

Generalizing the Dynamic Field Theory of the A-not-B Error Beyond Infancy: Three-Year-Olds' Delay- and Experience-Dependent Location Memory Biases

Anne R. Schutte and John P. Spencer

Thelen and colleagues recently proposed a dynamic field theory (DFT) to capture the general processes that give rise to infants' performance in the Piagetian A-not-B task. According to this theory, the same general processes should operate in noncanonical A-not-B-type tasks with children older than 12 months. Three predictions of the DFT were tested by examining 3-year-olds' location memory errors in a task with a homogeneous task space. Children pointed to remembered locations after delays of 0 s to 10 s. The spatial layout of the possible targets and the frequency with which children moved to each target was varied. As predicted by the DFT, children's responses showed a continuous spatial drift during delays toward a longer term memory of previously moved-to locations. Furthermore, these delay-dependent effects were reduced when children moved to an "A" location on successive trials, and were magnified on the first trial to a nearby "B" location. Thus, the DFT generalized to capture the performance of 3-year-old children in a new task. In contrast to predictions of the DFT, however, 3-year-olds' responses were also biased toward the midline of the task space—an effect predicted by the category adjustment (CA) model. These data suggest that young children's spatial memory responses are affected by delay- and experience-dependent processes as well as the geometric structure of the task space. Consequently, two current models of spatial memory—the DFT and the CA model—provide incomplete accounts of children's location memory abilities.

INTRODUCTION

Spatial memory has captured the interests of developmental psychologists for decades, largely driven by Piaget's demonstrations that children's understanding of space undergoes considerable change during the first few years of life (e.g., Piaget, 1954; Piaget & Inhelder, 1956). One of the most dramatic—and most studied—changes in early development occurs in the A-not-B situation. After repeatedly finding hidden toys at an "A" location, 8- to -10-month-old infants will reach back to A 3 to 5 s after seeing a toy hidden at a nearby "B" location. A few months later, infants begin responding correctly in this task at the same delay—they reach to B on the B trials (e.g., Diamond & Doar, 1989; Piaget, 1954; Smith, Thelen, Titzer, & McLin, 1999).

A variety of conceptual and mathematical models have been proposed to explain the A-not-B error and its disappearance in early development. In the past decade, at least four different models have been proposed: the memory + inhibition model (Diamond, Cruttenden, & Neiderman, 1994), the response vs. representation model (Marcovitch & Zelazo, 1999; Zelazo & Zelazo, 1998), the latent versus active memory model (Munakata, 1998; Munakata, McClelland, Johnson, & Siegler, 1997), and the dynamic field model (Smith et al., 1999; Thelen, Schöner, Scheier, & Smith, 2001). Each of these models gives a different explanation for why the A-not-B error goes away

around 12 months of age, ranging from changes in inhibitory processes (see Diamond et al., 1994) to changes in the ability to sustain location-related activation in working memory (see Thelen et al., 2001).

Although there are dramatic changes in children's spatial memory abilities around 12 months, development does not, of course, stop there. At approximately 2 years, children remember locations more accurately, even after manipulations of body position (Huttenlocher, Newcombe, & Sandberg, 1994; Newcombe, Huttenlocher, Drummey, & Wiley, 1998). This advance occurs, in part, because 2-year-olds are more adept at using distant landmark cues (Newcombe et al., 1998). Two-year-olds also show evidence of using boundaries and edges to help them categorize space. For instance, several studies have found that 2-year-olds' memory responses in a sandbox task are biased toward the center of the sandbox (Huttenlocher et al., 1994; Spencer, Smith, & Thelen, 2001). According to Huttenlocher and colleagues (1994), these memory errors reflect biases toward the center of a spatial category formed when children use the edges of the sandbox as category boundaries.

How are these spatial memory abilities that emerge around 2 years related to the advances taking place at 12 months in the A-not-B situation? Is there continu-

ity in the development of spatial memory abilities between infancy and early childhood? This issue has received relatively little attention in the literature. There are many models that capture the performance of infants in the canonical Piagetian A-not-B task. In contrast, different accounts capture young children's performance in tasks such as the sandbox task, most notably, the category adjustment (CA) model proposed by Huttenlocher, Hedges, and Duncan (1991).

From one perspective, the existence of a set of models to explain infants' performance in the A-not-B task and a separate set of models to explain 2-year-olds' performance in other tasks is not problematic. Rather, these separate accounts reflect the reality of development: different systems lead to different errors at different ages in different tasks. For instance, several researchers have proposed that the A-not-B error is caused by a bias in a "response" system built up by repeated hiding and finding events at A (Marcovitch & Zelazo, 1999; Newcombe & Huttenlocher, 2000). By contrast, 2-year-olds' performance in the sandbox task reflects the characteristics of a "place learning" system—the system involved in encoding locations relative to external landmarks (Newcombe & Huttenlocher, 2000; Newcombe et al., 1998). Thus, one can effectively study the development of the response and place-learning systems by examining the performance of children in these different tasks.

