Common Mechanisms for Working Memory and Attention: The Case of Perseveration with Visible Solutions

Jennifer Merva Stedron[1], **Sarah Devi Sahni**[2], **and Yuko Munakata**[3]

Abstract

■ Everyone perseverates at one time or another, repeating previous behaviors when they are no longer appropriate. Such perseveration often occurs in situations with working memory demands, and the ability to overcome perseveration has been linked to brain regions critical for working memory. Many theories thus explain perseveration in terms of working memory deficits. However, perseveration also occurs in situations without apparent working memory demands, in which the visible environment specifies appropriate behavior. Such findings appear to challenge working memory accounts of perseveration. To evaluate this challenge, a neural network model of a working memory account of perseveration was tested on tasks with visible solutions. With advances in the mechanisms that support working memory, networks became increasingly able to attend to relevant information in the environment. These developments led to improvements in performance on tasks with visible solutions, paralleling the developmental progression observed in infants. The simulations demonstrate how mechanisms of working memory can subserve perseveration and success on tasks with and without obvious memory demands. In both types of tasks, controlled processing occurs through the activation of task-relevant representations, which provide top–down biasing of other processing pathways. More generally, the simulations demonstrate how common mechanisms can support working memory and attention. ■

INTRODUCTION

An essential aspect of human cognition is the ability to behave flexibly, or adapt actions to changing environmental demands. However, under certain circumstances, infants, children, and adults are unable to behave flexibly and instead perseverate, repeating prepotent or habitual behaviors when they are no longer applicable (e.g., Morton & Munakata, 2002b; Zelazo, Frye, & Rapus, 1996; Dunbar & Sussman, 1995; Gerstadt, Hong, & Diamond, 1994; Lhermitte, 1986; Diamond, 1985; Stuss & Benson, 1984; Milner, 1963; Piaget, 1954). For example, after infants search for a hidden toy, they then often perseverate, searching the same hiding location even after seeing the toy hidden in a new location (making the A-not-B error, Diamond, 1985; Piaget, 1954). Three-year-olds can sort cards according to their color or shape, but they tend to perseverate with whatever sorting rule they used first, even after being asked to sort the cards according to a new rule (Zelazo et al., 1996). Similar profiles of perseveration are observed across adults with brain damage (e.g., Stuss & Benson, 1984; Milner, 1963), neurologically intact adults (Dunbar & Sussman, 1995), and other species (e.g., Diamond & Goldman-Rakic, 1989).

Such perseveration may result at least in part from limitations in working memory for currently relevant information (O'Reilly, Braver, & Cohen, 1999; Munakata, 1998; Roberts & Pennington, 1996; Cohen & Servan-Schreiber, 1992; Goldman-Rakic, 1987; Diamond, 1985). In all of these situations, one must remember the information specifying the correct response (e.g., the hidden toy's location, or the current rule for sorting cards). In addition, perseveration increases when working memory is taxed, as in dual-task conditions (Dunbar & Sussman, 1995). Moreover, the prefrontal cortex (PFC), which is critical for working memory (e.g., Braver et al., 1997; Miller, Erickson, & Desimone, 1996), is strongly implicated in overcoming perseveration (Casey, 2001; Bell & Fox, 1992; Diamond & Goldman-Rakic, 1989; Milner, 1963). Thus, task analyses and behavioral and brain evidence converge on the idea that working memory limitations contribute to perseveration.

However, perseveration also occurs in situations without apparent memory demands, in which the environment specifies the appropriate behavior (Berger, 2004; Aguiar & Baillargeon, 2000; Hauser, 1999; Hood, Hauser, Anderson, & Santos, 1999; Lockman & Pick, 1984; Diamond, 1981). For example, when faced with 2 cloths—one with a distant toy on it and one with a toy behind it

---

[1]University of Denver, [2]University of Wisconsin, [3]University of Colorado—Boulder

(Figure 1A, A trial)—infants pull the cloth with the toy on it. If the cloths are switched so that the cloth that was to the infants' left is now to the infants' right (Figure 1A, B trial), infants perseverate, continuing to pull the cloth on the same side as before although it does not yield the toy (Aguiar & Baillargeon, 2000). Such behaviors are not limited to infants, or to humans. Tamarin monkeys perseverate in a similar cloth-pulling task, by incorrectly pulling a cloth with a large piece of food next to it (which does not yield the food) rather than a cloth with a small piece of food on it (Hauser, 1999). Adults may similarly perseverate under circumstances without apparent memory demands, for example, reaching back to the usual location for an object (such as a remote control), even when it is visibly not there. Such behaviors seem to challenge the widely held view that working memory deficits play a large role in perseveration.

