Fall 2013

Closing the developmental loop on the behavioral and neural dynamics of flexible rule-use

Aaron Thomas Buss
University of Iowa

Recommended Citation

Buss, Aaron Thomas. "Closing the developmental loop on the behavioral and neural dynamics of flexible rule-use." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.
http://ir.uiowa.edu/etd/4949.

CLOSING THE DEVELOPMENTAL LOOP ON THE BEHAVIORAL AND NEURAL DYNAMICS OF FLEXIBLE RULE-USE

by

Aaron Thomas Buss

A thesis submitted in partial fulfillment
of the requirements for the
Doctor of Philosophy degree in Psychology
in the Graduate College of
The University of Iowa

December 2013

Thesis Supervisor: Professor John P. Spencer

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Aaron Thomas Buss

has been approved by the Examining Committee
for the thesis requirement for the Doctor of Philosophy
degree in Psychology at the December 2013 graduation.

Thesis Committee: __________________________________
John P. Spencer, Thesis Supervisor

____
Eliot Hazeltine

____
Vincent Magnotta

____
Bob McMurray

____
Larissa Samuelson

To my mom and dad.

The brain is wider than the sky, for, put them side by side, the one the other will contain with ease, and you beside.

Emily Dickenson

ACKNOWLEDGMENTS

I would like to thank the people and families who participated in the experiments reported here. I would especially like to thank Nicholas Fox for helping in all aspects of data collection, subject recruitment, and study management. This work was supported by National Science Foundation BCS-1029082 and National Institute of Health DA031583.

ABSTRACT

Executive function (EF) is a central aspect of cognition that undergoes significant changes in early childhood. Changes in EF in early childhood are robustly predictive of academic achievement and general quality of life measures later in adulthood. Here, I develop a dynamic neural field (DNF) model 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 5-year-olds, lack the flexibility to do so and perseverate on the first set of rules when instructed to switch. In *Study 1,* I use the DNF model to integrate behavioral and neural processes by simulating hemodynamics associated with the early emergence of flexible rule-use. I then test predictions of the model using near-infrared spectroscopy. In *Study 2,* I develop a DCCS task that can be used with adults that sheds light on key aspects of the task as they have been revealed with children. Using fMRI, a pattern of behavioral and neural effects shed light on the mapping of model components to neural regions. These two studies demonstrate that performance emerges as a property of system-wide interactions and that common neurocognitive effects are evident between childhood and adulthood.

TABLE OF CONTENTS

LIST OF TABLES..viii

LIST OF FIGURES...ix

CHAPTER

1. THE MULTIPLE COMPONENTS OF EXECUTIVE FUNCTION..........1

2. THE DCCS: BEHAVIORAL AND NEURAL EVIDENCE......................6

    Review of Behavioral Evidence..8
        The role of inhibition...8
        The role of working memory..9
        The role of task switching..9
        The role of feedback and demonstration cues...............................10
        The role of conflict...10
        The role of feature binding..11
        Review of neuroimaging studies..11
        Summary of behavioral and neural data.......................................12

3. THE DFT OF THE DCCS..15

    Current Theories of the Development of Executive Function.................17
        Cognitive complexity and control theory......................................17
        Other conceptual accounts..20
        Connectionist model...22
        Summary of theories...24
    Basic Concepts of DFT..25
    A DNF Model of object WM...27
    Autonomous Dimensional Attention and Rule-use..................................31
    The DNF model of and the DCCS task...34
        Summary of the model..39
    Simulating the Behavioral Dynamics of the DCCS Task.........................43
    Method...45
        Procedure..45
        Results..46
            Standard...46
            No-conflict Standard..47
            Total Change...48
        Summary of quantitative fits...48
    Conclusions...51

4. STUDY 1: USING NIRS TO PROBE MECHANISMS OF FLEXIBLE RULE-USE IN EARLY CHILDHOOD..61

Neural Dynamics of Flexible Rule-use in Early Childhood.....................61
Simulating Hemodynamics with Dynamic Neural Fields.........................65
Method: Simulating Moriguchi & Hiraki (2009).......................................68
Procedure...68
Results..69
Method: Generating Hemodynamic Predictions from DFT......................70
Procedure...70
Results..71
Testing Hemodynamic Predictions of the Model......................................73
Method..75
Participants..75
Stimuli and apparatus...75
Procedure..75
NIRS data collection..77
NIRS data processing...77
Behavioral results..79
fNIRS results...80
Discussion...82
Conclusions...87

5. STUDY 2: USING fMRI TO PROBE THE DYNAMICS OF FLEXIBLE RULE-USE IN ADULTHOOD...101

Neurocognitive dynamics of flexible rule-use in adulthood................102
Method...105
Participants..105
Behavioral procedure...106
fMRI data collection and analysis...106
Behavioral results..107
fMRI results...107
Discussion...109

6. GENERAL DISCUSSION..123

REFERENCES...129

LIST OF TABLES

Table

1. Developmental parameter differences in the dimensional attention system…...............89

2. Subject numbers and trial-counts across regions..90

3. Results of the omnibus ANOVA…..91

4. Factors contributing to DCCS conditions used in the fMRI study.............................114

5. Clusters of significantly different activation between switch and repeat trials...........115

6. Clusters of significantly different activation on switch trials across conditions.........116

7. Clusters of significantly different switch-effect activation across conditions…..........117

LIST OF FIGURES

Figure

1. Stimuli used in different versions of the DCCS task…..14

2. WM fields for the feature binding model…...52

3. The object WM model with the dimensional attention system....................................53

4. The sequence of events as the model sorts cards in the DCCS...................................54

5. A zoom-in on the critical dynamics of the attentional system on a rule-switch trial....55

6. Schematic diagram depicting the simplified processes of boosting and shifting.........56

7. H-boost distribution used for 3- and 4-year-old models...57

8. Shift-value distributions for 3- and 4-year-old models……..58

9. Model results………………………...59

10. Test of a behavioral prediction from the DNF model……..60

11. Hemodynamics of the DCCS...92

12. Model hemodynamic predictions…...93

13. Conditions used in the fNIRS experiment…...94

14. Depiction of the experimental display used in the NIRS study.................................95

15. fNIRS probe………...96

16. Behavioral data from fNIRS study……..97

17. Oxy by condition interactions……………...98

18. Oxy main-effects……………………..99

19. Comparing switchers and perseverators..100

20. Design of fMRI experiment…………...118

21. Stimulus dimensions and sequence of events used in fMRI study ………………..119

22. Behavioral data from the fMRI study……...120

23. Switch - Repeat contrasts for each condition...121

24. Switch-effect contrasts between conditions……..122

CHAPTER 1

THE MULTIPLE COMPONENTS OF EXECUTIVE FUNCTION

Early childhood is a time of rapid change in the organization of cognition. The period between 2 and 5 years is particularly dramatic, including the transition into formal schooling, the acquisition of language and mathematical abilities, learning to take the perspective of others in social interactions, and learning to appropriately adapt behavior across different contexts (Bull & Scerif, 2001; Frye, Zelazo, & Palfai, 1995; Kochanska, Coy, & Murray, 2001; Mazzocco & Kover, 2007; Samuelson & Smith, 1999; Zelazo, Muller, Frye, & Marcovitch, 2003). This developmental period is also marked by dramatic changes in executive function (EF). EF is an umbrella term that refers to the processes that allow individuals to rise above the exigencies of the environment, habits, or internally prepotent behaviors to behave in a contextually appropriate and goal-driven manner.

EF is an important topic of study in early childhood because it has widespread influences on the organization of behavior and behavioral control. For instance, improvements in EF have a positive impact on language development, and deficits in executive control have been linked to specific language impairment (Im-Bolter, Johnson, & Pascual-Leone, 2006; McEvoy, Rogers, & Pennington, 1993). Further, high levels of EF confer an initial advantage in mathematical and reading proficiency that has a facilitative effect on development through the early school years (Bull & Scerif, 2001; Mazzocco & Kover, 2007). Aspects of EF have also been linked to theory of mind and perspective taking which require children to suppress their own perspective and adopt the perspective of others (Carlson, Moses, & Breton, 2002; Frye et al., 1995; Hughes & Ensor, 2011). Finally, children with ADHD and autism show deficits in various aspects of EF, displaying weaker inhibitory control, a poorer ability to maintain information in WM, and greater difficulty switching tasks (Corbett, Constantine, Hendren, Rocke, & Ozonoff, 2009; Happé, Booth, Charlton, & Hughes, 2006; McEvoy et al., 1993;

Pennington & Ozonoff, 1996), The exact role of EF deficits in these pathologies, however, remains debated (Liss et al., 2001).

Critically, individual differences in EF early in development can produce long-lasting effects. Data show that enhancing EFs in early development can enhance school performance and reduce the prevalence of psychopathology (Diamond & Lee, 2011; Liss et al., 2001; Pennington & Ozonoff, 1996). Indeed, data suggest that EFs are more important for school readiness than IQ (Blair & Razza, 2007), in part, because EFs predict math and reading competence throughout the school years (Gathercole, Pickering, Knight, & Stegmann, 2004). EFs remain important into adulthood, predicting career and marriage satisfaction and positive mental and physical health (Dunn, 2010; Eakin et al., 2004; Prince et al., 2007). Reversely, children 3-11 years with poorer EFs have worse health, earn less, and commit more crimes as adults, even when controlling for IQ, gender, and social class (Moffitt et al., 2011).

Initial theories of EF proposed that the emergence of cognitive flexibility reflected developmental changes in a central executive system—a central resource that controls other aspects of cognition (Baddeley, 1986; J Duncan, Emslie, Williams, Johnson, & Freer, 1996; John 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, a large region anterior to the precentral sulcus. Lateral pre-frontal cortex is one of the slowest developing brain regions (Giedd et al., 1999) and evidence from patient populations (Baddeley, Della Sala, Papagno, & Spinnler, 1997; Milner, 1963) and single-unit neurophysiology (Asaad, Rainer, & Miller, 2000; Rao, Rainer, & Miller, 1997) showed that impairments of PFC leads to behaviors that mimic the performance of young children (Dempster, 1992; Diamond, 2002).

Factor analytic models suggest, however, that cognitive control and flexibility do not reflect the operation of a single resource. Rather, EF has multiple distinct components. Generally, EF is thought to involve the *inhibition* or suppression of irrelevant information or inappropriate actions, the stable maintenance or representation of information in *working memory* over time in a way that prevents interference or distraction, and the flexible updating or *switching* of cognitive processes to meet new goals (Collette et al., 2005; Davidson, Amso, Anderson, & Diamond, 2006; Garon, Bryson, & Smith, 2008; Lehto, Juujarvi, Kooistra, & Pulkkinen, 2003; Miyake et al., 2000). Factor analytic approaches have supported this multi-component view of EF, showing that different EF tasks load on inhibition, working memory, and task switching in different ways (Huizinga, Dolan, & van der Molen, 2006; Lehto et al., 2003; Miyake et al., 2000; however, see Wiebe et al., 2011). The review of the development of EF below is organized around these three component processes. Note that these labels are used simply to distinguish different aspects of functionality that executive control achieves (Zelazo et al., 2003), rather than making strong claims that these processes capture all aspects of EF. Ultimately, the real challenge in explaining EF is to understand the numerous ways cognitive control can emerge in specific tasks from a complex, multi-component system.

The multi-component nature of EF revealed by behavioral studies is also reflected in neural evidence. Data from neuroimaging studies have shown an extensive network of regions within frontal and posterior cortical areas that are involved in EF (Fair et al., 2007, 2008; Morton, 2010; Postle, 2006). Cognitive control and flexibility are thought to emerge from interactions within this system-wide network. One useful tool in understanding the neural basis of EF is resting state connectivity analysis, in which the endogenous fluctuations in baseline activity are correlated across voxels to reveal functionally connected regions. Fair and colleagues (Dosenbach et al., 2007; Fair et al., 2007, 2008) have used this approach to identify distinct networks involved in different aspects of cognitive control. One network is composed of frontoparietal connections across regions such as dorsolateral prefrontal cortex, intraparietal sulcus, and precuneus.

This network is hypothesized to be involved in trial-to-trial adaptation, task-initiation, and error adjustment. The other network is composed of cinguloopercular connections across regions thought to be involved in the stable maintenance of task-sets, including anterior prefrontal cortex, anterior cingulate cortex, anterior insula, and ventral prefrontal cortex, along with sensory areas in occipital and temporal cortex (Dosenbach et al., 2007).

Although neuroimaging data have shed important light on the neural systems involved in EF, they have also revealed new complexities. For instance, it is not clear how the different component processes of EF identified in the behavioral literature map onto the functional networks identified using neuroimaging techniques. The story gets even more complex when considering development. Each component process of EF has its own developmental trajectory with complex interactions among the processes.

In this thesis, I present a new theory of the development of executive function—Dynamic Field Theory (DFT)—that spans 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. I 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 thesis is divided into 4 chapters. In the next chapter, I review the behavioral and neural literature on the DCCS. In Chapter 3, I present a dynamic field theory (DFT) of executive function, discussing how this theory captures performance in key versions of the DCCS task. Moreover, I discuss behavioral predictions of this model which have been successfully tested and presented elsewhere (Buss & Spencer, in press).

Chapter 4 shows that the DNF model can capture aspects of developmental changes in frontal cortex activation between 3 and 5 years. To further examine the neural processes of the model, I also present data from an experiment designed to test hemodynamic predictions of the DNF model with 3- and 5-year-olds. In Chapter 5, I probe the longer developmental timescale of changes in EF by using a DCCS task with adults and fMRI. The task mimics key properties of DCCS variants used with children. This study represents an initial step toward more fully exploring the neurocognitive dynamics involved in flexibly shifting dimensional attention and provides a means to anchor DFT to the gold-standard of neuroimaging—fMRI. The thesis concludes with a General Discussion in Chapter 6 that evaluates the DNF model and describes key challenges moving forward toward an integrated theory of the development of executive function at both behavioral and neural levels of analysis.

CHAPTER 2

THE DCCS: BEHAVIORAL AND NEURAL EVIDENCE

The DCCS task has been extensively used over the last two decades to probe the development of cognitive flexibility in early childhood. Given the important and predictive role EF plays in early childhood (Blair & Razza, 2007; Diamond & Lee, 2011; Liss et al., 2001; Moffitt et al., 2011; Pennington & Ozonoff, 1996), this task provides an ideal starting point to develop a neural process account of EF. In this task, children are instructed to 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 for 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 early development of EF for three reasons. First, performance in the DCCS involves not only ‘switching’, but also builds upon ‘inhibition’ and ‘working memory’ which develop in some forms before flexible rule-use (Carlson, 2005; Garon et al., 2008). In particular, this task requires inhibition to suppress processing of 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). Further, ERP data suggest that multiple control processes unfold over the course of a trial, suggesting that the task does not tap into only a switching component (Waxer & Morton, 2011).

Second, this task reveals rapid and dramatic changes in children’s executive function in early development. Although 5-year-olds have little trouble switching rules, the majority of 3-year-olds (typically around 70%) 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. The dramatic nature of this developmental shift has led to an intensive investigation of why young children perseverate, creating a vast literature on the DCCS task. This literature has revealed complex and subtle aspects of children's rule-use, explaining why single-cause explanations (e.g., inhibitory control) fall short of capturing the full range of behavioral effects and why a formal, computational approach is warranted. Simply put, children fail to switch rules in very specific ways, and a successful integrative theory must capture these details. For example, 3-year-olds do not have trouble using the initial set of rules. Rather, difficulty only arises in particular circumstances during the post-switch phase. Further, children's post-switch performance is typically all-or-none—that is, they either get all of the trials correct or all of the trials incorrect. In addition, although perseveration is robust in many variants of the task, 3-year-olds are able to switch rules under specific conditions. Ultimately, the pattern of success and failure across different manipulations reveal critical details that a theory must capture.

The third reason why the DCCS task provides an ideal foundation for developing a theory of EF harkens back to the long history of trying to integrate behavioral and neural data in the EF literature. In addition to a rich behavioral data set in the DCCS literature, there is an emerging neuroimaging literature that has shown rapid and dramatic changes in children's neural dynamics in early development. This work has 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, Bosma, & Ansari, 2009). Thus, this simple task can provide insight into the wide spectrum of EF processes *and* associated neural changes in early childhood.

Below, I review the literature on the DCCS task, starting with a review of behavioral studies organized around the different components of executive function as well as other themes highlighted in the literature. I then discuss neuroimaging studies on the early development of EF. Finally, I survey existing theories of the DCCS task.

Review of Behavioral Evidence

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.

In a Negative Priming version, the features that were *relevant* for the pre-switch phase (e.g., color) are changed for the post-switch phase (see Figure 1B). With the pre-switch features changed, the stimuli no longer afford being sorted by the pre-switch values; thus, it is not required that children inhibit those values or rules during the post-switch phase. However, around 60% of 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, Dick, Gela, Overton, & Zelazo, 2006; Zelazo et al., 2003). That is, 3-year-olds appear to have too much inhibition of the post-switch dimension (e.g., shape) that they cannot overcome. In a Partial-Change version (Zelazo et al., 2003), the reverse issue is probed, that is, the features that were *irrelevant* during the pre-switch are changed in the post-switch phase (e.g., shape; 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, more than half of 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 post-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 persistent perseveration across these two conditions leaves one wondering whether *any* changes in the features of the cards can improve performance in this task.

Results from a Total-Change version speak to this issue: when the features of both dimensions are changed, 3-year-olds can reliably switch to sorting by the other dimension (Zelazo et al., 2003; see Figure 1D). 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 or rules 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 game while familiar features and labels are used for the post-switch game. In either case, it appears that an informative label can be maintained in working memory 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 a card in the post-switch (Deák, Ray, & Pick, 2002), 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 Negative Priming 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 1E-F). 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 neuroimaging studies

Research on the DCCS task has focused on early childhood, typically the age range from 3 to 5 years. This is time period during which it is difficult to collect neuroimaging data. This is because the gold standard of neuroimaging methods—fMRI—requires participants to lie still for extended durations, making it unsuitable for 3-year-olds. Consequently, researchers have turned to a new neuroimaging method—fNIRS—to study the early development of EF. In the one study of this type, Moriguchi and Hiraki (2009) used fNIRS to examine changes in frontal cortex activation as young children engaged in the standard DCCS task. They found increases in inferior pre-frontal cortex activity between 3 and 5 years associated with rule-shifting in the DCCS. In particular, 3-year-olds and 5-year-olds that were able to switch rules also showed significantly stronger frontal activation compared to 3-year-olds who perseverated. This suggests that changes in the engagement of the frontal cortex support the early emergence of flexible rule-use.

Another way to examine the neural processes that underlie performance in the DCCS task is to use fMRI with older children and adults. Morton and colleagues (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. Further, between adolescence and adulthood, there are increases in activation in superior parietal cortex, superior frontal sulcus, and fusiform gyrus. In addition to this study, Nagahama et al. (2001) used fMRI with adults and compared dimensional shifts of attention with simple response reversals for the features (akin to the ‘silly’ version reviewed above; Kloo, et al., 2008). Shifting rules, regardless of whether attention was also being shifted, activated inferior frontal sulcus. However, middle frontal gyrus was selectively activated when shifts of attention were required.

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 brain regions 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 studies examining the neural dynamics of perseveration and switching in the DCCS have shown increases in frontal activation when 3- and 5-year-olds switch rules. In addition, fMRI data with adults show 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 chapter, I describe a dynamic field theory of the development of executive function and flexible rule-use. I show how the theory captures the many nuances evident in children's behavioral data and discuss specific behavioral predictions that have been successfully tested elsewhere (Buss & Spencer, in press).

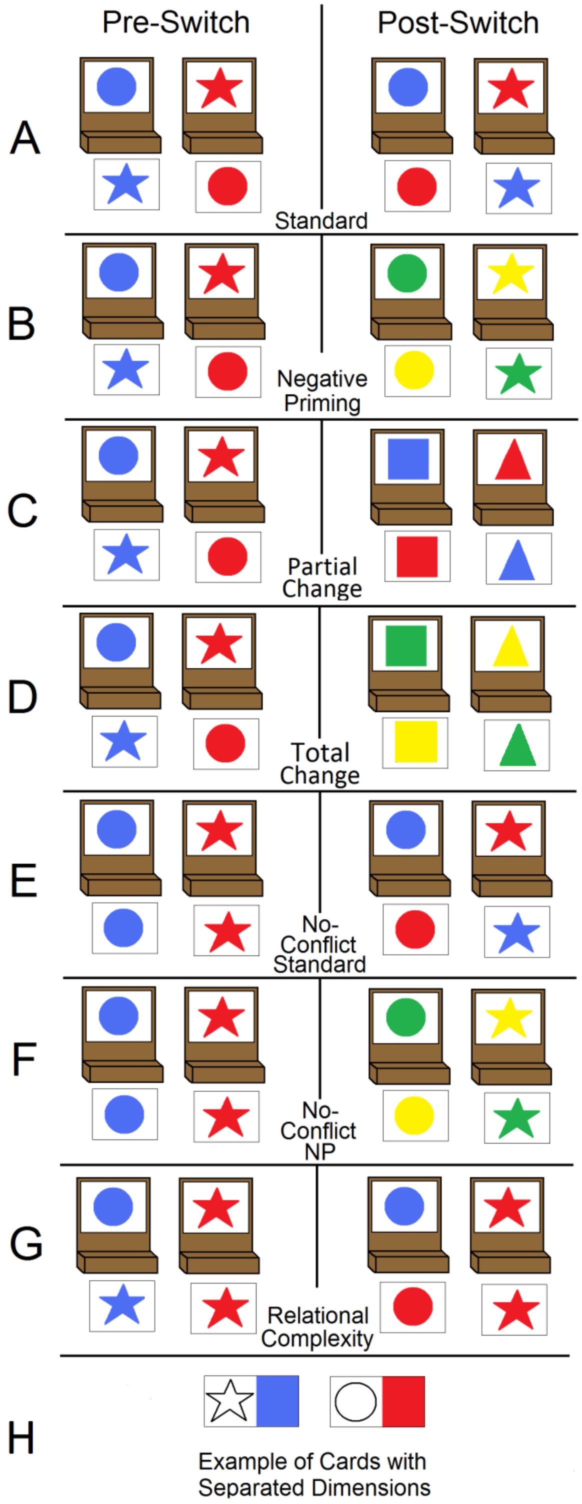

Figure 1: Target and test cards used in various versions of the DCCS.

CHAPTER 3

THE DFT OF THE DCCS

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, Michael Turvey and others (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 topics 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 spatial category biases to changes in the metric precision of spatial working memory from childhood to adulthood (Schutte, Spencer, & Schoner, 2003; Simmering, Schutte, & Spencer, 2008). More recently, this theoretical framework has 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; Spencer, Schneegans, & Schoner, in press), 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, Smith, Perry, & Spencer, 2011).

Importantly, the model presented here was not derived solely as a model of rule-use; rather, it is an extension of an object WM model (Johnson et al., 2008; Samuelson et al., 2011; Spencer et al., in press) to which I add an autonomous dimensional attention system. Rule-use, then, emerges in the neural interaction between dimensional attention nodes (a ‘shape’ or ‘color’ node) and an object WM system that localizes features in space and makes decisions about where objects should be placed. I show how this interactive system captures young children’s performance in the DCCS task. I then use the 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 5 years.

In the following sections, I first set the stage for this new theory by describing current theories of children’s performance in the DCCS task. Next, I describe the DNF model. It is important to note that terms such as neural activation, cortical fields, and excitatory/inhibitory neural interactions are referring to the simulated neural dynamics of the model. However, in Chapter 4, I build upon DFT’s long history of interfacing directly with neural data (Bastian, Riehle, Erlhagen, & Schöner, 1998; Bastian, Schoner, & Riehle, 2003; Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999; Jancke et al., 1999; Markounikau, Igel, Grinvald, & Jancke, 2010; McDowell, Jeka, Schoner, & Hatfield, 2002; Spencer, Barich, Goldberg, & Perone, 2012) by simulating hemodynamic data in early childhood and testing neural predictions from the model.

To describe the DNF model, I begin with an overview of several central concepts in DFT, including a discussion of neural population dynamics within multi-layered cortical fields, interactions between different cortical fields, as well as the variant of Hebbian learning used to capture changes in neural dynamics over a trial-to-trial timescale (Faubel & Schöner, 2008; Lipinski, Spencer, & Samuelson, 2010; Samuelson et al., 2011; Simmering et al., 2008). Next, I describe the object WM model (Johnson et al., 2008; Spencer et al., in press; see also, Samuelson et al., 2011). This is followed by an overview of the dimensional attention system that is couple to this model, as well as a discussion of the developmental hypotheses explored in the present study. I then 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, I present quantitative fits across different variations of the DCCS task, capturing performance across situations where 3-year-olds either perseverate or pass in the post-switch phase.

Current Theories of the Development of Executive Function

The review of the empirical literature on the DCCS task in Chapter 2 reveals a complicated pattern of results where *everything seemingly matters*. The DCCS, then, presents a formidable theoretical 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). Below, I 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 consciously 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 phase, 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 in some way. The active nature of this decision was captured in a subsequent revision of CCC theory, CCC-r, which incorporated the dynamics of activation and inhibition of rule sets (Zelazo et al., 2003). Specifically, one set of rules is used in the pre-switch phase 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. Thus, 3-year-olds 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, they lack the ability to integrate both branches of rules under the setting conditions at a second level in the hierarchy. 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).

Using these concepts of rule-representation and activation, CCC-r 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 actively selecting a set of rules against a competing alternative. Thus, in the no-conflict Standard (Figure 1D) and No-Conflict Negative Priming 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/CCC-r 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. Three-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/CCC-r theory generalizes to a broad range of effects, it has critical limitations as well. For example, CCC/CCC-r theory does not specify how children learn to construct complex rule-hierarchies over development. Although this theory proposes that rule-use is grounded in emergence from more basic process, the theory is focused on the conscious decisions made by the child and does not attempt to explain the real-time process of forming and engaging rules. In this sense, although CCC/CCC-r 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, 2006). Further, CCC/CCC-r falls short of accounting for different aspect of EF because it explains performance in the DCCS using a specialized rule-switching system. Finally, there are empirical data that highlight a fundamental limitation of an abstract rule representation framework. Specifically, 3-year-olds can more easily switch to attentionally salient, distinctive features than less salient, perceptually similar features (Fisher, 2011; Honomichl & Chen, 2010). In this case, the rule-hierarchy is the same, but performance is influenced by the attentional salience of the features involved in the different tasks. Thus, CCC/CCC-r theory does not specify how the salience or similarity of features, or attentional manipulations in general, influences reflection on the rule-structure, the representation of rule pairs, or the activation of rule pairs.

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 5-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 and colleagues (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 of 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 model

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, exerting a stronger influence on the decision-making process and facilitating 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 et al. (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. Finally, the studies examining the influence of dimensional labels or novel features reviewed above (Yerys & Munakata, 2006) were designed to probe the active/latent memory distinction in the model. Specifically, when the pre-switch game is simply called a 'sorting' game, or when unfamiliar features are used for the pre-switch game, weaker latent memories are established which poses less competition with the active memories during the post-switch phase.

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/CCC-r 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 Negative Priming 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 of theories

The theoretical perspectives reviewed here offer differing views of what a 'rule' is and how the ability to use rules changes over development. CCC/CCC-r 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 data. 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.

A primary aim in developing the DNF model described below is to achieve extensive theoretical coverage of young children's performance in the DCCS task early in development while adhering to, and anchoring our concepts in, formalized neural principles. This brain-behavior focus is inspired by 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 new ways to test and constrain developmental theories 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 simulated neurons 'tuned' to continuous metric dimensions (e.g., space, color, orientation; for related ideas, see, e.g., Georgopoulos, Schwartz, & Kettner, 1986). 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 (for discussion, see Spencer, Austin, & Schutte, 2012). 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., 2009; Edin, Macoveanu, Olesen, Tegnér, & Klingberg, 2007; Johnson, Spencer, & Schöner, 2009).

In this thesis, I use two-layered neural fields composed of an excitatory layer of neurons (called a working memory, WM, field) coupled to a layer of inhibitory interneurons (Inhib field). When activation reaches an activation threshold (>0), the WM field becomes self-excitatory—that is, activated neurons share excitation with close neighbors. Also, the activated neurons pass excitation to the Inhib field. Once these inhibitory interneurons are activated at above-threshold levels (>0), they 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.

