end{aligned} \quad 1$$

The first part of Equation 1, underlined with a dotted line, captures the neural resting level ($h_w<0$), stimulus input at particular locations presented at specific times in an experiment ($S(x,t)$), and a stabilization term ($-w(x,t)$). The resting level determines how far a field is from the activation threshold. The stabilizing term serves to maintain activation around an attractor state. That is, as the system is pushed away from its stable state, the rate of change goes in the opposite direction of the perturbation. Note that the tau parameter, τ_w, captures the timescale at which activation approaches the attractor state.

The remainder of the equation specifies the excitatory (section underlined with a dashed line) and inhibitory neural interactions (portion underlined with a solid line). Neural interactions within a field are determined by the convolution of a sigmoidal threshold function and a Gaussian projection. The term $\Lambda(w(x,t))$ is the sigmoided value of activation at each location in the WM field (w) used for the self-excitatory projection, while $\Lambda(v(x,t))$ is the sigmoided value of activation at each point in the Inhib field (v) used for the inhibitory projection into the WM field. The sigmoid function is given by Equation 2.

$$\Lambda \, (w) = \frac{1}{1 + exp \, [-\beta w]} \quad 2$$

The β parameter defines the slope of the sigmoid function as field activation is converted into neural output. Here, the activation threshold is represented by the point where the sigmoided output reaches 0.5, which is 0. With a large β value, the slope is steeper and there is a more abrupt transition in the sigmoided output. Thus, weak levels of activation in the field contribute relatively little to activation and peak formation, while strong levels of activation in the fields engage in robust interactions with associated neurons.

The spread and strength of neural interactions is determined by a Gaussian interaction kernel which is generically defined in Equation 3.

$$c\left(x-x^{'}\right)=c\,exp\left[\frac{\left(x-x^{'}\right)^{2}}{2\sigma^{2}}\right] \quad 3$$

The parameter denoted by *c* scales the strength of the projection while the width of the interaction kernel is given by σ. Finally, these terms are integrated to combine the contributions at each location, *x*, from all other locations, *x’* (see Equation 1).

The equation for the inhibitory layer takes the same general form as the equation for the WM field and is given by:

$$\tau_v\ \dot{v}\ (x,t)=-\ v\ (x,t)\ +h_v+\int c_{vw}\left(x-x^{'}\right)\Lambda\left(w\left(x^{'},t\right)\right)dx^{'} \quad 4$$

The timescale of activation in this field is given by τ_v ($\tau_v = 40$) while its resting level is denoted h_v. In this equation, input to the inhibitory layer is the integration of above-threshold activation within the WM field, $\Lambda(w(x,t))$, with the spread and strength of this projection given by the Gaussian interaction kernel, $c_{vw}(x–x')$.

A further contribution to the dynamics in the WM field comes from a Hebbian layer (HL) which is now added in Equation 5.

$$\begin{aligned} t_w\ \dot{w}\ (x,t)=&-\ w\ (x,t)\ +h_w+S\ (x,t)\ +\int c_{ww}\left(x-x^{'}\right)\Lambda\left(w\left(x,t\right)\right)dx^{'} \\ &+\int c_{wv}\left(x-x^{'}\right)\Lambda\left(v\left(x,t\right)\right)dx^{'}+\int c_{HL}\left(x-x^{'}\right)w_{HL}\left(x,t\right)dx^{'} \end{aligned} \quad 5$$

The Gaussian interaction kernel, $c_{HL}(x–x')$, determines the strength and width of the projection from the HL into the WM field. The dynamics of the HL (w_{HL}; Equation 6.0) are divided into two components (Equations 6.1 and 6.2) that capture the build and decay dynamics of HL separately:

$$\dot{w}_{HL}\left(x,t\right)=\dot{w}_{HLbuild}\left(x,t\right)+\dot{w}_{HLdecay}\left(x,t\right) \quad 6.0$$

$$\tau_{build}\dot{w}_{HLbuild}\left(x,t\right)=\left[-w_{HL}\left(x,t\right)+\Lambda\left(w\left(x,t\right)\right)\right]\cdot\theta\left(w\left(x,t\right)\right) \quad 6.1$$

$$\tau_{decay}\dot{w}_{HLdecay}\left(x,t\right)=-\ w_{HL}\left(x,t\right)\cdot\left[1-\theta\left(w\left(x,t\right)\right)\right] \quad 6.2$$

The shunting term, θ, gates activation into the HL from the WM field ($\theta = 1$ when $w(x,t)>0$, and $\theta = 0$ otherwise). With $\theta = 1$, Equation 6.1 is engaged and drives the accumulation of activation in the HL at sites associated with above-threshold activation in the WM field. By contrast, when $\theta = 0$, equation 6.2 is engaged and activation levels in the HL decay. Separating the build and decay mechanisms approximates accumulation and depression of synaptic change (Deco & Rolls, 2005). The build timescale (e.g., $\tau_{build} = 200$) is shorter than the decay timescale (e.g., $\tau_{decay} = 1{,}000$) which makes activation in the HL build more quickly relative to the rate of decay; however, both of these are significantly slower than the timescale for the WM field (i.e., $\tau = 80$). Thus, as inputs are presented to the WM field and peaks of activation are built, activation accumulates slowly in the HL. This accumulated activation acts as an input to the WM field which can have various influences on the stability properties of this field. This source of activation can make particular modes of behavior more stable than others, build up biases, or make peaks of activation build more quickly.

The Object WM Model

In this section we describe the equations governing activation in the full object representation model used to quantitatively fit 3- and 4-year-olds behavior in the DCCS. The parameters for the excitatory and inhibitory field dynamics is given in Table 2. The parameters for the spatial interactions among the WM fields is given in Table 3. The spatial WM field is given by Equation 7.

$$\begin{aligned}\tau_{w_s}\dot{w}_s(x,t) = {} & -w(x,t) + h_s + \int c_{ww_s}\left(x-x'\right)\Lambda\left(w_s\left(x',t\right)\right)dx' + \int c_{w_x v_s}\left(x,x'\right)\Lambda\left(v_s\left(x',t\right)\right)dx' \\ & + \int c_{w_s w_{f1}}\left(x-x'\right)\Lambda\left(w_{f1}\left(x',t\right)\right)dx' + \int c_{w_s w_{f2}}\left(x-x'\right)\Lambda\left(w_{f2}\left(x',t\right)\right)dx' \\ & + \int c_{w_s w_{HL}}\left(x-x'\right)w_{HL}\left(x',t\right)dx' + \int c_{w_s noise}\left(x-x'\right)\Gamma(x,t)\,dx'\end{aligned} \quad 7$$

$\dot{w}_s(x,t)$ specifies the rate of change of the field at every location, x, along the spatial dimension as a function of time. The constant τ_{w_s} determines the time scale of the dynamics (Erlhagen & Schöner, 2002), while $-w(x,t)$ is the stabilizing term that serves to drive activation toward h_{w_s}, the resting-level of the field. Next, the field is influenced by the self-excitatory and laterally-inhibitory projections, $\int c_{ww_s}\left(x-x'\right)\Lambda\left(w_s\left(x',t\right)\right)dx'$ and $\int c_{w_s v_s}\left(x-x'\right)\Lambda\left(v_s\left(x',t\right)\right)dx'$, respectively. Projections between fields are defined as the convolution of a sigmoid function (e.g., $\Lambda(w_s(x',t))$) with the Gaussian interaction kernel (e.g., $c_{ww_s}(x-x')$).

The next terms in Equation 7 are the projections of activation from w_{f1} and w_{f2} (the color and shape WM fields, respectively) into w_s. These are followed by the influence from accumulated HTs, $\int c_{w_s w_{HL}}\left(x-x'\right)w_{HL}\left(x',t\right)dx'$, and spatially-correlated noise added to the field.

The inhibitory layer for the spatial system is given by:

$$\tau_v \dot{v}_s\left(x,t\right)=-v_s\left(x,t\right)+h_{v_s}+\int c_{v_s w_s}\left(x-x^{'}\right)\Lambda\left(w\left(x^{'},t\right)\right)dx^{'}+\int c_{v_s noise}\left(x-x^{'}\right)\Gamma\left(x,t\right)dx^{'} \qquad 8$$

As with the equation for the WM field, $\dot{v}_s(x,t)$ is the rate of change of neural activity in the inhibitory field at every location, x, along the spatial dimension as a function of time. The term $-v_s(x,t)$ is a stabilizing influence that drives activation toward the resting level of the field, h_{v_s}. Input to this field comes from the convolution of the sigmoid of activation in the WM field, $\Lambda(w(x',t))$, with the Gaussian interaction kernel, $c_{v_s w_s}(x-x')$. This field also has a spatially-correlated stochastic influence. The dynamics of the HL for the spatial WM field are given by equations 6.0-6.2 above.

The feature WM field (in this example, it is the color WM field) is given by:

$$\begin{gathered}\tau\dot{w}_{f1}\left(x,y,t\right)--w_{f1}\left(x,y,t\right)+h_{w_{f1}}+S_{w_{f1}}\left(x,y,t\right)\\+\iint c_{ww_{f1}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_{f1}\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\\+\iint c_{wv_{f1}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(v_{f1}\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\\+\int c_{w_{f1}w_s}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_s\left(x^{'},t\right)\right)dx^{'}\\+\iint c_{w_{f1}w_{f2}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_{f2}\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\\+\iint c_{w_{f1}w_{HL_{f1}}}\left(x-x^{'}.y-y^{'}\right)w_{HL_{f1}}\left(x^{'},y^{'},t\right)dx^{'},dy^{'}\\+c_{wi_{ColorColor}}\Lambda\left(i_{Color}\left(t\right)\right)+c_{wi_{ColorShape}}\Lambda\left(i_{Shape}\left(t\right)\right)\\+\iint c_{w_{f1}noise}\left(x-x^{'},y-y^{'}\right)\Gamma\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\end{gathered} \qquad 9$$

The color WM fields have five excitatory inputs: a self-excitatory projection $\iint c_{ww_{f1}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_{f1}\left(x^{'},y^{'},t\right)dx^{'}dy^{'}\right)$, a projection from the spatial field $\int c_{w_{f1}w_s}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_s\left(x^{'},t\right)\right)dx^{'}\right)$, a projection from the second feature field $\iint c_{w_{f1}w_{f2}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_{f2}\left(x^{'},y^{'},t\right)dx^{'}dy^{'}\right)$, a projection from its associated HL $\iint c_{w_{f1}w_{HL_{f1}}}\left(x-x^{'},y-y^{'}\right)w_{HL_{f1}}\left(x^{'},y^{'},t\right)dx^{'}dy^{'}\right)$ and inputs from the 'shape' and 'color' nodes $(+c_{wi_{ColorColor}}\Lambda(i_{Color}(t))+c_{wi_{ColorShape}}\Lambda(i_{Shape}(t)))$. The projection from its associated inhibitory field is given by $\left(\iint c_{wv_{f1}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(v_{f1}\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\right)$. Finally, this field is influenced by 2-dimensional spatially-correlated noise.

The associated 2-dimensional inhibitory field is given by Equation 10.

$$\begin{gathered}\tau v_{f1}\dot{v}_{f1}\left(x,y,t\right)=-v_{f1}\left(x,y,t\right)+h_{v_{f1}}+\iint c_{ww_{f1}}\left(x-x^{'},y-y^{'}\right)\Lambda\left(w_{f1}\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\\+\iint c_{v_{f1}noise}\left(x-x^{'},y-y^{'}\right)\Gamma\left(x^{'},y^{'},t\right)\right)dx^{'}dy^{'}\end{gathered} \qquad 10$$

The input to this field is the convolution of the sigmoid of activation in w_{f1} with a Gaussian interaction kernel $\left(\iint c_{ww_{f1}}\left(x-x',y-y'\right)\Lambda\left(w_{f1}\left(x',y',t\right)\right)dx'dy'\right)$. The last term in this equation is the influence from 2-dimensional spatially-correlated noise.

The dynamics of HL in the 2D fields are governed by Equations 11.0-11.2. The parameters for these equations are given in Table 2.

$$\dot{w}_{HL_{f1}}(x,y,t)=\dot{w}_{HLbuild_{f1}}(x,y,t)+\dot{w}_{HLdecay_{f1}}(x,y,t) \qquad 10.0$$

$$\tau_{build}\dot{w}_{HLbuild_{f1}}(x,y,t)=\left\lfloor w_{HL_{f1}}(x,y,t)+\Lambda\left(w_{f1}(x,y,t)\right)\right\rfloor\cdot\theta\left(w_{f1}(x,y,t)\right) \qquad 11.1$$

$$\tau_{decay}\dot{w}_{HLdecay_{f1}}(x,y,t)=w_{HL_{f1}}(x,y,t)\cdot\left\lfloor 1-\theta\left(w_{f1}(x,y,t)\right)\right\rfloor \qquad 11.2$$

Note that the same equations govern activation in the shape WM, inhibitory, and HL fields (w_{f2}).

Finally, the dynamics of the ‘shape’ and ‘color’ nodes are given by Equations 12 and 13. The parameters relevant for these nodes are given in Table 4.

$$\begin{aligned}\tau\dot{i}_{shape}(t)=&-i_{shape}(t)+h_{i_{shape}}+S_{i_{shape}}(t)+c_{ii_{excite}}\Lambda\left(i_{shape}(t)\right)-c_{ii_{inhib}}\Lambda\left(i_{color}(t)\right)\\&+c_{iw_{ShapeColor}}\sum\sum\Lambda\left(w_{f1}(t)\right)+c_{iw_{ShapeShape}}\sum\sum\Lambda\left(w_{f1}(t)\right)\end{aligned} \qquad 12$$

$$\begin{aligned}\tau\dot{i}_{color}(t)=&-i_{color}(t)+h_{i_{color}}+S_{i_{color}}(t)+c_{ii_{excite}}\Lambda\left(i_{color}(t)\right)-c_{ii_{inhib}}\Lambda\left(i_{shape}(t)\right)\\&+c_{iw_{ColorShape}}\sum\sum\Lambda\left(w_{f2}(t)\right)+c_{iw_{ColorColor}}\sum\sum\Lambda\left(w_{f1}(t)\right)\end{aligned} \qquad 13$$

References

Aggelopoulos NC, Rolls ET. Scene perception: inferior temporal cortex neurons encode the positions of different objects in the scene. European Journal of Neuroscience. 2005; 22:2903–2916. [PubMed: 16324125]

Amari S. Dynamics of pattern formation in lateral-inhibition type neural fields. Biological Cybernetics. 1977; 27:77–87. [PubMed: 911931]

Amari S. Topographic organization of nerve fields. Bulletin of Mathematical Biology. 1980; 42:339–364. [PubMed: 6246997]

Amari, S.; Arbib, MA. Competition and cooperation in neural nets.. In: Metzler, J., editor. Systems Neuroscience. Academic Press; New York: 1977.

Andersen RA. Encoding of intention and spatial location in the posterior parietal cortex. Cerebral Cortex. 1995; 5:457–469. [PubMed: 8547792]

Ashby FG, Waldschmidt JG. Fitting computational models to fMRI data. Behavior Research Methods. 2008; 40(3):713–721. [PubMed: 18697666]

Baddeley, AD. Working Memory. Oxford University Press; 1986.