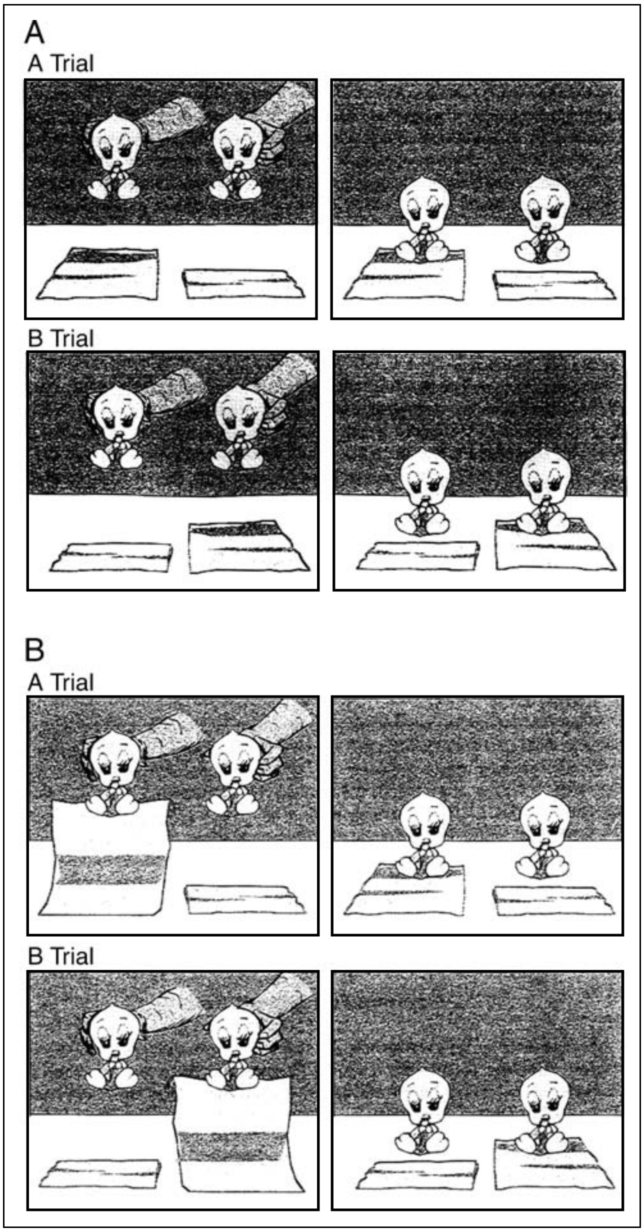

Figure 1. A and B trials in toy-unattached (A) and toy-attached (B) versions of the cloth-pulling task (Aguiar & Baillargeon, 2000). On A trials, infants pull the towel with the toy on it. On B trials, infants reach perseveratively to the prior location, although the correct towel is specified by the visible environment. Reprinted from Aguiar and Baillargeon (2000), with permission from Elsevier.

In theory however, the same mechanisms that support working memory could play a role in such situations, by guiding attention appropriately to relevant information in the environment. Common mechanisms appear to support working memory for nonvisible stimuli and attention to visible stimuli (Oh & Kim, 2004; Woodman & Luck, 2004; Curtis & D'Esposito, 2003; Engle, 2002; De Fockert, Rees, Frith, & Lavie, 2001; Kastner & Ungerleider, 2001; Miller & Cohen, 2001; Chafee & Goldman-Rakic, 2000; Desimone & Duncan, 1995; Fuster, Bauer, & Jervey, 1985). For example, the same neurons in PFC that participate in working memory appear to bias activity in posterior neurons toward relevant stimuli in the environment (Chafee & Goldman-Rakic, 2000; Desimone & Duncan, 1995; Fuster et al., 1985). In addition, visual search tasks (which require attention to visible stimuli) interfere with spatial working memory tasks and vice versa (Oh & Kim, 2004; Woodman & Luck, 2004). Increasing working memory load in humans leads to more processing of irrelevant stimuli and to increased neural activity in brain regions for the irrelevant stimuli (De Fockert et al., 2001). This suggests that mechanisms supporting working memory direct attention to relevant stimuli, and limitations in these mechanisms lead to greater intrusion of irrelevant stimuli. Thus, even when the visible environment specifies appropriate behaviors, the same mechanisms that support working memory may serve to guide attention to relevant information in the environment in the face of prepotent responses. Limitations in these mechanisms could lead one to fall back on habitual or prepotent responses, and improvements in these mechanisms could support correct responses. In theory then, the same mechanisms that support working memory could play a role in perseveration with both visible and nonvisible solutions.

To investigate this possibility, we tested a neural network model of a working memory account of perseveration (Morton & Munakata, 2002a; Munakata, 1998; Cohen & Servan-Schreiber, 1992; Cohen, Dunbar, & McClelland, 1990) on tasks with fully visible solutions. In the model, working memory takes the form of sustained activations of processing units, corresponding to the continued firing of neurons in PFC. These sustained activations are supported by self-excitatory recurrent connections, which can be modified to simulate changes in working memory. Such working memory representations support task-relevant information and can contribute to controlled processing through the top–down activation of other processing pathways. Prepotent responses are subserved by strong connection weights between processing units, which bias the network toward certain representations and behaviors. Variants of this model have simulated perseveration on a number of tasks, including the A-not-B (Munakata, 1998) and card-sorting (Morton & Munakata, 2002a) tasks described above.[1] Success on such tasks requires remembering information that was presented but is no longer available, such as the location of a toy that was presented and then hidden or the rule for sorting cards that was presented transiently. Perseveration results when the connection weights supporting prepotent responses are stronger than the working memory representations of currently relevant information. Flexible behavior results when working memory representations are stronger than prepotent connection weights, allowing controlled processing through top–down activation of task-relevant information.