Our DNF models also use a variant of Hebbian learning which allows neural populations to learn as they acquire a history in the task (see, e.g., Clearfield, Dineva, Smith, Diedrich, & Thelen, 2009; Lipinski, Simmering, Johnson, & Spencer, 2010; Perone, Simmering, & Spencer, 2011). These Hebbian traces vary in strength from 0 to 1, much like a “weight” or synaptic connection in a connectionist model. As Hebbian traces increase in strength, the affected neurons acquire a resting level closer to the activation threshold of 0 (i.e., the resting level becomes less negative). Thus, Hebbian traces 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 I use was initially developed to capture neural activation patterns within visual cortex (Amari, 1977; see also, Compte et al., 2000; Wilson & Cowan, 1972). 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 et al., 1998, 2003; Erlhagen et al., 1999; Jancke et al., 1999). A related approach has also been used with voltage-sensitive dye imaging (VSDI; see Markounikau et al., 2010). These approaches enable researchers to directly test DNF models using neurophysiological measures in non-human animals. Other efforts have shown that DNF models can be tested with humans using event-related potentials (see, Spencer et al., 2012). For instance, McDowell and colleagues (2002) examined the prediction of a DNF model of motor planning that movement direction and response probability should be interdependent. They tested this by measuring both reaction times and the amplitude of the P300 event-related potential. Both types of measures showed the predicted interactions. In summary, then, the differential equations used in the present report capture real, observable aspects of activation dynamics within populations of cortical neurons.

A DNF Model of Object WM

The model I propose here is based, in part, on the model of object working memory shown in Figure 2 (see Buss & Spencer, in press for full model equations). At the top of this figure, I show a picture of the brain, with several highlighted cortical regions whose function is captured by aspects of the model. The bottom shows a simulation of the model at different points in time as it 'binds' object features present in the visual display.

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; Mishkin, Ungerleider, & Macko, 1983): a dorsal ('where') pathway that is primarily concerned with encoding the spatial locations of objects (Andersen, 1995) and a ventral ('what') pathway that is composed of cortical fields that encode different object features such as color or orientation in distinct neural populations (Desimone & Gross, 1979; Xiao, Wang, & Felleman, 2003). This creates a 'binding' problem in vision (Treisman & Gelade, 1980; Treisman, 1996): 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 2 shows a concrete example of this challenge: when shown the blue square and the yellow diamond in the display in Figure 2A, how does the brain know that the blue hue (represented by the dark grey shading of the object in the Display panel) should be linked to the square shape given that these neural representations 'live' in different parts of the brain? Concretely, how does the brain know that the 'blue' neurons in the fusiform area (see light grey circle on brain image; Simmons et al., 2007) should be coupled to the 'square' neurons in the lateral occipital complex (see medium grey circle on brain image; 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 (Aggelopoulos & Rolls, 2005; DiCarlo & Maunsell, 2003; Kravitz, Vinson, & Baker, 2008; Op De Beeck & Vogels, 2000). Thus, it might be possible to anchor the featural representations in both cortical areas to a precise representation of spatial locations in, for instance, parietal cortex (see dark grey circle on brain image; Andersen, 1995; for a review, see Silver & Kastner, 2009). This is the solution that the DNF model implements.

The sequence of simulations in Figure 2 illustrates how the DNF model 'binds' visual features together. The top row of this figure shows the display of inputs that is presented to the model at three different time-points: at the onset of the display (A), after the objects have been consolidated or 'bound' in working memory (B), and after the display has been turned off and the trial has ended (C). The next row shows a one-dimensional spatial field that captures aspects of neural processing in the parietal cortex. This field consists of a layer of neurons with receptive fields sensitive to variations in the horizontal positions of objects in the display. Neurons that 'prefer' leftward locations in the display are on the left side of the spatial field, while neurons that 'prefer' rightward locations in the display are on the right side of the spatial field. The y-axis in this figure shows the activation of each neuron in the field. For instance, in panel A, the two bumps of activation show the initial response of spatially-tuned neurons to the presentation of an item to the left of center (the square) and a second item to the right of center (a diamond).

Below the spatial field in Figure 2A is a two-dimensional color-space field (which I will refer to as the 'color' field). This layer of neurons has receptive fields sensitive to the combination of spatial position and color (hue), like neural populations in the fusiform area (Simmons et al., 2007). Considered together, the entire two-dimensional field of neurons can represent any collection of colors at any location in the task space. The spatial tuning of the neurons is again shown along the x-axis, while the y-axis now displays the color tuning of the neurons (represented by the grey scale to the left of the field which shows the mapping of the RGB color spectrum to the y-axis). The level of activation of neurons in the color field is depicted by the grey shading, with higher levels of activation indicated by lighter shading (see activation scale at the bottom of the figure). As can be seen in the color field in Figure 2A, there is an activation bump (light grey oval) in the lower left region of the field that reflects the initial encoding of the blue (dark grey) square to the left, and a second bump of activation in the upper right region of the field that reflects the initial encoding of the yellow (light grey) diamond to the right.

Finally, the bottom panels in Figure 2 show the activation within a two-dimensional shape-space field (called the 'shape' field). Here, these neurons have receptive fields that are sensitive to spatial position and shape information as in the lateral occipital complex (Kourtzi et al., 2003). Like the color field, the x-axis shows the spatial tuning of the neurons while the y-axis shows the tuning along an abstract shape-similarity dimension. As before, the level of activation for these neurons is depicted by the grey shading (see activation scale at bottom). Thus, the activation bump in the upper left region of the field (light grey oval) reflects the initial encoding of the square on the left, and the activation bump to the lower right reflects the initial encoding of the diamond on the right.

In Figure 2A, all of the activation profiles are sub-threshold (i.e., activation < 0) because the objects were just presented in the display. Thus, strong neural interactions have not yet become engaged, and the model has not formed robust WM representations of the objects. Critically, the coarse spatial encoding properties of the feature-space fields produce overlap along the spatial dimension for the feature inputs. Left alone, these bimodal cortical fields would have difficulty binding the correct colors to the correct shapes (for a demonstration of this difficulty, see Johnson et al., 2008).

Figure 2B shows the model resolving this ambiguity through spatial coupling. Since all of the WM fields share a common spatial dimension, activation is coupled along this dimension (see bi-directional arrows between dorsal and ventral cortical areas in the brain image). As activation rises, spatial information is passed back and forth among the WM fields. Because the spatial field has precise spatial information, this helps resolve any spatial conflict in the feature-space fields. This is shown in Figure 2B which shows the fields after peaks have emerged. As can be seen, the model correctly binds ‘blue’ and ‘square’ on the left and ‘yellow’ and ‘diamond’ on the right. In particular, there are robust peaks to the left and right of the spatial field (second panel), at the blue hue value to the left and yellow hue value to the right in the color field (third panel), and at the square shape to the left and diamond shape to the right in the shape field (bottom panel). *Note that the shared spatial excitation can be seen in the vertical ridges within the feature-space fields.* Although activation is passing back-and-forth at the left and right spatial positions, this does not lead to explosive excitation at every sight along these ridges; rather, the lateral or surround inhibition associated with each peak in the feature-space fields keep excitation locally contained around the consolidated object feature values. This spatial coupling and associated feature binding will play a critical role in the simulations of the DCCS discussed below.

Figure 2C shows the fields after the inputs have been turned off and the “trial” is over. This reveals the contribution from the Hebbian process that operates in the model: there is now slightly elevated activation corresponding to the blue square that was on the left and the yellow diamond that was on the right. As can be seen, these memories are very coarse, elevating broad regions around sites associated with the WM peaks in Figure 2B. What is the consequence of the broad memory traces? When these or similar objects are presented again, peaks will build more readily at these particular feature-space conjunctions given the slightly elevated activation levels.

To summarize, in the DNF model, neural populations that code for different object features are bound together by a common spatial dimension anchored to the precise spatial representations found in the parietal cortex. Note that this solution to the binding problem shares elements with Feature Integration Theory proposed by Treisman and Gelade (1980). Further, as the model builds integrated WM representations across the different cortical fields, a Hebbian learning process operates to bring the activated sites closer to their activation threshold. This creates a type of priming effect when these objects are re-presented, causing neurons in the fields to become activated more quickly on subsequent presentations.

Autonomous Dimensional Attention and Rule-use

The model shown in Figure 2 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, I need to add an additional concept—dimensional attention. Figure 3 shows the complete model architecture which adds a frontal dimensional attention system to the object binding model show in Figure 3C. The brain image in Figure 3A adds a black circle highlighting frontal cortical regions implicated in executive function and the control of dimensional attention (see, e.g., Morton et al., 2009). This frontal system is reciprocally coupled to cortical fields in the ventral and dorsal pathways as indicated by the bi-directional arrows (Crone, Wendelken, Donohue, & Bunge, 2006; Fair et al., 2007).

The architecture of the dimensional attention system in the model is shown in Figure 3B. This architecture 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, such as shape or color. These nodes also learn as they are repeatedly activated in a task. In particular, the frontal system uses the same type of Hebbian process 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 on subsequent trials.

The information that each node actively represents is reflected in the pattern of connectivity between the frontal nodes and the object WM system. As is shown in the figure, the ‘color’ node has strong excitatory connections to the color field (solid bi-directional line) and weaker excitatory connections to the shape field (dashed bi-directional line). 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 object WM system based on color information. Reversely, the ‘shape’ node has strong excitatory connections to the shape field (solid bi-directional line) and weaker excitatory connections to the color field (dashed bi-directional line). 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 object WM system based on shape information. This pattern of connectivity is akin to connection weights in a PDP network. I assume that the system begins with randomized weights which become selectively mapped as the ‘shape’ dimensional neuron becomes associated with the shape field and the ‘color’ dimensional neuron becomes associated with the color field. Note that the pattern of connectivity in the model is fully reciprocal; thus, the growth of neural activity in the posterior neural fields (e.g., the shape and color fields) can also impact the activation of the frontal system. The activation dynamics within the model are consistent with neurophysiological evidence (Egner & Hirsch, 2005; Lepsien & Nobre, 2007; Zanto, Rubens, Bollinger, & Gazzaley, 2010) which has demonstrated boosts in baseline activity for cortical areas processing task-relevant information under situations with high executive demands.

In the present work, I probed whether specific changes in the connectivity pattern between the frontal and posterior cortical fields were sufficient to capture the emergence of rule-use over development in the DCCS task. In particular, I 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 weakly boosted the shape field. By contrast, the 5-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 Buss & Spencer, in press for details). As I discuss below, these changes in connectivity were critical in capturing key aspects of young children’s performance in the DCCS task. Conceptually, such changes are akin to changing the weight matrix in a PDP model over learning such that ‘shape’ and ‘color’ nodes become associated with a specific feature dimension.

In addition to refining the connectivity between the frontal and posterior neural systems in the model, I also implemented a specific developmental hypothesis central to other work using DFT—the spatial precision hypothesis (Schutte & Spencer, 2009; Spencer et al., 2009). According to this hypothesis, excitatory and inhibitory neural interactions become stronger over development. Thus, the strength of self-excitation and lateral inhibition for the ‘shape’ and ‘color’ dimensional units was increased. These changes create stronger and more selective activation of the 'shape' and 'color' units. This addresses the question of whether changes within the dimensional attention system are sufficient to capture the developmental pattern of behavior in the DCCS task.

Note that, although both developmental changes in the model—changes in the pattern of connectivity between frontal and posterior systems and changes in the strength of neural interactions in the dimensional attention system—were implemented 'by hand', other work suggests that these changes could emerge via an autonomous learning mechanism (see, Perone & Spencer, 2013). For instance, patterns of connectivity between nodes and fields can be learned by standard forms of Hebbian learning as the system learns the correlation between the use of the word 'color' and the importance of task-relevant hue values in the task space (Faubel & Schöner, 2008; Sandamirskaya & Schöner, 2010).

The DNF Model of the DCCS Task

Figure 4 shows the DNF 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 for the young model (i.e., 3-year-olds). There are 12 periods where the dimensional nodes become more active (creating the box-car profile) corresponding to the 6 pre-switch and 6 post-switch trials. The middle panels show the object WM system at particular points during the task indicated by the time labels above each panel (which correspond to marked times on the top panel). The bottom panel shows the activation of the dimensional nodes for the older model (i.e., 5-year-olds). As is clear in this panel, the older model shows more robust activation of the dimensional attention system (i.e., greater activation values), with a clear separation in the activity of the 'color' and 'shape' nodes during the pre-switch and post-switch phases.

To understand how the model works in detail, it is important to dissect the neural interactions captured in the middle panel. Figure 4A shows the *target input* presented to the model. This input to each field in the object WM system reflects the presence of the target cards and trays used in the DCCS, for instance, the blue-circle on the left and the red-star on the right (see the display at the top of Figure 4A). In the spatial field, this input takes the form of 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.

Figure 4B shows the object WM system just after an input for a red circle test card has been given. This *test input* is presented to the feature fields as horizontal 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. Notice that there is spatial conflict between the two feature fields for the features presented on the test card in Figure 4B. The test input for the blue feature overlaps with the target input at the left location. The test input for the star feature, on the other hand, overlaps with the target input at the right location in the shape field. Thus, there is not enough information in these inputs alone to determine where the test input should be localized. The input for the blue feature will try to build a peak at the leftward location, while the input for the star feature will try to build a peak at the rightward location.

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. This can be seen in the slight boost in activation for the ‘color’ node (light grey line) in the top panel at timestep 300. As a test card is presented in panel B, the feature-space fields build-up excitation from the inputs. This sends activation to the dimensional nodes, boosting their activity (see the increase in activation in the top panel at timestep 600).

The dimensional attention system, in turn, sends activation back to the feature-space fields, with a slight boost to the color field since the ‘color’ node is more active. Figure 4C shows that this subtle bias in the activation of the ‘color’ node is enough to tip the balance in the direction of a color-based decision: the model has formed a peak at the left location in the spatial field, binding the blue and star features to this location in the task space. Note that this spatial binding is reflected in the vertical ridges of subthreshold excitation at the leftward location in the feature-space fields as spatial information is passed back-and-forth among the object WM fields. Figure 4D shows the consequence of this sorting decision: there is a Hebbian memory in the color field that associates the blue feature and the left sorting location; similarly, there is a Hebbian memory in the shape field that associates the star feature and the left sorting location (see circled region of each field).

The same sequence of events play out as the model sorts the second type of test card highlighted in panels E-G. In Figure 4E, the model is shown a red circle to sort. This boosts the activity of the dimensional nodes (see top panel). The small bias to the ‘color’ node is amplified, leading the model to sort the red circle based on its color. Consequently, the model forms a peak of activation to the right side of the color field, binds the circle feature to the right as well, and effectively sorts the red circle to the right (see Figure 4F).

Figure 4G 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 Hebbian traces of where the test cards were sorted on the previous trials (see highlighted circles). Consider the consequences these memories will have as the model is now told to switch and play the shape game. In Figure 4G, the Hebbian traces are highlighted with white circles, while the target inputs are highlighted with white ovals. As can be seen in the figure, this results in *cooperation* within the color field—the Hebbian traces (circles) overlap with the target inputs (ovals). Consequently, activation is likely to build quickly at these primed locations. By contrast, there is *competition* within the shape field—the Hebbian traces (circles) associate the star shape with a leftward location and the circle shape with a rightward location, while the target inputs (ovals) indicate a circle on the left and a star on the right. Consequently, activation in this field will build slowly as multiple locations within the field compete through lateral inhibition. Considered together, then, the pre-switch dimension (color) is primed and has a competitive advantage going into the post-switch phase.

What happens when the model transitions to the post-switch phase? As can be seen in the top panel at timestep 25,000, 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; thus, the shape node (solid line) is now more strongly activated than the color node (grey line). Note, however, that the difference in the activation strengths of the two nodes is quite small—smaller, in fact, than at the start of the pre-switch phase (see timestep 300). This reflects the influence of Hebbian learning in the dimensional attention system. Specifically, Hebbian traces built up for the ‘color’ node during the pre-switch phase. As a result, the ‘color’ node is primed going into the post-switch phase, leading to a weaker differential in the activation of the two nodes.

At timestep 27,000, the model is given the first test card in the post-switch phase. In particular, Figure 4H shows a red circle test card that is presented to the object WM system. Figure 4I shows that the young model perseveratively sorts this red circle to the right—the model sorts by color and not shape. This occurs for two reasons: (1) there is only a small boost to the shape field from the dimensional attention system, and (2) there is cooperation in the color field along with competition in the shape field.

How is the model, like 5-year-olds, able to overcome these biases and switch rules? As discussed above, I made two changes to the model over development. First, the 5-year-old model has stronger interactions between the ‘shape’ and ‘color’ nodes—self-excitation and lateral inhibition are both stronger. This implements the spatial precision hypothesis probed in previous work using DFT (Schutte & Spencer, 2009; Spencer et al., 2009). Second, the connectivity pattern between the frontal and posterior systems is stronger and more selective. For example, the ‘shape’ node now passes stronger excitation to and from the shape field and has relatively minimal interaction with the color field. Thus, the 5-year-old model has a more refined understanding of what ‘shape’ means.

The bottom panel of Figure 4 shows how these developmental changes impact the activation of the dimensional attention system across the pre- and post-switch phases. Here, the relevant node for each phase achieves robust activation and strongly suppresses the irrelevant node. As a consequence, the 5-year-old model is now able to switch rules during the post-switch phase. This is shown in Figure 4J. In this figure, the 5-year-old model correctly sorts the red circle to the left—sorting by shape instead of color. The older model correctly sorts by shape because the ‘shape’ node sends a strong boost of activity to the shape field and suppresses the activity of the ‘color’ node (see timestep 27,500 in the bottom panel). With a stronger boost, the shape field can resolve the competition between the Hebbbian memory and the test input, successfully sorting the red circle to the left.

To convey more clearly the differences in the neural dynamics between the young and the old model, Figure 5 zooms in on these neural dynamics during the critical moments when the model is making its initial sorting decision on the first post-switch trial. The top panel plots the activation of the dimensional attention system for the old (solid lines) and young (dashed lines) model. The middle and bottom panels show the object WM system for the young and old models, respectively, at critical time-points during the decision-making process. Because the spatial field is receiving the summed activation over space from the different feature-space fields, this field reflects the relative amount of competition or cooperation in the feature-space fields as a decision is being made.

Figure 5A shows the fields before the first post-switch test card is presented. Note the cooperation within the color field and competition within the shape field for both models. In Figure 5B, the red circle test card is presented to the model. At the corresponding timestep (26,975), the dimensional attention system is just becoming engaged (see top panel). In Figure 5C, the inputs from the test card have built close to threshold and spatial activation is being shared throughout the network. At this point in time in the spatial field, the left and right locations are roughly equally activated in both the young and old model.

Consequently, it is not clear how the object WM system will sort the card at this point in time. Looking at the dimensional attention system, however, differences between the old and young models are evident (see timestep 27,000). Specifically, the ‘shape’ and ‘color’ neurons for the old model have achieved a much larger separation in activation relative to the young model. Thus, the old model is attending to, or boosting, the shape field more strongly and selectively than the young model. The consequences of this difference are starting to become evident in Figure 5D. The young model has stronger overall activation at the rightward location due to the cooperation in the color field, while the old model has stronger overall activation at the leftward location due to the input from the ‘shape’ node to the shape field. Finally, in Figure 5E, both models have resolved the spatial conflict and made a sorting decision. The young model perseverated and sorted the red circle to the right, while the older model successfully switched and sorted the test card to the left.

Summary of the model

The primary aim in developing the DNF model described below is to achieve extensive theoretical coverage of young children’s performance in the DCCS task early in development while adhering to, and anchoring our concepts in, formalized neural principles. I have presented a DNF model that 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 and between these populations, stable peaks of activation can emerge that capture decisions to sort objects (i.e., cards) to the left or right location in the DCCS task. Hebbian processes operating within these fields serve to accumulate memories over the course of the task which can influence subsequent decisions.

To sort cards in a rule-like fashion, I 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). In addition, I implemented a specific type of developmental change in the dimensional attention system to capture the emergence of rule-use between 3 and 5 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 robust activation of the ‘shape’ node; consequently, cooperation in the color field dominated the sorting decision. By contrast, the 5-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.

The DNF model is clearly a model of children’s performance in the DCCS task. But is it more than that? At face value, the answer to this question is certainly ‘yes’. Recall that the DNF model builds on an earlier model of how adults ‘bind’ visual features together (see, Johnson et al., 2008; Spencer et al., in press). Thus, the model speaks to issues that extend beyond the purview of the DCCS task. Moreover, the model implementation makes specific claims about the general nature of rule-use and executive control. For instance, there was no explicit representation of the rules in the model. Rather, rule-use emerged from specific associations of features with spatial locations in an object WM system and the associations of dimensional attention nodes with cortical populations in the object system. This means that the ‘rule’ does not reside in the frontal system—that system knows nothing about the specific test card presented or where the target cards are in the task space. These details are ‘known’ in the object system; however, the object system knows nothing about the meaning of the ‘color’ or ‘shape’ game. Thus, rule-use—and the ability to switch rules—emerges from interactions across multiple neural populations.

One can also ask how the different components of EF—inhibitory control, working memory, and task-switching—are captured by the model. For instance, how does the model inhibit irrelevant information? This is partly the job of the dimensional attention system, for instance, boosting the excitation of the ‘color’ node on the first pre-switch trial in the color game. But inhibitory processes also play a role: when the ‘color’ node was boosted, lateral inhibitory interactions suppressed the activity of the ‘shape’ node. Inhibitory processes were also at work in the posterior neural fields. Here, the build-up of activation in the color field at the left location increased the activation of neurons on the left side of the shape field. As this activation increased in the shape field, laterally inhibitory interactions suppressed activation associated with the star feature at the rightward location, binding the test card features to the left location. Thus, inhibitory control—a key cognitive concept in early development—emerges from excitatory neural processes (e.g., boosting the ‘color’ node) and associated laterally inhibitory consequences (e.g., suppressing the ‘shape’ node; for related ideas, see Morton & Munakata, 2002; Roberts, Hager, & Heron, 1994; Roberts & Pennington, 1996; Stedron, Sahni, & Munakata, 2005).

Similar processes have been used to explain inhibitory control in the Simon task. In this task, the spatial location of the stimulus can either be congruent or incongruent with the spatial response required for the stimulus. To simulate this task, Erlhagen and Schöner (2002) gave a DNF model stimuli that either overlapped (as in a congruent trial) or conflicted with one another (as in an incongruent trial). When an incongruent stimulus was presented, it was inhibited through lateral inhibition generated by the build-up of excitation associated with the relevant response location. Thus, activation peaks and inhibitory processes go hand-in-hand.

What about a second component of EF—where does ‘switching’ live in the DNF model? There were many components of the model’s dynamics that produced rule-switching. This required a strong representation of the relevant dimension brought about through robust activation of the relevant dimensional node. This, in turn, produced a strong boost to a specific feature-space field, helping this field resolve the conflict between the task input at one location and the Hebbian trace at another location. But the full resolution of this conflict required building excitation at the ‘correct’ feature value in the feature-space field and suppressing the incorrect feature via lateral inhibition. Switching, then, requires a confluence of robust neural activation peaks and lateral inhibitory processes.

These examples are informative in that they highlight the challenge of relating cognitive concepts such as ‘inhibitory control’ and ‘task-switching’ to mechanisms in a neural model, and the utility of having a neural model to bridge between these two levels of description. Simulations show that the DNF model can produce functional behavioral outcomes central to cognitive control in early development such as changes in the ability to switch tasks and inhibit a prepotent response. But these outcomes do not map one-to-one onto specific neural mechanisms. That is, there is not an inhibitory module or a task-switching module in the model; rather, the model helps explain how functional behaviors labeled as, for instance, ‘inhibitory control’ can arise from neural mechanisms in a complex neural system.

For this link between cognitive concepts and neural mechanisms to have meaning, however, the utility of the theory must first be established. I do this in the following section by quantitatively simulating a host of behavioral effects in the literature. Next, I discuss a set of novel behavioral predictions that was successfully tested with 3-year-old children (see Buss & Spencer, in press). Critically, these predictions are not consistent with any other theory of children's performance in the DCCS task.

Simulating the Behavioral Dynamics of the DCCS Task

In this section, I present quantitative simulations of a diverse array of findings from the literature on children's performance in the DCCS task. 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 this section, I will focus on simulations of key conditions that will be relevant for the tests of the DNF model in Chapters 4 and 5. Note, however, that the model has been used to simulate a total of 14 different variations of the task (Buss & Spencer, in press).

In the previous section, I 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, I 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, I 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 4, but retained the key consequence of these neural dynamics—a boost in the baseline level of activation in a particular feature field.

I used two parameters to specify the activation passed from the dimensional nodes to the feature fields. First, I 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, I 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 6 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), I 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, I used the two *boost* distributions shown in Figure 7. As this figure shows, the 3-year-old model had a mean boost value of 0.35 and the 5-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. I also used two *shift* distributions (see Figure 8). The 3-year-old distribution encompassed a broad range of values centered on 0.5 while the 5-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 4). 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.

Method

Procedure

Simulations were conducted in Matlab 7.5.0 (Mathworks, Inc.) on a PC with an Intel® i7™ 3.33 GHz quad-core processor. *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, the resting-level of the WM fields was decreased 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 Buss and Spencer (in press). 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 5-year-olds' performance. First, because I 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, I show that the model is able to quantitatively capture these details.

Standard. The first issue I probed was whether the model quantitatively captures children's performance in the Standard DCCS task. Figure 9A shows the rates of perseveration and switching for the model and literature sampled for 3- and 5-year-olds (Halford et al., 2007; Müller et al., 2006; Zelazo et al., 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 5-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 top panel in Figure 9A explains why the 3-year-old model perseverates.

Going into the post-switch phase, there is cooperation of Hebbian traces (HTs) and target input in the pre-switch field (see color field in Figure 9A). 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 9A). 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 5-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. Zelazo et al. (2003) demonstrated that 3-year-olds are able to switch rules when the pre-switch test cards match the target cards along both color and shape dimensions. How does eliminating this conflict impact on the performance of the model? Figure 9B 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 top panel of Figure 9B, no-conflict cards during the pre-switch phase establish cooperation 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.

Total change. 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 top panel in Figure 9C, this version eliminates both the cooperation 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, and these inputs do not overlap the HTs from the pre-switch phase. Thus, the model does not need to inhibit these HTs. Figure 9C shows that both 3-year-olds and the model tend to switch rules 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.

Summary of quantitative fits

Buss and Spencer (in press) showed that the DNF model was able to quantitatively simulate the performance of 3- and 5-year-olds in a diverse array of DCCS tasks. Specifically, the model fit 37 data points across 14 different versions with a RMSE of 0.06. The model proved robust in accounting for performance under various manipulations to the task, such as changing different features between the pre- and post-switch phases, manipulating the attentional properties of the dimensions, or manipulating the conflict present during different phases of the task. Importantly, the model also satisfied various constraints on performance, such as only having difficulty during the post-switch phase and the all-or-none nature of children’s switching.

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, 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. Thus, 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.

Although these features of the theory are compelling, the true test of any theory is its ability to generate novel predictions. A behavioral prediction of the model has already been tested regarding the role of feature binding through space and the spatial layout of the target cards between the pre- and post-switch phases (Buss & Spencer, in press). This was done in the context of a Negiative Priming and No-Conflict Negative Priming versions. In both of these versions, the features that are relevant for the pre-switch phase are changed in the post-switch phase (see top panel of Figure 10). If children sorted red and blue stars and circles by color during the pre-switch phase, then they would be instructed to sort green and yellow stars and circles by shape during the post-switch phase. The difference between these two conditions is the test cards used during the pre-switch phase.

Typically, 3-year-olds perseverate in the Negative Priming version, but switch rules in the No-Conflict Negative Priming version (see plots at bottom of Figure 10). DFT explains these results, as with the previous results, based on the layout of spatially specific HTs within the feature-space fields. As can be seen in the middle panels of Figure 10, the Negative Priming version has competition within the post-switch field, but the No-Conflict Negative Priming version has cooperation within the post-switch field. If spatial cooperation and competition are the underlying factors that explain 3-year-olds’ behavior, then this pattern should be reversed by simply swapping the locations of the target cards before the start of the post-switch phase. Figure 10 illustrates the influence of this manipulation in the model. In an NNP-SpaceSwap condition, this spatial manipulation creates conflict in the shape field and the model which used to switch rules now perseverates (see blue bars in bottom panel of Figure 10). In the NP-SpaceSwap version, the spatial manipulation creates cooperation in the shape field and the model which used to perseverate now switches rules (see blue bars in bottom panel of Figure 10).