Baddeley AD, Della Sala S, Papagno C, Spinnler H. Dual-task performance in dysexecutive and nondysexecutive patients with a frontal lesion. Neuropsychology. 1997; 11:187–194. [PubMed: 9110326]

Baird AA, Kagan J, Gaudette T, Walz KA, Hershlag N, Boas DA. Frontal lobe activation during object permanence: Data from near-infrared spectroscopy. NeuroImage. 2002; 16(4):1120–1126. [PubMed: 12202098]

Barnea-Goraly N, Menon V, Eckert M, Tamm L, Bammer R, Karchemskiy A, Dant CC, Reiss AL. White matter development during childhood and adolescence: a cross-sectional diffusion tensor imaging study. Cerebral Cortex. 2005; 15(12):1848–1854. [PubMed: 15758200]

Bastian A, Riehle A, Erlhagen W, Schöner G. Prior information preshapes the population representation of movement direction in motor cortex. Neuroreport. 1998; 9:315–319. [PubMed: 9507975]

Bastian A, Schöner G, Riehle A. Preshaping and continuous evolution of motor cortical representations during movement preparation. European Journal of Neuroscience. 2003; 18:2047–2058. [PubMed: 14622238]

Bohlman NL, Fenson L. The effects of feedback on children's perseverative errors. Journal of Cognition and Development. 2005; 6(1):119–131.

Brace JJ, Morton JB, Munakata Y. When actions speak louder than words: improving children's flexibility in a card-sorting task. Psychological Science. 2006; 17(8):665–669. [PubMed: 16913947]

Brooks PJ, Hanauer JB, Padowska B, Rosman H. The role of selective attention in preschoolers' rule use in a novel dimensional card sort. Cognitive Development. 2003; 18:195–215.

Bull R, Scerif G. Executive function as a predictor of children's mathematical ability: inhibition, switching, and working memory. Developmental Neuropsychology. 2001; 19(3):273–293. [PubMed: 11758669]

Bunge SA. How we use rules to select actions: a review of evidence from cognitive neuroscience. Cognitive, Affective, & Behavioral Neuroscience. 2004; 4(4):564–579.

Bunge SA, Wallis JD, Parker A, Brass M, Crone EA, Hoshi E, Sakai K. Neural circuitry underlying rule use in humans and nonhuman primates. The Journal of Neuroscience. 2005; 9:10347–10350. [PubMed: 16280570]

Bunge SA, Zelazo PD. A brain-based account of the development of rule use in childhood. Current Directions in Psychological Science. 2006; 15(3):118–121.

Carlson SM. Developmentally sensitive measures of executive function in preschool children. Developmental Neuropsychology. 2005; 28(2):595–616. [PubMed: 16144429]

Collette F, Van der Linden M, Laureys S, Delfiore G, Degueldre C, Luxen A, Salmon E. Exploring the unity and diversity of the neural substrates of executive functioning. Human Brain Mapping. 2005; 25(4):409–423. [PubMed: 15852470]

Compte A, Brunel N, Goldman-Rakic PS, Wang X-J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cerebral Cortex. 2000; 10:910–923. [PubMed: 10982751]

Cowan N, Elliott EM, Saults JS, Morey CC, Mattox S, Hismjatullina A, Conway ARA. On the capacity of attention: its estimation and its role in working memory and cognitive aptitudes. Cognitive Psychology. 2005; 51(1):42–100. [PubMed: 16039935]

Cragg L, Nation K. Go or no-go? Developmental improvements in the efficiency of response inhibition in mid-childhood. Developmental Science. 2008; 11(6):819–827. [PubMed: 19046150]

Crone EA, Donohue SE, Honomichl R, Wendelken C, Bunge SA. Brain regions mediating flexible rule use during development. The Journal of Neuroscience. 2006; 26(43):11239–11247. [PubMed: 17065463]

Crone EA, Wendelken C, Donohue SE, Bunge SA. Neural evidence for dissociable components of task-switching. Cerebral Cortex. 2006; 16:475–486. [PubMed: 16000652]

Crone EA, Wendelken C, Donohue SE, van Leijenhorst L, Bunge SA. Neurocognitive development of the ability to manipulate information in working memory. Proceedings of the National Academy of Sciences. 2006; 103(24):9315–9320.

Cui X, Bray S, Bryant DM, Glover GH, Reiss AL. A quantitative comparison of NIRS and fMRI across multiple cognitive tasks. NeuroImage. 2011; 54(4):2808–2821. [PubMed: 21047559]

Deák GO, Ray SD, Pick AD. Effects of age, reminders, and task-difficulty on young children's rule-switching flexibility. Cognitive Development. 2004; 19:385–400.

Davidson NC, Amso D, Anderson LC, Diamond A. Development of cognitive control and executive functions from 4 to 13 years: evidence from manipulations of memory, inhibition, and task switching. Neuropsychologia. 2006; 44(11):2037–2078. [PubMed: 16580701]

Deco G, Rolls ET. A neurodynamical cortical model of visual attention and invariant object recognition. Vision Research. 2004; 44(6):621–642. [PubMed: 14693189]

Dempster FN. The rise and fall of the inhibitory mechanism: toward a unified theory of cognitive development and aging. Developmental Review. 1992; 12(1):45–75.

Desimone R, Gross CG. Visual areas in the temporal cortex of the macaque. Brain Research. 1979; 178(2-3):363–380. [PubMed: 116712]

Diamond, A. Normal development of prefrontal cortex from birth to young adulthood: cognitive functions, anatomy, and biochemistry.. In: Stuss, DT.; Knight, RT., editors. Principles of Frontal Lobe Function. Oxford University Press; Oxford: 2005. p. 406-503.

Diamond A, Carlson SM, Beck DM. Preschool children's performance in task switching on the dimensional change card sort task: separating the dimensions aids the ability to switch. Developmental Neuropsychology. 2005; 28(2):689–729. [PubMed: 16144433]

Diamond A, Kirkham NZ. Not quite as grown up as we like to think: parallels between cognition in childhood adulthood. Psychological Science. 2005; 16(4):291–297. [PubMed: 15828976]

DiCarlo JJ, Maunsell JH. Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. Journal of Neurophysiology. 2003; 89:3264–3278. [PubMed: 12783959]

Dosenbach NUF, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RAT, Fox MD, Snyder AZ, Vincent JL, Raichle ME, Schlagger BL, Petersen SE. Distinct brain networks for adaptive and stable task control in humans. Proceedings of the National Academy of Sciences. 2007; 104(26): 11073–11078.

Douglas, R.; Martin, K. Neocortex.. In: Shepherd, GM., editor. The Synaptic Organization of the Brain. Oxford University Press; New York: 1998.

Dowsett SM, Livesey DJ. The development of inhibitory control in preschool children: effects of “executive skills” training. Developmental Psychobiology. 2000; 36(2):161–174. [PubMed: 10689286]

Drucker DM, Aguirre GK. Different spatial scales of shape similarity representation in lateral and ventral LOC. Cerebral Cortex. 2009; 19:2269–2280. [PubMed: 19176637]

Duncan J, Emslie H, Williams P, Johnson R, Freer C. Intelligence and the frontal lobe: the organization of goal-directed behavior. Cognitive Psychology. 1996; 30:257–303. [PubMed: 8660786]

Duncan J, Johnson R, Swales M, Freer C. Frontal lobe deficits after head injury: unity and diversity of function. Cognitive Neuropsychology. 1997; 14(5):713–741.

Durstewitz D, Seamans JK, Sejnowski TJ. Neurocomputational models of working memory. Nature Neuroscience Supplement. 2000; 3:1184–1191.

Durston S, Thomas KM, Yang Y, Ulug AM, Zimmerman RD, Casey BJ. A neural basis for the development of inhibitory control. Developmental Science. 2002; 5(4):F9–F16.

Edin F, Macoveanu J, Olesen P, Tegnér J, Klingberg T. Stronger synaptic connectivity as a mechanism behind development of working-memory related brain activity during childhood. Journal of Cognitive Neuroscience. 2007; 19(5):750–760. [PubMed: 17488202]

Edin F, Klingberg T, Johansson P, McNab F, Tegnér J, Compte A. Mechanism for top-down control of working memory capacity. PNAS. 2009; 106:6802–6807. [PubMed: 19339493]

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

Erlhagen W, Bastian A, Jancke D, Riehle A, Schöner G. The distribution of neuronal population activation (DPA) as a tool to study interaction and integration of cortical representations. Journal of Neuroscience Methods. 1999; 94:53–66. [PubMed: 10638815]

Erlhagen W, Schöner G. Dynamic field theory of movement preparation. Psychological Review. 2002; 109(3):545–572. [PubMed: 12088245]

Fair DA, Dosenbach NUF, Church JA, Cohen AL, Brahmbhatt S, Miezin FM, Barch DM, Raichle ME, Petersen SE, Schlaggar BL. Development of distinct control networks through segregation and integration. Proceedings of the National Academy of Sciences. 2007; 104(33):13507–13512.

Faubel C, Schöner G. Learning to recognize objects on the fly: a neurally based dynamic field approach. Neural Networks. 2008; 21(4):562–576. [PubMed: 18501555]

Fisher AV. Automatic shifts of attention in the Dimensional Change Card Sort task: subtle changes in task materials lead to flexible switching. Journal of Experimental Child Psychology. 2011

Fisher AV, Sloutsky VM. When induction meets memory: evidence for gradual transition from similarity-based to category based induction. Child Development. 2005; 76(3):583–597. [PubMed: 15892780]

Frye D, Zelazo PD, Palfai T. Theory of mind and rule-based reasoning. Cognitive Development. 1995; 10:483–527.

Garon N, Bryson SE, Smith IM. Executive function in preschoolers: a review using an integrative framework. Psychological Bulletin. 2008; 134(1):31–60. [PubMed: 18193994]

Gerardi-Caulton G. Sensitivity to spatial conflict and the development of self-regulation in children 24-36 months of age. Developmental Science. 2000; 3(4):397–404.

Giedd JN, Blumenthal J, Jeffries NO, Castellanos FX, Liu H, Zijdenbos A, Paus T, Evans AC, Rapoport JL. Brain development during childhood and adolescence: a longitudinal MRI study. Nature Neuroscience. 1999; 2(10):861–863.

Gogtay N, Giedd JN, Lusk L, Hayashi KM, Greenstein D, Vaituzis AC, Nugent TF, Herman DH, Clasen LS, Toga AW, Rapoport JL, Thompson PM. Dynamic mapping of human cortical development during childhood through early adulthood. Proceedings of the National Academy of Sciences. 2004; 101(21):8174–8179.

Halford GS, Bunch K, McCredden JE. Problem decomposability as a factor in complexity of the dimensional change card sort. Cognitive Development. 2007; 22(3):384–391.

Haxby JV, Grady CL, Horwitz B, Ungerleider LG, Mishkin M, Carson RE, Herscovitch P, Schapiro MB, Rapoport SI. Dissociation of object and spatial visual processing pathways in human extrastriate cortex. Proceedings of the National Academy of Sciences. 1991; 88:1621–1625.

Honomichl RD, Chen Z. Relations as rules: The role of attention in the Dimensional Change Card Sort task. Developmental Psychology. 2011; 47(1):50–60. [PubMed: 20954760]

Hwang K, Velanova K, Luna B. Strengthening of top-down frontal cognitive control networks underlying the development of inhibitory control: A functional magnetic resonance imaging effective connectivity study. The Journal of Neuroscience. 2010; 30(46):15535–15545. [PubMed: 21084608]

Im-Bolter N, Johnson J, Pascual-Leone J. Processing limitations in children with specific language impairment the role of executive function. Child Development. 2006; 77(6):1822–1841. [PubMed: 17107463]

Isaacs EB, Vargha-Khadem F. Differential course of development of spatial and verbal memory span: a normative study. British Journal of Developmental Psychology. 1989; 7(4):377–380.

Jancke D, Erlhagen W, Dinse HR, Akhavan AC, Giese M, Steinhage A, Schöner G. Parametric population representation of retinal location: neuronal interaction dynamics in ca primary visual cortex. The Journal of Neuroscience. 1999; 19(20):9016–9028. [PubMed: 10516319]

Johnson JS, Spencer JP, Schöner G. Garzòn F, Laakso A, Gomila T. Moving to higher ground: the Dynamic Field Theory and the dynamics of visual cognition. Dynamics and Psychology [special issue]. *New Ideas in Psychology*.. 2008

Jordan PJ, Morton JB. Flankers facilitate 3-year-olds’ performance in a card-sorting task. Developmental Psychology. 2008; 44(1):263–274.

Kelly AMC, Di Martino A, Uddin LQ, Shehzad Z, Gee DG, Reiss PT, Margulies DS, Castellanos FX, Milham MP. Development of anterior cingulate functional connectivity from late childhood to early adulthood. Cerebral Cortex. 2009; 19:640–657. [PubMed: 18653667]

Kelso JAS, Scholz JP, Schöner G. Dynamics govern switching among patterns of coordination in biological movement. Physics Letters A. 1988; 134:8–12.

Kirkham NZ, Cruess L, Diamond A. Helping children apply their knowledge to their behavior on a dimension-switching task. Developmental Science. 2003; 6:449–467.

Klingberg T, Forssberg H, Westerberg H. Increased brain activity in frontal and parietal cortex underlies the development of visuospatial working memory capacity during childhood. Journal of Cognitive Neuroscience. 2002; 14(1):1–10. [PubMed: 11798382]

Kloo D, Perner J. Disentangling dimensions in the dimensional change card sort task. Developmental Science. 2005; 8(1):44–56. [PubMed: 15647066]

Kloo D, Perner J, Kerschhuber A, Dabernig S, Aichhorn M. Sorting between dimensions: conditions of cognitive flexibility in preschoolers. Journal of Experimental Child Psychology. 2008; 100:115–134. [PubMed: 18241879]

Kopecz K, Schöner G. Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. Biological Cybernetics. 1995; 73(1):49–60. [PubMed: 7654850]

Kourtzi Z, Erb M, Grodd W, Bülthoff HH. Representation of the perceived 3-D object shape in human lateral occipital complex. Cerebral Cortex. 2003; 13:911–920. [PubMed: 12902390]

Kravitz DJ, Vinson LD, Baker CI. How position dependent is visual object recognition? Trends in Cognitive Sciences. 2008; 12(3):114–122. [PubMed: 18262829]

Lehto JE, Juujärvi P, Kooistra L, Pulkkinen L. Dimensions of executive functioning: evidence from children. British Journal of Developmental Psychology. 2003; 21(1):59–80.

Lenroot RK, Giedd JN. Brain development in children and adolescents: insights from anatomical magnetic resonance imaging. Neuroscience & Biobehavioral Review. 2006; 30(6):718–729.