RESULTS AND DISCUSSION

The model (Figure 2) was tested with 2 versions of the cloth-pulling task described above, in which the solution is fully visible (Aguiar & Baillargeon, 2000). In the ''toy-attached'' version, one cloth has a toy attached to it and another cloth has a toy behind it (Figure 1B). Seven-month-olds perseverate in this toy-attached condition when the toy and cloth locations are switched, whereas 9-month-olds succeed. In the ''toy-unattached'' condition, one cloth has a toy on (but not attached to) it and another cloth has a toy behind it (Figure 1A). Nine-month-olds perseverate in this version, whereas 11-month-olds succeed. The tasks were presented to the model as patterns of activity that specified the relevant information for determining an appropriate response, so that success did not require remembering information that was no longer available. The model, like infants, thus faced a task without apparent memory demands. To assess potential contributions of working memory mechanisms to tasks with visible solutions, the model was tested at different ''ages,'' that is, with different strengths of recurrent connections supporting sustained activations. This single manipulation in the present model was the same manipulation that subserved increased working memory in prior models, on tasks that required working memory for where a toy was hidden (Munakata, 1998) or for what rule was presented for sorting cards (Morton & Munakata, 2002a).

Percent Correct Response

With changes to recurrent connections alone, the model simulated the developmental progression observed in infants from 7 to 11 months on these tasks with fully visible solutions (Figure 3). Networks succeeded on A trials at each age tested. On B trials, the youngest networks perseverated in both the toy-attached and toy-unattached conditions. Older networks succeeded in the toy-attached condition but perseverated in the toy-unattached condition. Finally, networks that were even older succeeded in both the toy-attached and toy-unattached conditions. How did changes to recurrent connections alone lead to this developmental progression?

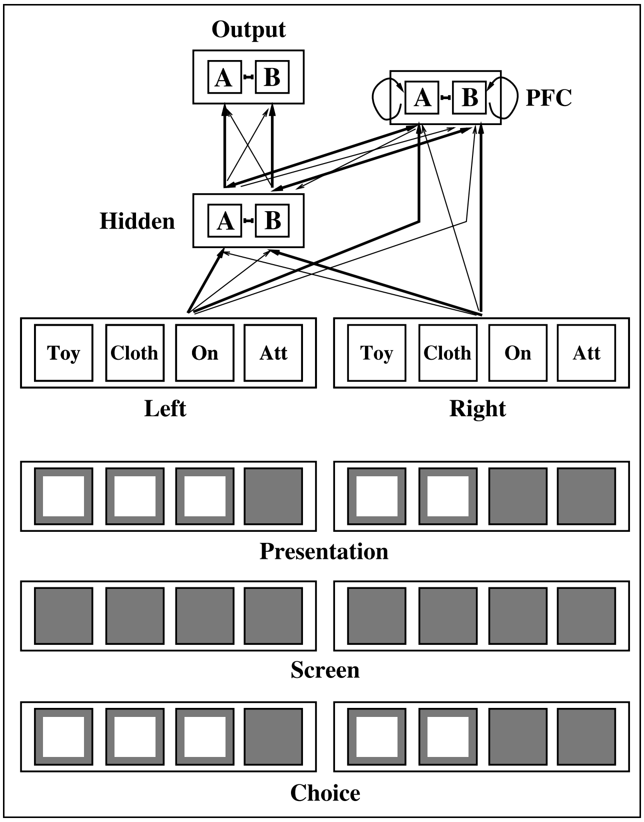

Figure 2. Neural network model of working memory account of perseveration, with sequence of inputs corresponding to one trial in a towel-pulling task: The model was composed of input, hidden, PFC, and output layers. Separate input layers represented information about the left and right sides of the environment. ''On'' and ''attached'' information was included because infants clearly perceive this information and use it to guide their reaching, as evidenced by their behavior on the cloth-pulling tasks as well as other tasks (Diamond, Churchland, Cruess, & Kirkham, 1999). Connectivity included an initial bias (shown with thicker connection lines) to accurately encode the location of various aspects of the display, so that the hidden, PFC, and output layers could represent and respond to the location of a toy on a cloth preferentially over a toy behind a cloth. These connections changed via a Hebbian learning rule, which increases the strength of connections between units that are simultaneously active. Inhibitory connections within layers supported competition between representations. Excitatory recurrent connections at the PFC layer (shown with curved arrows) supported sustained activations of these units. These recurrent connections were manipulated to simulate changes in working memory mechanisms.

Networks at all ages succeeded on A trials because of the greater input activation associated with the correct side (the correct side had 3 or 4 units activated whereas the incorrect side had only 2 units activated). Through