We tested these predictions with a group of 76 3-year-olds between 38 and 46 months of age (see Buss & Spencer, in press). These children were randomly assigned to one of four conditions (NNP, NP, NNP-SpaceSwap, and NP-SpaceSwap). As predicted by the model, children showed a higher rate of perseveration in the NNP-SpaceSwap version compared to the NNP version ($\chi 2(1)= 3.995$, $p= 0.04$; see red bars in Figure 10). Conversely, children showed a higher rate of switching in the NP-SpaceSwap version compared to the NP version ($\chi 2(1)= 4.48$, $p= 0.03$; see red bars in Figure 10). Thus, by simply swapping the spatial locations of the target cards before the post-switch phase, children’s performance was effectively reversed. Note that because the rule structure and the features were the same across the standard and space-swap conditions, this pattern of results is not easily explained by any other current theories. For example, according to CCC/CCC-r theory, children perseverate in the Negative Priming version due to the suppression of the post-switch rules, not due to any learning about where the post-switch features are sorted. Thus, it should not matter where the post-switch features are positioned in the task space. By contrast, the PDP model proposed by Morton and Munakata (2002) only learns about the relevant dimension on the pre-switch trials, that is, the feature dimension which is most strongly activated. Consequently, there are no weight changes associated with the post-switch features that could guide differential performance across the standard and space-swap conditions. Moreover, the spatial positions of the cards are not captured in the PDP model in any grounded way; rather, the mapping between the output units and space is built-in *a priori*. As such, there is no clear way to implement our manipulation in this model. I return to the issue of prediction in Chapter 4 where I focus on a critical feature of DFT—the ability to simulate real-time hemodynamics.

Conclusions

In the remainder of this thesis, I address two key challenges for the DFT of EF. The previous sections demonstrated that DFT can simulate various details of rule-switching, as well as generate novel behavioral predictions. The next goal is to use the DFT to gain greater insight into the development of brain-behavior interactions. *Study 1* will address this challenge by testing hemodynamic predictions about the mechanisms of flexible rule-use associated with ‘easy’ versions of the DCCS task with 3- and 5-year-olds.

A second challenge arises from the emphasis on the early development of EF. In particular, DFT has only been used to capture task-switching in early childhood. Important developmental changes in both brain and behavior, however, also occur through adolescence into adulthood (Morton et al., 2009; Crone, Donohue, et al., 2006). *Study 2* will use a task inspired by DCCS task variants used with young children to study the neural dynamics that underlie adults’ performance in rule-switching tasks. The goal is to examine whether key components of the DNF model—the dimensional attention system and the posterior feature-space fields—can be localized using the common subtraction logic from the fMRI literature. This would extend our probes of the neural bases of the DNF model to the gold-standard of neural imaging—fMRI. This would represent a first step toward clarifying the neural systems that underlie DCCS performance and would set up future work extending the model into adulthood, further refining and constraining the architecture and dynamics of the model.

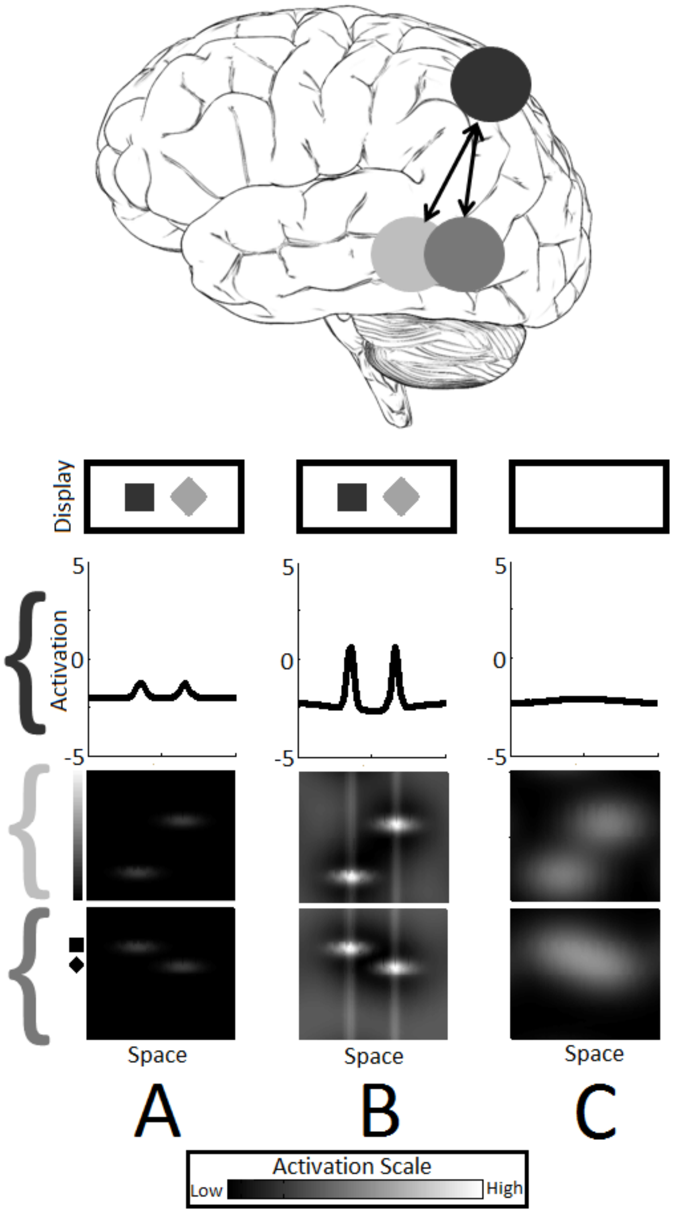

Figure 2: WM fields for the feature binding model.

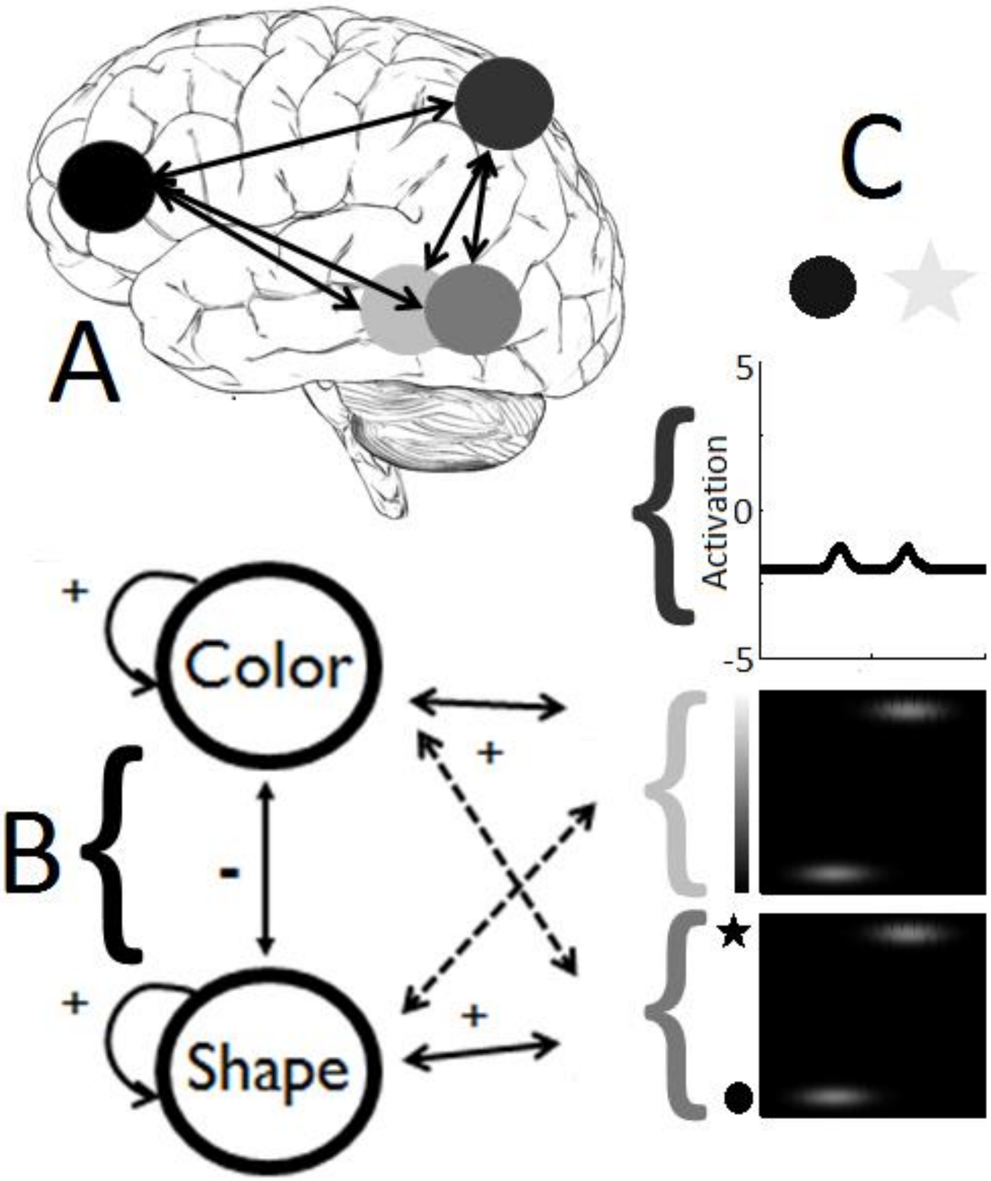

Figure 3: The object WM model with the dimensional attention system.

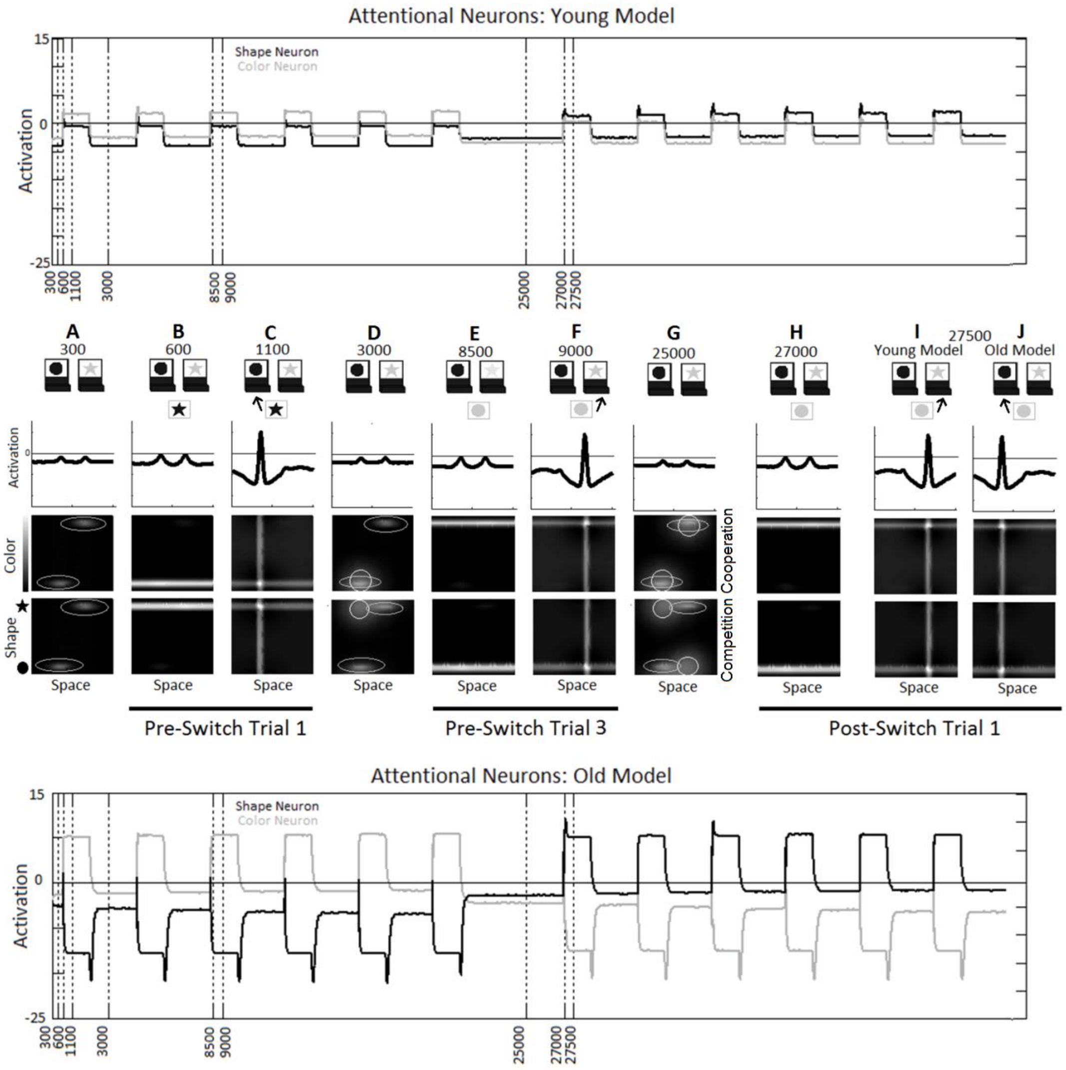

Figure 4: The sequence of events as the model sorts cards in the DCCS.

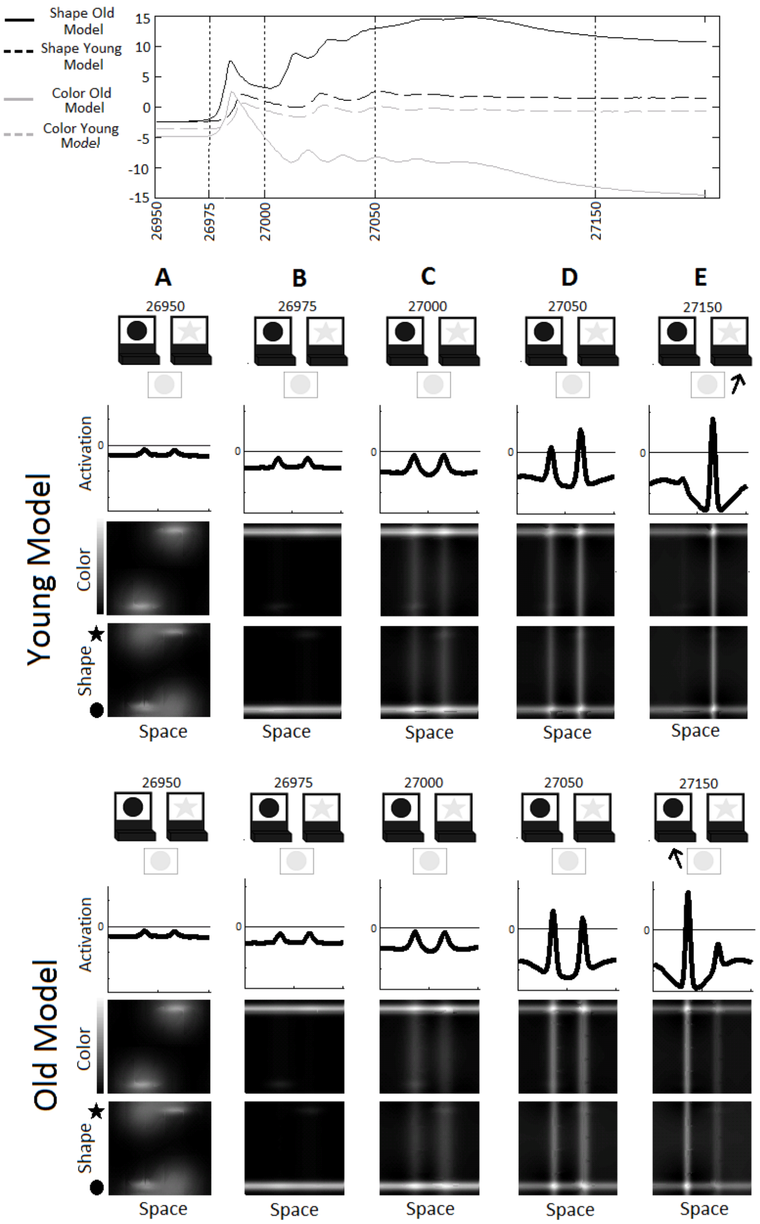

Figure 5: A zoom-in on the critical dynamics of the attentional system on a rule-switch trial.

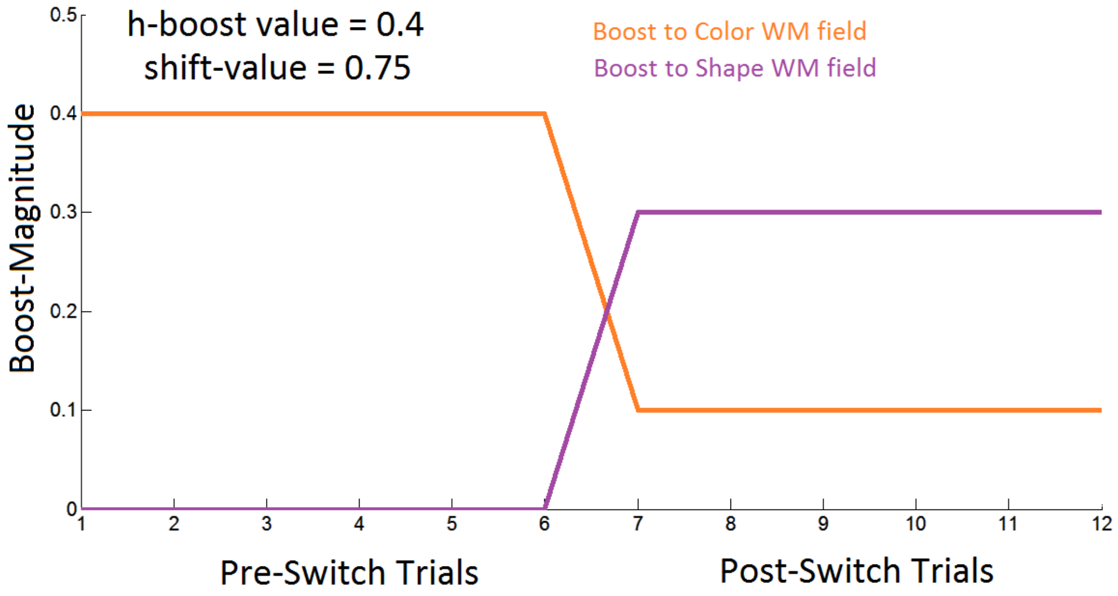

Figure 6: Schematic diagram depicting the simplified processes of boosting and shifting.

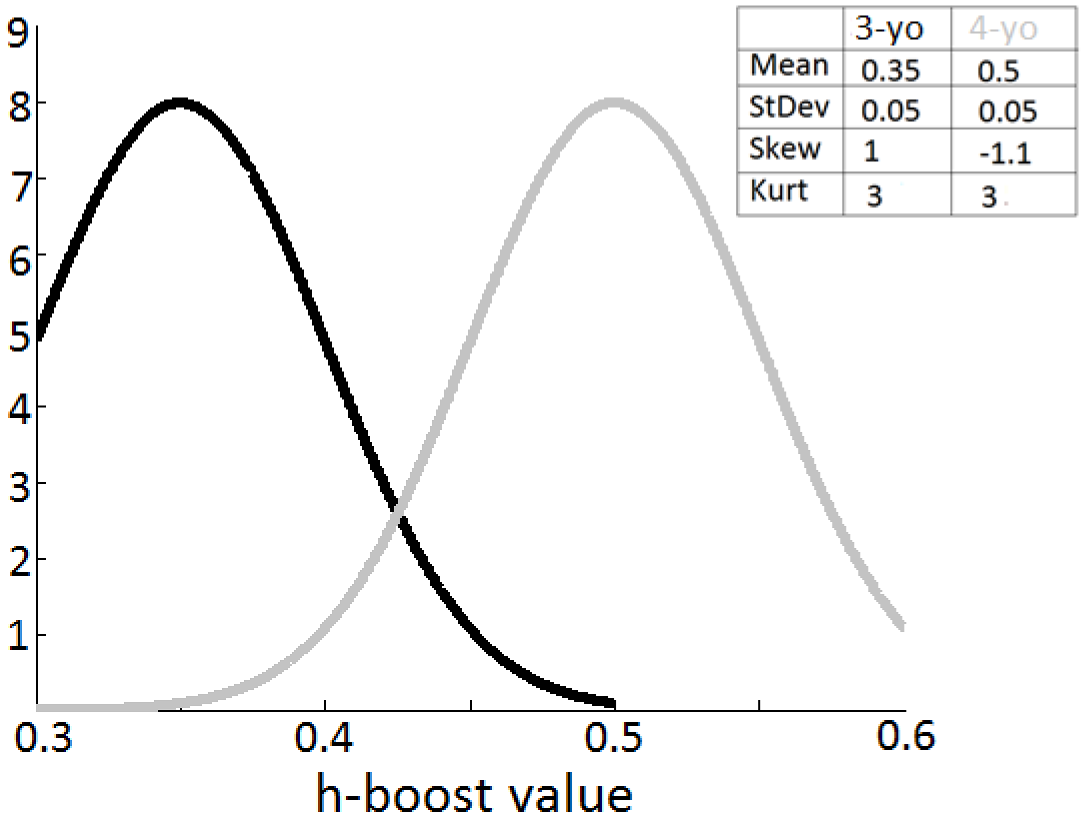

	3-yo	4-yo
Mean	0.35	0.5
StDev	0.05	0.05
Skew	1	-1.1
Kurt	3	3

Figure 7: H-boost distributions used for 3- and 4-year-old models.

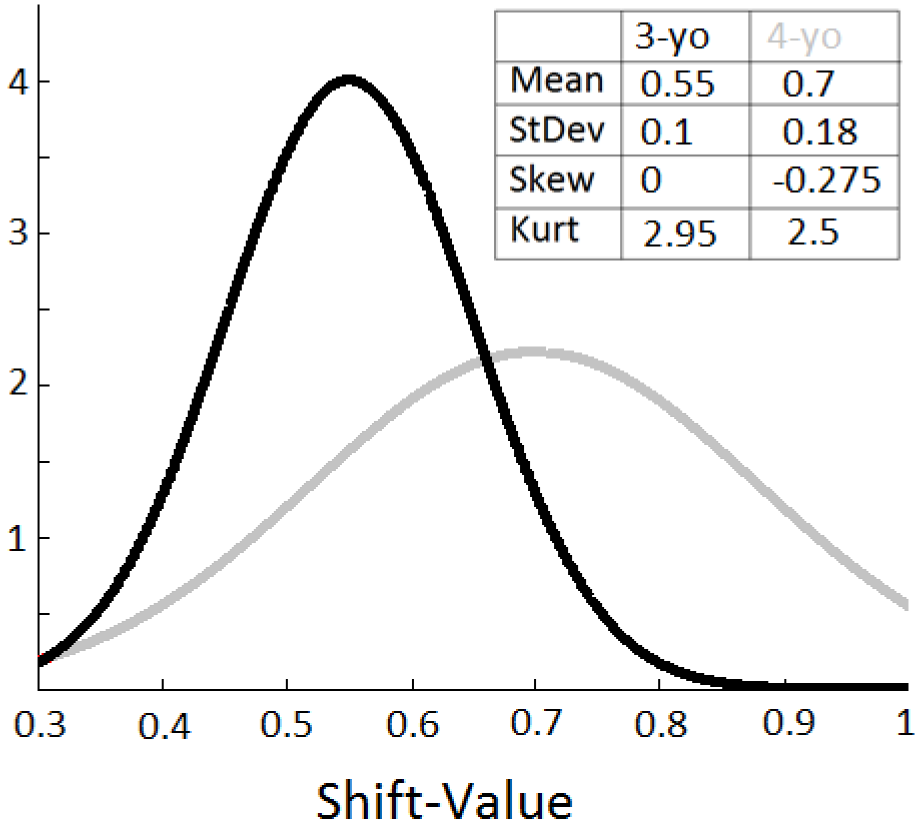

	3-yo	4-yo
Mean	0.55	0.7
StDev	0.1	0.18
Skew	0	-0.275
Kurt	2.95	2.5

Figure 8: Shift-value distributions for 3- and 4-year-old models.

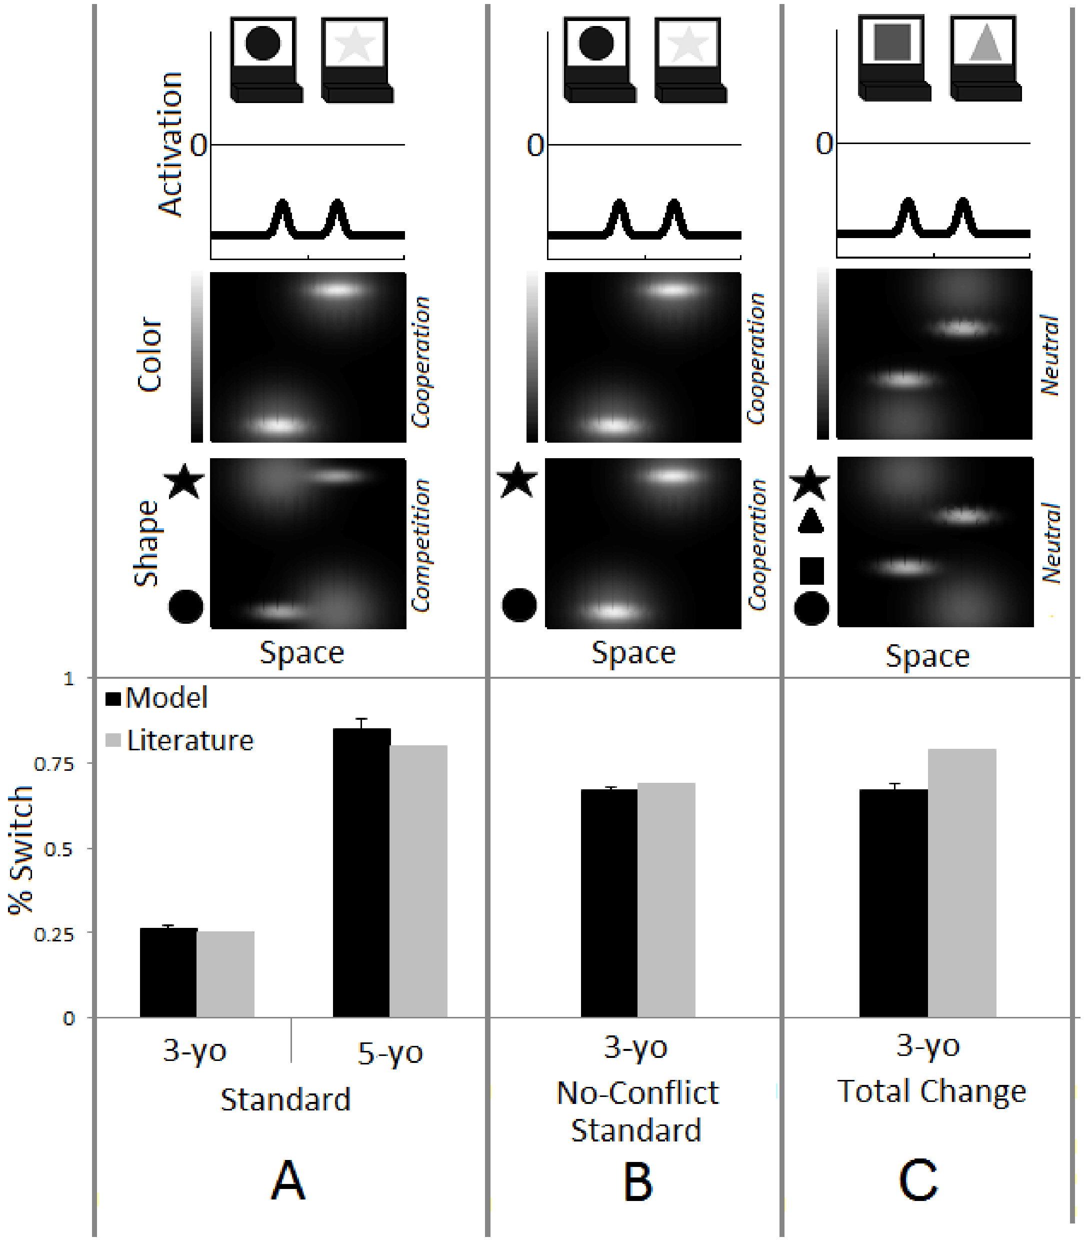

Figure 9: Model results.