Lepsien J, Nobre AC. Attentional modulation of object representations in working memory. Cerebral Cortex. 2007; 17:2072–2083. [PubMed: 17099066]

Lipinski J, Spencer JP, Samuelson LK. Biased feedback in spatial recall yields a violation of delta rule learning. Psychonomic Bulletin and Review. 2010; 17:581–588. [PubMed: 20702881]

Liss M, Fein D, Allen D, Dunn M, Feinstein C, Morris R, Waterhouse L, Rapin I. Executive functioning in high-functioning children with autism. Journal of Child Psychology and Psychiatry. 2001; 42(2):261–270. [PubMed: 11280422]

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. Nature. 2001; 412:150–157. [PubMed: 11449264]

Logie RH, Pearson DG. The inner eye and the inner scribe of visuo-spatial working memory: evidence from developmental fractionation. The European Journal of Cognitive Psychology. 1997; 9(3): 241–257.

Luck, SJ.; Vecera, SP. Attention.. In: Pashler, H.; Yantis, S., editors. Stevens' Handbook of Experimental Psychology: Vol. 1. Sensation and Perception. 3rd ed.. Wiley; New York: 2002. p. 235-286.

Mack W. Improving postswitch performance in the dimensional change card-sorting task: the importance of the switch and of pretraining by redescribing the test cards. Journal of Experimental Child Psychology. 2007; 98:243–251. [PubMed: 17599343]

Marcovitch S, Zelazo PD. The A-not-B error: evidence from a logistic meta-analysis. Child Development. 1999; 70(6):1297–1313.

McEvoy RE, Rogers SJ, Pennington BF. Executive function and social communication deficits in young autistic children. Journal of Child Psychology and Psychiatry. 1993; 34(4):563–578. [PubMed: 7685360]

McNab F, Leroux G, Strand F, Thorell L, Bergman S, Klingberg T. Common and unique components of inhibition and working memory: An fMRI within-subject investigation. Neuropsychologia. 2008; 46:2668–2682. [PubMed: 18573510]

Milner B. Effects of different brain lesions on card sorting: the role of the frontal lobes. Archives of Neurology. 1963; 9(1):90–100.

Minati L, Visani E, Dowell NG, Medford N, Critchley HD. Variability comparison of simultaneous brain near-infrared spectroscopy (NIRS) and functional MRI (fMRI) during visual stimulation. Journal of Medical Engineering Technology. 2011; 35:370–376. [PubMed: 21780948]

Miyaki A, Friedman NP, Emerson MJ, Witzki AH, Howerter A. The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: a latent variable analysis. Cognitive Psychology. 2000; 41:49–100. [PubMed: 10945922]

Moffitt TE, Arseneault L, Belsky D, Dickson N, Hancox RJ, Harrignton H, Houts R, Poulton R, Roberts BW, Ross S, Sears MR, Thomson WM, Caspi A. A gradient of childhood self-control predicts health, wealth, and public safety. Proceedings of the National Academy of Sciences. 2011; 108(7):2693–2698.

Moriguchi Y, Hiraki K. Neural origins of cognitive shifting in young children. Proceedings of the National Academy of Sciences. 2009; 106(14):6017–6021.

Morton JB. Understanding genetic, neurophysiological, and experiential influences on the development of executive function: The need for developmental models. Wiley Interdisciplinary Reviews: Cognitive Science. 2010; 1(5):709–723.

Morton JB, Bosma R, Ansari D. Age-related changes in brain activation associated with dimensional shifts of attention: An fMRI study. Neuroimage. 2009; 46:249–256. [PubMed: 19457388]

Morton JB, Munakata Y. Active versus latent representations: a neural network model of perseveration and dissociation in early childhood. Developmental Psychobiology. 2002; 40:255–265. [PubMed: 11891637]

Müller U, Dick AS, Gela K, Overton WF, Zelazo PD. The role of negative priming in preschoolers' flexible rule use on the dimensional change card sort task. Child Development. 2006; 77(2):395–412. [PubMed: 16611180]

Müller U, Zelazo PD, Imrisek S. Executive function and children's understanding of false-belief: how specific is the relation? Cognitive Development. 2005; 20(2):173–189.

Müller U, Zelazo PD, Lurye LE, Liebermann DP. The effect of labeling on preschool children's performance in the Dimensional Change Card Sort Task. Cognitive Development. 2008; 23:395–408.

Norman, DA.; Shallice, T. Attention to action: willed and automatic control of behavior.. In: Davidson, RJ.; Schwartz, GE.; Shapiro, D., editors. Consciousness and self-regulation: Advances in Research and Theory. Plenum Press; New York: 1986. p. 1-18.

Nosofsky RM. Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory, & Cognition. 1984; 10:104–114.

Op De Beeck H, Vogels R. Spatial sensitivity of macaque inferior temporal neurons. Journal of Comparative Neurology. 2000; 426:505–518. [PubMed: 11027395]

Perner J, Lang B. What causes 3-year-olds' difficulty on the dimensional change card sort task? Infant and Child Development. 2002; 11(2):93–105.

Pennington BF, Ozonoff S. Executive functions and developmental psychopathology. Journal of Child Psychology and Psychiatry. 1996; 37(1):51–87. [PubMed: 8655658]

Rennie DAC, Bull R, Diamond A. Executive functioning in preshcoolers: reducing the inhibitory demands of the dimensional change card sort task. Developmental Neuropsychology. 2004; 26(1):423–443. [PubMed: 15276903]

Rueda MR, Fan J, McCandliss BD, Halparin JD, Gruber DB, Lercari LP, Posner MI. Development of attentional networks in childhood. Neuropsychologia. 2004; 42(8):1029–1040. [PubMed: 15093142]

Samuelson LK, Smith LB, Perry LK, Spencer JP. Grounding word learning in space. PLoS One. in press.

Samuelson LK, Schutte AR, Horst JS. The dynamic nature of knowledge: insights from a dynamic field model of children's novel noun generalization. Cognition. 2009; 110(3):322–345. [PubMed: 19131050]

Sandamirskaya Y, Schöner G. An embodied account of serial order: How instabilities drive sequence generation. Neural Networks. 2010; 23(10):1164–1179. [PubMed: 20800989]

Sandhofer CM, Smith LB. Learning color words involves learning a system of mappings. Developmental Psychology. 1999; 35(3):668–697. [PubMed: 10380858]

Sandhofer CM, Smith LB. Why children learn color and size words so differently: evidence from adults' learning of artificial terms. Journal of Experimental Psychology: General. 2001; 130(4):600–620. [PubMed: 11757871]

Schroeter ML, Zysset S, Wahl M, von Cramon DY. Prefrontal activation due to Stroop interference increases during development—an event-related fNIRS study. Neuroimage. 2004; 23(4):1317–1325. [PubMed: 15589096]

Schöner G, Kelso JAS. Dynamic pattern generation in behavioral and neural systems. Science. 1988; 239:1513–1520. [PubMed: 3281253]

Schutte AR, Spencer JP, Schöner G. Testing the dynamic field theory: working memory for locations becomes more spatially precise over development. Child Development. 2003; 74(5):1393–1417. [PubMed: 14552405]

Schutte AR, Spencer JR. Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. Journal of Experimental Psychology: Human Perception and Performance. 2009; 35(6):1689–1725.

Silver MA, Kastner S. Topographic maps in human frontal and parietal cortex. Trends in Cognitive Science. 2009; 13(11):488–495.

Simmering, VA. Developing a magic number: The dynamic field theory reveals why visual working capacity estimates differ across tasks and development. University of Iowa; 2008.

Simmering VA, Schutte AR, Spencer JP. Becker S. Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. Computational Cognitive Neuroscience [special issue]. Brain Research. 2008; 1202:68–86. [PubMed: 17716632]

Simmons WK, Ramjee V, Beauchamp MS, McRae A, Martin A, Barsalou LW. A common neural substrate for perceiving and knowing about color. Neuropsychologia. 2007; 45:2802–2810. [PubMed: 17575989]

Smith LB, Jones SS, Landau B. Naming in young children: a dumb attentional mechanism? Cognition. 1996; 60(2):143–171. [PubMed: 8811743]

Sowell ER, Trauner DA, Garnst A, Jernigan TL. Development of cortical and subcortical brain structures in childhood and adolescence: a structural MRI study. Developmental Medicine & Child Neurology. 2002; 44(1):4–16. [PubMed: 11811649]

Spencer, JP.; Perone, S.; Johnson, JS. Dynamic field theory and embodied cognitive dynamics.. In: Spencer, JP.; Thomas, M.; McClelland, JL., editors. Toward a unified theory of development: Connectionism and dynamic systems theory re-considered. Oxford University Press; New York: 2009.

Stevens MC, Skudlarski P, Pearlson GD, Calhoun VD. Age-related cognitive gains are mediated by the effects of white matter development on brain integration. Neuroimage. 2009; 48:738–746. [PubMed: 19577651]

Stevens MC, Pearlson GD, Caloun VD. Changes in the interaction of resting-state neural networks from adolescence to adulthood. Human Brain Mapping. 2009; 30:2356–2366. [PubMed: 19172655]

Thelen E, Schöner G, Scheier C, Smith L. The dynamics of embodiment: A field theory of infant perseverative reaching. Behavioral and Brain Sciences. 2001; 24:1–86. [PubMed: 11515285]

Thelen, E.; Smith, LB. A Dynamic Systems Approach to Cognition and Action. MIT Press; Cambridge, MA: 1994.

Thomason ME, Race E, Burrows B, Whitfield-Gabrieli S, Glover GH, Gabrieli JDE. Development of spatial and verbal working memory capacity in the human brain. Journal of Cognitive Neuroscience. 2009; 21(2):316–332. [PubMed: 18510448]

Towse JN, Redbond J, Houston-Price CMT, Cook S. Understanding the dimensional change card sort: perspectives from task success and failure. Cognitive Development. 2000; 15:347–365.

Treisman AM. The binding problem. Current Opinions in Neurobiology. 1996; 6(2):171–178.

Treisman AM, Gelade G. A feature-integration theory of attention. Cognitive Psychology. 1980; 12:97–136. [PubMed: 7351125]

Tsujii T, Yamamoto E, Sayako M, Watanabe S. Longitudinal study of spatial working memory development in young children. Learning and Memory. 2009; 20(8):759–763.

Tsujimoto S. The prefrontal cortex: Functional neural development during early childhood. The Neuroscientist. 2008; 14:345–358. [PubMed: 18467667]

Turvey, MT.; Shaw, RE. Toward an ecological physics and a physical psychology.. In: Solso, R.; Massaro, D., editors. The Science of the Mind: 2001 and beyond. Oxford University Press; Oxford: 1995. p. 144-169.

Ungerleider, LG.; Mishkin, M. Two cortical visual systems.. In: Ingle, DJ.; Goodale, MA.; Mansfield, RJW., editors. Analysis of Visual Behavior. MIT Press; Cambridge, MA: 1982.

Velanova K, Wheeler ME, Luna B. The maturation of task set-related activation supports late developmental improvements in inhibitory control. The Journal of Neuroscience. 2009; 29(40): 12558–12567. [PubMed: 19812330]

Wilson HR, Cowan JD. Excitatory and inhibitory interactions in localized populations of model neurons. Biophysical Journal. 1972; 12:1–24. [PubMed: 4332108]

Xiao, Y-p.; Wang, Y.; Felleman, DJ. A spatially organized representation of color in macaque cortical area V2. Nature. 2003; 421:532–539.

Yerys B, Munakata Y. When labels hurt but novelty helps: children's perseveration and flexibility in a card-sorting task. Child Development. 2006; 77(6):1589–1607. [PubMed: 17107448]

Zanto TP, Rubens MT, Bollinger J, Gazzaley A. Top-down modulation of visual feature processing: The role of the inferior frontal junction. NeuroImage. 2011; 53:736–745. [PubMed: 20600999]

Zanto TP, Rubens MT, Thangavel A, Gazzaley A. Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. Nature Neuroscience. 2011; 14:656–663.

Zelazo PD. The development of conscious control in childhood. Trends in Cognitive Sciences. 2004; 8(1):12–17. [PubMed: 14697398]

Zelazo PD, Frye D, Rapus T. An age-related dissociated between knowing rules and using them. Cognitive Development. 1996; 11:37–63.

Zelazo PD, Reznick JS, Piñon DE. Response control and the execution of verbal rules. Developmental Psychology. 1995; 31(3):508–517.

Zelazo PD, Müller U, Frye D, Markovitch S. The development of executive function in early childhood. Monographs for the Society for Research in Child Development. 2003; 68(3) Serial No. 274.

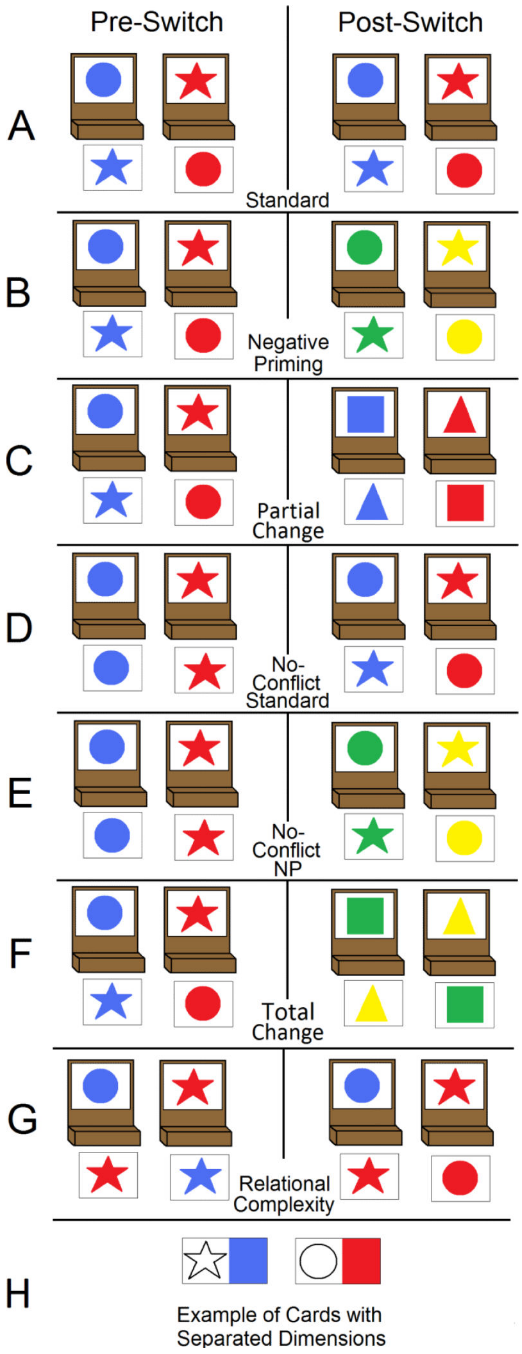

Figure 1.

Target and test cards used in various versions of the DCCS. For all of the examples shown ***color*** is the pre-shape dimension and ***shape*** is the post-switch dimension.