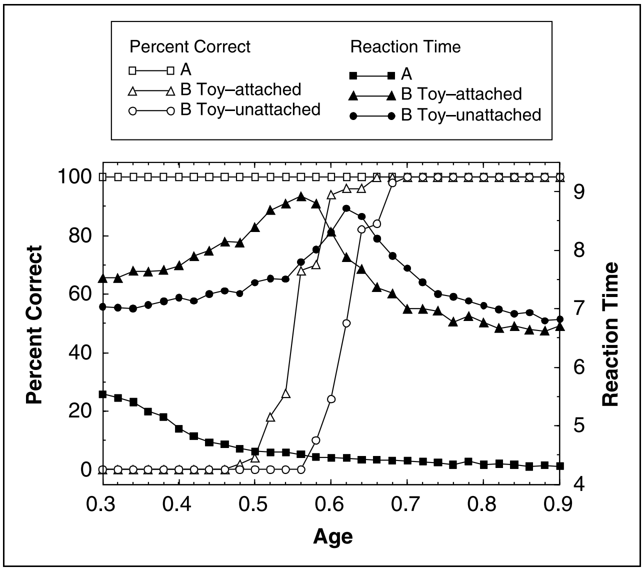

Figure 3. Percent correct and reaction time results from networks at different "ages" (with different strengths of recurrent connections supporting sustained activations), on A trials (shown for just the toy-attached condition, with similar results in the toy-unattached condition) and B trials. Networks showed the developmental progression observed in infants. Changes in recurrent connections alone allowed networks to progress from perseverating to succeeding on B trials, first in the toy-attached and then in the toy-unattached condition. Thus, performance improved through the same changes that increased working memory abilities in prior models, despite the lack of apparent memory demands in these tasks. Reaction times were slowest around the transitions from perseverating to succeeding.

feedforward weights, the greater input activity on the A side led the A hidden, PFC, and output units to win the competition for activation over the B units, resulting in correct performance on the A trials. Further, with each successful A trial, the weights connecting the A units were strengthened, laying down a bias for networks to respond to the A location.

During the B trials, greater input activation was again associated with the correct (B) side. However, this input had to compete with the bias to the A location. With lower recurrent weights, the networks' representation of the B location in the PFC layer was not strong enough to compete with the bias in the A input-to-hidden and hidden-to-output weights; thus, networks reached to the A location. Higher recurrent weights served to amplify the networks' representation of the B location in the PFC layer. This representation of task-relevant information provided top–down activation of the B location in the hidden layer, resulting in this representation winning out over the bias to A and leading to correct performance. In this way, the same manipulation that improved working memory in prior simulations improved performance on these tasks without apparent memory demands, by guiding attention to relevant information in the environment. Thus, greater activation of task-relevant representations can improve controlled processing, both for information that is present and must be attended to (as in tasks with visible solutions, such as this towel-pulling task) and for information that is absent and must be remembered (as in the A-not-B and card-sorting tasks).

Why did networks succeed on the B trials in the toy-attached condition before the toy-unattached condition? The strength of the networks' representations of the correct location increased both with age and with increases in the input activity for that location. On B trials, the toy-attached condition had greater input activity for the B location during the presentation phase (4 units activated) than in the toy-unattached condition (3 units activated). This greater input activity increased the networks' representation of the correct location, increasing its ability to compete with a bias for the A location, leading to earlier success in the toy-attached condition.

This neural network account corresponds to the following explanation of infants' behavior. Seven-, 9-, and 11-month-old infants have a basic understanding that pulling a cloth with a toy on it will allow them to retrieve the toy, and as a result, a toy supported by a cloth is more prominent or salient to them than a toy behind a cloth. This greater salience allows infants to actively represent, attend to, and retrieve the toy on A trials, with repeated trials strengthening infants' bias for this side. On B trials, the toy supported by the cloth is again more salient than the toy behind the cloth, allowing infants to attend to these stimuli. However, now attention to the B side must compete with a bias toward the A side. Younger infants have a limited ability to attend to the B side in the face of the strong bias toward A, for the same reasons they have limited working memory abilities; they cannot strongly activate task-relevant representations. For older infants, the more developed mechanisms that support their greater working memory abilities also support greater attention to B; they can strongly activate task-relevant representations, allowing controlled processing through top–down activation of other processing pathways. Infants succeed on different versions of the cloth-pulling task at different ages because of the differential salience between the sides. In the toy-attached condition, the correct and incorrect cloths differ in 2 important features (the toy is both on and attached to the correct cloth) and this helps infants attend to the correct side in the face of a bias toward the incorrect side. The difference is less salient in the toy-unattached condition, because the correct and incorrect cloths differ in only one feature. As a result, infants cannot attend as readily to the correct side in the face of a bias toward the incorrect side in the toy-unattached condition and require stronger abilities to amplify their representations to direct their attention to the correct side.

This account may also explain behavior in another variant of the cloth-pulling task, the "gap" task, in which one of the cloths is actually 2 cloths separated in the middle by a gap (Figure 4, Aguiar & Baillargeon, 2000). Eleven-month-old infants succeed on A trials but perseverate on B trials. In the model, the gap condition could be represented with lower input activation on the ‘‘cloth’’ unit for the cloth with a gap in the middle, with the stimuli on both sides activating the ‘‘toy’’ and ‘‘on’’ units (Figure 5). Networks could succeed on A trials because of the correct side’s greater input activation. However, on B trials, the slight difference in input activations might not overcome the bias built up toward A, unless there were very strong abilities to support attention to the correct location. That is, the less salient the difference between the correct and incorrect cloths, the stronger the abilities required to direct attention to the correct side on B trials.[2]

Reaction Time

On B trials in both the toy-attached and toy-unattached conditions, the model produced a developmental inverted U-shaped reaction time curve, in the number of processing cycles required before networks settled on a stable response (Figure 3). Networks settled more quickly when they were (a) relatively young and perseverating and (b) relatively old and succeeding. In contrast, networks responded relatively slowly during transitional ages, before and after first success. These differences in reaction time resulted from differences in the degree of competition between the PFC representation for the current location of the cloth supporting the toy and the bias for the old location. Specifically, the more imbalanced this competition, the faster the competition was resolved and the faster the reach. These results suggest that infants should be fastest on B trials when they are either really perseverating or really succeeding (i.e., when they are far from the transition period from perseveration to success), and they should be slowest in the transition period.