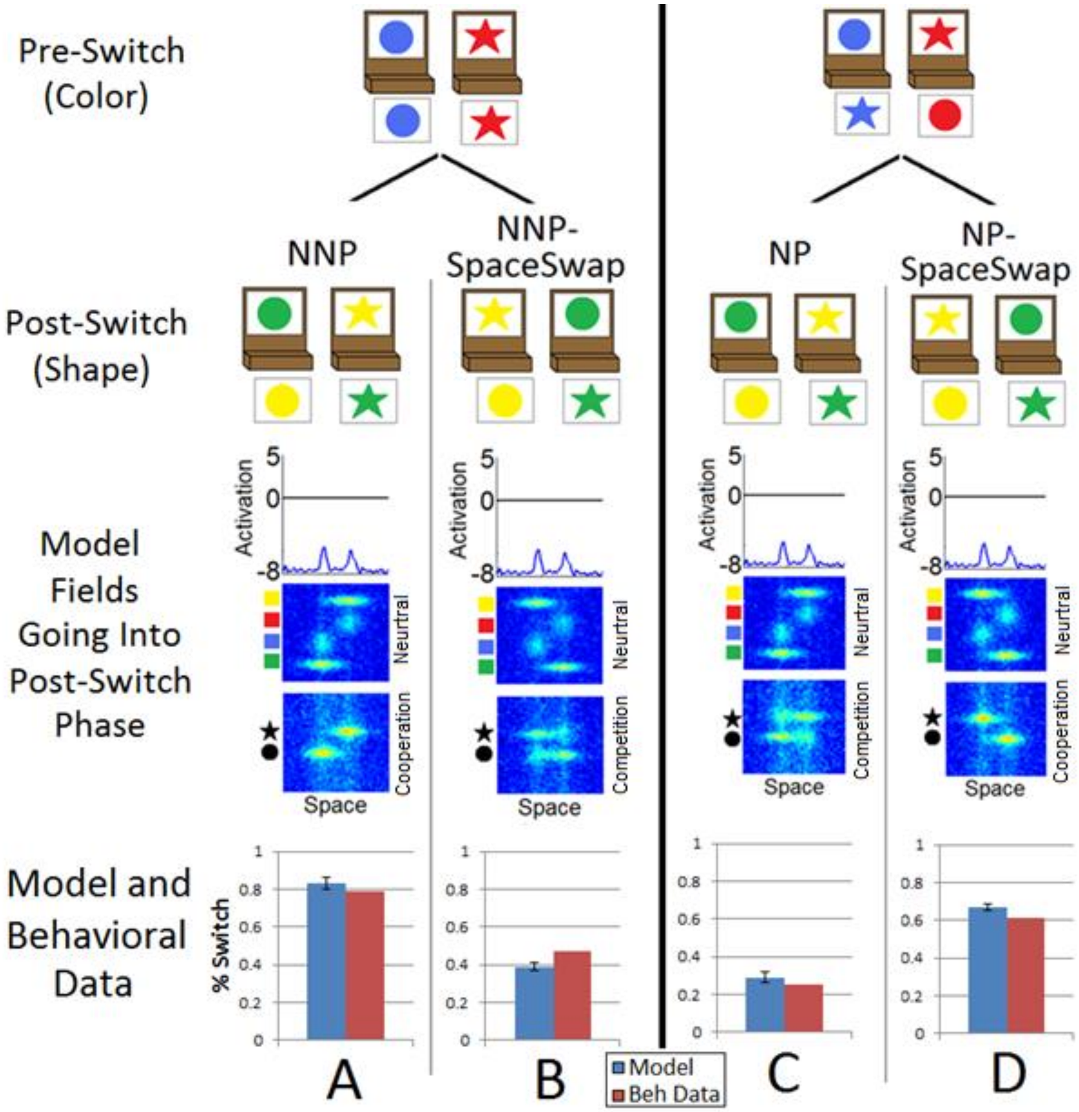

Figure 10: Test of a behavioral prediction from the DNF model.

CHAPTER 4

STUDY 1: USING NIRS TO PROBE MECHANISMS OF FLEXIBLE RULE-USE IN EARLY CHILDHOOD

Chapter 3 demonstrated that DFT provides a robust account of the behavioral literature on the DCCS, quantitatively capturing behavioral data with 3- and 5-year-olds and generating novel behavioral predictions that no other theory can explain. In this chapter, I ask whether the DNF model also offers an account of children's performance at the level of the brain. First, I discuss data from the literature pertaining to the neural dynamics of flexible rule-use in early childhood. Next, I describe a method used to translate real-time neural dynamics in the DNF model to simulated hemodynamic data. I then demonstrate the viability of this method by simulating the NIRS data reported by Moriguchi and Hiraki (2009). Building from this, I generate hemodynamic predictions on 'easy' versions of the DCCS task where young children typically succeed. Finally, I present an experiment using NIRS with 3- and 4-year-olds that tests these hemodynamic predictions of the DNF model.

Neural Dynamics of Flexible Rule-use in Early Childhood

This chapter aims to elucidate the role that frontal cortex plays in flexible rule-use in early childhood. Different areas of frontal cortex have been shown to be associated with different functional roles. With adults, shifting rules is associated with activation in dorsolateral PFC, pre-supplementary motor area, and inferior frontal junction. Further, shifting rules, regardless of whether attention is also being shifted, activates inferior frontal sulcus. However, middle frontal gyrus is selectively activated when shifts of attention are required (Nagahama et al., 2001). The ventro-lateral prefrontal cortex and rostro-lateral prefrontal cortex 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 are more active when switching rules or reconfiguring a rule-set (Crone, Wendelken, et al., 2006).

Generally, these frontal areas are implicated in the top-down organization of task sets or attention. Research has further revealed differences in these neural dynamics over development. Specifically, stronger activation is seen in superior frontal sulcus between adolescence and adulthood (Morton et al., 2009). Additionally, 8- to 12-year-olds display less-differentiated frontal activation compared to adults. These children activate the pre-supplementary motor area for rule-representation, not simply for rule-switching as with adults. Further, while adults show heightened ventro-lateral pre-frontal cortex 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, Honomichl, Wendelken, & Bunge, 2006).

This line of work using fMRI is promising, but fMRI is not suitable for use with very young children due to the requirements to lie still for extended periods of time in a noisy, confined space. However, recent advances in neuroimaging technology have allowed for investigation of the role of frontal cortex in rule-switching in early childhood. Functional Near-Infrared Spectroscopy (fNIRS) provides a measure of blood flow similar to fMRI using the optical properties of hemoglobin. This technology provides a more child-friendly experience than fMRI because participants are fitted with a lightweight cap that allows for freedom of movement during data collection.

fNIRS has revealed a close association between rule-switching performance and frontal activation in early childhood. Moriguchi and Hiraki (2009) used NIRS with 3-year-olds and 5-year-olds in the DCCS task. Children were fitted with a NIRS probe containing lights and detectors at left and right frontal areas. The top panel of Figure 11 shows the hemodynamics (relative concentration of oxy-hemoglobin) recorded during the pre-switch and post-switch phases. Three-year-olds who perseverated showed significantly weaker bilateral frontal activation during both phases of the task compared to 3-year-olds who were able to switch rules and 5-year-olds (all of whom switched rules). This demonstrates a close association between activity of the frontal cortex and cognitive flexibility—with strong frontal activation, children are able to switch rules; with weak frontal activation, children are not able to switch rules.

Based on these data, Moriguchi and Hiraki (2009) suggest that activation of the frontal cortex is a developmental achievement which is necessary in order to flexibly switch rules. But this raises a key question: If frontal cortex activation is necessary to switch rules, then how are 3-year-olds able to switch rules in ‘easier’ versions of the DCCS? For example, 3-year-olds have little trouble switching rules in a No-Conflict version. This version is the same as the standard version except the test cards match the target cards along both dimensions during the pre-switch phase (see Figure 1D). If 3-year-olds are generally unable to activate the frontal cortex when asked to switch rules, then what neural mechanism underlies their successful performance in this version? Are 3-year-olds able to somehow activate frontal cortex in this condition, or are other regions compensating for the lack of frontal activation? If flexible rule-use also depends on activation in other regions, then how is frontal cortex activation influenced by the activity of this more global network?

Research using fMRI has indicated that regions in temporal and parietal cortex are also activated when switching rules. For example, Morton and colleagues (2009) found stronger switch-related activation in superior parietal cortex and fusiform gyrus between adolescence and adulthood. Over development, there are also changes in the connectivity within distinct cognitive control networks. One such network encompasses frontoparietal connections between dorsolateral prefrontal cortex, intraparietal sulcus, and precuneus. This network is hypothesized to be involved in trial-to-trial adaptation, task-initiation, and error adjustment. Another network is composed of cinguloopercular connections across regions thought to be involved in the stable maintenance of task-sets, including anterior prefrontal cortex, anterior cingulate cortex, anterior insula, ventral prefrontal cortex, along with sensory areas in occipital and temporal cortex (Dosenbach et al., 2007; Fair et al., 2009; Fair et al, 2007). However, the fNIRS data from early childhood reported by Moriguchi and Hiraki (2009) only examined frontal cortex. Thus, nothing is known about the development of these networks in early childhood and whether other regions might be involved in the developmental transition from perseveration to switching in the DCCS task.

Dynamic Field Theory (DFT) is well positioned to provide insight into the neural network underlying the development of rule-use. Buss & Spencer (in press) presented a dynamic neural field model that reproduced the performance of 3 and 5-year-olds across 14 different variations of the DCCS (Buss & Spencer, in press). A central feature of this model is the reciprocal excitatory loop between a dimensional attention system and an object representation system. This implements a form of dimensional attention: activation from the object representation system feeds into the dimensional attention system; once the dimensional attention system is activated, it feeds excitation back into the object representation system. The input from the dimensional attention system biases activation within the object representation system leading to the formation of rule-like decisions based on particular visual dimensions.

As reviewed in Chapter 3, in addition to addressing the developmental transition between 3- and 5-year-olds, DFT is able to explain the improved performance of 3-year-olds in the No-Conflict version. The mechanism by which the model achieves this stems from the spatial layout of Hebbian traces (HTs) and target inputs. When no-conflict cards are sorted during the pre-switch phase, the HTs accumulated during the pre-switch phase overlap with the target inputs to the feature-space field that are relevant for the post-switch phase. As illustrated in Figure 9B, if the pre-switch dimension is color, then the memories for star and circle will overlap with the locations of those features in the task space. By eliminating this form of feature-space conflict, the young-model, like 3-year-olds, is able to switch rules.

Can this model be used to gain insight into the neural mechanisms that underlie successful performance of 3-year-olds in the No-Conflict version? One way to achieve this is to compare model neural dynamics to the primary measure in cognitive neuroscience—the hemodynamic response. Different components of the model can be localized to particular neural regions based on their function. This can provide a starting point for comparing model dynamics to brain dynamics.

Figure 3A provided such a mapping between the DNF architecture and cortical regions involved in EF. Specifically, the dimensional attention system is mapped onto processing within frontal cortex because it implements a form of abstract dimensional attention that provides biases to other regions. The feature fields of the object representation system can be localized to processing within object regions in the ‘what’ visual pathway in temporal cortex. The spatial field can be localized to processing within the ‘where’ visual pathway in parietal cortex (Haxby et al., 1991; Mishkin et al., 1983). The model, then, implements an architecture that can potentially capture the interaction of different brain regions in flexible rule-use tasks.

By simulating hemodynamic data with DFT, the mechanisms and processes implemented by the model can be tested to determine whether they produce the same pattern of hemodynamic data from neuroimaging studies. First, though, the relationship between model activity and the hemodynamic response measured with fMRI/fNIRS must be formalized. In the next section, I discuss a recent approach developed by Buss, Spencer, and colleagues (Buss, Magnotta, Schoner, & Spencer, submitted; Buss, Wifall, Hazeltine, & Spencer, in press) to map neural activity in dynamic field models to simulated hemodynamics. I then use this approach to examine whether the DNF model captures existing hemodynamic data from the study by Moriguchi and Hiraki (2009).

Simulating Hemodynamics with Dynamic Neural Fields

Simulating the real time dynamics of both neural and behavioral data has been a challenging endeavor (Ashby & Waldschmidt, 2008). To achieve this using dynamic field theory, I adapt an approach based on biophysical work exploring the neural basis of the BOLD signal. In particular, Logothetis, Pauls, Augath, Trinath, and Oeltermann

(2001) recorded single- and multi-unit data along with local field potentials (LFP) and the BOLD signal in visual cortex of macaques while they were shown checkerboard displays. An LFP is a measure of dendritic activity over a localized population of neurons, accounting for changes in both inhibitory and excitatory ion flow. This provides a measure of the input to, and local processing within, a region of cortex. Logothetis et al. (2001) reported that the LFP was most strongly correlated with the BOLD response. Further, by convolving the LFP with an impulse response function (specifying the time course of the slow blood flow response to neural activity), the authors were able to reproduce the BOLD signal. This suggests that the LFP is a strong contributor to the neural signal driving the BOLD response.

Deco, Rolls and Horwitz (2004) expanded upon this work using an integrate-and-fire network to simulate the hemodynamics underlying visual working memory. They estimated an LFP by summing the absolute value of the total synaptic current flow into the cells, including excitatory components (implemented through NMDA and AMPA receptors) and inhibitory components (GABA). According to the model, these components are the primary contributors to the rate of change in membrane potential, excluding the stability term and factors that influence the resting potential. By convolving this estimated LFP with an impulse response function, the authors were able to simulate the BOLD response in a visual working memory task.

Here, I adapt the approach of Deco et al. (2004) for use with DNFs (see, Buss, et al., submitted; Buss, Wifall, et al., in press). In particular, I created a DFT-LFP measure by summing the absolute value of all terms contributing to the rate of change in activation within each component of the model, excluding the stability term and the two factors that impact the neuronal resting level—a resting-level parameter and the HTs which modulate the baseline resting levels over learning. The included terms reflect excitatory and inhibitory interactions within each component of the model, excitation passed between components of the model, and noise. Note that I excluded the stimulus input in the LFP computation because the inputs were applied directly to the model rather than implementing these in a more neurally realistic manner (e.g., by using simulated input fields as in Spencer et al., in press).

To apply this method, an LFP (which I call a DFT-LFP) can be generated by computing the sum of the absolute value of excitatory and inhibitory interactions for specific parts of the model over the course of the task. This is illustrated in the following example with the equations for the dimensional attention system given by Equations 1 and 2.

$$\begin{aligned}\tau \dot{i}_{shape}(t) &= -i_{shape}(t) + h_{i_{shape}} + S_{i_{shape}}(t) + c_{ii_{excite}} \Lambda(i_{shape}(t)) - c_{ii_{inhib}} \Lambda(i_{color}(t)) \\ &+ c_{iw_{ShapeColor}} \iint \Lambda(w_{f1}(x,y,t))dxdy + c_{iw_{ShapeShape}} \iint \Lambda(w_{f2}(x,y,t))dxdy\end{aligned} \tag{1}$$

$$\begin{aligned}\tau \dot{i}_{color}(t) &= -i_{color}(t) + h_{i_{color}} + S_{i_{color}}(t) + c_{ii_{excite}} \Lambda(i_{color}(t)) - c_{ii_{inhib}} \Lambda(i_{shape}(t)) \\ &+ c_{iw_{ColorShape}} \iint \Lambda(w_{f2}(x,y,t))dxdy + c_{iw_{ColorColor}} \iint \Lambda(w_{f1}(x,yt))dxdy\end{aligned} \tag{2}$$

The terms in these equations contributing to the rate of change in activation were included in the calculation of the DFT-LFP (Edin et al., 2007). For example, from Equation 1, this includes the self-excitatory projection ($c_{ii_{excite}} \Lambda(i_{shape}(t))$), the mutually inhibitory projections ($c_{ii_{inhib}} \Lambda(i_{color}(t))$), and the inputs from the color WM field ($c_{iw_{ShapeColor}} \iint \Lambda(w_{f1}(x,y,t))dxdy$) and shape WM field ($c_{iw_{ShapeShape}} \iint \Lambda(w_{f2}(x,y,t))dxdy$). In particular, the DFT-LFP for the dimensional nodes was defined by Equations 3 and 4.

$$\begin{aligned}LFP_{ishape}(t) &= | c_{ii_{excite}} \Lambda(i_{shape}(t)) | + | c_{ii_{inhib}} \Lambda(i_{color}(t)) | \\ &+ | c_{iw_{ShapeColor}} \iint \Lambda(w_{f1}(x,y,t)) | + | c_{iw_{ShapeShape}} \iint \Lambda(w_{f2}(x,y,t)) |\end{aligned} \tag{3}$$

$$\begin{aligned}LFP_{icolor} &= | c_{ii_{excite}} \Lambda(i_{color}(t)) | - | c_{ii_{inhib}} \Lambda(i_{shape}(t)) | \\ &+ | c_{iw_{ColorShape}} \iint \Lambda(w_{f2}(x,y,t)) | + | c_{iw_{ColorColor}} \iint \Lambda(w_{f1}(x,yt)) |\end{aligned} \tag{4}$$

This measure of real-time neural activity can then be convolved with a general impulse response function given by Equation 5 (Deco & Rolls, 2005; n_1=7.000, t_1=1.2, n_2=12.250, t_2=1.22, $a_2=1.6\text{x}10^{-12}$).

$$h(t) = c_1 t^{t_1} e^{-t/t_1} - a_2 c_2 t^{n_2} e^{-t/t_2}, c_i = \max(t^{n_i} e^{-t/t_i}) \quad 5$$

Equations 3-5 provide a formal mapping between simulated real-time neural dynamics in the DNF model and simulated hemodynamic data. The sections below examine whether this approach to simulating hemodynamics can capture functional neuroimaging data from children. First, I provide an initial demonstration that DFT is able to capture hemodynamic data in early childhood by simulating the fNIRS data presented by Moriguchi and Hiraki (2009). Second, I examine what the model predicts should happen at the level of neural dynamics from 'easy' (No-Conflict) versus 'hard' (Standard) versions of the DCCS task. I then test these novel *neural* predictions in a study using fNIRS with 3- and 4-year-old children.

Method: Simulating Moriguchi & Hiraki (2009)

Procedure

Simulations were conducted in Matlab 7.5.0 (Mathworks, Inc.) on a PC with an Intel® i7™ 3.33 GHz quad-core processor. The 'young' and 'old' autonomous models (see Buss & Spencer, in press for model details and parameters) were given the Standard version of the DCCS task. A DFT-LFP was tracked from each component of the model. Table 1 lists the parameter differences in the dimensional attention system between models. There were 4 total parameter differences between the 'old' and 'young' models. Two of these differences corresponded to changes in the excitatory (Self-Excitation) and inhibitory (Lateral Inhibition) connectivity between the dimensional attention nodes. Together, these create strong suppression of the irrelevant dimensional attention node and stronger activation of the selected dimensional attention node. Two other parameter differences corresponded to the strength of activation from a dimensional attention node to its associated feature-space field and the strength of the reciprocal activation from a feature-space field to its associated dimensional attention node. These parameter changes create a stronger boost of the relevant feature-space field when a dimensional node is activated (and vice versa due to the reciprocal coupling).

The general procedure with the model followed that of the simulations presented in Chapter 3. The model was given 5 trials during 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. 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.

The DFT-LFP corresponding to the dimensional nodes was tracked throughout the course of the pre- and post-switch phases. The DFT-LFP was first convolved with the impulse response function to generate a hemodynamic time course. The data were divided into pre- and post-switch time courses which were aligned with the onset of the first test card during each phase. The resulting data were then analyzed by initializing the hemodynamic response at the start of each sorting phase to a 0 baseline value. Note that the fNIRS data reported by Moriguchi and Hiraki (2009) was initialized to a baseline level determined while children sorted blank cards. The simulated hemodynamic data, then, were normalized by calculating a z-score across runs and ages. This is comparable to the analysis conducted by Moriguchi and Hiraki.

Results

This simulated hemodynamic response during the pre- and post-switch phases is plotted in the bottom panels of Figure 11 along with the data reported by Moriguchi and Hiraki (2009) in the top panels. The model showed the same relationship between frontal activity and rule-switch as was reported by Moriguchi and Hiraki (2009). The ‘young’ model that tends to perseverate (plotted in green) showed significantly weaker hemodynamic activity from the dimensional attention system compared to the ‘old’ model that tends to switch rules (plotted in blue). Thus, by mapping the dimensional attention system to frontal cortex, the same developmental mechanisms that produce correct switching in the ‘old’ model also produce the associated changes in frontal cortex activation.

Method: Generating Hemodynamic Predictions from DFT

The previous section demonstrated that the DNF model captures the key neural effect associated with the developmental transition from perseverating to switching in the DCCS task. Having established the viability of the method for simulating hemodynamics, I now turn to simulations that predict hemodynamic patterns for the ‘young’ and ‘old’ models in ‘easy’ (No-Conflict version) and ‘hard’ (Standard version) variants of the DCCS task. The study by Moriguchi and Hiraki (2009) only recorded data from frontal cortex, limiting model comparisons to the dimensional attention system. In this section, however, I expand the presentation of results to also include simulated hemodynamics from the spatial field and feature-space fields from the object representation system.

Procedure

Simulations were conducted in Matlab 7.5.0 (Mathworks, Inc.) on a PC with an Intel® i7™ 3.33 GHz quad-core processor. The ‘young’ and ‘old’ models were given the Standard version of the DCCS task while a DFT-LFP was tracked from each component of the model. The general procedure with the model followed that of the simulations in previous sections. The model was given 5 trials during 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 target 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. 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.

The DFT-LFP corresponding to the dimensional attention nodes, spatial field, and feature-space fields were tracked throughout the course of the pre- and post-switch phases. The DFT-LFP was first convolved with the impulse response function to generate a hemodynamic timecourse. The data were normalized by computing the mean hemodynamic response for each run and dividing by the maximum value across runs. This was done separately for each component of the model yielding a localized normalization similar to what is done in fMRI.

To facilitate comparison with fNIRS data presented in the following section, data are presented using a measure of proportion signal change. Further, the data are presented as an event-related time course. The relevant time window was time-locked to the onset of each trial and each trial was normalized to begin at 0 baseline value. The average hemodynamic response was then calculated across sorting phase and condition.

Results

Figure 12 shows the average hemodynamic response from the ‘young’ and ‘old’ models for the Standard and No-Conflict conditions. Significant differences between conditions are marked with a double asterisks ($p<.05$) and statistical trends are marked with a single asterisk ($p<.1$). In the dimensional attention system of the ‘young’ model (top left panel), a larger hemodynamic response is seen for the No-conflict version (in which the model correctly switches) compared to the Standard version (in which the model perseverates; $t(9)=41.72, p<.05$). This difference is driven by the same mechanism that produces correct switching behaviorally—the HTs and target inputs overlap in the post-switch feature WM field. This overlap results in stronger activation within this feature WM field and sends a stronger signal to the dimensional attention nodes. Importantly, no differences were seen in the hemodynamics of the dimensional attention system of the ‘old’ model ($t(9)=0.69$, *ns*).

Figure 12 also presents hemodynamic data from the spatial field and the feature WM fields of the object representation system. Here, there are also differences between the Standard and No-Conflict conditions. For the 'young' model, both the spatial WM field and the feature WM field that is relevant for the post-switch phase show a stronger hemodynamic response on No-Conflict trials when sorting correctly compared to Standard trials when perseverating ($t(9)=20.96$, $p<.05$; $t(9)=87.93$, $p<.05$). Finally, The only difference between conditions for the 'old' model was a larger response for the No-Conflict condition in the pre-switch feature field ($t(9)=1.83$, $p=.07$).

In summary, the DNF model predicts a stronger hemodynamic response in the No-Conflict version of the DCCS task relative to the Standard version for 3-year-old perseverators. Critically, this difference across conditions should extend beyond frontal cortex into parietal and temporal cortical regions associated with the spatial and feature-space fields. Finally, the model predicts few condition differences for switchers, with the exception of a stronger hemodynamic response in temporal cortex in the No-Conflict version.

The prediction of a robust frontal response in the No-Conflict condition for the young model is striking in that Moriguchi and Hiraki (2009) showed little frontal activation in their study. It is worth noting that this increased frontal activity is not the cause of correct switching; rather, it is a consequence of the object representation system which itself is driving correct rule switching. That is, the model succeeds in the No-Conflict version due to the cooperation, or overlap, of feature-space memories and task inputs. Note that there is no frontal difference between conditions with the 'old' model because the activation of the dimensional attention system is robust enough that the 'bottom-up' influences from the feature-space fields are swamped by the strong local dynamics. The DNF model, then, makes a detailed set of hemodynamic predictions over development. I now turn to an fNIRS experiment designed to test these predictions with 3- and 4-year-olds.

Testing Hemodynamic Predictions of the Model

The present study is aimed at testing the hemodynamic predictions of the DNF model between conditions of the DCCS task. One way to accomplish this is to build off of the method used by Moriguchi and Hiraki (2009); however, this study only included a single condition (the Standard condition). Having children participate in multiple versions of the DCCS has only been done in a limited fashion in the literature (Kloo et al., 2008; Perner & Lang, 2002). Moreover, there are constraints imposed by using fNIRS. The hemodynamic response is slow, so a design is needed that can optimize the time between trials and allow for the relevant neural signal to be pulled out for different trials. Furthermore, a larger number of trials are needed to obtain a robust estimate of the hemodynamic response for each individual. This represents new territory with many challenges. The approach I used here is to create a design that interleaves different sorting phases in order to create a continuous, event-related design. The series of phases was designed to set up the target conditions of interest: the Standard and No-Conflict conditions. That is, Standard and No-Conflict switches were interleaved with conditions that would provide the necessary pre-switch phases.

An example sequence of phases is shown in Figure 13. Each sorting phase is composed of three trials. The experiment starts, for example, by telling children to sort red and blue stars and circles by color. Next, children are told to sort by shape with the same stimuli. The phase labeled B constitutes a Standard switch based on the relationship with phase A. *This is a target phase in which 3-year-olds are expected to perseverate.* From here, children would be told to keep sorting by shape using new features and no-conflict cards. This is called a Total-Change-No-Switch (TCNS) switch (phase C). During this phase, it is expected that 3-year-olds should sort correctly because no-conflict test cards are used and the dimension does not change. Next, they would be told to switch and sort by color with the same stimuli using test cards that contain conflict (phase D). Based on the relationship between phase C and phase D, this would constitute a No-

Conflict switch. *This is the target phase in which we expect 3-year-olds to correctly switch rules.* After this, children would be given new features and would be told to sort by shape with conflict cards (phase E). This would constitute a Total-Change switch. During this phase, it is expected that 3-year-olds should sort correctly based on previous studies with this condition (Zelazo et al., 2003). From there, children would be told to sort by color with the same stimuli (phase F). This returns the design to a Standard switch phase. For each session, the task would unfold in this sequence of phases for a total of 3 iterations (providing 3 phases for each condition) with shape and color dimensions counterbalanced across iterations.

The target conditions for the hemodynamic predictions occur during phases B (Standard) and D (No-Conflict). In order to analyze these phases, however, it is necessary that children sort correctly during the just previous phases (A for Standard and C for No-Conflict) to ensure that children are actually switching rules. These phases, then, were preceded by conditions in which it was expected that 3-year-olds would sort correctly. For the No-Conflict condition, this is not necessarily a concern since the pre-switch phase involves no-conflict test cards. However, to help children sort correctly before the Standard condition, this phase was preceded with a Total-Change condition which is an easier version for 3-year-olds (Zelazo et al., 2003).

This design represents a novel approach to using the DCCS task. Typically, children perform a single set of pre- and post-switch phases, although there have been efforts to administer multiple conditions with the same children (see Kloo et al., 2008; Perner & Lang, 2002). This continuous, event-related design involves a series of interleaved switches, allowing for a comparison of performance across conditions within-subjects. If successful, this type of design could be a major innovation, providing a new way to assess performance in early childhood across a range of conditions within-subjects, and opening up a new way to measure individual differences across children.

Method

Participants

Forty children between the ages of 40- and 56-months participated in this study (*M* age = 4.03 yrs, *SD* age = 0.51 yrs; 19 males, 21 females). An additional 10 children were recruited but were not included in analysis because these children did not complete both sessions (*N*=5) or had NIRS data with excessively weak and noisy signals (*N*=5).

Stimuli and apparatus

The task was explained to children using physical trays and cards. The cards were laminated 8” by 5” flash cards which contained red or blue stars or circles. The task proper was completed on a 46” LCD television monitor that was connected to a PC running E-Prime 2.0 software (Psychology Software Tools, Pittsburgh, PA). Children were seated approximately 25” from the screen. The stimulus set was composed of 13 different sets of shapes and colors which were presented on virtual cards on a grey background. Importantly, when the colors or shapes changed between phases these features never repeated. That is, each time a set of features changed, new values were introduced rather than returning to values used in previous phases. Images of trays (subtending 15 cm by 19 cm) were displayed below the horizontal center of the screen which contained images of target cards (subtending 10.5 cm x 10.5 cm). Test card images were presented just above the center of the screen (subtending 10.5 cm x 10.5 cm; see Figure 14).