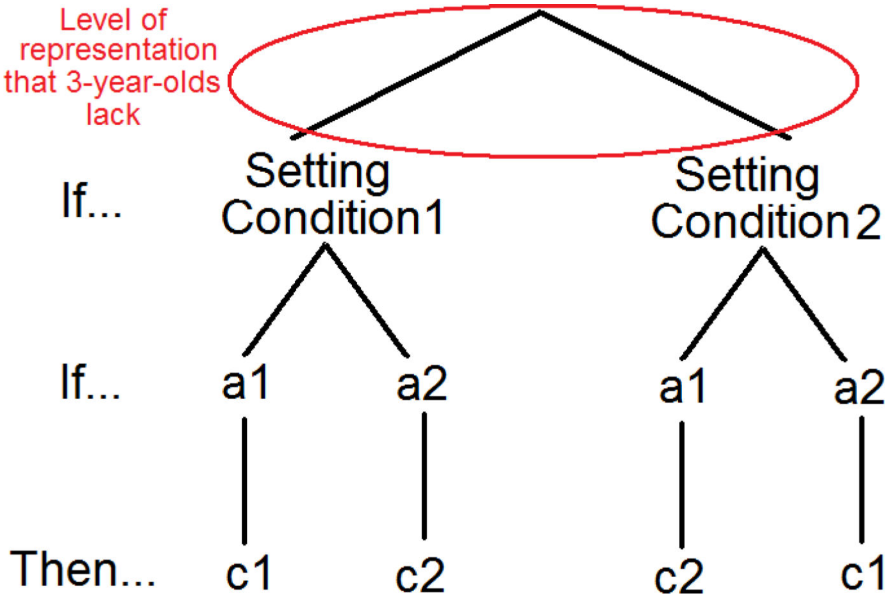

Figure 2.

The rule-hierarchy of the Cognitive Complexity and Control (CCC) Theory (Zelazo et al., 2003). Setting conditions determine if shape or color rules are to be used. The *a*'s are different features of each dimension (*a1* goes with *a1* and *a2* with *a2* under each setting condition as the combination of features on the test cards). The *c*'s are the different decisions to be made for each feature (i.e., where it is to be placed).

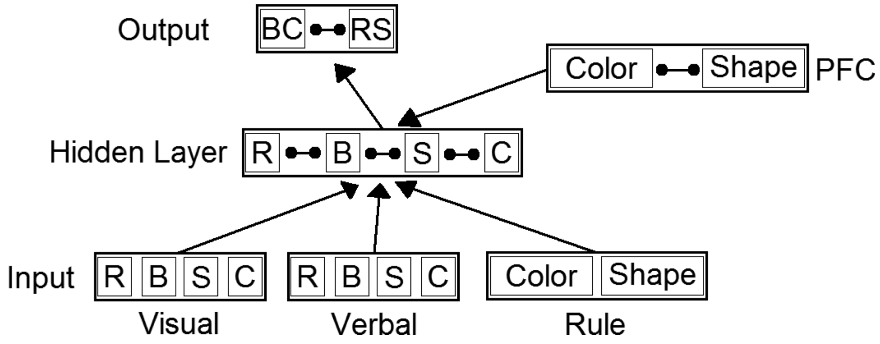

Figure 3.

The PDP model proposed by Morton & Munakata (2002). Visual, verbal, and rule inputs are fed through a hidden layer and a PFC layer (which also feeds into the hidden layer and modifies the strength of the connections between the hidden layer and the output layer). Decisions are made in the output layer for which target card the test card should be matched.

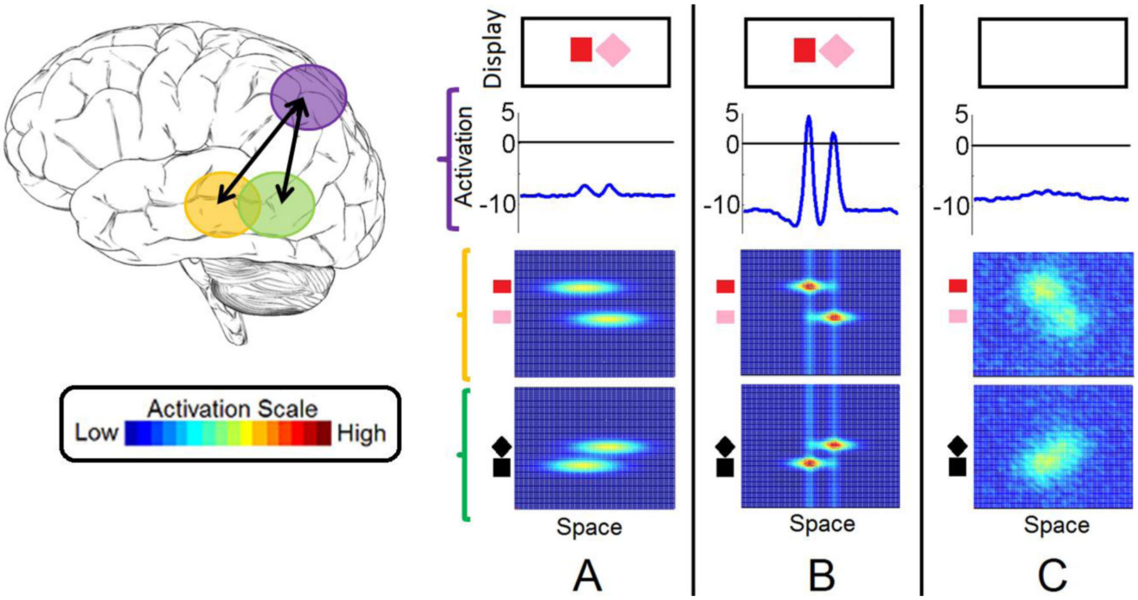

Figure 4.

WM fields for the feature binding model. Panel A depicts the model just after the inputs have been turned on. Panel B shows the WM fields after the inputs have reached threshold. Within fields, the neural interactions have been engaged to form a peak. Between fields, spatial activation is being shared (visible in the feature WM fields as the vertical ridge of activation) to anchor features together in the representation of an object with a particular shape and color. In Panel C the WM fields are shown after the inputs have turned off and excitation has relaxed to resting level. The contribution from LTM is now visible. The left panel shows the putative mapping of the different fields to cortical locations.

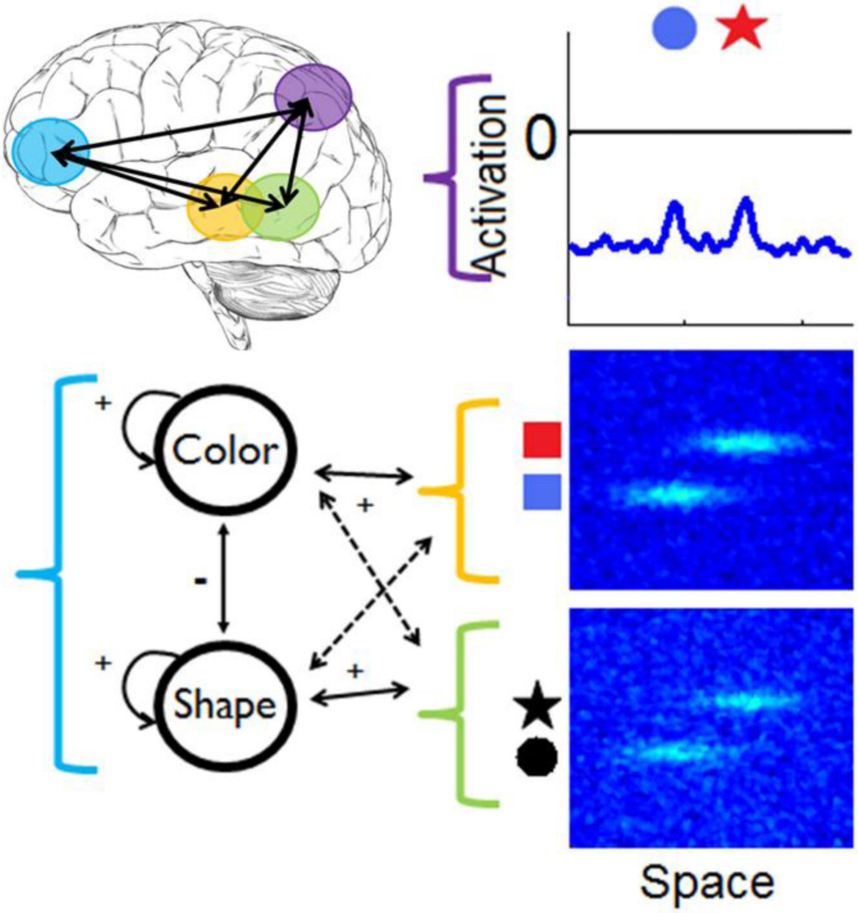

Figure 5.
The object WM model, the dimensional attention architecture, and the putative mapping to cortex.

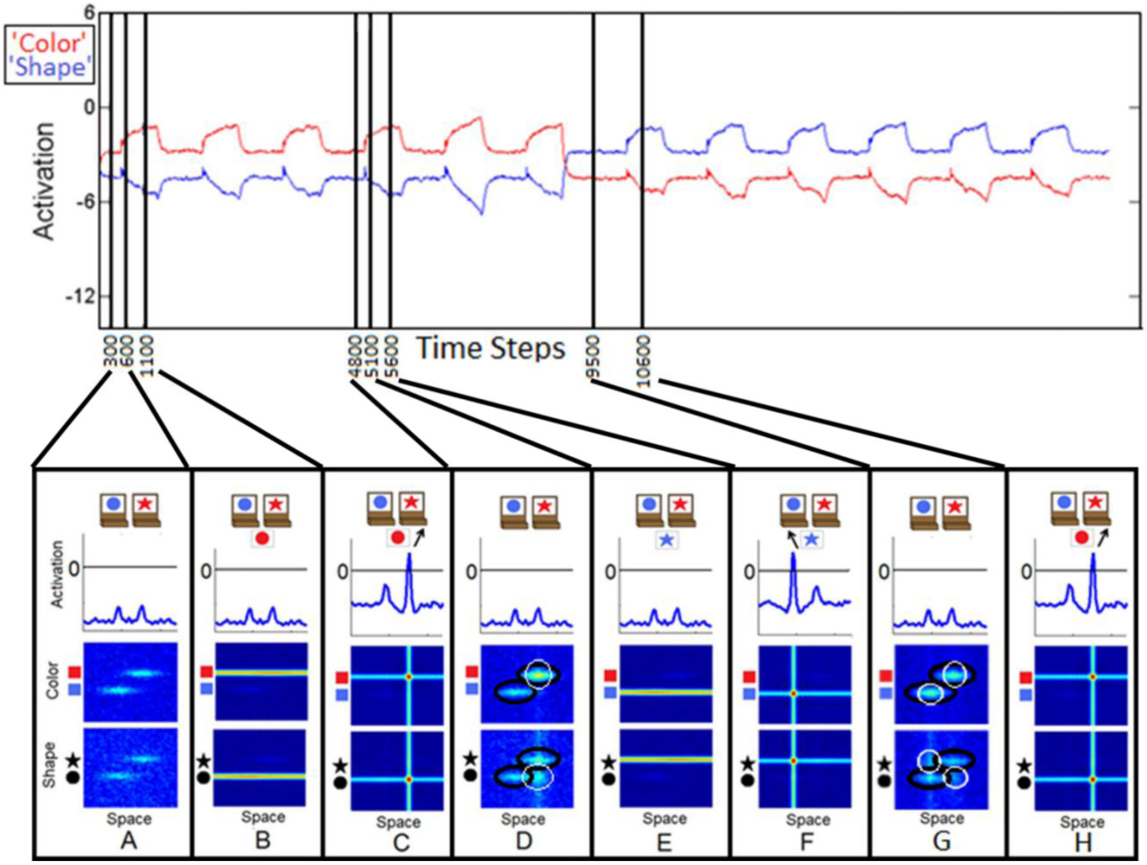

Figure 6.

The sequence of events as the model sorts cards in the DCCS. For simplicity, only the spatial, color, and shape WM fields are shown in each panel. The top panel depicts the activation of the dimensional attention nodes over the 6 pre- and 6 post-switch trials ('color' is presented in right, 'shape' is presented in blue). In this example, color is the relevant dimension for the pre-switch and shape is the relevant dimension for the post-switch. Panel A shows the inputs for the target cards and trays. Panel B shows the input for a red circle test card. Panel C shows the model sorting this card to the right. Panel D shows the formation of Hebbian traces from making that decision (target inputs are circled in black ovals and HTs are circled in white). Panels E and F show the inputs and decision being made for a blue star during the pre-switch phase. Panel G shows the WM fields of the model going into the post-switch phase after sorting during pre-switch phase. Black ovals outline the target inputs, while white circles outline the HTs. Panel H shows the model perseverating and soring the blue circle by color even though the 'shape' node is now more strongly activated.

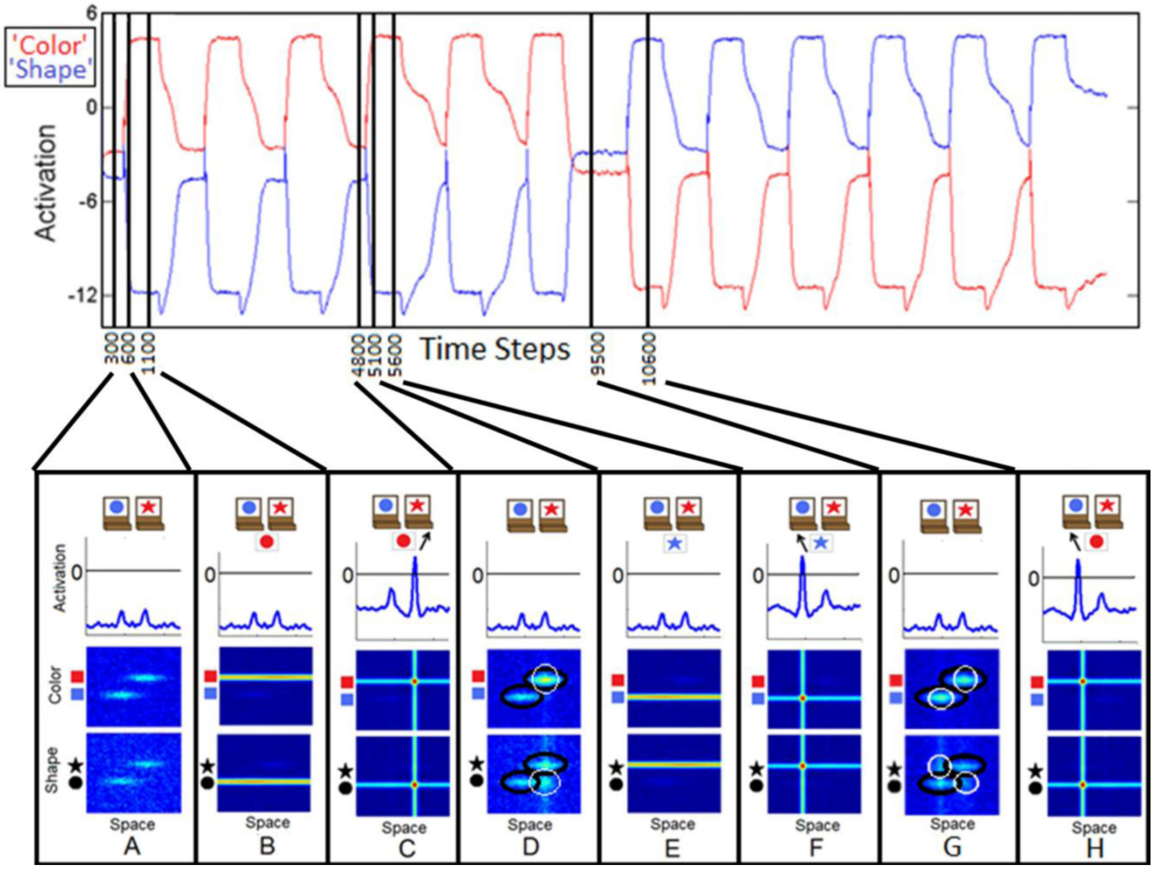

Figure 7.