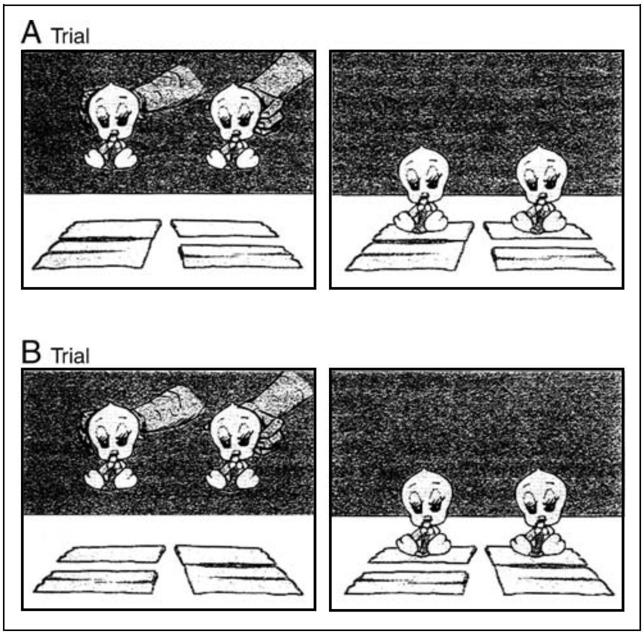

Figure 4. A and B trials in the gap version of the cloth-pulling task (Aguiar & Baillargeon, 2000): Cloths on both the left and right side support toys; however, one of the cloths is actually 2 cloths separated in the middle by a gap. Reprinted from Aguiar and Baillargeon (2000), with permission from Elsevier.

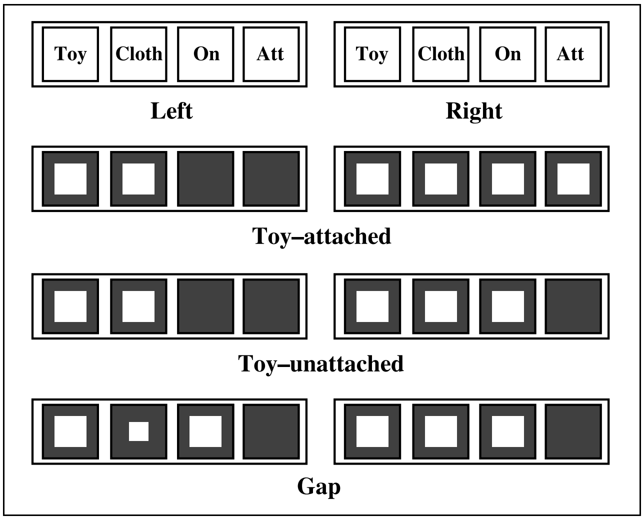

Figure 5. Input representations for B trials in the toy-attached, toy-unattached, and Gap conditions. In each condition, the correct response (on the right) has more input activity than the incorrect response (on the left). These differences in input activity would allow the network to succeed at all ages on preceding A trials (not shown). However, B trials require an ability to direct attention to the correct side, in the face of a bias to respond to the previous location. Because differences in input activity between the correct and incorrect sides decrease across the conditions as ordered, networks would require increasingly strong activation abilities to direct attention to the correct side across conditions.

Conclusions

This neural network model of a working memory account of perseveration simulated infants’ developmental progression on tasks without apparent memory demands. With weak recurrent connections in the PFC layer (the same manipulation that led to weak working memory abilities in prior models), networks tended to repeat inappropriate prior behaviors even when the visible environment specified alternate appropriate behaviors. With increases in these recurrent connections (the same manipulation that led to increased working memory abilities in prior models), networks became increasingly able to attend to relevant information in the environment to overcome their perseverative tendencies. Thus, this model demonstrates how the same mechanisms may support flexible behaviors in situations with nonvisible solutions (Morton & Munakata, 2002a;

Munakata, 1998; Cohen & Servan-Schreiber, 1992; Cohen et al., 1990) and visible solutions. Greater activation of task-relevant information supports working memory for nonvisible information and attention to visible information, both of which contribute to controlled processing. In this way, behaviors that appear to challenge the role of working memory mechanisms in perseveration may instead be quite compatible with working memory accounts.

Note that although we refer to PFC in general throughout the article, different areas within PFC may be specialized for different functions (e.g., Petrides, 2000) or for variations of a common function (Dias, Robbins, & Roberts, 1997; Goldman-Rakic, 1996). Thus, PFC should not be viewed as a unified entity. However, the points demonstrated in this article can apply across specializations within PFC. That is, regardless of where various functions are localized within PFC, the same mechanisms of activation of task-relevant representations and top–down biasing may subserve working memory and attention. Perseveration in the absence of working memory demands does not require a separate mechanism to be posited.