In the present study, we adopted a different perspective and asked whether the same processes involved in performance in the A-not-B situation affect the performance of older children in different tasks. In particular, we adopted the general processes perspective specified by one model of the A-not-B error—the *dynamic field theory* (DFT) proposed by Thelen and colleagues (2001). According to this model, the A-not-B error results from general processes involved in planning and remembering goal-directed actions to specific locations. Because these processes are general, they should operate in any task that requires children to plan and remember actions to specific locations. Thus, the model should, in theory, be able to capture both infants' performance in the A-not-B situation and young children's performance in other spatial memory tasks. This possibility was examined by generating a set of predictions from the dynamic field model about the types of spatial memory errors young children should make in a task that differs from the A-not-B situation. Specifically, we predicted that in the absence of salient location cues, young children's memory responses would show a continuous spatial "drift" during short-term delays toward locations moved to on previous trials.

The Dynamic Field Theory of the A-not-B Error

The DFT was proposed to explain how an infant's plan to move to a particular spatial location in the A-not-B task changes in the context of the different events that happen in the task (for a mathematical treatment of the theory, see Amari, 1977; Amari & Arbib, 1977; Schöner, Kopecz, & Erlhagen, 1997; Thelen et al., 2001). Given the prevalence of neural network models in the literature, it is useful to think about the DFT in neural network terms. As in most network models, the DFT specifies how "activation" changes in the context of different inputs. In the A-not-B situation, three inputs are relevant (Thelen et al., 2001). The task input captures the stable perceptual cues in the task space that designate the different behavioral alternatives—the lids that mark the hiding locations; the specific input captures the specification of the target location, for instance, the experimenter hiding a toy at the A or B location; and the memory input captures infants' *longer term* memory of where they have reached on previous trials.

These inputs are fed into a motor planning field. Unlike many network models, the neurons in the planning field have a spatial ordering—neurons are ordered from left to right to capture the continuous range of possible reachable locations from a far left location to a far right location. Stronger activation at particular points along this field indicates a stronger representation of the associated locations in space. Conceptually, this means that an infant is more "likely" to reach to locations associated with highly activated sites in the field. (Note that "likely" is used here in a conceptual way, and is not meant to refer to probability density functions.) Given the relationship between sites in the network and locations in space, patterns of activation in the model can be translated into a "reaching response" in a rather simple way—the most strongly activated site in the planning field at a particular moment in time determines where the model (i.e., the infant) will reach.

The central elements of the DFT are shown in Figure 1. This figure shows how the model behaves on the first B trial in the canonical A-not-B task (i.e., the first trial after pretraining to the A location and two A trials). The top panel shows a simulation of an 8- to 10-month-old infant's performance. The bottom panel shows a simulation of a 12-month-old infant's performance. The figures in the left column of each panel show the three inputs to the model, whereas the figures to the right in each panel (Figures 1d, and 1h) show the planning field. In each figure, the range of possible reaching locations are captured along the x-axis; time, from the start of a trial (0 s) to the end of




Figure 1 The dynamic field theory of the A-not-B error. The top panel shows time-dependent changes in a motor planning field (d) in the context of three inputs (a, b, c) during the first B trial for an 8- to 10-month-old infant. The bottom panel shows changes in a motor planning field (h) in the context of the three inputs (e, f, g) during the first B trial for older infants. White arrows in (b, f) indicate the appearance of the toy at 2 s. Black arrow in (d) indicates the increase in activation in the planning field as a result of this event. See text for further details.

a trial (10 s), is on the y-axis; and activation is on the z-axis.

Figure 1a shows the task input. In the canonical A-not-B task, there are two hiding locations covered by identical lids. This is captured in the figure by the two peaks of input activation centered over the A and B locations. These activation peaks are relatively small because the lids in the canonical A-not-B task are not very salient. Furthermore, because the two possible hiding locations are visible throughout the trial, the task input remains constant across the duration of the trial (from 0 s to 10 s). The specific input is shown in Figure 1b. At the start of the trial when the toy is not visible, the specific input is 0. Next, the experimenter waves the toy near the B location and hides it under the B lid. This event is captured by the strong input activation at the B location between 2 s and 4 s (see white arrow in Figure 1b). After the hiding event is over (after 4 s), the specific input is once again 0 (i.e., the toy is not visible). The final input, the memory input, is shown in Figure 1c. The memory input has activation centered at the A location. Recall that Figure 1 shows how the model behaves on the first B trial. Thus, the stronger activation at A in the memory input reflects the infant's longer term memory of the previous trials to A. Note that the static memory input used here is a simplification of the dynamically constructed memory input described in Thelen et al. (2001). Although the construction of the memory input is neglected for simplicity, our simulations show the same qualitative effects that are present in the full model.