Procedure

At the start of the experiment, children were shown the trays and target cards and were told that they were going to play a series of matching games. Children were instructed for either the shape or color game (whichever dimension they would start with in the computerized task). They were told, “You’re going to play a matching game. This is called the color game. In the color game, you are going to match by color. All of the blue ones go here and all of the red ones go there.” The child was then shown a demonstration of how to sort each test card. The experimenter then gave the child 5 cards to sort, one at a time, prompting the child by saying “Where does this one go in the color/shape game”. For all trials, the experimenter repeated the rules if the child sorted incorrectly.

After the initial warm up phase, the experimenter began the computerized version. The sequence of events for a trial is displayed in Figure 14. Each trial began with the presentation of the target cards and sorting trays. At this time, on the first trial of each phase, the experimenter told the child the rules for the game (e.g., “We’re going to play the color game. In the color game we sort by color. All of the red ones go here, but all of the blue ones go there.”). The rules were also repeated after a card was sorted incorrectly. The experiment then pressed a button to initialize the trial with an auditory prompt saying, “*Let’s play the color/shape game!*”. After 1500 ms, the test card appeared above the center of the screen. The test card remained on the screen until the child responded by pointing to the location where it should be sorted. The experimenter then recorded this response by pushing a left or right response key. After the response was recorded, the screen was blanked for a minimum inter-trial interval of either 2, 4, or 6 seconds. Additional time was allowed to orient the child toward the screen and prepare for the next trial. Once the child was ready, the experimenter pressed a button to bring up the next set of target cards and initiate the next trial (i.e., *“let’s play the color/shape game!”*).

In the analyses that follow, I only include data from ‘valid switch’ phases in which children sorted correctly on the just previous phase. A valid switch phase was defined as phases where children sorted at least 2 out of 3 cards correctly on the just previous phase. Participants completed 2 runs during each of 2 visits for a maximum of 36 total trials (3 trials per phase x 3 switch phases per run x 2 runs x 2 visits) for each condition.

NIRS data collection

NIRS data were collected at 25 Hz using a 24-channel TechEn CW6 system with wavelengths of 830 nm and 690 nm. Light was delivered via fiber optic cables that terminated in a customized cap (Figure 15) placed on the head with sources and detectors secured within six flexible plastic arrays. Each array contained 2 sources and 4 detectors placed 3 cm apart creating 6-channel arrays with foam padding underneath to ensure that they would rest comfortably on the head. These arrays were placed on the head relative to the 10-20 system with one array placed over the left frontal cortex (over F5-F7), a second array over the right frontal cortex (over F6-F8), a third array placed over the left parietal cortex (over P3-P5), a fourth array placed over the right parietal cortex (over P4-P6), a fifth array placed over the left temporal cortex (over T3-T5), and a sixth array placed over the right temporal cortex (over T4-T6). Figure 15 shows the placement of the arrays on a model head created by digitizing the optode positions on a sample participant using a Polhemus motion-tracking system. NIRS data were time-stamped at the onset of the auditory prompt saying (e.g., "*Let's play the color game!*") on each trial through synchronization with the computer running the experiment using E-Prime software.

NIRS data processing

One approach to analyzing NIRS data is to only include participants who contribute data for all NIRS channels and all conditions; examination of our data, however, revealed that this was not ideal. For instance, motion artifact appeared to impact signals locally (e.g., near F5) rather than globally. Thus, I opted to exclude different numbers of trials over different cortical regions for each participant. The NIRS data were split into 8 regions composed of the 3-channel sets depicted in Figure 15 and different participants were allowed to contribute data for each region. I will refer to these 3-channel sets using the 10-20 sites they were placed over and the channel number moving from the front of the head to the back. For instance, the 3-channel array near F5 will be referred to as F5-1, F5-2, and F5-3.

An initial behavioral criterion was set to include trials in the NIRS analyses. Specifically, each phase of 3 trials was categorized as a correct or incorrect phase. Each phase of 3 trials was categorized as a ‘correct’ phase if at least 2/3 were sorted correctly; phases were categorized as an ‘incorrect’ phase if at least 2/3 were sorted incorrectly. Only NIRS data from incorrect trials were included from ‘incorrect’ phases, and only NIRS data from correct trials were included from ‘correct’ phases. This was to ensure a cleaner estimate of the hemodynamic profile associated with different trial types.

Using HomER2 software (Huppert, Diamond, Franceschini, & Boas, 2009), data were first demeaned and converted into an optical density measure. Next, the data were band-pass filtered to remove frequencies slower than .016 Hz and faster than 2 Hz. Motion artifacts were then removed from each region by eliminating trials with a change in optical density larger than 0.3 absorbance units within the time window between 2 seconds before to 12 seconds after the onset of the dimensional cue. Motion artifacts were typically caused by large movements of the head and body. Data were then band-pass filtered again to retain only frequencies between .016 and .5 Hz. Concentration data were computed using the modified Beer-Lambert Law and the known extinction coefficients of oxygenated and deoxygenated hemoglobin (Boas et al., 2001). Finally, outlier trials were removed that contained amplitudes of oxy-Hb that were more than 2.5 standard deviations above or below a participant’s mean in each condition for 9 consecutive time-samples (a duration of 360 ms). Outliers were only removed from channels and conditions for which a robust estimate of the mean could be obtained. Thus, participants had to contribute a minimum of 5 trials in order to perform outlier removal for a given condition in a given region.

Children were then grouped based on their behavioral performance in the Standard version. Children that performed lower than 50% correct were categorized as *Perseverators* while children that performed above 50% correct were categorized as *Switchers.* The fNIRS analyses only focus on the Standard and No-Conflict conditions.

For both groups, the No-Conflict condition was composed of correct trials. However, the Standard condition was composed of incorrect trials for *Perseverators* and of correct trials for *Switchers.*

To be included in the following NIRS analyses, a participant needed to contribute at least 2 trials of data to the No-Conflict and Standard conditions within a given region. Table 2 displays the number of participants, average trial counts, and the percent of included trials for each condition and region. In all cases, the average number of trials that were retained was high, allowing for robust estimates of children's hemodynamic responses.

Given that the number of trials varied by participant in each region, a weighted mean was computed for each participant in which the average hemodynamic response (HbO_2, HbR) was weighted by the number of trials. This reduced the possibility that statistically significant effects might be driven primarily by participants with few trials. For statistical analyses, the mean weighted average was computed within an 8 second time window spanning 2 seconds after the dimensional cue to 10 seconds after dimensional cue. The average time between trials was 13.7 seconds (*SD*= 3.9 s) for *Perseverators* and 10.3 seconds (*SD*= 2.6 s) for *Switchers.* This time window, then, reflects a 2 second shift from the onset of the dimensional cue to the onset of the next trial for the faster group (i.e., Switchers). Importantly this time window should capture the peak response on each trial, which typically occurs 7-8 seconds post stimulus onset (Schroeter, Zysset, Wahl, & von Cramon, 2004) and is unlikely to be influenced by the rise in hemodynamic activity associated with the next trial.

Behavioral results

Figure 16 shows the accuracy across conditions for *Perseverators* (*M* age= 3.7 yrs, *SD* age= 0.4 yrs) and *Switchers* (*M* age=4.22 yrs, *SD* age= 0.5 yrs). *Switchers* showed a main effect of condition, $F(24,3)=16.67$, $p<.05$. Tukey HSD tests indicated that performance in the Total-Change-No-Switch (TCNS) condition was significantly better than in all other conditions (*p*<.05). *Switchers*, then, sorted equally well in the target conditions, the Standard and No-Conflict. *Perseverators* also showed a main effect of condition, *F*(12,3)=68.77, *p*<.05. Tukey HSD tests showed these children sorted significantly worse in the Standard condition compared to the No-Conflict condition (*p*<.05). Thus, the target pattern of behavior needed to test the hemodynamic predictions using fNIRS data was present. *Perseverators* also sorted significantly better in the TCNS condition than in all others (*p*<.05), and better in the No-Conflict version than in the Total Change condition (*p*<.05). The only conditions that were not significantly different for *Perseverators* were the Standard and Total Change.

fNIRS results

The weighted means from each channel were analyzed using a Condition (Standard, No-Conflict) x HbX (oxy-Hb, deoxy Hb) ANOVA. Oxy-Hb and deoxy-Hb typically show an inverse relationship (Cui, Bray, & Reiss, 2010). Thus, an HbX effect indicates a robust hemodynamic response with significant differences between the oxy-Hb and deoxy-Hb signals. Table 3 lists all channels that showed a significant HbX effect (or interaction) from the omnibus ANOVA.

Recall that the central question of this study was whether perseverators would show a stronger hemodynamic response on No-Conflict trials compared to Standard trials. To examine this, I first focus on channels that showed an HbX by Condition interaction. Activation time courses for these channels are plotted in Figure 17. *Perseverators* showed a robust interaction across six frontal channels. As can be seen in Figure 17, all of these interactions showed the predicted effect with a stronger hemodynamic response on the No-Conflict version compared to the Standard version.

In contrast to these findings with *Perseverators*, *Switchers* only showed HbX by Condition interactions on the temporal channels near T3 (T3-1, T3-2, and T3-3). As can be seen in Figure 17, these channels showed a stronger hemodynamic response in the No-

Conflict condition compared to the Standard condition. This is consistent with the marginal effect seen in the model in the pre-switch field hemodynamics.

Next, I assessed channels that showed general task-related activation as indexed by an HbX main effect (see Table 3). *Perseverators* showed an HbX effect on two frontal channels (F6-1 and F8-1) and one temporal channel (T6-3). The time courses for these channels are plotted in Figure 18. *Switchers* showed robust, task-related activation in one frontal channel (F5-1), one temporal channel (T6-1), and one parietal channel (P6-2). Recall that *Switchers* did not show any HbX by Condition interactions on frontal channels, but Figure 18 shows that they did show task-related frontal activation. This is consistent with the simulation results from the ‘old’ model showing no difference in hemodynamics on the dimensional attention system between the Standard and No-Conflict conditions. Note that *Switchers* also showed a condition effect on channel P6-2. Examination of this pattern showed a stronger response on the No-Conflict condition (*M* Hb = 0.159 μM) compared to the Standard condition (*M* Hb = -0.016 μM).

The previous analyses focused on condition-related effects to evaluate the DNF predictions for each group separately. In a final set of analyses, we examined whether there were any differences between groups. This is important because there was a clear increase in the strength of the hemodynamic response for the ‘old’ DNF model in both the dimensional attention system and the post-switch feature field (see Figure 12).

The first set of analyses examined whether there were group differences in the Standard condition. Recall that Moriguchi and Hiraki (2009) reported a difference in frontal activation between *Switchers* and *Perseverators* in the Standard condition. Did we replicate their effect? An HbX (oxy-Hb, deoxy-Hb) by Group (*Perseverator*, *Switcher*) ANOVA revealed no group-related differences in activation for any channel. To visualize the relationship between *Switchers* and *Perseverators*, Figure 19A plots NIRS data from the standard condition for these groups for channel F5-1, the common channel that showed frontal activation for both groups. As can be seen, *Switchers* show a robust early response while perseverators show no response. Recall that the time window for computing the average hemodynamic response extended from 2-10 s which clearly excluded a portion of the early hemodynamic response for switchers. Thus, I conducted a post-hoc HbX (oxy-Hb, deoxy-Hb) by Group (*Perseverator*, *Switcher*) ANOVA comparing *Switchers* and *Perseverators* in the Standard condition on channel F5-1 averaging over the time window from 1 to 5 seconds. This revealed a significant Group main effect, $F(1,38)= 4.33$, $p<.05$. This analysis suggests that in addition to the regional differences in activation described above, there might be differences in the timing and duration of the hemodynamic response between *Switchers* and *Perseverators*.

A further set of analyses were conducted on the No-Conflict version to determine if this condition revealed any group differences. This analysis revealed a single channel, F6-1, which showed a difference between groups, $F(1,36)=5.028$, $p<.05$. As can be seen in Figure 19B, *Switchers* showed a significantly larger response on the No-Conflict condition compared to *Perseverators*. This was primarily driven by the difference in the HbR signal, however. Note that the presence of a group difference in frontal activation is consistent with the DNF predictions, although there was no associated change in temporal cortex that would parallel the prediction from the post-switch feature field.

Discussion

The goal of the present study was to test a set of hemodynamic predictions of the DNF model across ‘easy’ and ‘hard’ versions of the DCCS task. Specifically, the model predicted that *Perseverators* should show a robust hemodynamic response in the No-Conflict version (in which they switch rules) compared to the Standard version (in which these children fail to switch rules). *Switchers*, by contrast, were predicted to show few differences across conditions with the exception of a stronger hemodynamic response in the No-Conflict condition in temporal cortex. To test this, I designed a novel DCCS paradigm that interleaved multiple switch phases in a continuous event-related format.

The design used in this study was innovative in that multiple conditions were interleaved to facilitate a within-subject comparison of performance across conditions. This shares some similarity to other studies that have used additional sorting phases. For example Hanania (2010) interleaved two different switch types to determine if children are perseverating on specific features or on dimensions. Specifically, the task was constructed with the standard pre- and post-switch phases. A third phase used the same dimension as the post-switch phase, but changed all of the features on the cards (similar to the Total Change version). Children who perseverate on specific features would then sort well during this phase, but children who are perseverating on dimensions would continue to sort by the pre-switch dimension. The design of the NIH Toolbox version of the DCCS (Zelazo et al., 2013) also includes multiple switch phases. This version is designed with the standard pre- and post-switch phases and then includes a phase where the shape and color games are randomly assigned within a block of trials.

The paradigm used in the current study successfully replicated important aspects of performance in the literature across different conditions. First, children were required to sort correctly during the pre-switch phases of the target conditions, Standard and No-Conflict. Second, one group of children (*Perseverators*) performed poorly on the Standard version but sorted significantly better in the No-Conflict version. Another group of children (*Switchers*) performed well in both the Standard and No-Conflict conditions. Thus, this paradigm elicited the target pattern of behavior which allowed for the direct comparison of performance across different conditions within-subjects. One interesting aspect of these data was that *Perseverators* also had difficulty in the Total Change version. This suggests this version may be more difficult in the context of the more demanding task design used in the current study (see also, Hanania, 2010). In the context of the DNF model, this seems plausible. In contrast to the No-Conflict version, the Total Change condition (see Figure 9C) does not have HTs in the post-switch for the features that are relevant in the post-switch phase. This is because all of the features are changed during this switch phase. Thus, this field is in a neutral state given the configuration of HT's and target inputs. The No-Conflict version, however, has HTs that overlap and act in cooperation with the target inputs in the post-switch feature field. Without this extra support from the HTs, the Total Change version may be at a disadvantage relative to the No-Conflict version.

This new task design also achieved another important goal of obtaining a robust number of trials for the target conditions to allow for averaging of the fNIRS data. The fNIRS data revealed task-related activity across a range of regions previously implicated in flexible rule-use. Since Moriguchi and Hiraki (2009) only recorded from frontal cortex in their fNIRS study, this is the first time that the DCCS task has been shown to activate regions outside of frontal cortex in early childhood.

Recall that the model made particularly striking predictions regarding the role of frontal cortex in flexible rule-use during early childhood. First, *Perseverators* should show stronger activation in the No-Conflict condition (when sorting correctly) compared to the Standard condition (when sorting incorrectly). This effect was most prominent in the dimensional attention system (corresponding to frontal cortex), but was also seen in the spatial field and post-switch feature field (corresponding to parietal and temporal regions). In the fNIRS data, this predicted effect was robust across six frontal channels. Note that some of these channels showed a negative-going HbO time course. This may reflect the fact that trial averages are all calculated as changes relative to time point 0. It could also be a 'sink' effect that results from oxy-hemoglobin being drawn to other nearby regions, creating a local decrease in oxy-hemoglobin concentration.

Beyond frontal regions, the condition effects in the fNIRS data were not as pervasive as in the model. The 'young' model showed larger hemodynamic responses on the No-Conflict condition compared to the Standard condition in the spatial field and the posts-switch feature field. However, there were no condition effects in parietal or temporal regions with *Perseverators*. This may reflect a real mismatch between the model and the data. It may also be a result of the lower trial and subject counts in these regions compared to frontal regions. Obtaining clean signals from these channels proved difficult due to the curvature of the head. Thus, a more detailed evaluation of activation in posterior brain regions in *Perseverators* will require a new data set with a better NIRS cap design. For now, the conclusions regarding parietal and temporal activation with *Perseverators* will have to remain tentative.

Regarding *Switchers*, the 'old' model did not show a difference between conditions on the dimensional attention system. This was consistent with the fNIRS data—*Switchers* showed task related activation on a common channel with *Perseverators,* channel F5-1, but no difference between conditions. The only difference between conditions in the model was an effect within the pre-switch feature field. In this case, there was a stronger hemodynamic response on the No-Conflict condition compared to the Standard condition. *Switchers* showed this pattern on 3 temporal channels (T3-1, T3-2, T3-3). Switchers also showed a condition effect on a parietal channel (P6-2); however, no corresponding difference was seen in the hemodynamics of the spatial field of the 'old' model.

The data from Moriguchi and Hiraki (2009) showed that *Switchers* activate frontal cortex more strongly than *Perseverators* when switching rules. This close alignment of neural and behavioral data suggests that the developmental state of the frontal cortex is critical for rule-switching between 3- and 5-years-old. By this view, 3-year-olds perseverate because they have an immature frontal cortex. In this regard, it is rather remarkable that *Perseverators* showed greater activation on the No-Conflict condition compared to the Standard condition.

In the model, this effect is driven by the state of the post-switch feature field which has increased excitation due to overlapping HTs and target inputs. Given the reciprocal excitatory connectivity between the dimensional attention system and the object representation system, this creates a stronger input into the dimensional attention system which results in a larger hemodynamic response in this condition. Thus, frontal activation in this condition is not the cause of correct rule switching; rather, it is a consequence of robust activation in the posterior cortical fields.

One concern with the current dataset is that there was not a robust replication of the pattern reported by Moriguchi and Hiraki (2009). There was a significantly larger response from *Switchers* compared to *Perseverators* in the Standard condition on channel F5-1 when the average time course was restricted to a 1-5 second time window. Additionally, *Switchers* showed a significantly larger response on channel F6-1 in the No-Conflict condition compared to *Perseverators*. Several factors might account for the weak differences in frontal activation between *Switchers* and *Perseverators*. First, Moriguchi and Hiraki (2009) used a block design over the course of a 25 second time window during the pre- and post-switch phases while the current study focused on a 10 second time window that was time locked to individual trials. Inspecting the time course of activation presented by Moriguchi and Hiraki (2009), there are clear upward trends in the magnitude of the hemodynamic response toward the end of each 25 second phase. Thus, the differences across studies might reflect the use of an event related design here and a block average design in Moriguchi and Hiraki (2009).

Although the replication of the Moriguchi and Hiraki (2009) result was tenuous, there are broad consistencies between the current findings and results from the fMRI literature. For instance, *Perseverators* showed more widespread frontal activation in the present study than *Switchers*. This is consistent with data showing a shift from diffuse to focused frontal activation over development (Crone, Donohue, et al., 2006). In addition, Morton et al. (2009) reported left lateralized activation in frontal cortex in the DCCS task, consistent with the activation from channel F5-1 in the current study. The most striking difference between *Switchers* and *Perseverators* was in the constellation of activation in parietal and temporal cortices. Specifically, *Perseverators* showed task-related activation on only a single channel outside of frontal cortex (T6-3) while

Switchers showed task related activation on a parietal channel (P6-2) as well as 4 temporal channels (T3-1, T3-2, T3-3, T6-1). This is also consistent with fMRI data reported by Morton et al. (2009) who found the most robust activation in parietal and temporal cortex when switching rules in the DCCS. Further, developmental differences were found in superior parietal cortex along with superior frontal sulcus. Thus, the developmental shift from perseveration to rule-switching may involve not only differential activation of frontal regions, but also greater activation of parietal and temporal areas.

Conclusions

Chapter 3 showed that the DNF model could explain developmental changes in EF in the context of DCCS task. The model captures behavioral data well and novel behavioral predictions have been successfully tested. The goal in the present study was to ask whether we could use the DNF model to explain not just behavioral data, but also neural data. Toward this end, I applied an approach that involved simulating LFPs from the model to generate hemodynamic time courses from different components of the model. I first tested whether this approach could capture data from the literature: the model successfully reproduced the developmental shift in frontal hemodynamics associated with the transition from perseverating to switching (Moriguchi & Hiraki, 2009). I then generated a pattern of novel neural predictions (Figure 11) regarding the relationship between ‘easy’ and ‘hard’ conditions, and tested these predictions using an innovative paradigm. Behavioral data from this paradigm produced the expected pattern of results. The key question was whether the neural data were consistent with the model predictions. The fNIRS data revealed broad activation in frontal, parietal, temporal regions which is consistent with the fMRI literature (Crone, Donohue, et al., 2006; Morton et al., 2009). Moreover, the NIRS data were consistent with the prediction that the No-Conflict condition should elicit a stronger hemodynamic response relative to the

Standard condition for *Perseverators*, and that this effect should only be seen in temporal cortex for *Switchers*.

The central issue moving forward is to determine how distributed these effects are across the frontal-parietal-temporal network, especially with *Perseverators*. Future work will be needed to assess that given limits in the recording techniques used here. Future work would also benefit from moving from a channel-wise analysis to an image-based analysis. Despite these limitations, the current results are a promising first step. The DNF model appears to be a useful tool for understanding the development of the neural mechanisms that underlie EF in early childhood.

Table 1. *Developmental parameter differences in the dimensional attention system.*

	‘Young’ Model	‘Old’ Model
Self-Excitation	0.2	3.0
Lateral Inhibition	0.75	20
Feature Field → Attention Neuron	0.285	1.000
Attention Neuron → Feature Field	0.5	5.0

Table 2. *Subject numbers and trial counts across regions.*

	Perseverators		Switchers	
Region	N	nTR	N	nTR
F5	14	9.7	26	14.3
F6	14	8.4	24	14
F7	6	7.5	16	8
F8	9	7	18	9.8
P3	10	5.1	18	8.3
P4	3	7.2	9	8.5
P5	9	6.1	12	9.8
P6	7	8.1	17	8.4
T3	7	6.4	16	10.7
T4	5	7.8	14	9.5
T5	3	6.2	16	6.9
T6	6	6.8	18	9.7

Table 3. *Results of the omnibus ANOVA*

Omnibus ANOVA Channel	**Perseverators**		**Switchers**	
	Effect	**F(dfeffect,dferror)**	**Effect**	**F(dfeffect,dferror)**
Frontal				
F5-1	HbX*Cond	*F*(1,13)=4.69	HbX	*F*(1,25)=10.83
F5-2	HbX*Cond	*F*(1,13)=14.79		
F5-3	HbX*Cond	*F*(1,13)=8.26		
F6-1	HbX	*F*(1,13)=7.94		
F6-3	HbX*Cond	*F*(1,13)=4.90		
F7-2	HbX*Cond	*F*(1,5)=8.54		
F7-3	HbX*Cond	*F*(1,5)=9.75		
F8-1	HbX	*F*(1,8)=6.34		
Parietal				
P6-2			HbX	*F*(1,16)=5.01
			Cond	*F*(1,15)=4.84
Temporal				
T3-1			HbX*Cond	*F*(1,15)=4.43
T3-2			HbX*Cond	*F*(1,15)=8.03
T3-3			HbX*Cond	*F*(1,15)=4.61
T6-1			HbX	*F*(1,17)=4.38
T6-3	HbX	*F*(1,5)=16.51		

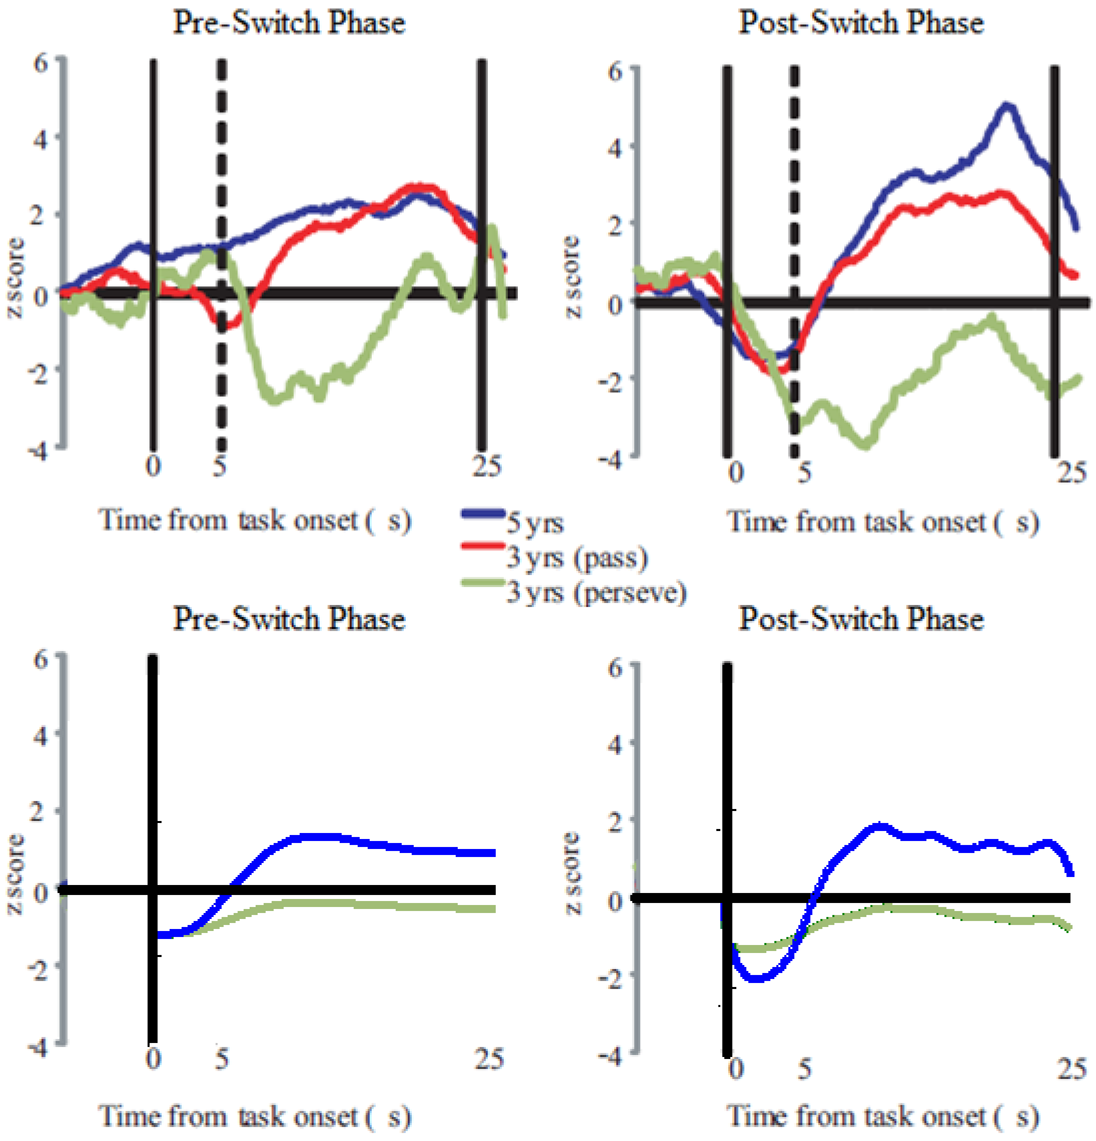
Pre-Switch Phase

Post-Switch Phase

Pre-Switch Phase

Post-Switch Phase

Figure 11: Hemodynamics of the DCCS.