The 4-year-old model sorting during the pre- and post-switch phases. This figures shows the same sequence of events as Figure 6. Critically, Panel H shows the model correctly switching and sorting the red circle by shape. This is due to the robust activation of the dimensional nodes during each trial which provides strong top-down modulation of activation to the relevant feature field.

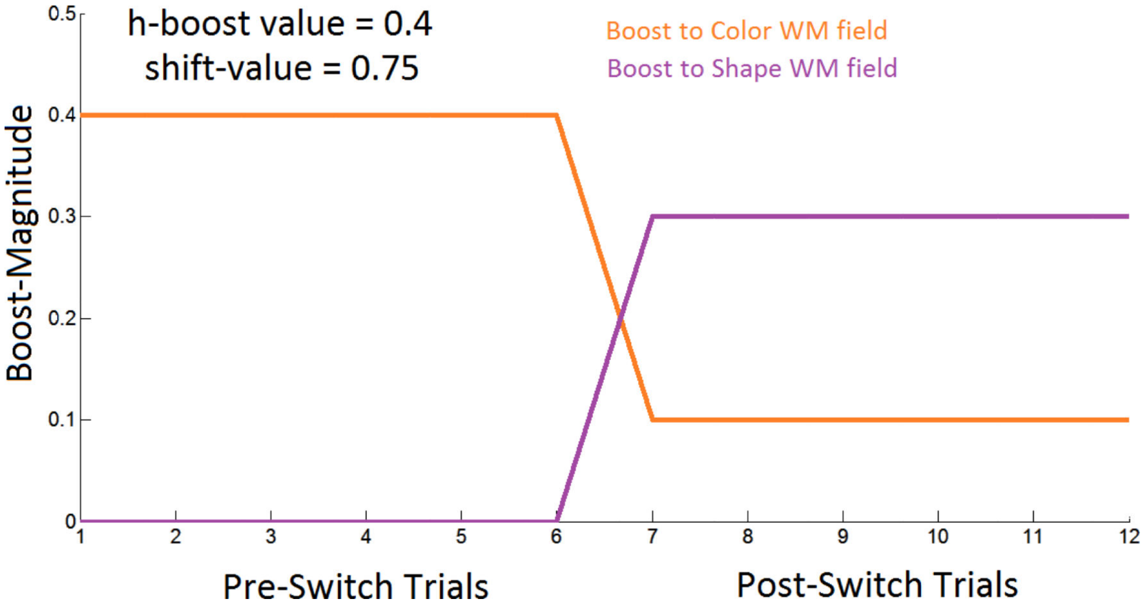

Figure 8.

Schematic diagram depicting the simplified processes of boosting and shifting. For trials 1-6, the boost (.4) is applied to the pre-switch field (the color WM field) is the full boost value. At trial 7 a portion of the boost is shifted to post-switch field (the shape WM field) determined by the shift value (0.75).

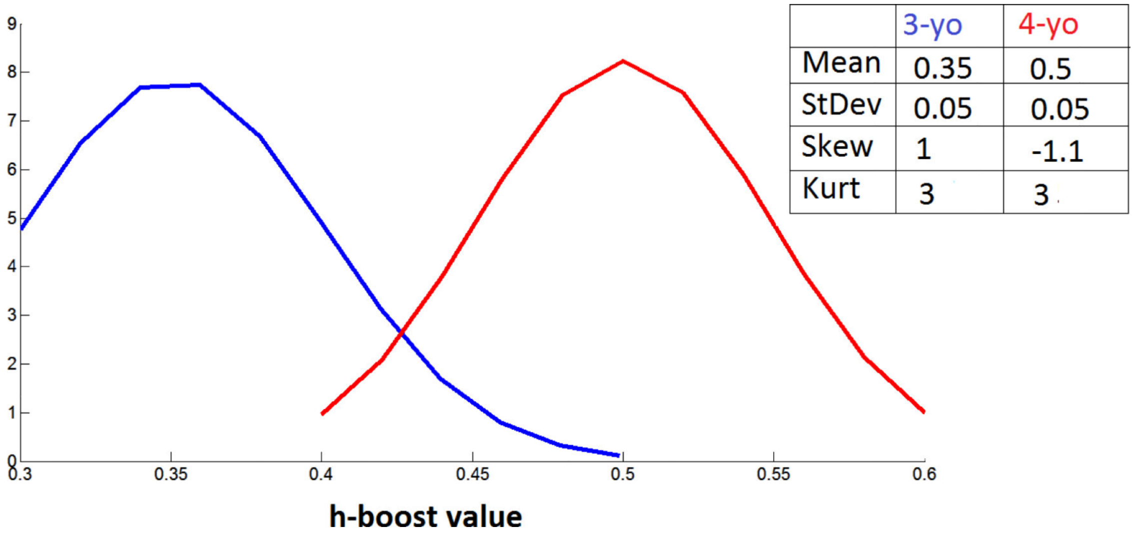

	3-yo	4-yo
Mean	0.35	0.5
StDev	0.05	0.05
Skew	1	-1.1
Kurt	3	3

Figure 9.

H-boost distributions for 3- and 4-year-old models. The distribution for 3-year-olds is dominated by lower h-boost values, while the distribution for 4-year-olds is dominated by higher h-boost values reflecting the stronger mapping of the dimensional nodes to their relevant feature fields.

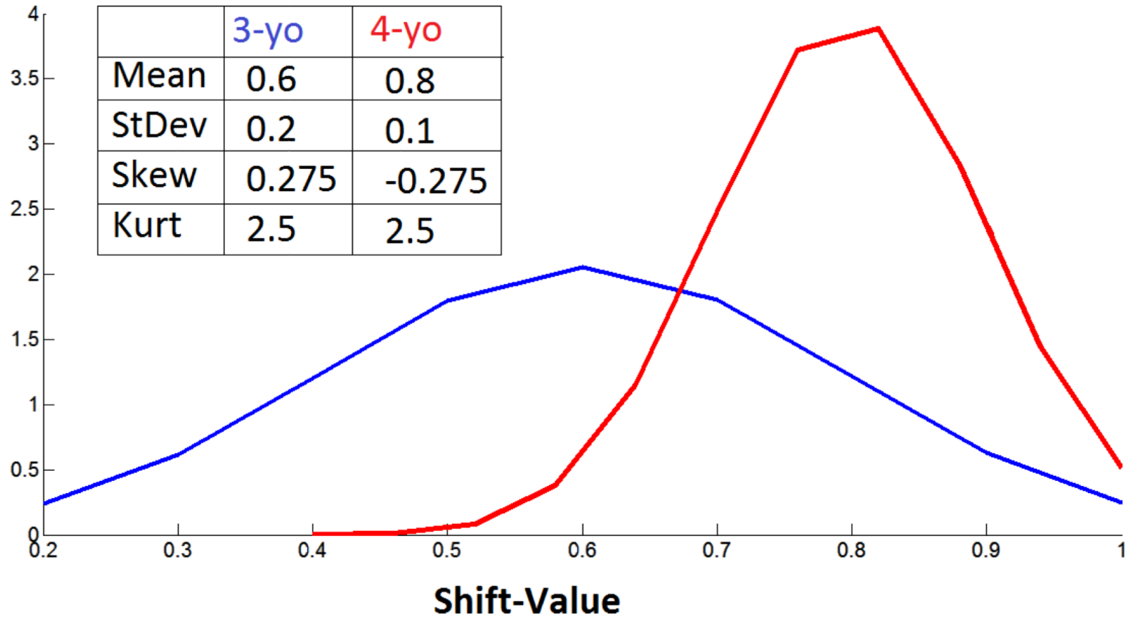

	3-yo	4-yo
Mean	0.6	0.8
StDev	0.2	0.1
Skew	0.275	-0.275
Kurt	2.5	2.5

Figure 10.

Shift-value distributions for 3- and 4-year-old models. The 3-year-old distribution covers a wide range of intermediary value while the 4-year-old distribution is skewed to values closer to 1 capturing the increasing fidelity with which the dimensional nodes are able to achieve robust selective activation.

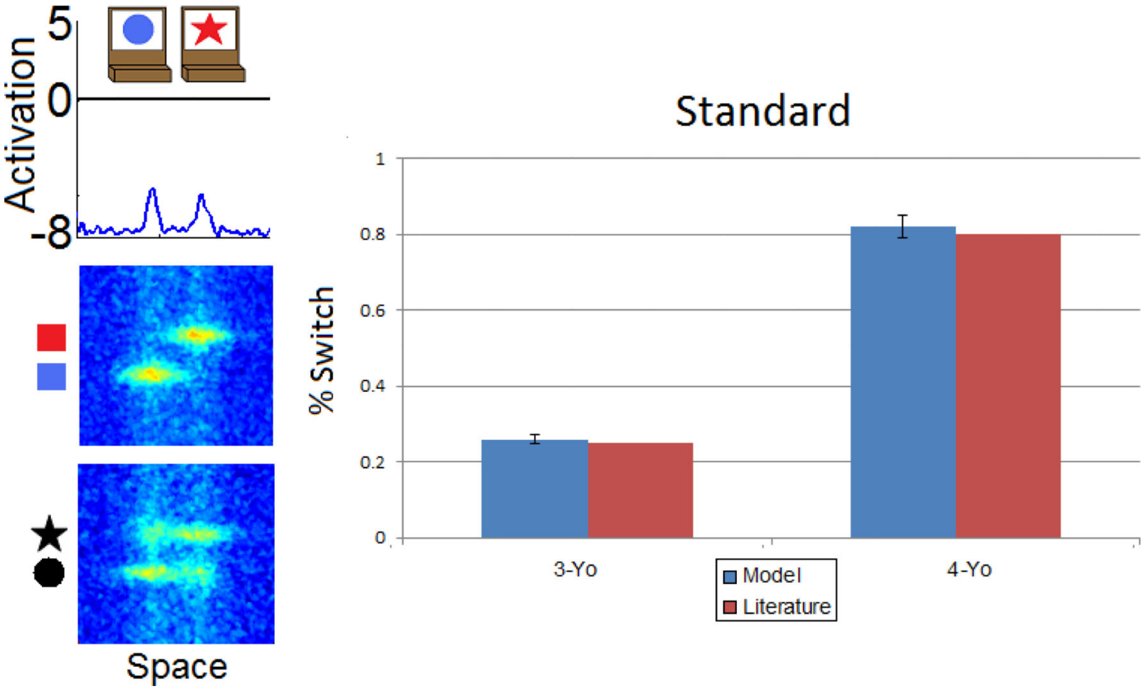

Figure 11.

WM fields of the model at the beginning of the post-switch phase and simulation results for the Standard version.

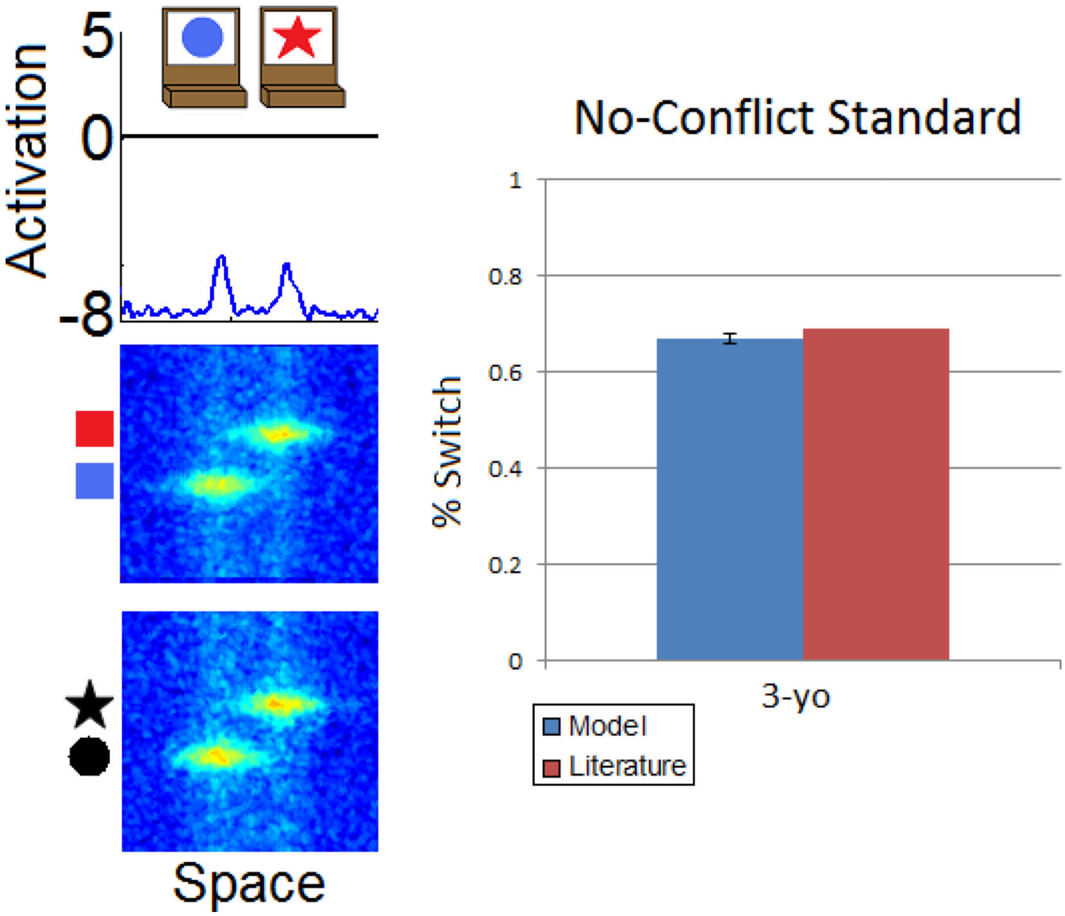

Figure 12.

WM fields of the model at the beginning of the post-switch phase and simulation results for the No-Conflict Standard version.