Other theories may also accommodate perseveration on tasks without apparent memory demands. For example, such perseveration might also reflect inhibitory deficits (Kirkham, Cruess, & Diamond, 2003) or the ''miscategorization'' of new tasks as old, leading to the repeating of previous solutions (Aguiar & Baillargeon, 2000). Such accounts could be instantiated and tested in computational models (e.g., by simulating development via increases in the strength of inhibitory, bottom–up, or top–down connections) and might lead to competing predictions about reaction times (e.g., that infants should respond quickly as long as they are perseverating). We focused on working memory accounts given the apparent challenge to such accounts from perseveration on tasks without apparent memory demands. These simulations provide an existence proof that this challenge is more apparent than real; changes in the same mechanisms that support working memory can improve performance on tasks with visible solutions, by supporting attention to relevant information in the environment.

Perseveration in the face of visible solutions is not the only apparent challenge to working memory accounts of flexible behavior. When individuals perseverate, they sometimes show dissociations in their behavior, indicating through other measures that they seem to know what they should be doing instead. For example, in the A-not-B task, infants occasionally look correctly to the new B hiding location while reaching perseveratively back to the old A location (Hofstadter & Reznick, 1996; Diamond, 1985; Piaget, 1954). Perhaps even more compelling, children can correctly answer questions about the new rule they should be using in sorting cards (such as where trucks should go in the shape game), even as they perseverate in sorting cards according to the previous rule (e.g., by color; Zelazo et al., 1996). Similarly, patients with frontal lobe damage occasionally report that they know they are (incorrectly) using the previous rule and can even state the rule they should be using instead (Milner, 1963). Neurologically intact adults show similar dissociations (Morton & Munakata, in preparation). Such dissociations may suggest that working memory for task-relevant information cannot be the sole problem in perseveration.

However, memory can be graded in nature, so that individuals may succeed on one measure of memory without ruling out a role for memory limitations in their perseveration. For example, limited working memory for a new rule may be sufficient for tasks that do not involve conflict, but insufficient for tasks that involve conflict. Sorting a card (e.g., a red truck) according to a new rule (e.g., shape) requires resolving the conflict between the new rule and the previous rule (color). In contrast, answering a standard verbal query about the new rule (e.g., ''Where do the trucks go in the shape game?'') does not involve conflict. Neural network models have demonstrated how stronger working memory for a new rule is required for tasks with greater conflict (Morton & Munakata, 2002a). This analysis correctly predicted that children should perseverate when conflict is introduced into verbal queries (e.g., ''Where do the red trucks go in the color game?'') (Munakata & Yerys, 2001; see also Morton & Munakata, 2002b). In addition, neural network models have demonstrated how stronger working memory for a hidden toy is required for behaviors with less frequent updating, such as reaching versus looking measures in the A-not-B task (Munakata, 1998). Thus, individuals may perseverate because of working memory limitations despite succeeding on some measures of memory.

In summary, the current model demonstrates how the same developments that decrease perseveration in tasks with obvious working memory demands can also decrease perseveration in tasks without obvious working memory demands. These findings provide a computational perspective on how common mechanisms may support working memory for nonvisible stimuli and attention to visible stimuli. In addition, the model's reaction time predictions provide a clear direction for further research to evaluate the role of working memory mechanisms in flexible behavior.

METHODS

Neural Network

Each input layer contained 4 units corresponding to a toy, a cloth, the toy's placement on the cloth, and the toy's attachment to the cloth. Lateral inhibitory connections within the PFC, hidden, and output layers, produced competition between units within layers.

The activation of units was determined using 2 primary equations, designed to capture electrophysiological properties of neurons (O'Reilly & Munakata, 2000). One equation computed the membrane potential of a unit as a function of its excitatory and inhibitory inputs:

$$V_m(t+1) = V_m(t) - dt_{vm}[g_e(t)\bar{g}_e(V_m(t) - E_e) + g_i(t)\bar{g}_i(V_m(t) - E_i) + g_l(t)\bar{g}_l(V_m(t) - E_l)] \tag{1}$$

where dt_{vm} is a time constant, set to 0.05, slowing the potential change and capturing the corresponding slowing of this change in a neuron, primarily as a result of the capacitance of the cell membrane. Units settled for 200 processing steps, sufficient for reaching equilibrium potential, expressed as:

$$V_m = \frac{g_e\bar{g}_e E_e + g_i\bar{g}_i E_i + g_l\bar{g}_l E_l}{g_e\bar{g}_e + g_i\bar{g}_i + g_l\bar{g}_l} \tag{2}$$

A second equation converted the membrane potential into an activation value using a thresholded sigmoidal function:

$$a_j = \frac{\gamma[V_m - \theta]_+}{\gamma[V_m - \theta]_+ + 1} \tag{3}$$

A small amount of Gaussian noise (variance 0.0001) was added to the membrane potential to allow variation in performance across the 50 networks tested. Network activations were cleared between trials to reflect the behavioral activities and associated processes occurring between trials (e.g., as infants played with retrieved toys).