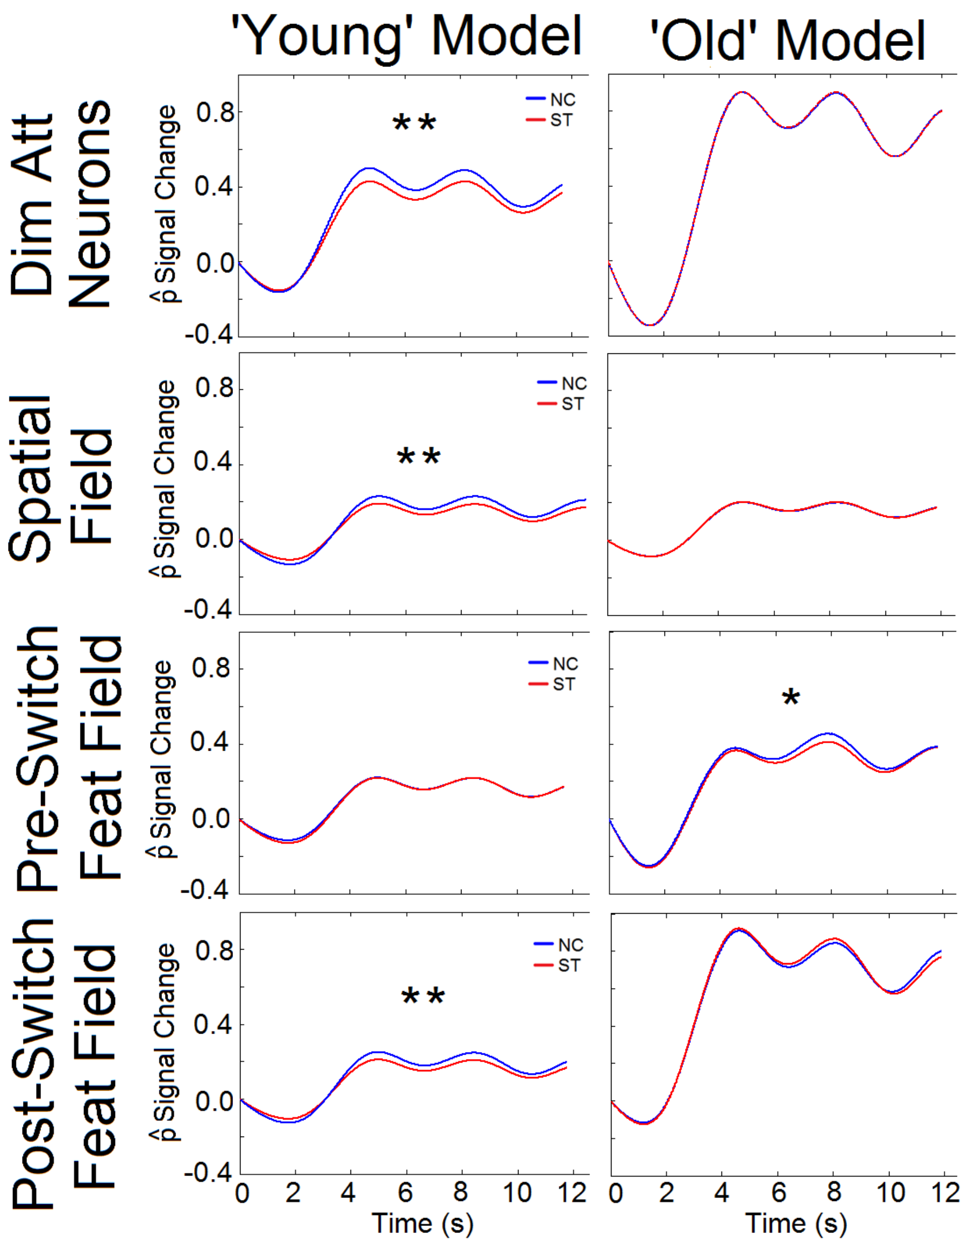

Figure 12: Model hemodynamic predictions.

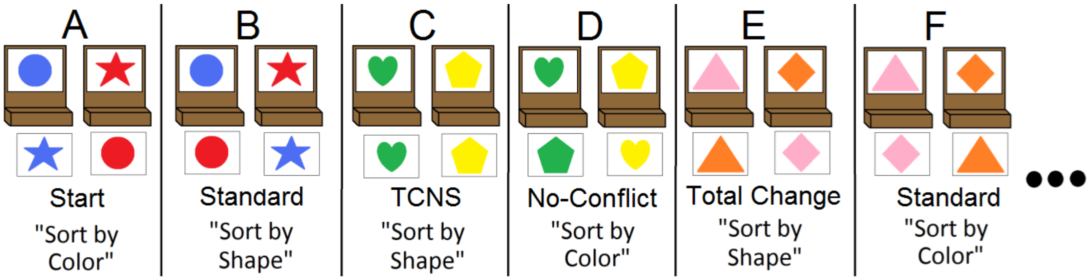

Figure 13: Conditions used in the fNIRS experiment.

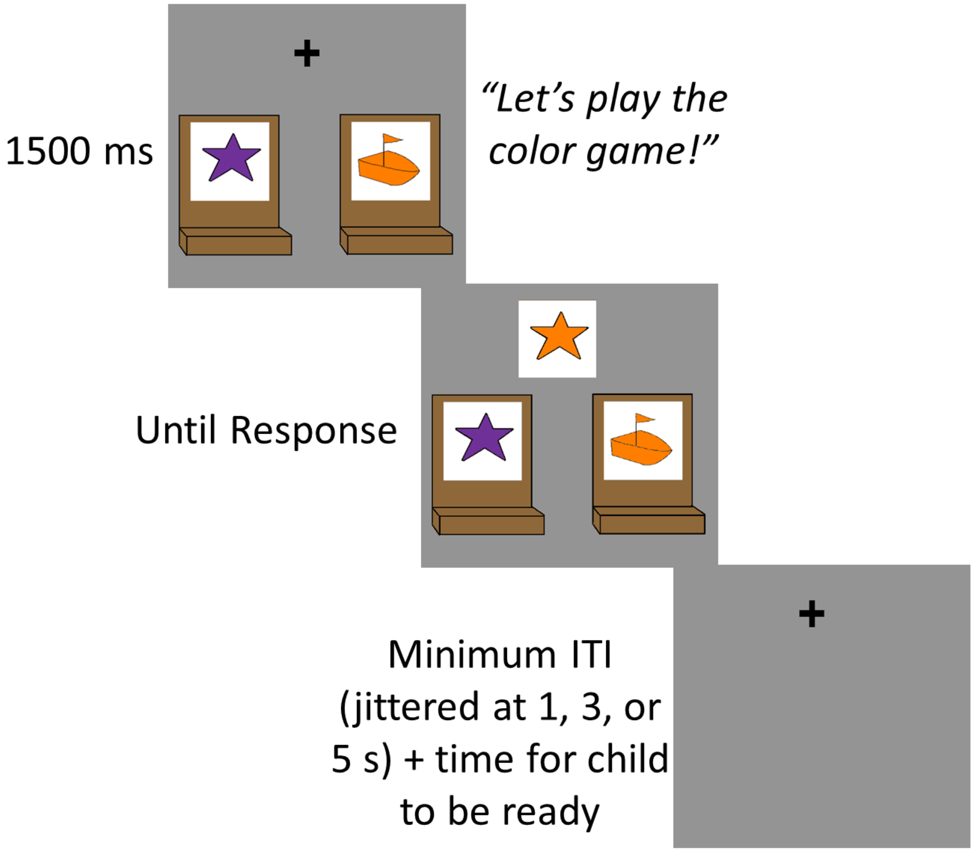

Figure 14: Depiction of the experimental display used in the NIRS study.

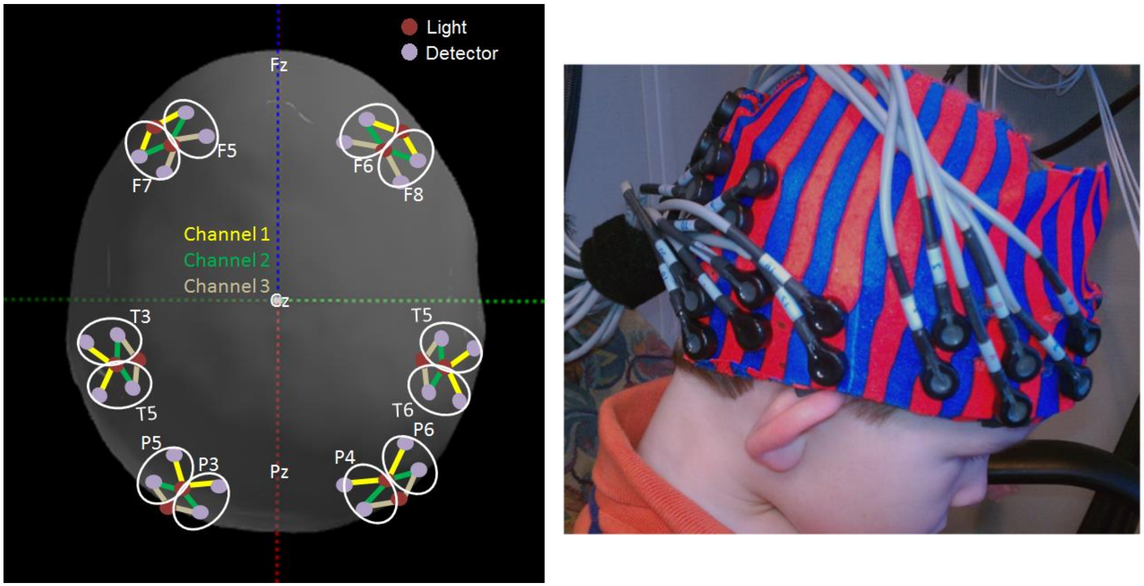

Figure 15: fNIRS probe.

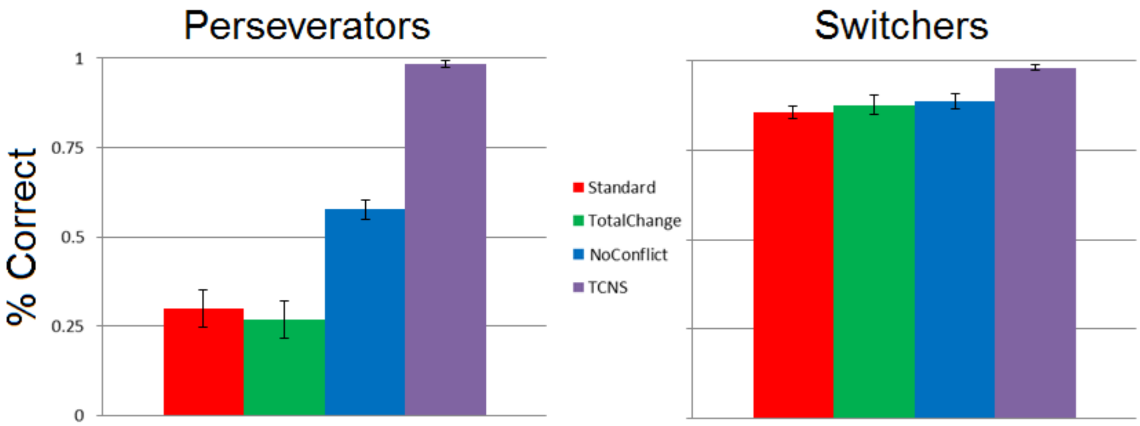

Figure 16: Behavioral data from fNIRS study.

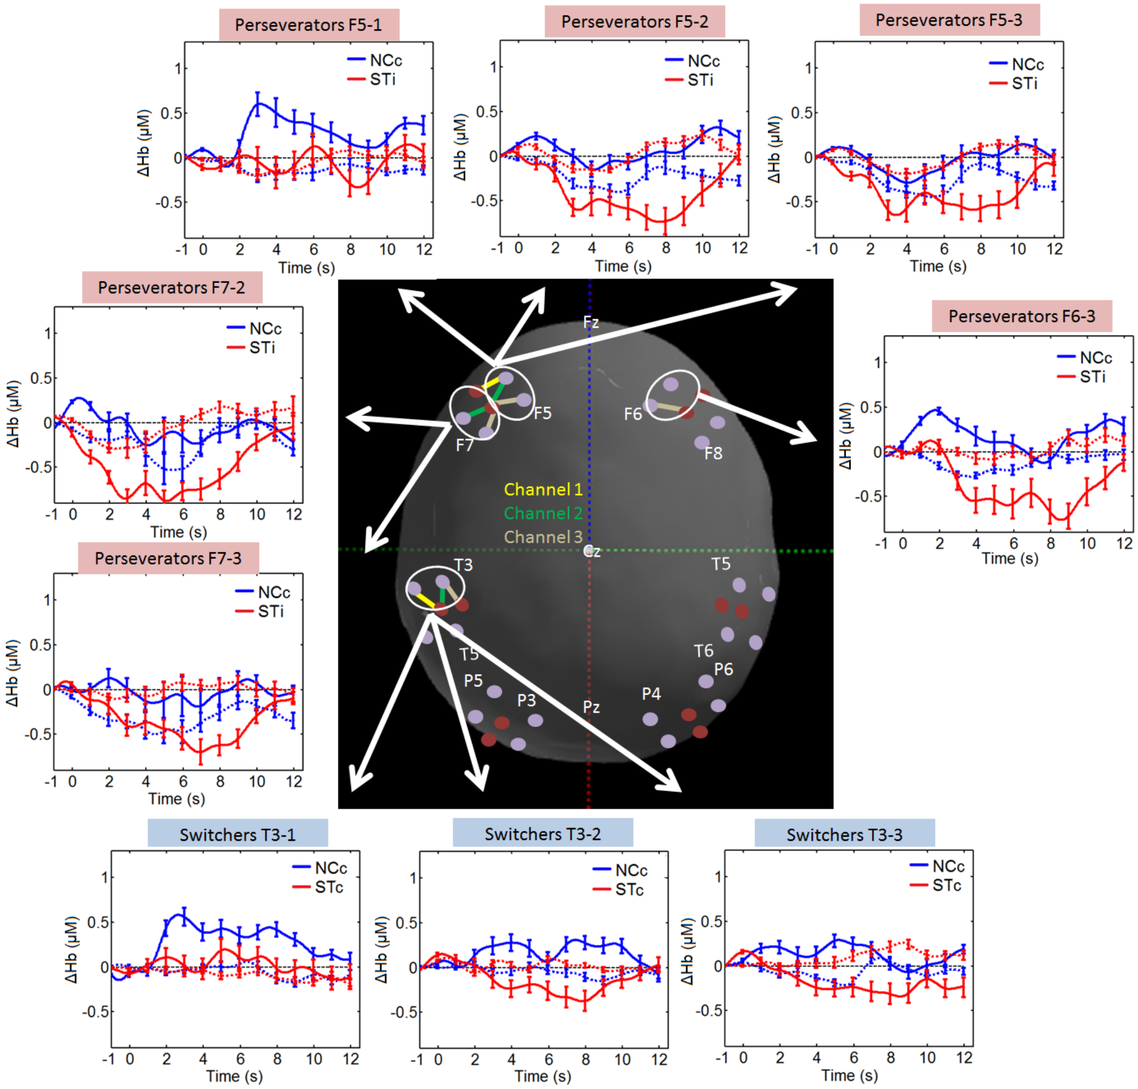

Figure 17: HbX by condition interactions.

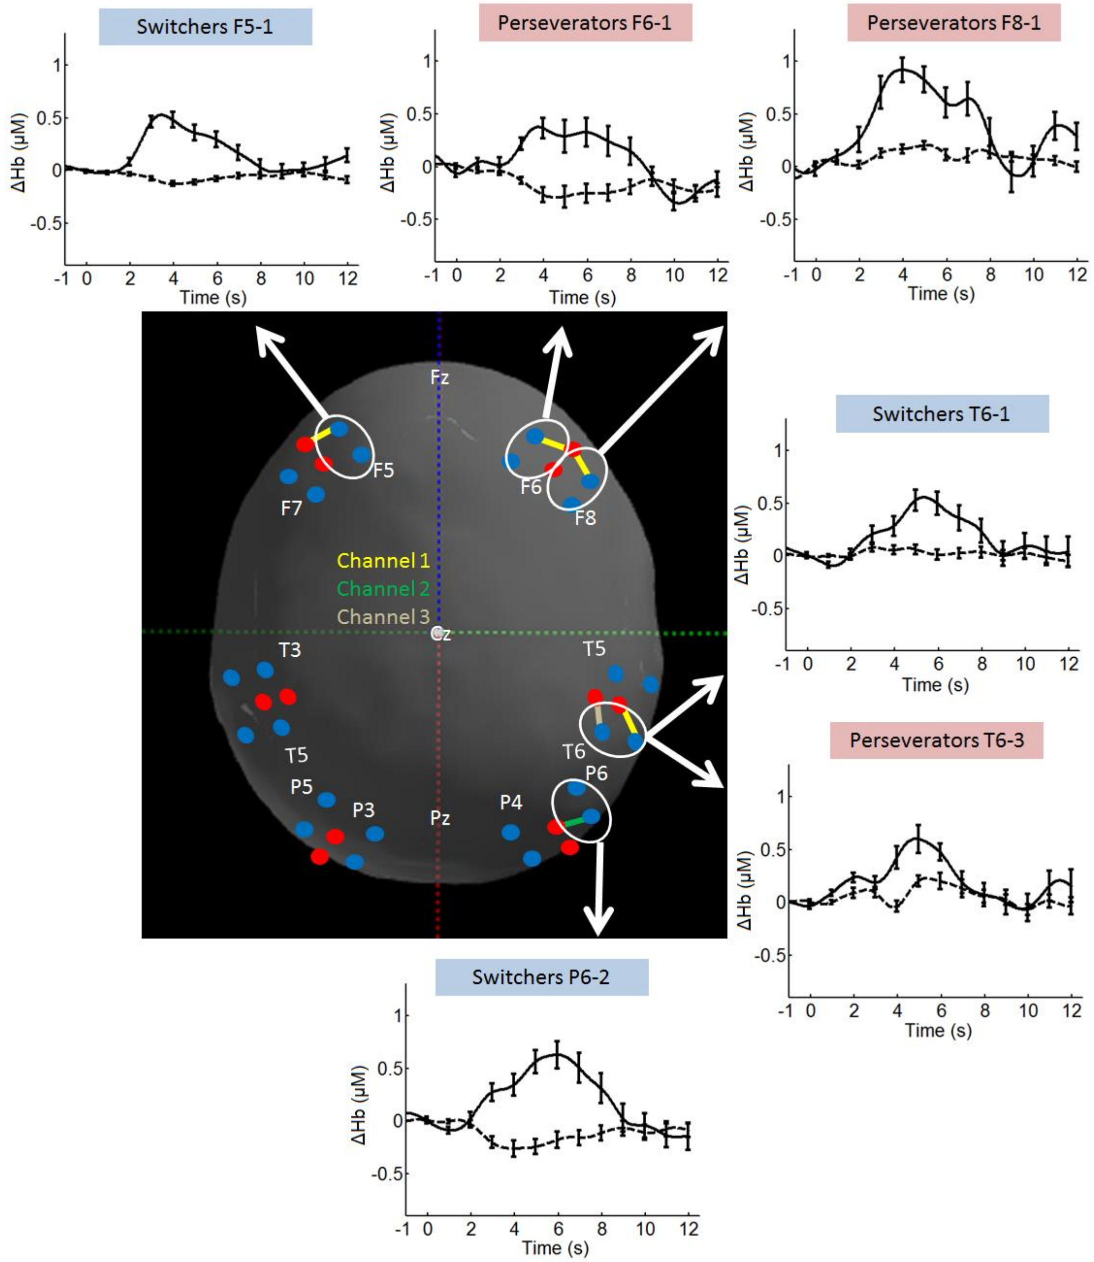

Figure 18: Oxy main-effects.

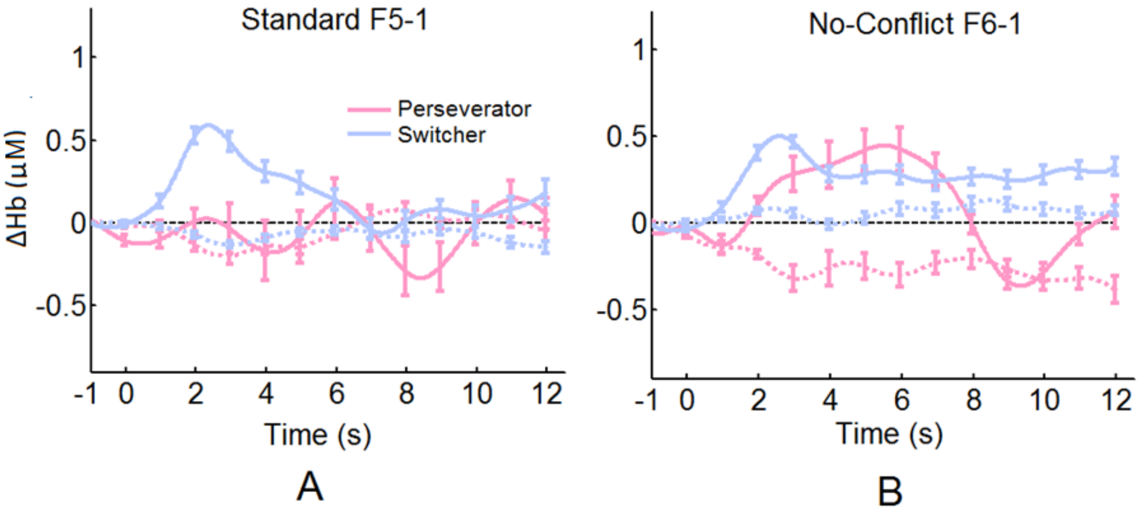

Figure 19: Comparing switchers and perseverators.

CHAPTER 5

STUDY 2: USING fMRI TO PROBE THE DYNAMICS OF FLEXIBLE RULE-USE IN ADULTHOOD

Chapter 4 focused on a key transition in EF in early development between 3- and 4-years. Although some of the most drastic changes in EF occur during this critical early window, changes in performance, neural activation, and neural connectivity continue into adulthood (Crone, Donohue, et al., 2006; Ezekiel, Bosma, & Morton, 2013; Fair et al., 2008; Morton et al., 2009). To speak to these issues, the DNF model of flexible rule-use needs to be generalized to performance in adulthood. This will provide an opportunity to test the model using new behavioral measures such as reaction times and quantitative switch costs as well as the gold-standard for neuroimaging—fMRI. Extension of these theoretical concepts to the adult literature will also help probe the limits of the developmental changes hypothesized to underlie the emergence of EF.

The long-term goal here is to be able to establish a mapping between model components and neural regions. Recall that rule-switching occurs in the dimensional attention system of the DNF model as the ‘color’ and ‘shape’ nodes compete, while post-switch conflict is localized to the feature-space fields. As an initial extension of the DNF model to later development, I adopted a common approach used in the fMRI literature. In particular, I identified two key factors that influence neural activation in the model—whether a dimension-switch is required and whether there is featural conflict along the post-switch dimension—and varied these factors in a two-by-two design in order to use a standard subtraction method to identify brain regions that are sensitive to these manipulations. The design of the fMRI study, then, sets up particular contrasts of interest in an effort to localize these two aspects of EF to neural systems in the brain. These would then be candidate neural sites for the dimensional attention system in the DNF model and the posterior feature-space fields. Note that the current study also opens the door to future work, building on the work from Chapter 4. For instance, a future goal will be to quantitatively simulate adults' behavioral data with the DNF model and then ask whether simulated model hemodynamics using the LFP approach can capture the fMRI findings.

In the section below, I first review what is currently known about the neural dynamics that underlie flexible rule-use in adulthood. I then describe the specific design used here to localize two key aspects of the DNF model using fMRI contrasts.

Neurocognitive Dynamics of Flexible Rule-use in Adulthood

Previous research has explored the DCCS task with adults. This work has demonstrated that adults are both significantly more error prone and slower to respond when switching dimensions compared to when the shape or color rules are repeated. Morton et al. (2009) used fMRI with 11- to 13-year-olds and adults and found greater activation on switch blocks compared to repeat blocks in right superior parietal and intraparietal sulcus, bilaterally in dorsolateral prefrontal cortex, right inferior frontal junction, and right pre-motor cortex. Although there were no differences in behavioral performance between these two age groups, there were developmental changes within a network of frontal, parietal, and temporal areas. Specifically, left superior parietal cortex and fusiform gyrus showed greater activity on switch blocks only for adults. However, right superior frontal sulcus showed greater activity on switch blocks only for children. This suggests that there may be decreases in activation strength in particular areas as functional specialization emerges (see also, Rubia et al., 2006). Finally, the thalamus, which is an intermediary between cortex and basal ganglia (Alexander, DeLong, & Strick, 1986) showed greater activity on switch blocks for adults but not children.

These data clearly show that there is a complex pattern of neural activity that underlies rule-switching in adolescence and adulthood, and that aspects of this network change across adolescence. But, to date, the theoretical understanding of these complex neural patterns has been limited. Part of the challenge is that there have been only a handful of fMRI or behavioral studies that have used rule-switching tasks like the DCCS with adults. This stands in sharp contrast to the vast literature on the DCCS task in early development. The other challenge is that the DCCS task is itself quite complex, requiring rule-switching, inhibition, and working memory.

The goal of the present study is to examine whether the DNF model can shed unique light on the neural dynamics of rule-switching in the DCCS task. Two aspects of the model are a focus for the current experiment. One aspect is dimensional switching which occurs within the dimensional attention system. When the model is required to make a rule switch, the neural activation of the post-switch node (e.g., ‘color’) must out-compete the neural activation of the pre-switch node (e.g., ‘shape’). This requires a boost to the post-switch node and strong inhibition to suppress the pre-switch node, both of which should produce a strong BOLD response which is sensitive to both excitatory and inhibitory processes. Manipulating the need to switch dimensions, then, should reveal differential activation of neural systems involved in dimensional attention when there is a change in dimensions compared to not. In Chapter 3, I hypothesized that the dimensional attention system in the DNF model is localized to areas of the frontal cortex that includes inferior frontal junction (IFJ; see Figure 3). Thus, I expect to see greater IFJ activation when a dimensional switch is required versus not.

A second aspect of the model is featural competition between target inputs and Hebbian traces (HTs) which occurs within the posterior feature-space fields. Manipulations to this factor by, for instance, introducing new features and eliminating conflict, should reveal differential activation of neural systems involved in feature-space conflict when there is competition compared to not. Competition in the feature-space fields should slow down the growth of neural activation. By contrast, the absence of competition should yield a fast rise in neural activation. In Chapter 3, I hypothesized that the feature-space fields are localized to ventral visual areas including the fusiform color area and the lateral occipital complex (see Figure 3). Thus, I expect to see greater temporal cortex activation when there is no conflict versus when conflict is present. Note that in the DNF model, competition can live in the pre-switch field or the post-switch field. Thus, I created a design that only manipulated competition in the post-switch field and maintained cooperation in the pre-switch field for all conditions.

Based on these two factors, the present design includes 4 conditions (see Table 4). Two conditions have a dimension change—Standard (ST) and Partial-Change (PC)—while the other two conditions—SpaceSwap (SS) and SpaceSwap Partial-Change (SSPC)—do not. Similarly, two conditions have post-switch feature conflict—ST and SS—while the other two conditions do not—PC and SSPC. Note that in the current design, the absence of post-switch feature conflict is created by changing the feature values of the target and test cards along the post-switch dimension. Thus, post-switch conflict and the absence/presence of a feature change are synonymous (see Table 4).

Figure 20 shows a sequence of stimuli that illustrate the different conditions that will be used for this study. Panels A-E represent individual phases of the experiment. Each phase is composed of 3 trials; only the switches between phases are shown in the figure for simplicity. The top of each panel shows an example test display (target images are below the test image which is presented at the center top) along with the condition name and the dimension to be used on each phase. The bottom of each panel shows a cartoon of the model illustrating the state of the feature-space fields at the start of each condition. Here, target inputs are represented by ovals and HTs are represented by shaded circles.

The example sequence starts with an initial trial (Start) to set up the pre-switch condition for the first phase change which occurs in Figure 20B. This first switch is to a Standard (ST) condition. This involves a dimension change and competition within the post-switch feature field (i.e., the shape field). Next, there is a SpaceSwap-Partial-Change condition (Figure 20C). This condition does not require switching dimensional attention but instead spatially swaps the locations of the target images. This version has a neutral post-switch feature-space field (i.e., the shape field) because the relevant features (i.e., the shapes) are changed. The third condition, a Partial-Change (PC; Figure 20D), involves a dimension change but a neutral post-switch field (because the features that are relevant for the post-switch phase—in this case, color—are changed). Fourth, there is a SpaceSwap (SS) condition (Figure 20E) which does not involve switching dimensions but does contain conflict within the post-switch field (i.e., the color field). Note that the SpaceSwap condition then serves as the pre-switch phase for a Standard condition (not shown), and the design then loops through the sequence of phases again, but with new stimuli (for details, see methods).