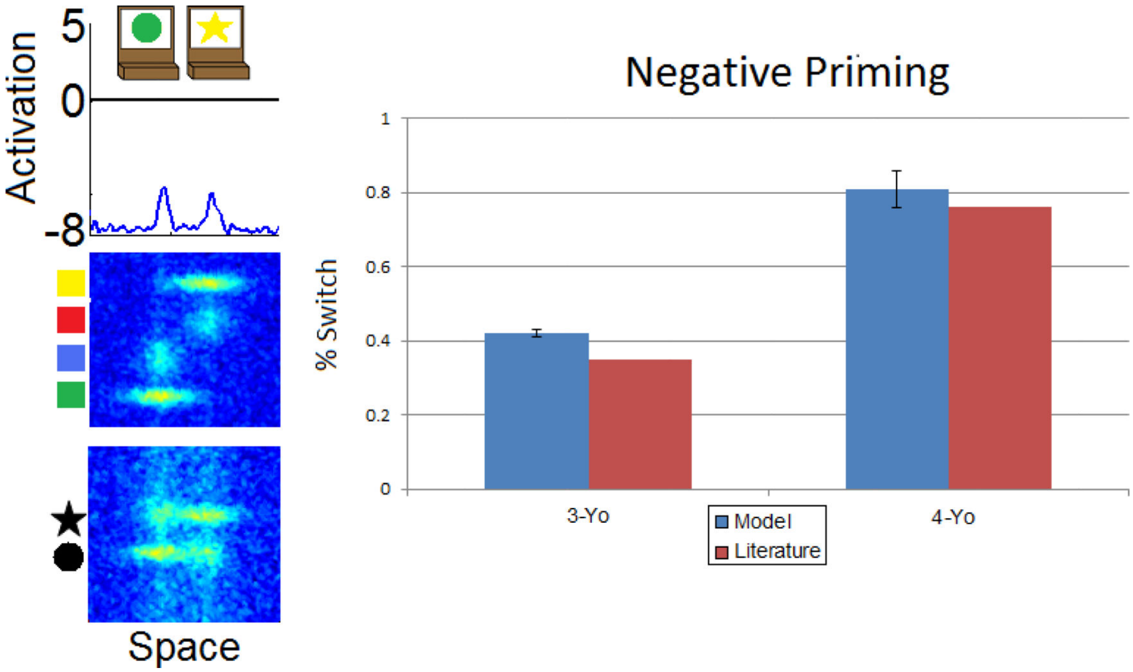

Figure 13.

WM fields for the model at the beginning of the post-switch phase and simulation results for the NP version.

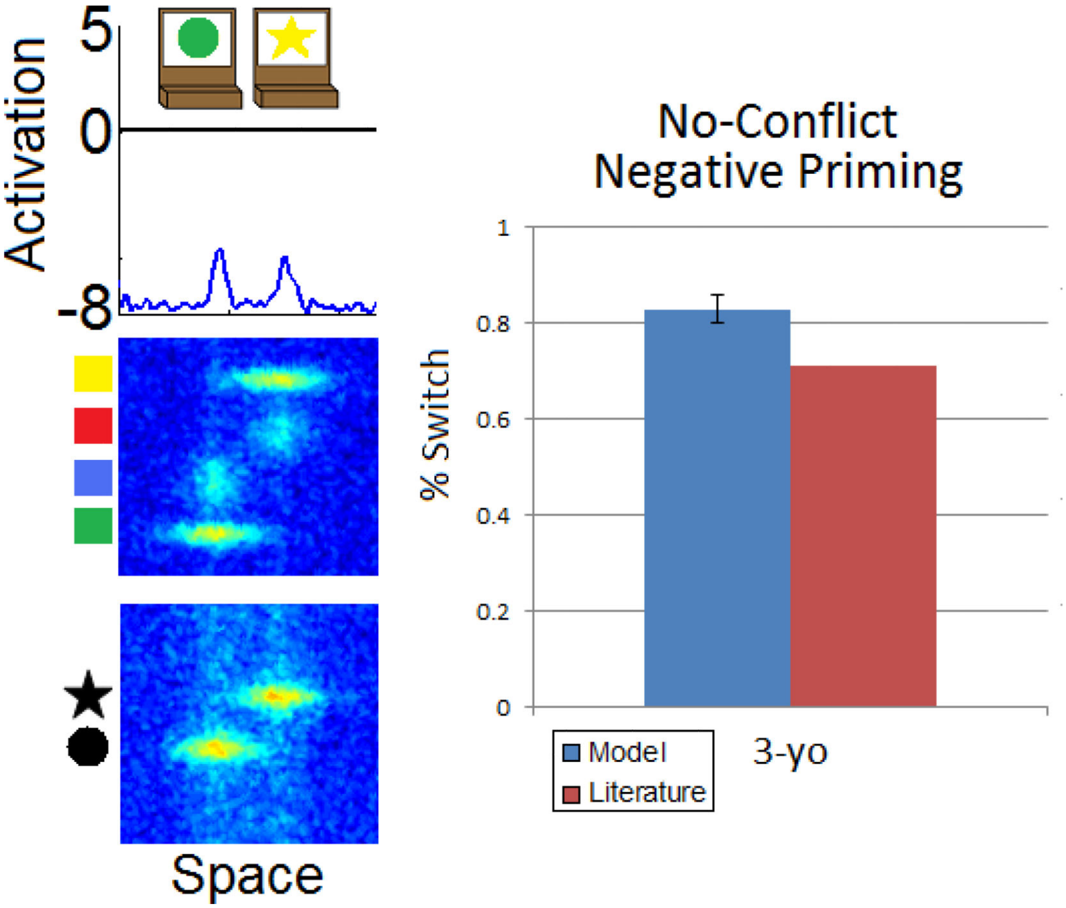

Figure 14.
WM fields for the model at the beginning of the post-switch phase and simulation results for the No-Conflict NP version.

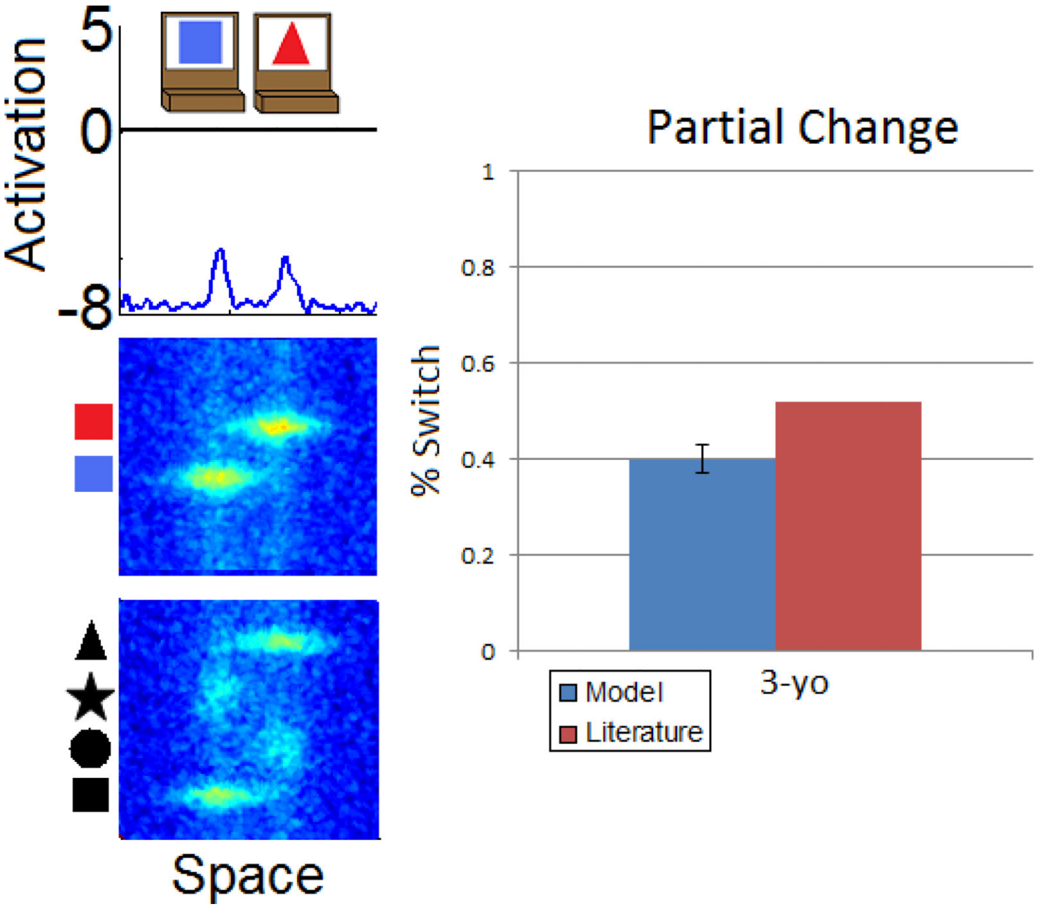

Figure 15.

WM fields for the model at the beginning of the post-switch phase and simulation results for the Partial-Change version.

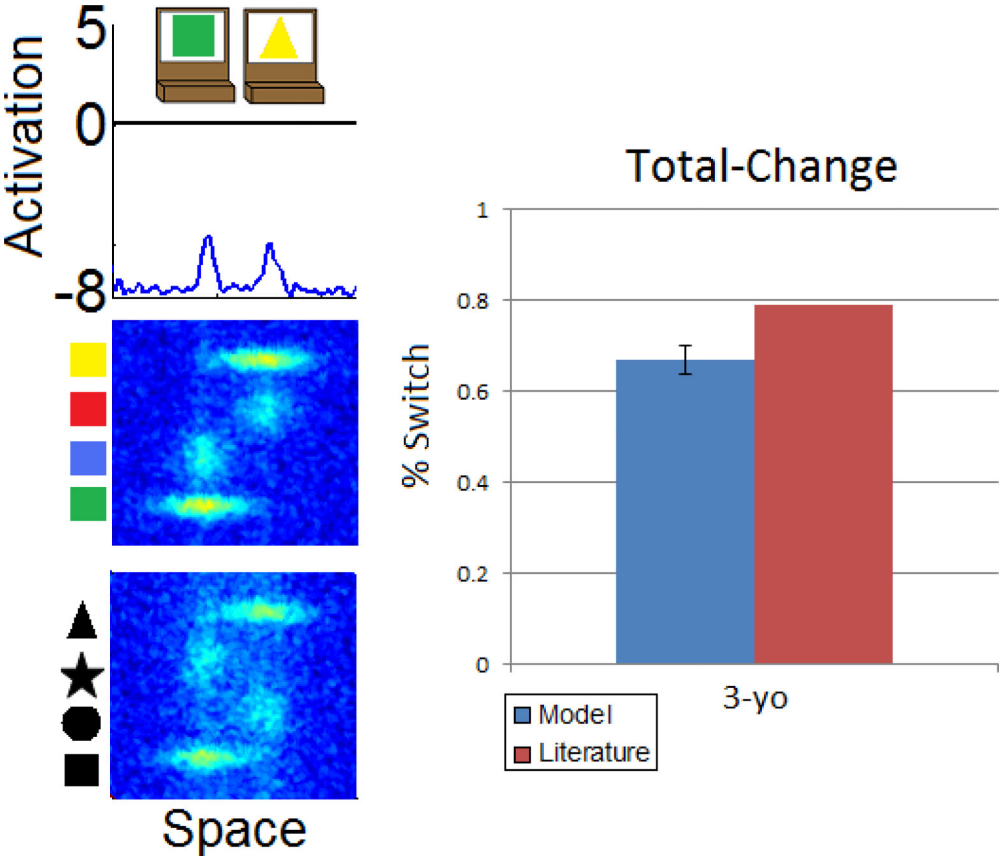

Figure 16.
WM fields for the model at the beginning of the post-switch phase and simulation results for the Total-Change version.

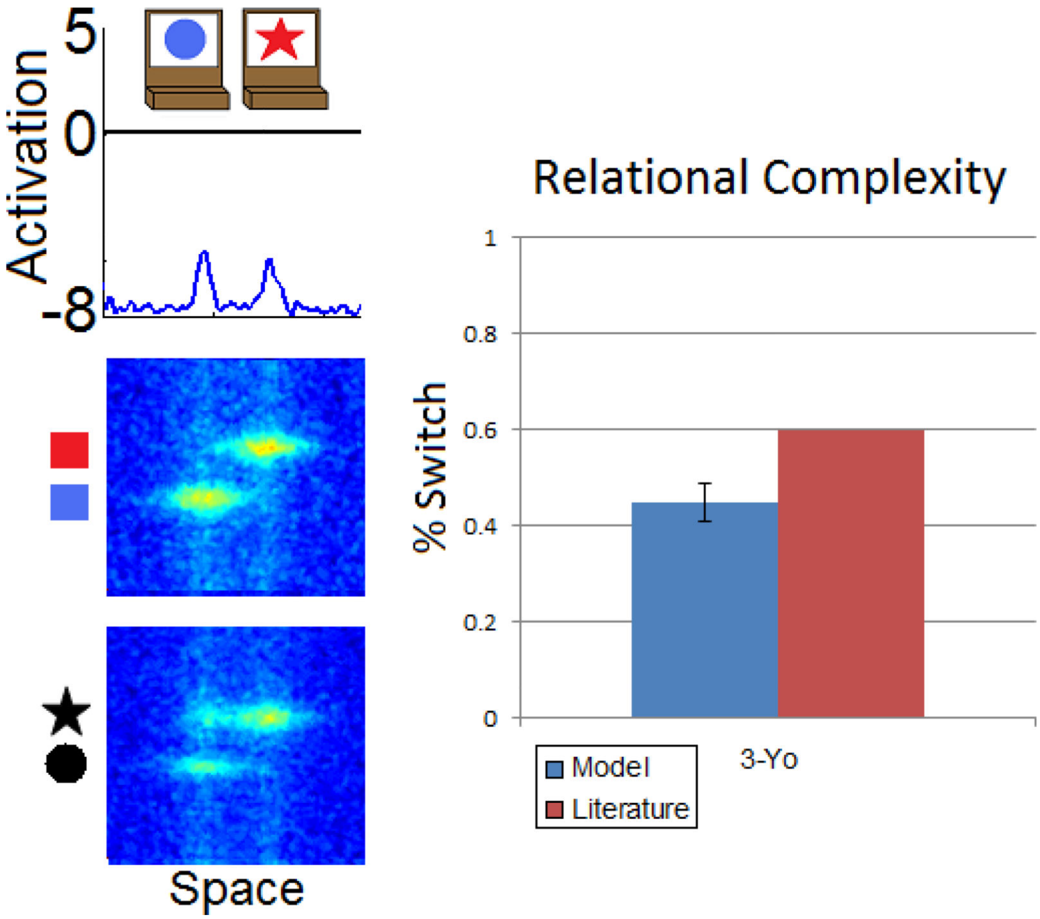

Figure 17.
WM fields for the model at the beginning of the post-switch phase and simulation results for the Relational Complexity version.