Connections between units were adjusted according to a Hebbian learning rule:

$$\Delta w_{ij} = \varepsilon a_j(a_i - w_{ij}) \tag{4}$$

where ε is the learning rate, set to 0.025 as in earlier simulations with this working memory model (Morton & Munakata, 2002a; Munakata, 1998). Feedforward connections to the PFC units learned at a much slower rate (1% the learning rate of other connections), such that the PFC part of the model was less susceptible to bias than the rest of the system. Other models have used such a reduced learning rate in the prefrontal system to allow it to overcome perseverative tendencies (Morton & Munakata, 2002a; O'Reilly, Noelle, Braver, & Cohen, 2002). This computational work thus provides one hypothesis about mechanisms in PFC that could aid in its specialization in supporting flexible behavior.

The network's initial connectivity included a bias to respond appropriately to location information. Input units and their corresponding hidden and PFC units had connections of 0.6, as did corresponding PFC, hidden, and output units. Connections between noncorresponding units (e.g., the left input units and the B hidden unit; the B hidden unit and the A output unit) were 1% the strength of those between corresponding units. Inhibitory connections within hidden, PFC, and output layers, which prevented runaway activity, were 0.6.

Stimuli

The model was presented with stimuli representing the toy-attached and toy-unattached conditions of the cloth-pulling task. Each condition consisted of 2 A trials and 1 B trial, reflecting the number of test trials in the behavioral studies (Aguiar & Baillargeon, 2000).[3] On A trials, the cloth supporting the toy was at one location (A) and on the B trial, the cloth supporting the toy was at the other location (B).

Each trial consisted of segments (Figure 2), corresponding to segments in the behavioral studies: the presentation of the toys and cloths at the A and B locations, a screen period in which reaching was prevented (simulating a brief period in behavioral studies when a screen was placed between infants and the cloths to prevent immediate reaching to the most recently fixated toy-cloth pairing), and a choice period when the cloths and toys were again visible. The screen stimulus might appear to create memory demands given that stimuli are not visible during this period; however, the relevant information about which toy was on its towel was presented again in the Choice period (the point where the model was tested). On each trial, the Presentation, Screen, and Choice inputs were presented for 2, 1, and 5 time steps, respectively (designed to capture the time period for each of these phases of a trial). Relevant input units were activated to 1 for the Presentation and Choice inputs, and to 0 for the Screen input. The network response on each trial was measured when it was first allowed to reach (the first time step in the Choice input), with a stable response assessed when the activation level of an output unit exceeded 0.5. The number of processing cycles to reach the stable response was divided by 5 to scale to infant reaction times. This scaling factor is arbitrary; the substantive finding is the overall pattern of increased reactions times around the transitions from perseveration to success.

Measures

The model was tested with different strengths of the recurrent excitatory connections at the PFC layer. Performance was evaluated with 2 measures: percent correct response and reaction time. The network's percent correct response was computed as the activation of the appropriate output unit divided by the sum of activation over both output units (e.g., on A trials, A/[A + B]). Reaction time was computed as the number of processing cycles required to settle on a stable response, to assess predictions from the model. Performance was averaged across 50 different networks.

Acknowledgments

This research was supported by grants from the NIH (F31 NS42968), NICHD (R01 HD37163), and NSF (IBN-9873492). We thank Randy O'Reilly, Michael Frank, Harlene Hayne, David Jilk, Mike Mozer, and members of the Cognitive Development Center and Computational Cognitive Neuroscience Lab for helpful feedback.

Reprint requests should be sent to Yuko Munakata, Department of Psychology, University of Colorado-Boulder, 345 UCB, Boulder, CO 80309-0345, or via e-mail: munakata@colorado.edu.

Notes

1. In some models, the working memory component has been referred to as ''active'' memory, distinguished from ''latent'' memory in the connection weights. Active memory is viewed as one component of working memory.

2. We have discussed 2 manipulations to affect the activation associated with a given stimulus: the activation of additional units (e.g., the ''on'' and ''attached'' units in our model) and the differential activation of shared units (e.g., the ''cloth'' unit in the proposed Gap condition model). Either or both of these could subserve perceptual sensitivity to differences in visual displays. For example, although our model implements the activation of additional units, this could also be simulated with differential activity of shared units. In addition, the current model simply reflects the perceptual sensitivities infants bring to the task. A different model might explore how infants' perceptual sensitivities emerge and why they emerge in the order they do.

3. The behavioral studies also included pretest trials, in which infants were presented with a single cloth supporting a toy. Only infants who retrieved the toy on pretest trials (demonstrating a basic understanding of the cloth-pulling task) proceeded to test. Because the network had a basic understanding of the cloth-pulling task built in via its initial weights, it was unnecessary to test the network with pretest trials. In addition, infants in the behavioral studies received repeated A trials until they succeeded on 2 consecutive A trials. Thus, the 2 A trials presented to the networks reflected infants' experience immediately preceding the B trial.

REFERENCES

Aguiar, A., & Baillargeon, R. (2000). Perseveration and problem solving in infancy. In H. Reese (Ed.), *Advances in child development and behavior* (vol. 27, pp. 135–180). New York: Academic Press.

Bell, M., & Fox, N. A. (1992). The relations between frontal brain electrical activity and cognitive development during infancy. *Child Development, 63,* 1142–1163.

Berger, S. E. (2004). Demands on finite cognitive capacity cause infants' perseverative errors. *Infancy, 5,* 217–238.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of frontal cortex involvement in human working memory. *Neuroimage, 5,* 49–62.