These conditions will allow for the identification of the neural networks involved in switching dimensional attention and resolving featural competition. Using standard fMRI subtraction techniques for comparing the neural activation associated with particular pairs of conditions, the neural systems involved in shifting dimensional attention can be identified by contrasting conditions that require shifting dimensional attention (i.e., the Standard or PC versions) with conditions that do not require shifting dimensional attention (i.e., the SS or SSPC conditions). Further, the neural systems involved in resolving feature-space competition can be revealed by contrasting conditions that contain feature competition (i.e., the Standard or SS conditions) with condition that do not (i.e., the PC or SSPC versions). In this way, these four conditions can isolate the dimensional attention system and the feature-space fields from the model. This research constitutes the first attempt to apply the DNF model of the DCCS task to the adult literature and fMRI and probes whether this theoretical account can provide clarity on the complex pattern of results evident in the fMRI literature on this task.

Method

Participants

Twenty subjects participated (*M* age =23.8 yrs, *SD*= 3.8 yrs; 9 males) for monetary compensation. All subjects were right-handed and reported normal or corrected-to-normal vision. The data were collected as part of a larger 6-session fMRI study. The DCCS sessions presented here were collected over 4 half-hour sessions.

Behavioral procedure

This study uses a continuous event-related design with interleaved switch phases, similar to the design used with children in Chapter 4. For this study, colors are sampled from CIE Lab color space (see Figure 21). Shape stimuli are generated from a continuous space defined by the phase angle of Fourier components (Drucker & Aguirre, 2009; Zahn & Roskies, 1972; see Figure 21). This allows the metric details of shape to be controlled in a similar fashion as color. The objects used in the task were approximately 60 pixels by 60 pixels and consisted of a solid black outline filled with color (see Figure 21).

Figure 21 shows the sequence of events on a given trial. Each trial began with an ‘S’ or ‘C’ presented at the center of the screen to indicate whether to match by shape or color on the upcoming trial. In addition, the target images were presented to the left and right of center. After 1000 ms, the dimensional cue was removed and replaced with an ‘X’. After 1500 ms, a test image appeared above the center of the screen. Participants were given 1500 ms to press a button with their right index or middle finger, mapping onto a leftward or rightward sorting response. If a response was not entered during this time window, a warning appeared on the screen telling the participant to respond more quickly. Each phase consisted of 3 trials and participants completed 10 phases of each condition (counterbalancing shape and color) in each run with 4 total runs.

fMRI data collection and analysis

The fMRI study used a 3T Siemens TIM Trio system with a 12 channel head coil. Anatomical T1 weighted volumes were collected using and MP-RAGE sequence. Functional BOLD imaging was acquired using an axial 2D echo-planar gradient echo sequence with the following parameters: TE=30ms, TR=2000ms, flip angle= 70°, FOV=240x240mm, matrix=64x64, slice thickness/gap=4.0/1.0mm, and bandwidth=1920Hz/pixel. Standard preprocessing was used that included slice timing correction, outlier removal, motion correction, and spatial smoothing (Gaussian FWHM=5mm). First level analysis was performed using multiple linear regression, which included regressors for trial type (switch and repeat) and condition (ST, SS, PC, and SSPC), as well as motion parameters as nuisance regressors. An initial contrast of switch versus repeat was performed for each condition. These activation maps were used to define a region of interest for comparing activation across conditions. Percent signal change was estimated for switch trials and repeat trials from the individual subjects using regions of interest defined at the group level using ANOVA.

Behavioral results

Response times were averaged separately for switch trials (the first trial of each 3-trial block) and for repeat trials (the remaining 2 trials in a 3-trial block). Figure 22 shows plots of average RT's for each condition and trial type. A Condition (ST, SS, PC, PCSS) x Repeat (switch, repeat) ANOVA revealed a main effect of Repeat, $F(1,19)=7.60$, $p=.013$, with participants performing slower on switch than repeat trials. The ANOVA also revealed a main effect of Condition, $F(1,19)=4.58$, $p=.016$. A Tukey HSD test revealed that responses in the Standard condition were significantly slower than responses in the SpaceSwap-Partial-Change condition ($p<.05$). No other contrasts reached statistical significance, which is not surprising given the gradual nature of the change in RTs across conditions.

fMRI results

As a first step in the fMRI analysis, I conducted individual contrasts comparing switch and repeat trials for each condition. Table 5 lists significant clusters of activation ($p < .05$, corrected) and Figure 23 shows clusters of activation in different brain regions for each condition. The Standard condition was associated with activation in left superior parietal, left middle frontal, and left inferior frontal gyrus. The Partial-Change version was associated with a broader network of regions, also including bilateral fusiform and middle occipital activations. Interestingly, there were no switch related areas associated with the SpaceSwap condition. The SpaceSwap-Partial-Change version elicited a pattern of switch-related activity in bilateral fusiform and left superior parietal cortex. Thus, conditions requiring a switch in dimensional attention (ST and PC) showed increased activation in middle frontal and inferior frontal areas on switch trials. Conditions that eliminated feature-space conflict, on the other hand, were associated with stronger activation in fusiform areas on switch trials. All conditions, except the SS condition which showed no switch-related activity, also activated parietal cortex on switch trials.

Next, a mask was made from the union of significant clusters in the switch versus repeat contrasts from each condition. This was used to define a region of interest in which to compare Switch trials across conditions ($p < .01$ cluster-wise threshold; see Table 6) and to compare switch-effect related activation (Switch – Repeat) across conditions ($p <$ .005 cluster-wise threshold; see Table 7). These two sets of contrasts yielded generally consistent results, but switch-effect comparisons revealed a broader network of robust effects; thus, I focus on these contrasts below. Figure 24 shows the clusters of differential activation revealed by the switch-effect contrasts across conditions. Note that this mask was also used to compare activation on Repeat trials across conditions; however, this did not reveal any differences in activation.

In the first set of contrasts, I examined the effect of switching dimensions. The Standard condition (dimension change + feature-space competition) showed significantly stronger activation on switch trials in inferior frontal cortex compared to the SpaceSwap condition (no dimension change + feature-space competition). The Partial-Change condition (dimension change + no feature-space competition), however, did not elicit significantly stronger activation in any regions compared to the SpaceSwap-Partial-Change condition (no dimension change + no feature-space conflict). The Standard and SpaceSwap conditions both involve feature-space competition, and it is in this context that a difference in switch-trial activation is seen. Activation within the dimension switching network, then, is sensitive to the influence or demands of feature-space competition.

Next, I examined the effect of changing features and eliminating feature-space competition. These contrasts should reveal the network involved in resolving feature-space competition. The Partial-Change condition showed greater switch-related activation within fusiform cortex compared to the Standard condition, and the SpaceSwap-Partial-Change condition also elicited stronger activation in fusiform cortex compared to the SpaceSwap version. Thus, a feature change which eliminates feature-space competition is associated with a larger hemodynamic response in fusiform areas compared to the presence of feature-space conflict.

Next, I compared conditions that differed along both factors. The contrast between the SpaceSwap-Partial-Change condition and the Standard condition revealed significantly stronger switch-related activity in fusiform cortex. Again, a feature change which eliminates feature-space competition is associated with stronger fusiform activation. Finally, there were particularly widespread differences between the Partial-Change and the SpaceSwap conditions across the entire rule-use network identified in the initial contrasts between Switch and Repeat trials. In this case, these effects are driven by a combination of factors since this contrast compares a situation with a dimensional switch and no feature-space competition with a condition that does not require a shift in dimensional attention and has feature-space competition.

Discussion

The present experiment was designed as an initial step toward a model-based approach to the DCCS with adults using fMRI. Although there is a vast literature on the DCCS task in early childhood, little is known about the behavioral and neural dynamics of performance in the DCCS in adulthood. The aim of this study was twofold: to explore whether the model can be used to gain insight into the neurocognitive dynamics of DCCS performance in adulthood, and to establish a behavioral and neural dataset that can be used in future work to parameterize a model of adult performance in the DCCS using both behavioral and neural measures.

Two factors were manipulated which are critical to the performance of the DNF model: dimension switching and the presence of feature-space competition within the post-switch feature field. In the behavioral data, as in previous studies of the DCCS with adults (e.g., Diamond & Kirkham, 2005; Morton et al., 2009), robust behavioral switch costs were observed—Switch trials were performed more slowly compared to Repeat trials. A main effect of condition was also found, and follow-up tests revealed that the Standard condition was performed more slowly than the SpaceSwap or SpaceSwap-Partial-Change conditions. This is the first study to reveal differential performance across variations of the DCCS task, suggesting that adult performance is susceptible to similar influences as children. In this case, the condition that involved both a change in dimension as well as feature-space competition was slower than both conditions that did not require a change in dimension, suggesting that switching dimensional attention is a primary contributor to performance. However, the Partial Change condition, which also requires changing dimensions, but does not contain feature-space competition, was not significantly different from conditions that did not require changing dimensions. This suggests that the effect of dimension switching is modulated by the presence of feature space competition.

As in previous fMRI studies using the DCCS task with adults (Ezekiel et al., 2013; Morton et al., 2009), Switch trials activated a network of frontal, parietal, and fusiform brain regions more strongly than Repeat trials. The central question here was whether there would be selective activation corresponding to the set of *a priori* predictions made based on the mapping of the DNF model to neural regions. Specifically, the dimensional attention system of the model corresponds to processes in frontal cortex. Thus, frontal cortex activation should reflect sensitivity to demands on the dimensional attention system. Further, the feature-space system of the model corresponds to visual object regions in temporal fusiform areas. Thus, these areas should reflect sensitivity to the presence or absence of feature-space competition.

These general predictions were supported by the fMRI data. One condition that required a shift in dimensional attention activated regions in frontal cortex more strongly than an associated condition that did not require this shift (St > SS). By contrast, there was no differential activation when a dimensional shift was required in the absence of post-switch conflict (PC = SSPC). This suggests that that dimensional attention system is sensitive to the presence of feature-space competition and responds more strongly when feature-space competition needs to be resolved. At face value, these findings appear consistent with the reciprocal coupling between the dimensional attention and feature-space systems in the DNF model. Exploring this relationship will be a central focus of future efforts to quantitatively simulate the fMRI data with the DNF model.

The second primary manipulation—manipulating post-switch conflict—revealed a pattern of results across conditions that were consistent with model predictions. Conditions that eliminated feature-space competition more strongly activated fusiform areas than conditions that had feature-space conflict (PC>St, SSPC > SS, SSPC>St). It is important to note that there were rather large anatomical differences in the location of fusiform activation across comparisons. This might reflect anatomical differences in ‘shape’ versus ‘color’ areas in temporal cortex. Recall that the present study was part of a larger, 6-session study examining executive function. As part of this larger effort, several tasks were included to isolate ‘shape’ and ‘color’ areas on a within-subject basis. This will provide a means of isolating shape and color areas within fusiform cortex in order to determine if the shifting fusiform effects are associated with particular dimensions.

Note that in the current design, feature-space competition was always eliminated through a change in the features used in the task. Thus, it is unclear at present whether the fusiform response reflects the introduction of a new feature or a change in feature conflict along the post-switch dimension. Introducing new features could yield a strong neural response because they capture attention. Alternatively, it is possible that the neural response to old features is subject to repetition suppression effects (Desimone, 1996). Teasing these effects apart in the context of the DCCS, however, is difficult. Specifically, it is not possible to change features and maintain post-switch competition—changing the features means that the inputs are necessarily moving away from HT's accumulated during the previous phase. Similarly, it is not possible to keep the same features and eliminate competition without introducing cooperation within the post-switch field. Cooperation within the post-switch field was not studied in the current experiment, and would thus require further testing to reveal unique neural signatures associated with this configuration of activation within the relevant feature-space field.

Note that the SpaceSwap condition (no dimension switch + post-switch conflict) was the only condition to not show any switch-related activity. This is not surprising given that dimension switching and an absence of feature-space conflict drove the pattern of significant contrasts across conditions. It is possible that adults found this task easy because they simply had to re-map the spatial responses and they were given plenty of time to do this before the test item appeared. Interestingly, children appear to have more difficulty with the space swap manipulation. In the case of the NNP-SpaceSwap version reviewed in Chapter 4, children robustly perseverated. Note, however, that this condition also required a switch in dimensions. Thus, the difference between children and adults across space swap versions likely reflects the additional demands placed on children.

Although the work presented here is merely a first step toward a model-based approach to fMRI with adults in the DCCS, the results are promising. The novel task design yielded reliable switch costs while incorporating multiple switch types in an event-related design. Further, the pattern of fMRI results mapped nicely to the *a priori* predictions based on localization of model components to neural regions. Future work with DFT will use this dataset to parameterize the model and quantitatively simulate the pattern of behavioral and neural data. With a parameterized model, DFT can then be used to generate and test specific predictions regarding the neural network involved in flexible rule-use with adults.

Table 4. *Factors contributing to DCCS conditions used in the fMRI study*

		Dimension Change			
		Yes	**No**		
Post-Switch Conflict	**Yes**	Standard (St)	SpaceSwap (SS)	**No**	**Feature Change**
	No	Partial-Change (PC)	SpaceSwap Partial-Change (SSPC)	**Yes**	

Table 5. *Clusters of significantly different activation between switch and repeat trials.*

Switch vs Repeat	Region	Hemisphere	X	Y	Z	K
St	Superior Parietal	L	14	74	54	1711
	Middle Frontal	L	28	4	42	1112
	Inferior Frontal	L	34	-28	20	312
PC	Superior Parietal	L	22	68	32	752
		L	46	46	46	380
	Fusiform	L	32	42	-24	593
		R	-42	54	-26	706
	Middle Occipital	L	32	86	16	610
	Middle Frontal	L	20	10	54	268
	Inferior Frontal	L	50	-12	28	346
SS	None					
SSPC	Fusiform	R	-38	72	-8	5011
		L	32	52	-18	3398
	Superior Parietal	L	28	48	50	810

Table 6. *Clusters of significantly different activation on switch trials across conditions*

Switch Trial x Condition	Region	Hemisphere	X	Y	Z	K
St > SS	Inferior Frontal	L	48	-30	16	73
PC > St	Fusiform	L	32	38	38	92
PC > SS	Parietal	L	26	60	48	121
	Inferior Frontal	L	50	-30	16	88
		L	40	2	24	56

Table 7. *Clusters of significantly different switch-effect activation across conditions*

Switch Effect x Condition	Region	Hemisphere	X	Y	Z	K
St > SS	Inferior Frontal	L	36	-28	22	114
PC ≠ SSPC	None					
PC > St	Fusiform	L	32	40	-26	172
SSPC > SS	Fusiform	L	44	80	-4	601
		L	-42	80	0	433
SSPC > St	Fusiform	L	28	100	-2	431
		R	-34	96	0	143
PC > SS	Medial Frontal	M	-6	-8	56	264
	Parietal	R	2	82	42	241
	Cuneus	L	12	78	6	170
	Thalamus	L	20	36	2	168
	Inferior Frontal	L	46	-36	12	168
	Fusiform	L	34	52	-20	132

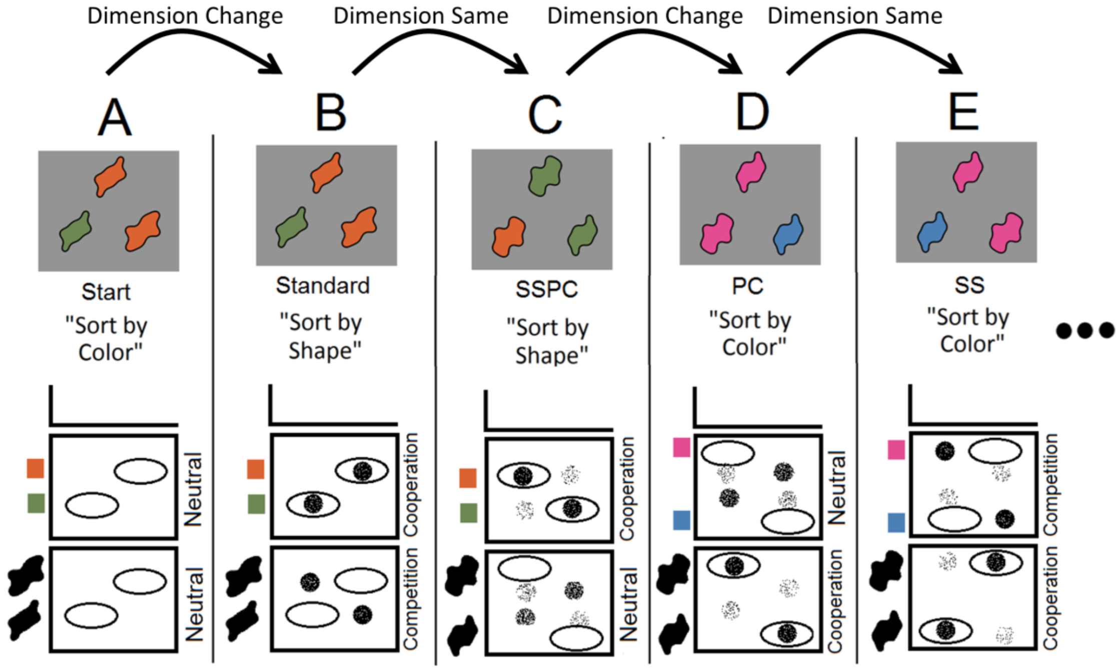

Figure 20: Design of fMRI experiment.

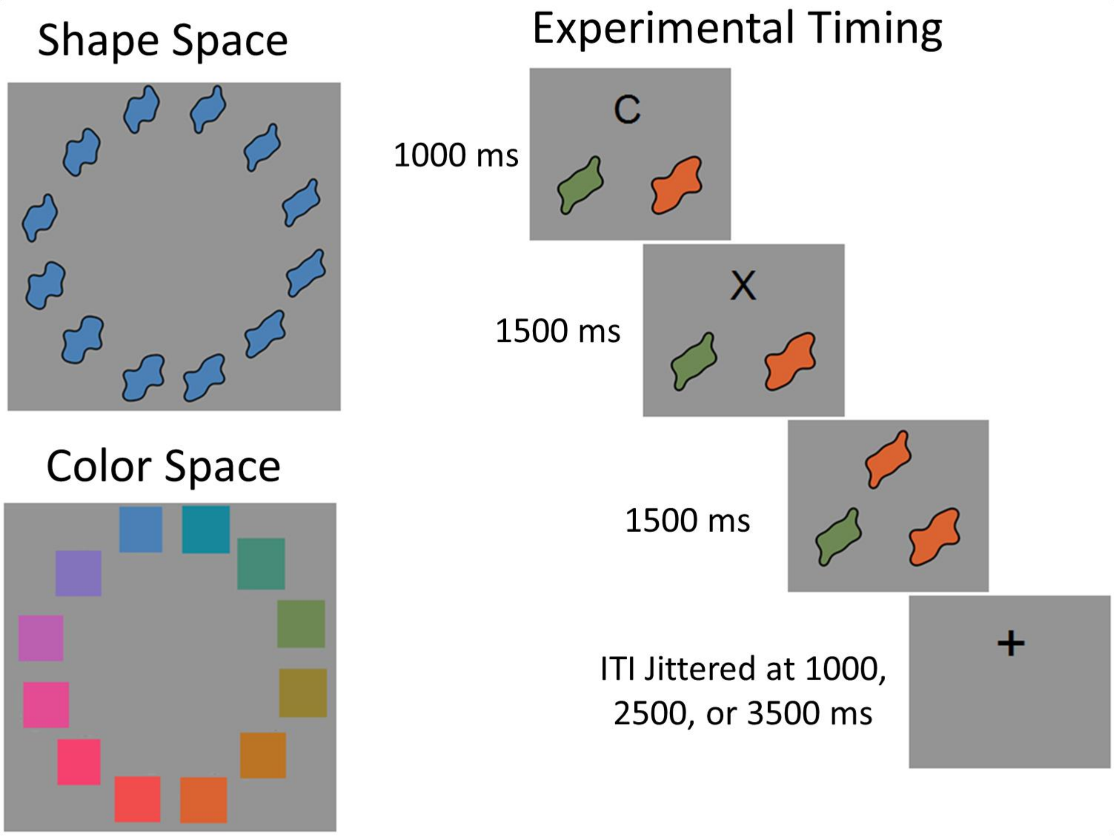

Figure 21: Stimulus dimensions and sequence of events used in the fMRI study.

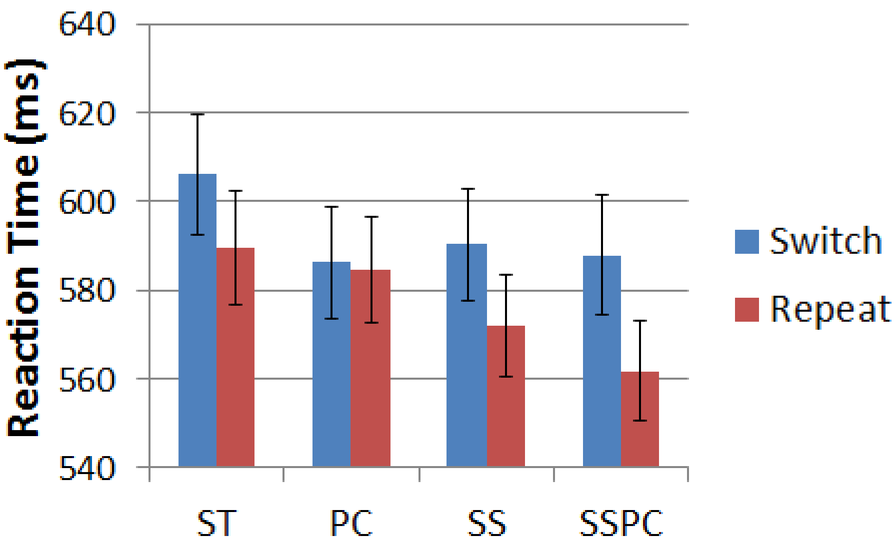

Figure 22: Behavioral data from the fMRI study.

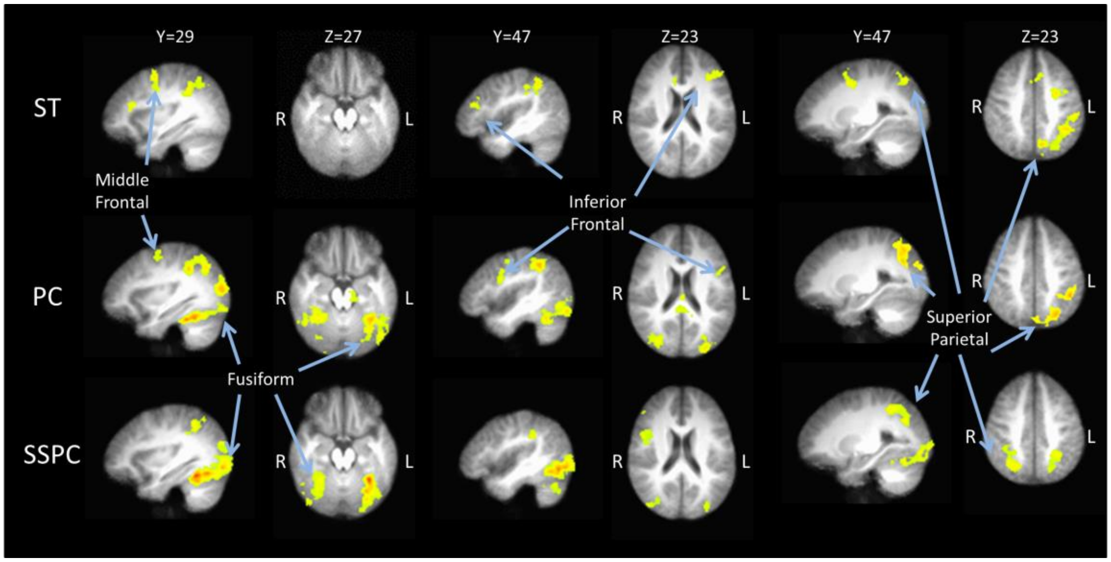

Figure 23: Switch - Repeat contrasts for each condition.

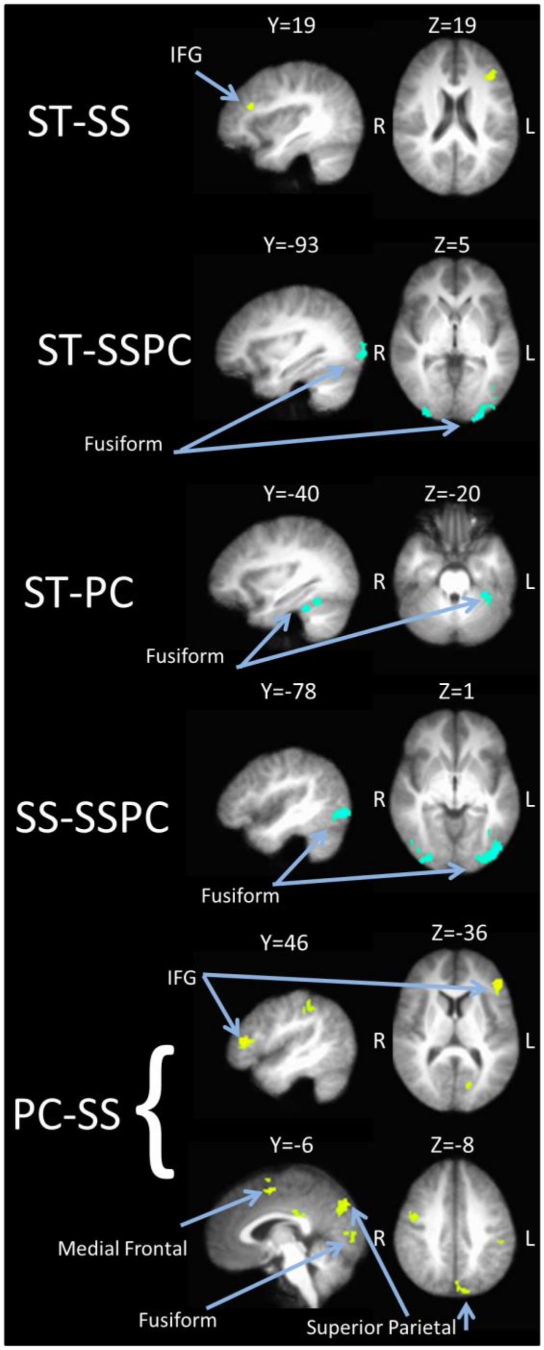

Figure 24: Switch-effect contrasts between conditions.

CHAPTER 6

GENERAL DISCUSSION

EF is an important aspect of cognition because control and regulation are central in numerous cognitive domains. It is an especially important topic of study in early childhood because it is during this period that EF begins to emerge and undergo dramatic and lasting changes. Moreover, measures of EF during early childhood robustly predict physical health, substance dependence, personal finances, and criminal offending outcomes nearly 3 decades later (Moffitt et al., 2011). Further, preschool interventions aimed at improving EF have produced significant increases in school achievement and behavioral function (Diamond & Lee, 2011). Thus, EF is a central aspect of cognition that undergoes foundational changes in early childhood.

The goal of the present work was to cultivate a DFT of flexible rule-use that can integrate behavioral and neural processes over development. Chapter 3 presented a computational framework that I used to capture the dynamics of flexible rule-use. The model is composed of a dimensional attention system that is coupled to an object representation system. These systems are reciprocally coupled and provide a mechanism for on-the-fly shifts in dimensional attention. Rule-use in this framework is a product of modulating dimensional attention and object representation processes while making spatial decisions about multi-feature objects. Developmental changes occur by increasing the strength of neural interactions within the dimensional attention system and improving the selectivity of the mapping from the dimensional attention system to the posterior feature fields which leads to a shift from perseverative responding to correct rule-switching.

In many ways, the model retains the strengths of previous accounts. For example, the dimensional attention system and the developmental mechanisms implemented in this system bear many similarities to the connectionist model proposed by Morton and Munakata (2002). CCC theory posits linguistically mediated reflection as a mechanism of cognitive flexibility. Likewise, the dimensional attention system in the DNF model is framed around ‘shape’ and ‘color’ labels which are associated with visual features. Importantly, DFT moves beyond these theories in specific ways by couching rule-use in a real-time, embodied neural architecture where dimensional attention modulates processes of perception and action.

The model has been used to account for a diverse array of findings in the literature—14 conditions in total with 3- and 5-year-olds. Further, behavioral predictions of the model regarding the role of spatial information in the DCCS task have been successfully tested with a group of 3-year-olds. Together, this establishes DFT as a framework that is rigorously grounded in the behavioral literature in early childhood. One aspect of the model that stands to be fleshed out is the role of space. In the versions of the DCCS that have been presented, the response space and the space of the objects were always confounded. Space as a response dimension and space as a feature binding dimension, however, may serve dissociable roles in rule-use. In the context of binding features, space plays a special role because objects are defined by the spatial overlap of features. In the context of responses, however, space plays a general role that can be interchanged with other dimensions, such as a vocal dimension if the task requires vocal responses.