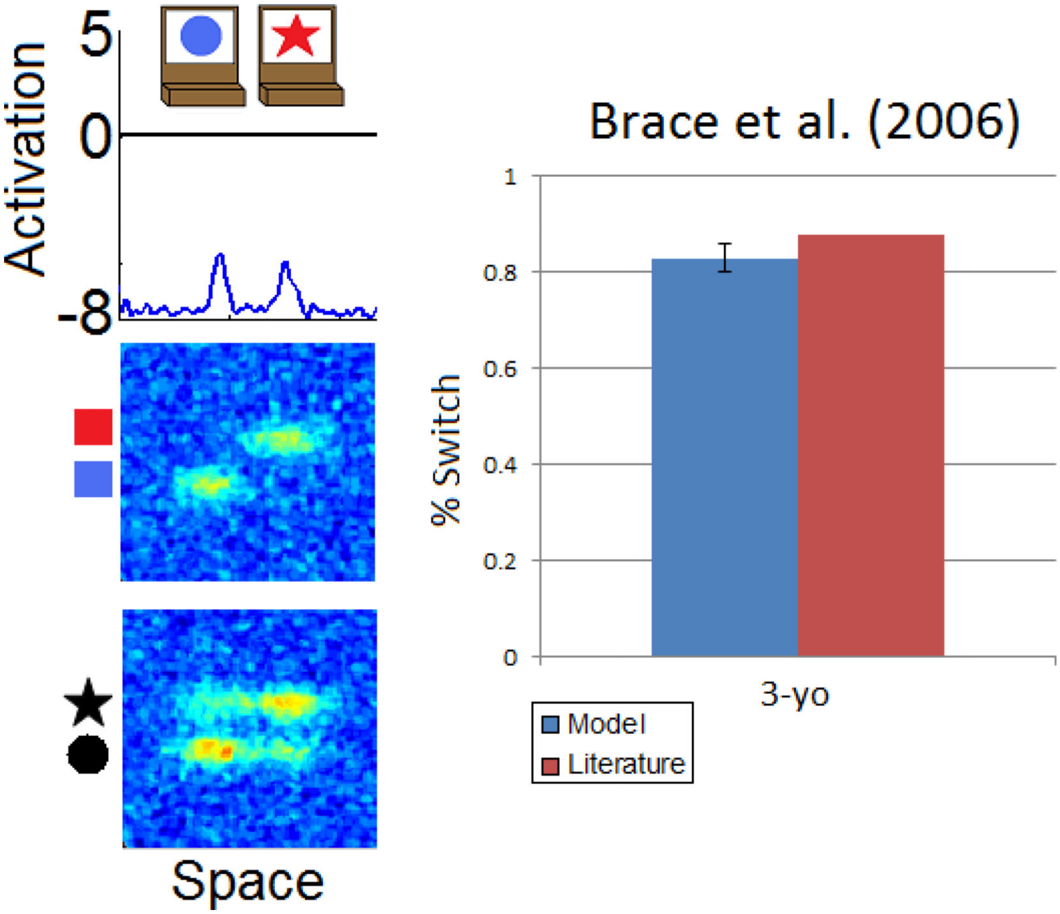

Figure 18.

WM fields for the model at the beginning of the post-switch phase and simulation results for the training study by Brace, Morton, & Munakata (2006).

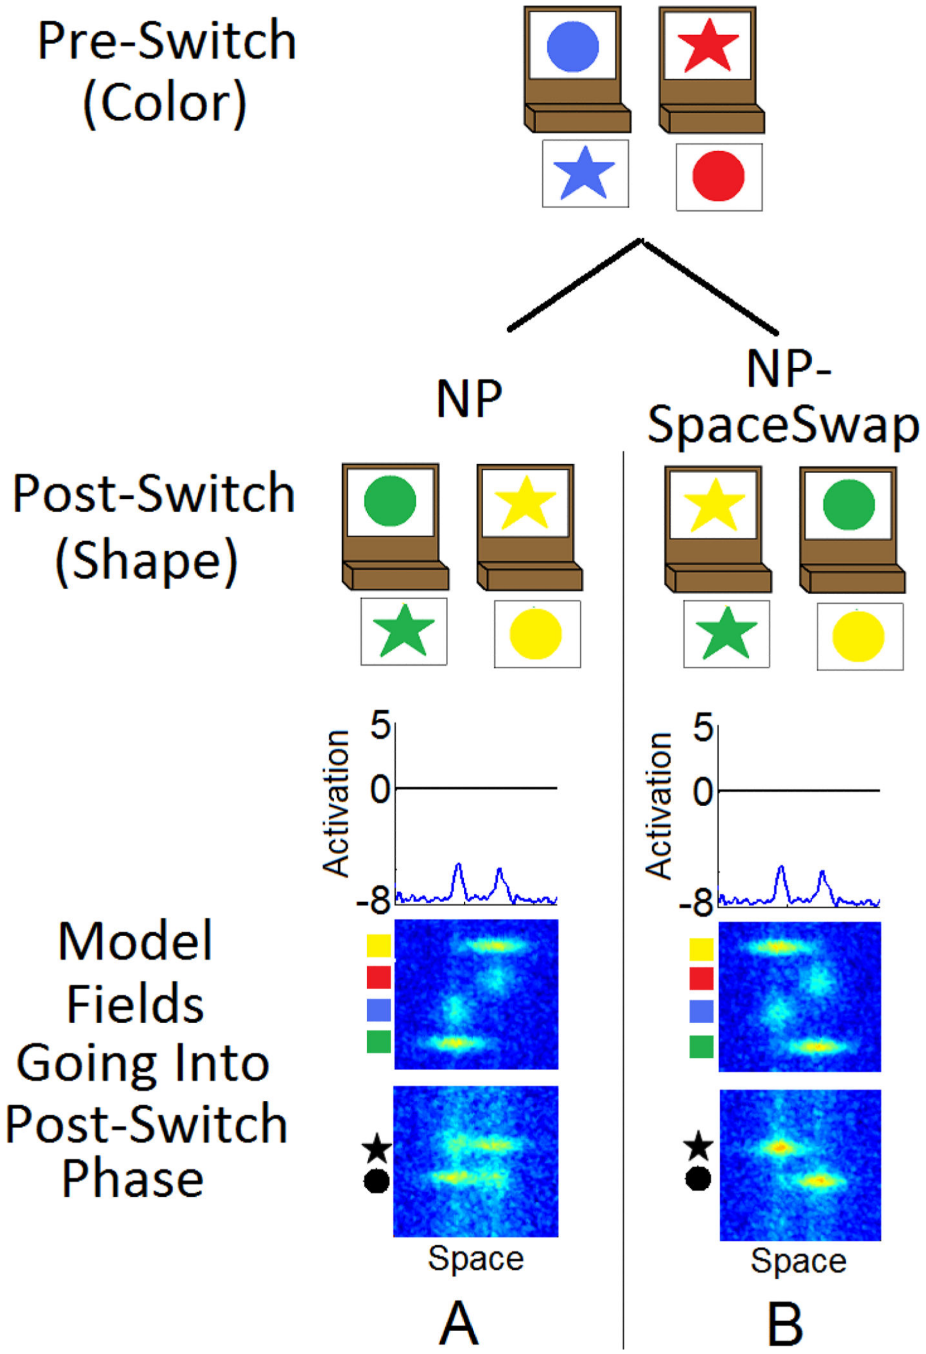

Figure 19.

Model at the start of the pre-switch phase for the NP and NPS versions. In this example, color is the pre-switch dimension and shape is the post-switch dimension. In panel B conflict in the post-switch (shape) WM field is eliminated by swapping the locations of the target cards.

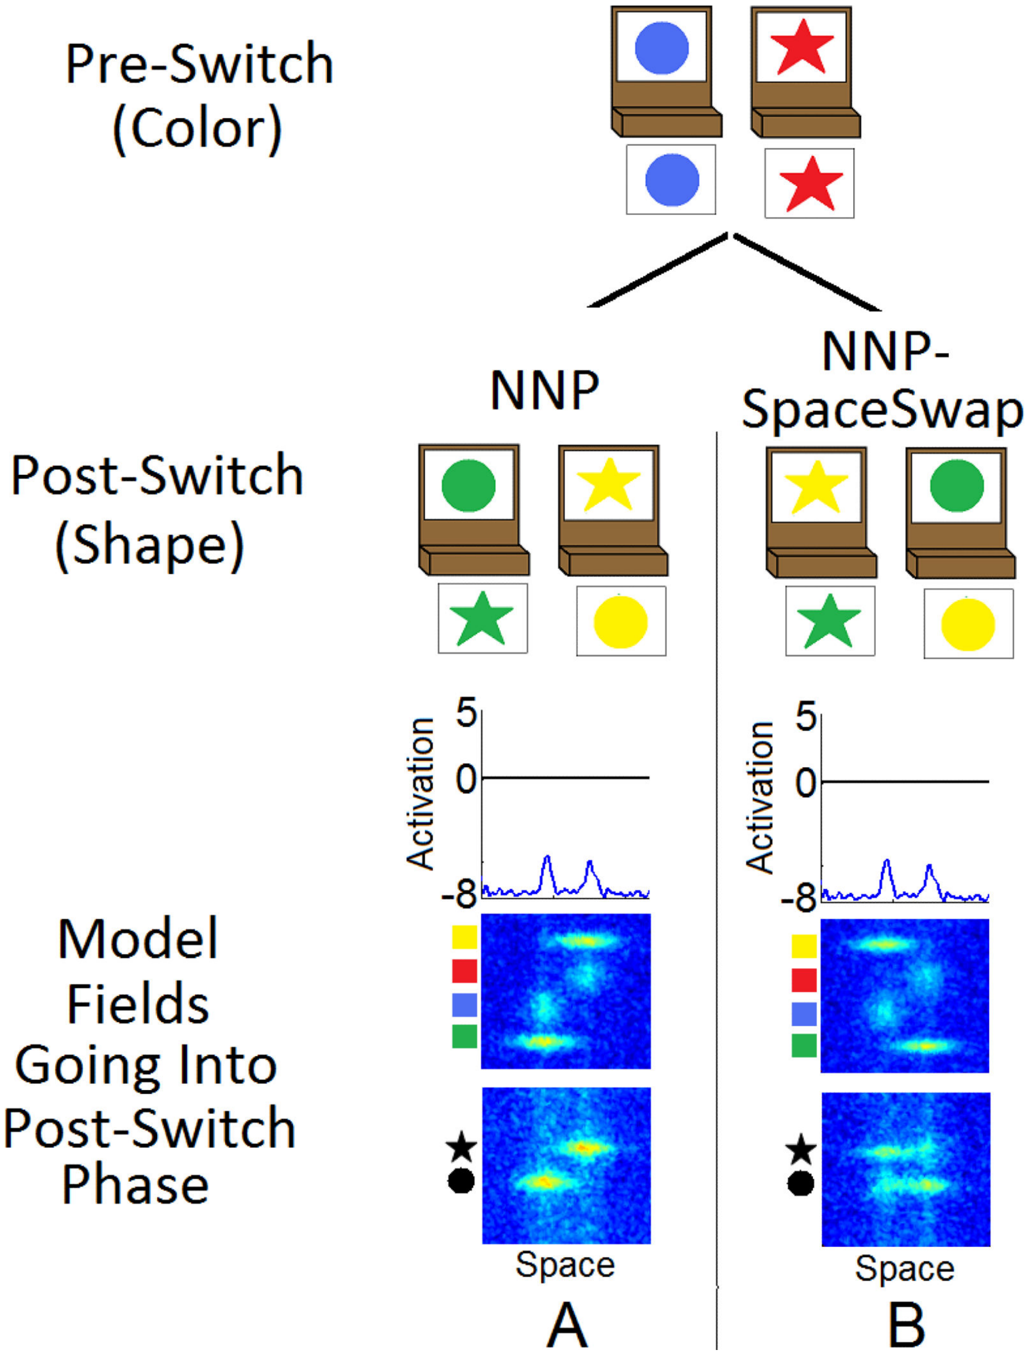

Figure 20.

Model at the start of the pre-switch phase for the NNP and NNPS versions. In this example, color is the pre-switch dimension and shape is the post-switch dimension. In panel B conflict is introduced in the post-switch (shape) WM field swapping the locations of the target cards.

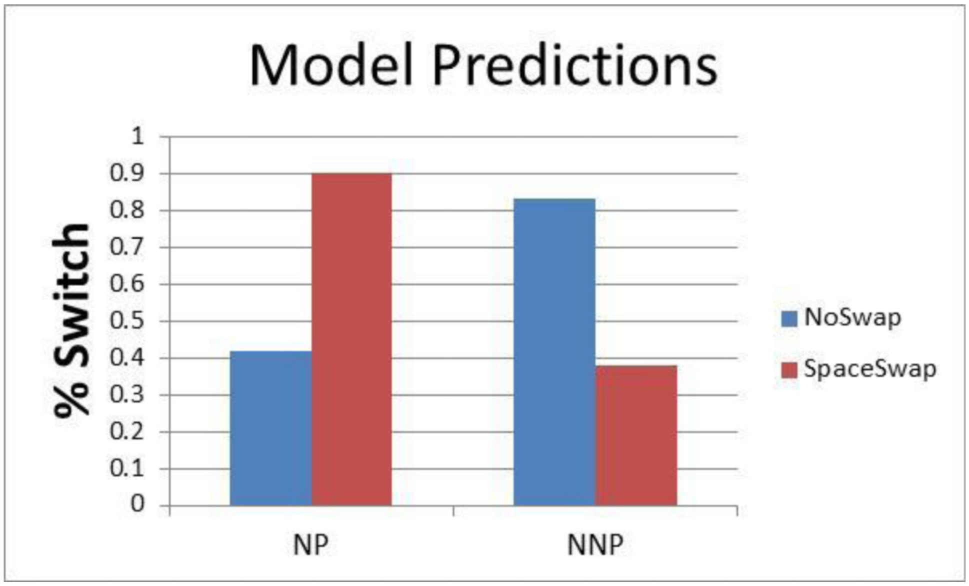

Figure 21.
Simulated predictions of 3-year-olds' performance in the SpaceSwap versions.

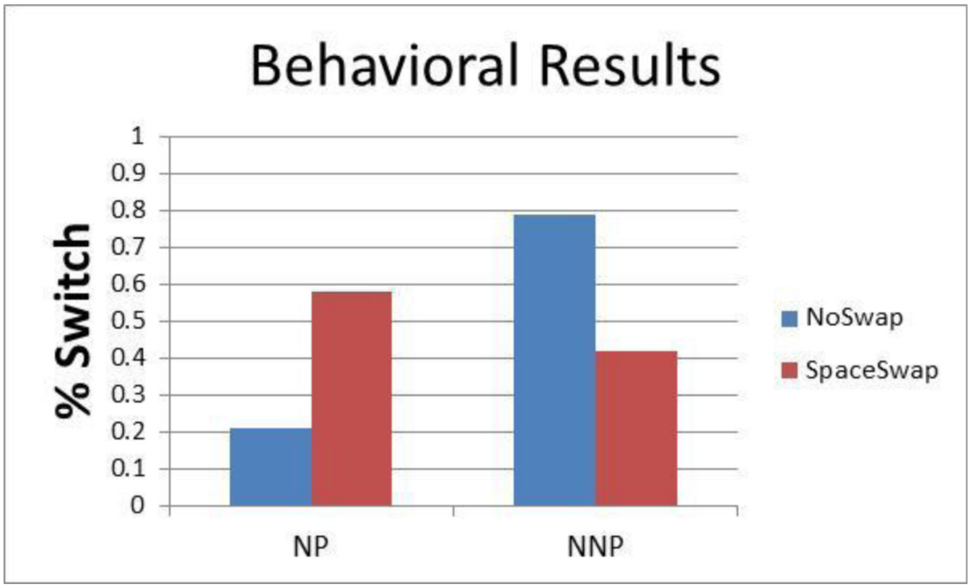

Figure 22.

Children's performance in the NP and NNP versions plotted as a function of whether the target cards swapped locations for the post-switch phase.