Casey, B. J. (2001). Disruption of inhibitory control in developmental disorders: A mechanistic model of implicated frontostriatal circuitry. In R. Siegler & J. McClelland (Eds.), *Mechanisms of cognitive development: The Carnegie Symposium on Cognition* (vol. 28) (pp. 327–349). Hillsdale, NJ: Erlbaum.

Chafee, M. V., & Goldman-Rakic, P. S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *Journal of Neurophysiology, 83,* 1550.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review, 97,* 332–361.

Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review, 99,* 45–77.

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences, 7,* 415–423.

De Fockert, J., Rees, G., Frith, C., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science, 291,* 1803–1806.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18,* 193–222.

Diamond, A. (1981). Retrieval of an object from an open box: The development of visual–tactile control of reaching in the first year of life. *Society for Research in Child Development Abstracts, 3,* 78.

Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on AB. *Child Development, 56,* 868–883.

Diamond, A., Churchland, A., Cruess, L., & Kirkham, N. (1999). Early developments in the ability to understand the relation between stimulus and reward. *Developmental Psychology, 35,* 1507–1517.

Diamond, A., & Goldman-Rakic, P. S. (1989). Comparison of human infants and rhesus monkeys on Piaget's AB task: Evidence for dependence on dorsolateral prefrontal cortex. *Experimental Brain Research, 74,* 24–40.

Dias, R., Robbins, T. W., & Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin Card Sort Test: Restriction to novel situations and independence from ''on-line'' processing. *Journal of Neuroscience, 17,* 9285–9297.

Dunbar, K., & Sussman, D. (1995). Toward a cognitive account of frontal lobe function: Simulating frontal lobe deficits in normal subjects. In J. Grafman, K. J. Holyoak, & F. Boller (Eds.), *Structure and functions of the human prefrontal cortex* (pp. 289–304). New York: New York Academy of Science Press.

Engle, R. W. (2002). Working memory capacity as executive attention. *Current Directions in Psychological Science, 11,* 19–23.

Fuster, J. M., Bauer, R. H., & Jervey, J. P. (1985). Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain Research, 330,* 299–307.

Gerstadt, C., Hong, Y., & Diamond, A. (1994). The relationship between cognition and action: Performance of 3.5–7 year old children on a Stroop-like day–night test. *Cognition, 53,* 129–153.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. *Handbook of Physiology—The Nervous System, 5,* 373–417.

Goldman-Rakic, P. S. (1996). The prefrontal landscape: Implications of functional architecture for understanding human mentation and the central executive. *Philosophical Transaction of the Royal Society of London, B, Biological Sciences, 29,* 1445–1453.

Hauser, M. D. (1999). Perseveration, inhibition and the prefrontal cortex: A new look. *Current Opinion in Neurobiology, 9,* 214.

Hofstadter, M. C., & Reznick, J. S. (1996). Response modality affects human infant delayed-response performance. *Child Development, 67,* 646–658.

Hood, B. M., Hauser, M., Anderson, L., & Santos, L. (1999). Gravity biases in a non-human primate. *Developmental Science, 2,* 35–41.

Kastner, S., & Ungerleider, L. G. (2001). The neural basis of biased competition in human visual cortex. *Neuropsychologia, 39,* 1263–1276.

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

Lhermitte, F. (1986). Human autonomy and the frontal lobes: Part II. Patient behavior in complex and social situations: The ''environmental dependency syndrome.'' *Annals of Neurology, 19,* 335–343.

Lockman, J. J., & Pick, H. L. (1984). Problems of scale in spatial development. In C. Sophian (Ed.), *Origins of cognitive skills* (pp. 3–26). Hillsdale, NJ: Erlbaum.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* 167–202.

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefontal cortex of the macaque. *Journal of Neuroscience, 16,* 5154–5167.

Milner, B. (1963). Effects of different brain lesions on card sorting. *Archives of Neurology, 9,* 90–100.

Morton, J. B., & Munakata, Y. (2002a). Active versus latent representations: A neural network model of perseveration and dissociation in early childhood. *Developmental Psychobiology, 40,* 255–265.

Morton, J. B., & Munakata, Y. (2002b). Are you listening? Exploring a knowledge action dissociation in a speech interpretation task. *Developmental Science, 5,* 435–440.

Morton, J. B., & Munakata, Y. (in preparation). The appearance (and disappearance) of knowledge-action dissociations in adults.

Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP model of the AB task. *Developmental Science, 1,* 161–184.

Munakata, Y., & Yerys, B. E. (2001). All together now: When dissociations between knowledge and action disappear. *Psychological Science, 12,* 335–337.

Oh, S. H., & Kim, M. S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin and Review, 11,* 275–281.

O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 375–411). New York: Cambridge University Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain.* Cambridge: MIT Press.

O'Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex, 12,* 246–257.

Petrides, M. (2000). The role of the mid-dorsolateral prefrontal cortex in working memory. *Experimental Brain Research, 133,* 44.

Piaget, J. (1954). *The construction of reality in the child.* New York: Basic Books.

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

Stuss, D., & Benson, D. (1984). Neuropsychological studies of the frontal lobes. *Psychological Bulletin, 95,* 3–28.

Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin and Review, 11,* 269–274.

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