The two studies presented in this thesis sought to extend this account in two directions. First, *Study 1* proposed a method for simulating hemodynamics with DFT and tested novel neural predictions from the model. *Study 2* looked beyond childhood and proposed an fMRI study to provide a starting point for understanding the neural dynamics that underlie flexible rule-use in adulthood.

The question then becomes how to link diverse processes such as these, which are emergent from the same dynamic interactions, to neural processes in the brain. *Study 1* applied an LFP approach for simulating hemodynamic data. This method proved fruitful, capturing developmental increases in the strength of frontal activation associated with the emerging ability to switch rules reported by Moriguchi and Hiraki (2009). Building from this, I successfully tested hemodynamic predictions of the model and showed that ‘easy’ versions of the task are associated with a stronger hemodynamic response compared to the Standard version.

This method points toward the development of a cognitively functional account of neural activation. This means going beyond localization of function to account for how neural activation relates to cognitive processes. This is illustrated in the simulations of the DCCS fNIRS data. The goal of this simulation project was not just to localize dimensional attention processes, but to identify hemodynamic signatures of the underlying dynamics that give rise to dimensional attention. Based on the reciprocal connectivity of the dimensional attention and feature-space fields, the model accurately predicted a stronger hemodynamic response in the frontal system in a No-Conflict version compared to the Standard version. This stronger activation, however, is not necessarily the cause of correct rule-switching since that would be accomplished based on the cooperation of Hebbian traces (HTs) and target inputs within the post-switch feature field. In this case, differences in activation are a byproduct of the coupling within the neural system rather than a source of behavioral differences. Frontal activation is both the cause and the consequence of rule-switching and is open to influences from other neural regions.

Study 2 provided a rich behavioral and neural dataset that can shed light on the neural and cognitive processes involved in flexible rule-use in adulthood. This study manipulated the need to switch dimensions between phases and the presence of feature-space conflict in a 2 x 2 design. Using standard subtraction methods, it was revealed that frontal regions (corresponding to the dimensional attention system) were sensitive to the demands on dimension switching, and fusiform areas (corresponding to feature-space fields) were sensitive to the absence/presence of feature-space conflict. Thus, in contrast to previous studies (see Morton et al., 2009), DFT offers an account of why temporal and parietal areas are activated on different conditions and the function that these areas serve in flexible rule-use. This points to an embodied account where perception and action and key aspects of flexible rule-use.

The results from *Study 2* set up future work that will use the DNF model to not only link brain and behavior, but also link these systems at different developmental time points. 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 architecture by isolating different functional components that are associated with activity in different frontal regions (such as superior frontal and inferior frontal regions highlighted in the current data set). In this way, fMRI can be used 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.

Although I focused on the DCCS task in the work reported here, the present findings provide 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, DFT demonstrates how these components can emerge from system-wide interactions in a complex 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 as the system attempts to switch rules. Thus, all three processes—working memory, inhibition, and task-switching—conspire to enable flexible behavior and to explain the myriad findings in the literature. Within the DFT framework, the functional aspects of ‘inhibitory control’, ‘working memory’, and ‘switching’ all emerged from general neural dynamics and the structure of inputs for different versions of the task.

I have taken additional steps to explore the generality this framework by using a dynamic field architecture similar to the one reported here to capture aspects of dual-task performance including changes in behavioral (reaction times) and fMRI data over learning (Buss et al., in press). The dual-task model has posterior neural populations that associate visual or auditory stimuli with manual or vocal responses. These two-dimensional neural fields are coupled to a dimensional attention system which modulates the resting level of the posterior fields when a stimulus is presented, similar to how the dimensional attention system work in the DCCS model.

I used this model to capture a behavioral and neural dataset reported by Dux et al. (2009). This study exposed participants to dual- and single-task conditions over 8 days of practice with fMRI conducted at the beginning, middle, and end of training. These researchers identified a region in inferior frontal cortex that closely followed dual-task costs over practice. That is, early in learning, inferior frontal cortex showed a large response that reduced to single-task levels by the end of practice. In this case, processing within inferior frontal areas becomes more efficient over learning and requires less activation to execute both tasks.

The model showed these same neural signatures due to Hebbian learning in both the modality-specific, posterior fields and the dimensional attention system (see Buss et al., in press). Most critically, after exposing the model to the training regimen used by Dux et al. (2009), I found that changes in the hemodynamic response generated from the dimensional attention system closely paralleled changes in inferior frontal cortex activation over learning. And the same model produced reaction times that quantitatively matched the empirical data from both single and dual-task conditions over learning.

This dataset provides an opportunity to probe another time scale of EF that was not previously addressed: learning. In the case of dual-task performance, learning has a behavioral and neural signature that can be probed with the model. Dual-task RT’s reduce to near single-task levels, and activation in IFJ reduces to single-task levels. Using the model, I probed the locus of learning and efficiency by ‘turning off’ learning in either the stimulus-response fields or in the dimensional attention system. While the fMRI data appear to suggest that efficiency is restricted to IFJ, the simulation results showed that learning in all aspects of the model contributed to the behavioral and neural effects observed over practice as there were significant reductions in the fits when learning was restricted to only one part of the model.

This thesis presented a theory of EF and flexible rule-use that was developed in the context of a particular cognitive flexibility task, the DCCS. This model was most thoroughly applied and developed in the context of the early developmental transition from perseveration to flexible rule-switching. Here, the model simulated behavioral and neural data, as well as made behavioral and neural predictions that were successfully tested. The results of this work shed new light on the cognitive dynamics involved in this task, demonstrating how a simple dimensional attention mechanism combined with dynamics involved in object representation, could account for a wealth of behavioral data. The results of this work also shed new light on the functional role that neural activation plays in flexible rule-use. The NIRS experiment provided evidence suggesting that frontal cortex activation is influenced by other regions and is both the cause and the consequence of rule-switching. An fMRI study was then used to probe whether the model could also provide insight on the neurocognitive dynamics of flexible rule-use in adulthood. Here, I manipulated the demands on the dimensional attention system and the object representation system. This revealed not only differential performance across conditions, but also revealed that manipulations to these components differentially activated target neural regions based on the conceptual mapping of these model components. Future work with DFT is positioned to further develop this theory beyond flexible rule-use to provide a full account of EF.

REFERENCES

Aggelopoulos, N. C., & Rolls, E. T. (2005). Scene perception: inferior temporal cortex neurons encode the positions of different objects in the scene. *The European Journal of Neuroscience*, *22*(11), 2903–16.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*, 357–81.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*(2), 77–87.

Andersen, R. A. (1995). Encoding of intention and spatial location in the posterior parietal cortex. *Cerebral Cortex*, *5*(5), 457–69.

Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology*, *84*(1), 451–9.

Ashby, F. G., & Waldschmidt, J. G. (2008). Fitting computational models to fMRI data. *Behavior Research Methods*, *40*(3), 713–721.

Baddeley, A. D. (1986). *Working Memory*. Oxford University Press.

Baddeley, A. D., Della Sala, S., Papagno, C., & Spinnler, H. (1997). Dual-task performance in dysexecutive and nondysexecutive patients with a frontal lesion. *Neuropsychology*, *11*(2), 187–94.

Bastian, Riehle, A., Erlhagen, W., & Schöner, G. (1998). Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport*, *9*(2), 315–9.

Bastian, Schoner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, *18*(7), 2047–2058.

Blair, C., & Razza, R. P. (2007). Relating effortful control, executive function, and false belief understanding to emerging math and literacy ability in kindergarten. *Child development*, *78*(2), 647–63.

Boas, D. A., Gaudette, T., Strangman, G., Cheng, X., Marota, J. J., & Mandeville, J. B. (2001). The accuracy of near infrared spectroscopy and imaging during focal changes in cerebral hemodynamics. *NeuroImage*, *13*(1), 76–90.

Bohlmann, N. L., & Fenson, L. (2005). The Effects of Feedback on Perseverative Errors in Preschool Aged Children. *Journal of Cognition and Development*, *6*(1), 119–131.

Brooks, P. J., Hanauer, J. B., Padowska, B., & Rosman, H. (2003). The role of selective attention in preschoolers' rule use in a novel dimensional card sort. *Cognitive Development*, *18*(2), 195–215.

Bull, R., & Scerif, G. (2001). Executive functioning as a predictor of children's mathematics ability: inhibition, switching, and working memory. *Developmental neuropsychology*, *19*(3), 273–93.

Buss, A. T., Magnotta, V. A., Schoner, G., & Spencer, J. P. (n.d.). Model-based fMRI reveals the structure-function mappings that underlie visual working memory.

Buss, A. T., & Spencer, J. P. (n.d.). The emergent executive: A dynamic neural field theory of the development of executive function. *Monographs of the Society for Research in Child Development*.

Buss, A. T., Wifall, T., Hazeltine, E., & Spencer, J. P. (n.d.). Integrating the behavioral and neural dynamics of reponse selection in a dual-task paradigm: A dynamic neural field model of Dux et al. (2009). *Journal of Cognitive Neuroscience*.

Carlson, S. M. (2005). Developmentally sensitive measures of executive function in preschool children. *Developmental Neuropsychology*, *28*(2), 595–616.

Carlson, S. M., Moses, L. J., & Breton, C. (2002). How Specific is the Relation between Executive Function and Theory of Mind? Contributions of Inhibitory Control and Working Memory. *Infant and Child Development*, *11*, 73–92.

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

Collette, F., Van der Linden, M., Laureys, S., Delfiore, G., Degueldre, C., Luxen, A., & Salmon, E. (2005). Exploring the unity and diversity of the neural substrates of executive functioning. *Human brain mapping*, *25*(4), 409–23.

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*(9), 910–23.

Corbett, B. A., Constantine, L. J., Hendren, R., Rocke, D., & Ozonoff, S. (2009). Examining executive functioning in children with autism spectrum disorder, attention deficit hyperactivity disorder and typical development. *Psychiatry Research*, *166*(2-3), 210–22.

Crone, E. A., Donohue, S. E., Honomichl, R., Wendelken, C., & Bunge, S. A. (2006). Brain regions mediating flexible rule use during development. *The Journal of* *neuroscience : the official journal of the Society for Neuroscience*, *26*(43), 11239–47.

Crone, E. A., Wendelken, C., Donohue, S. E., & Bunge, S. a. (2006). Neural evidence for dissociable components of task-switching. *Cerebral cortex (New York, N.Y. : 1991)*, *16*(4), 475–86.

Cui, X., Bray, S., & Reiss, A. L. (2010). Functional near infrared spectroscopy (NIRS) signal improvement based on negative correlation between oxygenated and deoxygenated hemoglobin dynamics. *NeuroImage*, *49*(4), 3039–46.

Davidson, M. C., Amso, D., Anderson, L. C., & Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia*, *44*(11), 2037–2078.

Deák, G. O., Ray, S. D., & Pick, A. D. (2002). Matching and naming objects by shape or function: Age and context effects in preschool children. *Developmental Psychology*, *38*(4), 503–518.

Deco, G., Rolls, E. T., & Horwitz, B. (2004). “What” and “where” in visual working memory: A computational neurodynamical perspective for integrating FMRI and single-neuron data. *Journal of Cognitive Neuroscience*, *16*(4), 683–701.

Dempster, F. N. (1992). The rise and fall of the inhibitory mechanism: Toward a unified theory of cognitive development and aging. *Developmental Review*, *12*(1), 45–75.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, *93*(24), 13494–13499.

Desimone, R., & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. *Brain Research*, *178*(2), 363–380.

Diamond, A. (2002). Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In D. T. Stuss & R. T. Knight (Eds.), *Principles of Frontal Lobe Function* (pp. 465–503). New York: Oxford University Press.

Diamond, A., Carlson, S. M., & Beck, D. M. (2005). Preschool children’s performance in task switching on the Dimensional Change Card Sort task: Separating the dimensions aids the ability to switch. *Developmental Neuropsychology*, *28*(2), 689–729.

Diamond, A., & Kirkham, N. Z. (2005). Not Quite as Grown-Up as We Like to Think: Parallels between cognition in childhood and adulthood. *Psychological Science*, *16*(4), 291–297.

Diamond, A., & Lee, K. (2011). Interventions shown to aid executive function development in children 4 to 12 years old. *Science*, *333*(6045), 959–64.

DiCarlo, J. J., & Maunsell, J. H. R. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *Journal of Neurophysiology*, *89*(6), 3264–78.

Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., … Petersen, S. E. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(26), 11073–8.

Drucker, D. M., & Aguirre, G. K. (2009). Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cerebral Cortex*, *19*(10), 2269–80.

Duncan, J, Emslie, H., Williams, P., Johnson, R., & Freer, C. (1996). Intelligence and the frontal lobe: the organization of goal-directed behavior. *Cognitive psychology*, *30*(3), 257–303.

Duncan, John, Johnson, R., Swales, M., & Freer, C. (1997). Frontal lobe deficits after head injury: Unity and diversity of function. *Cognitive Psychology*, *14*(5), 713–741.

Dunn, J. R. (2010). Health behavior vs the stress of low socioeconomic status and health outcomes. *The Journal of the American Medical Association*, *303*(12), 1199–200.

Eakin, L., Minde, K., Hechtman, L., Ochs, E., Krane, E., Bouffard, R., … Looper, K. (2004). The marital and family functioning of adults with ADHD and their spouses. *Journal of Attention Disorders*, *8*(1), 1–10.

Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegner, J., & Compte, A. (2009). Mechanism for Top-down control of working memory capacity. *Procedings of the National Academy of Sciences*, *106*(16), 6802–6807.

Edin, F., Macoveanu, J., Olesen, P., Tegnér, J., & Klingberg, T. (2007). Stronger synaptic connectivity as a mechanism behind development of working memory-related brain activity during childhood. *Journal of cognitive neuroscience*, *19*(5), 750–60.

Egner, T., & Hirsch, J. (2005). Where memory meets attention: neural substrates of negative priming. *Journal of cognitive neuroscience*, *17*(11), 1774–84.

Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schöner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, *94*(1), 53–66.

Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*(3), 545–572. Retrieved from http://doi.apa.org/getdoi.cfm?doi=10.1037/0033-295X.109.3.545

Ezekiel, F., Bosma, R., & Morton, J. B. (2013). Dimensional Change Card Sort performance associated with age-related differences in functional connectivity of lateral prefrontal cortex. *Developmental Cognitive Neuroscience*, *5*, 40–50.

Fair, D. A., Cohen, A. L., Dosenbach, N. U. F., Church, J. A., Miezin, F. M., Barch, D. M., … Schlaggar, B. L. (2008). The maturing architecture of the brain's default network. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(10), 4028–32.

Fair, D. A., Dosenbach, N. U. F., Church, J. a, Cohen, A. L., Brahmbhatt, S., Miezin, F. M., … Schlaggar, B. L. (2007). Development of distinct control networks through segregation and integration. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(33), 13507–12.

Faubel, C., & Schöner, G. (2008). Learning to recognize objects on the fly: a neurally based dynamic field approach. *Neural networks : the official journal of the International Neural Network Society*, *21*(4), 562–76.

Frye, D., Zelazo, P. D., & Palfai, T. (1995). Theory of mind and rule-based reasoning. *Cognitive Development*, *10*(4), 483–527.

Garon, N., Bryson, S. E., & Smith, I. M. (2008). Executive function in preschoolers: a review using an integrative framework. *Psychological bulletin*, *134*(1), 31–60.

Gathercole, S. E., Pickering, S. J., Knight, C., & Stegmann, Z. (2004). Working Memory Skills and Educational Attainment: Evidence from National Curriculum Assessments at 7 and 14 Years of Age. *Applied Cognitive Psychology*, *18*(1), 1–16.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416–1419.

Giedd, J. N., Blumenthal, J., Jeffries, N. O., Castellanos, F. X., Liu, H., Zijdenbos, A., … Rapoport, J. L. (1999). Brain development during childhood and adolescence: a longitudinal MRI study. *Nature neuroscience*, *2*(10), 861–3.

Halford, G., Bunch, K., & Mccredden, J. (2007). Problem decomposability as a factor in complexity of the dimensional change card sort task. *Cognitive Development*, *22*(3), 384–391.

Hanania, R. (2010). Two types of perseveration in the Dimension Change Card Sort task. *Journal of Experimental Child Psychology*, *107*(3), 325–36.

Happé, F., Booth, R., Charlton, R., & Hughes, C. (2006). Executive function deficits in autism spectrum disorders and attention-deficit/hyperactivity disorder: examining profiles across domains and ages. *Brain and Cognition*, *61*(1), 25–39.

Haxby, J. V, Grady, C. L., Horwitz, B., Ungerleider, L. G., Mishkin, M., Carson, R. E., … Rapoport, S. I. (1991). Dissociation of object and spatial visual processing pathways in human extrastriate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *88*(5), 1621–5.

Hughes, C., & Ensor, R. (2011). Individual differences in growth in executive function across the transition to school predict externalizing and internalizing behaviors and self-perceived academic success at 6 years of age. *Journal of experimental child psychology*, *108*(3), 663–76.

Huizinga, M., Dolan, C. V, & van der Molen, M. W. (2006). Age-related change in executive function: developmental trends and a latent variable analysis. *Neuropsychologia*, *44*(11), 2017–36.

Huppert, T. J., Diamond, S. G., Franceschini, M. A., & Boas, D. A. (2009). HomER: a review of time-series analysis methods for near-infrared spectroscopy of the brain. *Applied Optics*, *48*(10), 1–33.

Im-Bolter, N., Johnson, J., & Pascual-Leone, J. (2006). Processing limitations in children with specific language impairment: the role of executive function. *Child development*, *77*(6), 1822–41.

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

Johnson, J. S., Spencer, J. P., & Schöner, G. (2008). Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. *New ideas in psychology*, *26*(2), 227–251.

Johnson, J. S., Spencer, J. P., & Schöner, G. (2009). A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain Research*, *1299*, 17–32.

Kelso, J. A. S., Scholz, J. P., & Schöner, G. (1988). Dynamics governs switching among patterns of coordination in biological movement. *Physics Letters A*, *134*(1), 8–12.

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

Kloo, D., Perner, J., Kerschhuber, A., Dabernig, S., & Aichhorn, M. (2008). Sorting between dimensions: conditions of cognitive flexibility in preschoolers. *Journal of experimental child psychology*, *100*(2), 115–34.

Kochanska, G., Coy, K. C., & Murray, K. T. (2001). The development of self-regulation in the first four years of life. *Child Development*, *72*(4), 1091–111.

Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. *Biological Cybernetics*, *73*(1), 49–60.

Kourtzi, Z., Erb, M., Grodd, W., & Bülthoff, H. H. (2003). Representation of the perceived 3-D object shape in the human lateral occipital complex. *Cerebral cortex (New York, N.Y. : 1991)*, *13*(9), 911–20.

Kravitz, D. J., Vinson, L. D., & Baker, C. I. (2008). How position dependent is visual object recognition? *Trends in Cognitive Sciences*, *12*(3), 114–22.

Lehto, J. E., Juujarvi, P., Kooistra, L., & Pulkkinen, L. (2003). Dimensions of executive functioning: Evidence from children. *British Journal of Developmental Psychology*, *21*, 59–80.

Lepsien, J., & Nobre, A. C. (2007). Attentional modulation of object representations in working memory. *Cerebral cortex (New York, N.Y. : 1991)*, *17*(9), 2072–83.

Lipinski, J., Simmering, V. R., Johnson, J. S., & Spencer, J. P. (2010). The role of experience in location estimation: Target distributions shift location memory biases. *Cognition*, *115*(1), 147–153.

Lipinski, J., Spencer, J. P., & Samuelson, L. K. (2010). Biased feedback in spatial recall yields a violation of delta rule learning. *Psychonomic Bulletin & Review*, *17*(4), 581–8.

Liss, M., Harel, B., Fein, D., Allen, D., Dunn, M., Feinstein, R. M., … Rapin, I. (2001). Predictors and Correlates of Adaptive Functioning in Children with Developmental Disorders. *Journal of Autism and Developmental Disorders*, *31*(2), 219–230.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–7.

Mack, W. (2007). 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*, *98*(4), 243–51.

Markounikau, V., Igel, C., Grinvald, A., & Jancke, D. (2010). A dynamic neural field model of mesoscopic cortical activity captured with voltage-sensitive dye imaging. *PLoS Computational Biology*, *6*(9).

Mazzocco, M. M. M., & Kover, S. T. (2007). A longitudinal assessment of executive function skills and their association with math performance. *Child Neuropsychology*, *13*(1), 18–45.

McDowell, K., Jeka, J. J., Schoner, G., & Hatfield, B. D. (2002). Behavioral and electrophysiological evidence of an interaction between probability and task metrics in movement preparation. *Experimental Brain Research*, *144*(3), 303–313.

McEvoy, R. E., Rogers, S. J., & Pennington, B. F. (1993). Executive Function and Social Communication Deficits in Young Autistic Children. *Journal of Child Psychology and Psychiatry*, *34*(4), 563–578.

Milner, B. (1963). Effects of Different Brain Lesions on Card Sorting: The Role of the Frontal Lobes. *Archives of Neurology*, *9*(1), 90–100.

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences*, *6*(null), 414–417.

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, a H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex “Frontal Lobe” tasks: a latent variable analysis. *Cognitive psychology*, *41*(1), 49–100.

Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., … Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(7), 2693–8.

Moriguchi, Y., & Hiraki, K. (2009). Neural origin of cognitive shifting in young children. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(14), 6017–21.

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

Morton, J. B., Bosma, R., & Ansari, D. (2009). Age-related changes in brain activation associated with dimensional shifts of attention: an fMRI study. *NeuroImage*, *46*(1), 249–56.

Morton, J. B., & Munakata, Y. (2002). Active Versus Latent Representations: A Neural Network Model of Perseveration, Dissociation, and Decalage. *Developmental Psychobiology*, *40*(3), 255–265.

Müller, U., Dick, A. S., Gela, K., Overton, W. F., & Zelazo, P. D. (2006). The role of negative priming in preschoolers' flexible rule use on the dimensional change card sort task. *Child development*, *77*(2), 395–412.

Nagahama, Y., Okada, T., Katsumi, Y., Hayashi, T., Yamauchi, H., Oyanagi, C., … Shibasaki, H. (2001). Dissociable mechanisms of attentional control within the human prefrontal cortex. *Cerebral cortex (New York, N.Y. : 1991)*, *11*(1), 85–92.

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, D. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory.* (4th ed., pp. 1–18). New York: Plenum.

Op De Beeck, H., & Vogels, R. (2000). Spatial sensitivity of macaque inferior temporal neurons. *The Journal of Comparative Neurology*, *426*(4), 505–18.

Pennington, B. F., & Ozonoff, S. (1996). Executive Functions and Developmental Psychopathology. *Journal of Child Psychology and Psychiatry, 37*(1), 51–87.

Perner, J., & Lang, B. (2002). What Causes 3-year-olds ' Difficulty on the Dimensional Change Card Sorting Task ? *Infant and Child Development*, *11*, 93–105.

Perone, S., Simmering, V. R., & Spencer, J. P. (2011). Stronger neural dynamics capture changes in infants' visual working memory capacity over development. *Developmental Science*, *14*(6), 1379–92.

Perone, S., & Spencer, J. P. (2013). Autonomous visual exploration creates developmental change in familiarity and novelty seeking behaviors. *Frontiers in Psychology*, *4*, 648.

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*(1), 23–38.

Prince, M., Patel, V., Saxena, S., Maj, M., Maselko, J., Phillips, M. R., & Rahman, A. (2007). No health without mental health. *Lancet*, *370*(9590), 859–77.

Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science (New York, N.Y.)*, *276*(5313), 821–4.

Rennie, D. A. C., Bull, R., & Diamond, A. (2004). Executive functioning in preschoolers: reducing the inhibitory demands of the dimensional change card sort task. *Developmental Neuropsychology*, *26*(1), 423–43.

Roberts, R. J., Hager, L. D., & Heron, C. (1994). Prefrontal cognitive processes: Working memory and inhibition in the antisaccade task. *Journal of Experimental Psychology: General1*, *123*(4), 374–393.

Roberts, R. J., & Pennington, B. F. (1996). An interactive framework for examining prefrontal cognitive processes. *Developmental Neuropsychology*, *12*(1), 105–126.

Rubia, K., Smith, A. B., Woolley, J., Nosarti, C., Heyman, I., Taylor, E., & Brammer, M. (2006). Progressive increase of frontostriatal brain activation from childhood to adulthood during event-related tasks of cognitive control. *Human Brain Mapping*, *27*(12), 973–93.

Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. (J. Wiles, Ed.)*PloS One*, *6*(12), e28095.

Samuelson, & Smith, L. B. (1999). Early noun vocabularies: do ontology, category structure and syntax correspond? *Cognition*, *73*(1), 1–33.

Sandamirskaya, Y., & Schöner, G. (2010). An embodied account of serial order: how instabilities drive sequence generation. *Neural networks : the official journal of the International Neural Network Society*, *23*(10), 1164–79.

Schöner, G., & Kelso, J. A. (1988). Dynamic pattern generation in behavioral and neural systems. *Science (New York, N.Y.)*, *239*(4847), 1513–20.

Schroeter, M. L., Zysset, S., Wahl, M., & von Cramon, D. Y. (2004). Prefrontal activation due to Stroop interference increases during development--an event-related fNIRS study. *NeuroImage*, *23*(4), 1317–25. doi:10.1016/j.neuroimage.2004.08.001

Schutte, A. R., Spencer, J. P., & Schoner, G. (2003). Testing the Dynamic Field Theory: Working Memory for Locations Becomes More Spatially Precise Over Development. *Child Development*, *74*(5), 1393–1417.

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*(11), 488–495.

Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. *Brain research*, *1202*, 68–86.

Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A., & Barsalou, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, *45*(12), 2802–10.

Spencer, Barich, K., Goldberg, J., & Perone, S. (2012). Behavioral dynamics and neural grounding of a dynamic field theory of multi-object tracking. *Journal of Integrative Neuroscience*, *11*(3), 339–62.

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

Spencer, J. P., Schneegans, S., & Schoner, G. (n.d.). Integrating “what” and “where”: Visual working memory for objects in a scene. In G. Schoner & J. P. Spencer (Eds.), *Dynamic Thinking--A Primer on Dynamic Field Theoryi*.

Stedron, J. M., Sahni, S. D., & Munakata, Y. (2005). Common mechanisms for working memory and attention: the case of perseveration with visible solutions. *Journal of cognitive neuroscience*, *17*(4), 623–31.

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

Towse, J. N., Redbond, J., Houston-price, C. M. T., & Cook, S. (2000). Understanding the dimensional change card sort: Perspectives from task success and failure. *Cognitive Development*, *15*, 347–365.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, *6*(2), 171–178.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

Turvey, M. T., & Shaw, R. E. (1995). Toward an ecological physics and a physical psychology. In R. Solso & D. Massaro (Eds.), *The Science of the Mind: 2001 and beyond.* (pp. 144–169). Oxford: Oxford University Press.

Waxer, M., & Morton, J. B. (2011). The development of future-oriented control: an electrophysiological investigation. *NeuroImage*, *56*(3), 1648–54.

Wiebe, S. A., Sheffield, T., Nelson, J. M., Clark, C. a C., Chevalier, N., & Espy, K. A. (2011). The structure of executive function in 3-year-olds. *Journal of experimental child psychology*, *108*(3), 436–52.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24.

Xiao, Y., Wang, Y., & Felleman, D. J. (2003). A spatially organized representation of colour in macaque cortical area V2. *Nature*, *421*(6922), 535–9.

Yerys, B. E., & Munakata, Y. (2006). When Labels Hurt but Novelty Helps: Children’s Perseveration and Flexibility in a Card-Sorting Task. *Child Development*, *77*(6), 1589–1607.

Zahn, C. T., & Roskies, R. Z. (1972). Fourier Descriptors for Plane Closed Curves. *IEEE Transactions on Computers*, *C-21*(3), 269–281.

Zanto, T. P., Rubens, M. T., Bollinger, J., & Gazzaley, A. (2010). Top-down modulation of visual feature processing: the role of the inferior frontal junction. *NeuroImage*, *53*(2), 736–45.

Zelazo, P. D., Anderson, J. E., Richler, J., Wallner-Allen, K., Beaumont, J. L., & Weintraub, S. (2013). NIH Toolbox cognition battery (CB): Measuring executive function and attention. *Monographs of the Society for Research in Child Development*, *78*(4), 16–33.

Zelazo, P. D., Muller, U., Frye, D., & Marcovitch, S. (2003). The development of executive function in early childhood. *Monographs of the Society for Research in Child Development2*, *68*, 1–137.