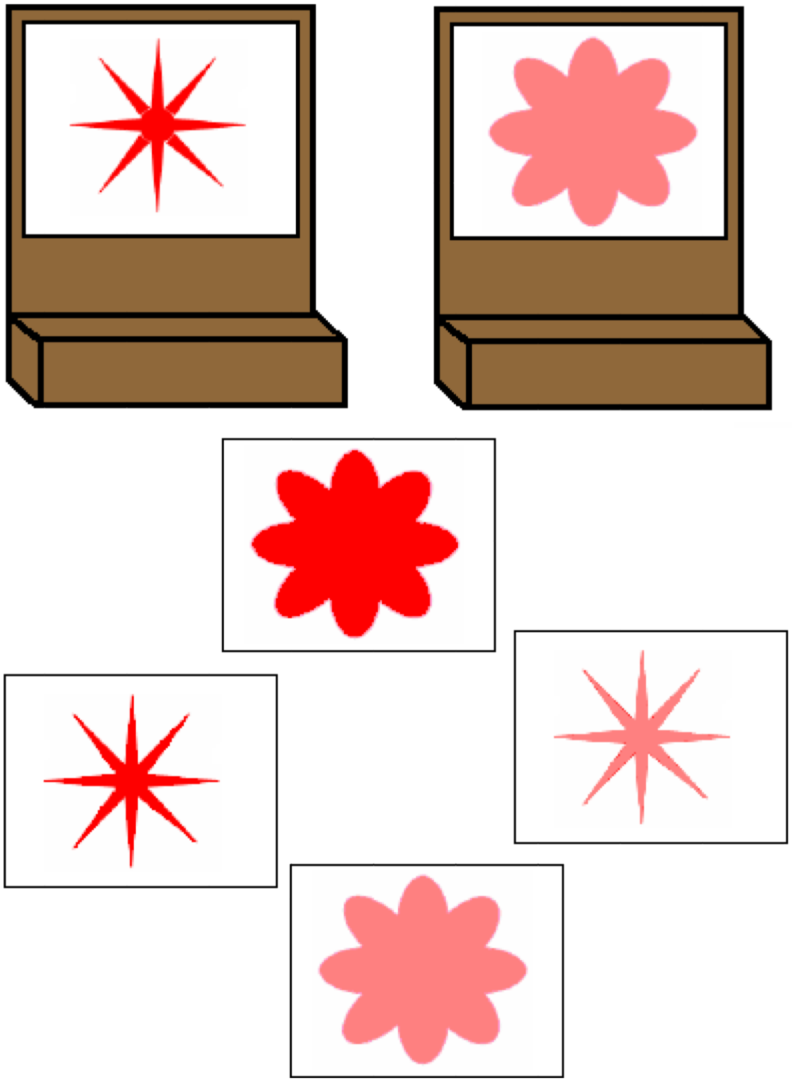

Figure 23.
Cards used in the feature-saliency version by Fisher (2011).

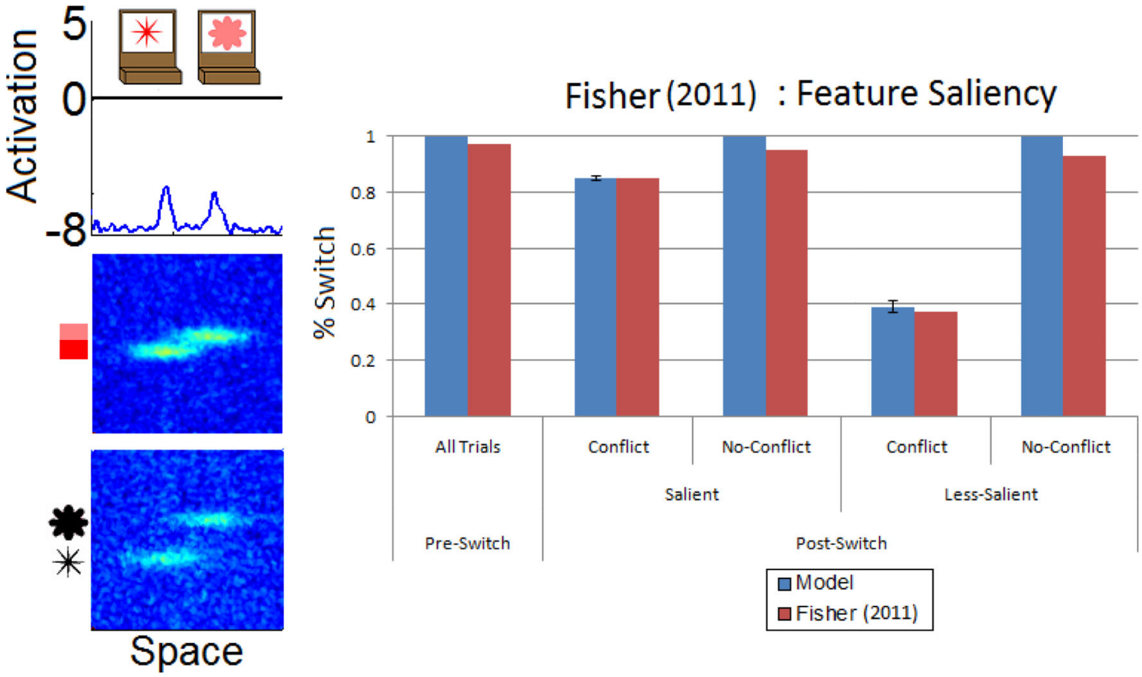

Figure 24.

Inputs to the object WM model and quantitative fits of the feature-saliency versions in Fisher (2011).

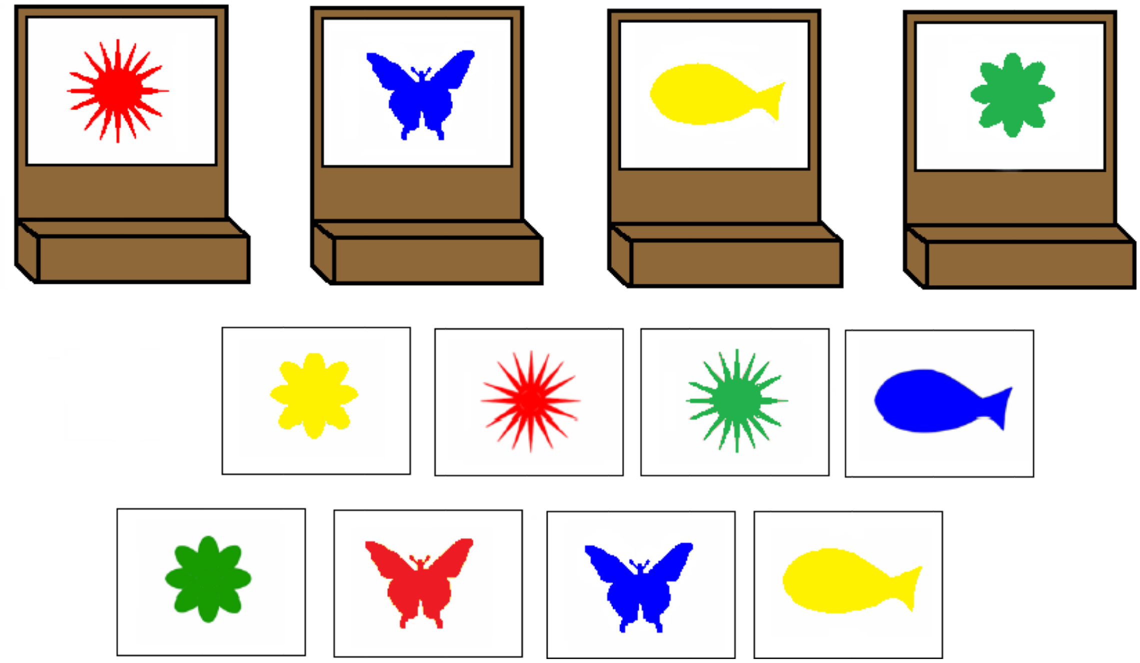

Figure 25.
Cards used in the feature-weights version by Fisher (2011).

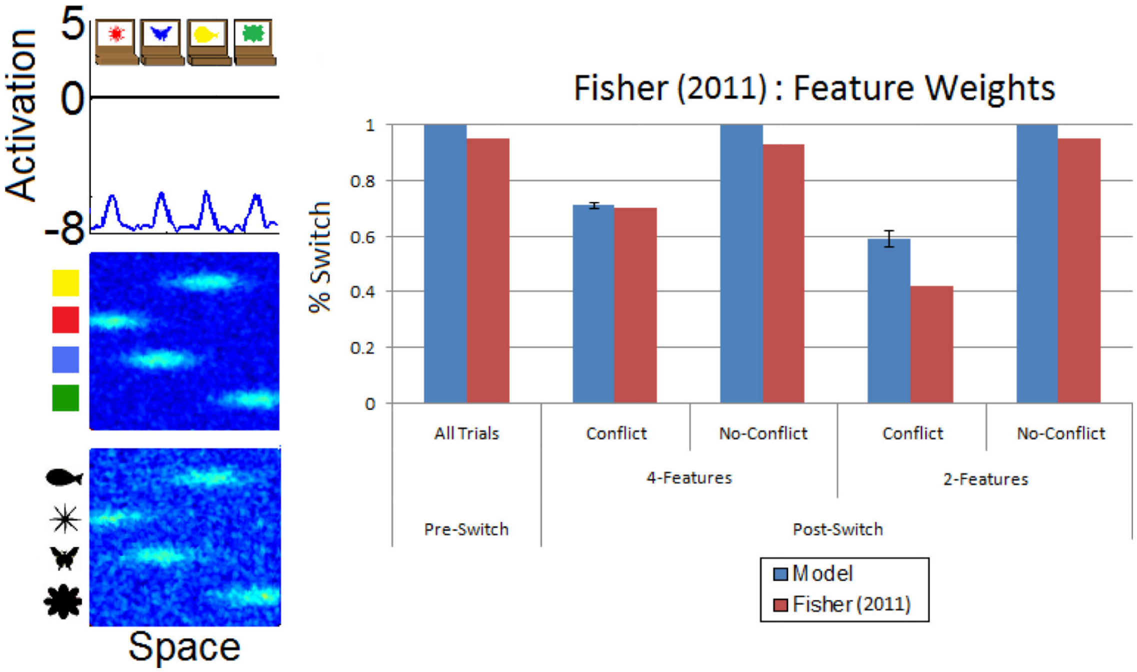

Figure 26.

Inputs to the object WM model and quantitative fits of the attentional-weights versions in Fisher (2011).

Table 1

Summary of Data Points and Number of Trials Simulated For Each Version of the DCCS.

Condition	Data points	Trials	Total Trials
Standard 3-yo	2	12	3600
Standard 4-yo	2	12	3600
NST	2	12	3600
NP 3-yo	2	12	3600
NP 4-yo	2	12	3600
NNP	2	12	3600
PC	2	12	3600
TC	2	12	3600
RC	2	12	3600
Training	3	18	5400
2-Features	4	16	4800
4-Feautres	4	16	4800
To-Salient	4	16	4800
To-Less-Salient	4	16	4800
Total	**37**	**190**	**57000**

Table 2

Field Parameters and Lateral Interaction Strengths

Resting Level	Lateral Excitation	Lateral Inhibition	Global Inhibition
$h_{ws} = -4$	$c_{ww} = 0.9$	$c_{wvs} = 1.75$	$c_{wvsg} = 0.4$
$h_{vs} = -4$	$c_{vws} = 1.35$		
$h_{wf1} = -7$	$c_{wwf1} = 0.7$	$c_{wvf1} = 0.3$	$c_{wvf1g} = 0.5$
$h_{vf1} = -4$	$c_{vwf1} = 0.65$		
$h_{wf2} = -7$	$c_{wwf2} = 0.7$	$c_{wvf2} = 0.3$	$c_{wvf2g} = 0.5$
$h_{vf2} = -4$	$c_{vwf2} = 0.65$		

Note: The time constant for all WM fields (w) is $\tau = 40$ while the time constant for all inhibitory fields (v) is $\tau = 5$. The lateral excitation width are $\sigma_{wws} = \sigma_{wwf1} = \sigma_{wwf2} = 3$ and the lateral inhibition width are $\sigma_{wvs} = \sigma_{wvf1} = \sigma_{wvf2} = 20$. The sigmoid steepness for all if the local-excitatory and laterally-inhibitory interactions within the fields is $\beta_{wws} = \beta_{wvs} = \beta_{vws} = \beta_{wwf1} = \beta_{wvf1} = \beta_{vwf1} = \beta_{wwf2} = \beta_{wvf2} = \beta_{vwf2} = 5$. The parameters for the Hebbian layer are: $HL_{build} = 500$, $HL_{decay} = 2000$, $c_{HL} = 0.065$, $\sigma_{HLfeature} = 5$, $\sigma_{HLspace} = 10$

Table 3

Parameters of Spatial Interactions between Fields

Excitatory Strength	Kernel Width
$c_{\mathrm{wswf1}} = 0.2$	$\sigma_{\mathrm{wswf1}} = 2$
$c_{\mathrm{wswf2}} = 0.2$	$\sigma_{\mathrm{wswf2}} = 2$
$c_{\mathrm{wf1ws}} = 0.1$	$\sigma_{\mathrm{wf1ws}} = 2$
$c_{\mathrm{wf2ws}} = 0.1$	$\sigma_{\mathrm{wf2ws}} = 2$
$c_{\mathrm{wf1wf2}} = 0.35$	$\sigma_{\mathrm{wf1wf2}} = 2$
$c_{\mathrm{wf2wf1}} = 0.35$	$\sigma_{\mathrm{wf2wf1}} = 2$

Note: The sigmoid steepness for all interactions between fields is $\beta_{\mathrm{wsf1}} = \beta_{\mathrm{wsf2}} = \beta_{\mathrm{wf1ws}} = \beta_{\mathrm{wf2ws}} = \beta_{\mathrm{wf1wf2}} = \beta_{\mathrm{wf2wf1}} = 1$.

Table 4

Parameters for the Dimensional Attention Nodes

3-yo Model	**4-yo Model**
$h_\text{i} = -4$	$h_\text{i} = -4$
$c_\text{ii\_excite} = 1$	$c_\text{ii\_excite} = 5$
$c_\text{ii\_inhib} = 5$	$c_\text{ii\_inhib} = 20$
$c_\text{iw\_colorcolor} = 0.005$	$c_\text{iw\_colorcolor} = 0.3$
$c_\text{iw\_colorshape} = 0.0025$	$c_\text{iw\_colorshape} = 0.001$
$c_\text{iw\_shapeshape} = 0.005$	$c_\text{iw\_shapeshape} = 0.3$
$c_\text{iw\_shapecolor} = 0.0025$	$c_\text{iw\_shapecolor} = 0.001$
$c_\text{wi\_colorcolor} = 1$	$c_\text{wi\_colorcolor} = 3$
$c_\text{wi\_colorshape} = .5$	$c_\text{wi\_colorshape} = .1$
$c_\text{wi\_shapeshape} = 1$	$c_\text{wi\_shapeshape} = 3$
$c_\text{wi\_shapecolor} = .5$	$c_\text{wi\_shapecolor} = .1$