Figure 1d shows the planning field for a younger infant (8–10 months). This field integrates the task, specific, and memory inputs as the different events in the trial unfold. At the start of the trial, there is stronger activation in the planning field at A than at B. This is due to the stronger memory input at A. From 2 s to 4 s, the experimenter waves the toy at the B location. The strong specific input that captures this event builds a peak of activation in the planning field at the B location (see black arrow in Figure 1d). At 4 s, the toy is hidden. In the absence of strong specific input, the activation peak at B in the planning field decays. Consequently, by the time the infant is allowed to reach (at 10 s), activation at A—driven by the memory input—dominates and the infant reaches to A. That is, the infant makes the A-not-B error.

The bottom panel of Figure 1 shows a simulation of the model that captures the performance of older infants (12 months). Notice, first, that the inputs in the left column (Figures 1e, 1f, and 1g) are identical to the inputs shown in the top panel. The planning field, however, behaves quite differently (see Figure 1h). At the

start of the trial, there is stronger activation at A than at B. As with the younger infant in Figure 1d, this reflects the stronger memory input at A. Next, the toy is held up and hidden at the B location. This leads to strong specific input at the B location that is amplified in the planning field. When the toy is hidden at 4 s, the specific input goes to 0 (see Figure 1f). Nevertheless, activation at the B location in the planning field is maintained during the delay. Thus, even in the absence of input, the planning field is able to retain an accurate memory of the hiding event at B. Consequently, after the delay (at 10 s), the older infant reaches correctly to the B location.

What accounts for the qualitatively different patterns of activation in the two planning fields shown in Figure 1; that is, why does the model of the older infant stop making the A-not-B error? Thelen et al. (2001) proposed that neurons in the planning field go from being weakly “cooperative” to being strongly cooperative between 10 and 12 months. Cooperativity refers to how neurons “work with” one another in the planning field: when a neuron is activated, it increases the activation of nearby “neighbors” (local excitation) and suppresses the activity of far away neighbors (lateral inhibition). Weak cooperativity means that neurons do not have a strong affect on one another. Instead, neuronal activity is largely determined by input. Strong cooperativity, by contrast, means that neuronal activity can take on a life of its own; even after strong input has disappeared, neurons can continue to excite one another within the local region initially stimulated by input. This self-sustaining stimulation is what maintains the activation peak at the B location in Figure 1h after the specific input is set to 0.

It is important to note two characteristics of the developmental account proposed by Thelen et al. (2001). First, it only takes a small, quantitative parameter change to switch the dynamic field model from weak to strong cooperativity. Thus, as is the case with many dynamic systems accounts of developmental phenomena, a small change in the parameters of the model can lead to qualitatively different behaviors over development (e.g., Hartelman, van der Maas, & Molenaar, 1998; van der Maas & Molenaar, 1992; Van Geert, 1998). Second, although changes in cooperativity can be realized with a small parameter change, there are also other ways to create self-sustaining local excitation. For instance, self-sustaining activation can be created by very strong input, the type of input that might result from perceptual learning. Thus, development in the model should not be considered an all-or-none developmental event. Rather, this type of developmental change is likely context and experience dependent (Smith et al., 1999; Thelen et al., 2001).

The DFT Is a General Processes Account

Although the DFT was proposed to explain the performance of 8- to 12-month-old infants in the A-not-B task, the processes formalized in the field theory are general processes that make goal-directed actions to remembered locations. Thus, although the details of the inputs in Figure 1 are specific to the A-not-B situation—two lids mark the possible hiding locations, and there is a relatively strong memory of the A location on the first B trial—the processes that govern activation in the field are not. Consequently, these processes should operate in tasks other than the canonical Piagetian A-not-B task and should capture the behavior of children older than 12 months.

Spencer et al. (2001) examined the generality of the model by investigating 2-year-olds' location memory responses in a sandbox task. In this task, children watched as a toy was buried in a long, rectangular sandbox. After a short delay, children were allowed to search for the toy (see also, Huttenlocher et al., 1994; Newcombe et al., 1998). The toy was hidden at an A location several times in a row. These A trials were followed by several trials to a B location 8 inches to 10 inches (20.3 cm to 25.4 cm) from A. Two-year-olds' responses on the A trials were accurate. Nevertheless, on the B trials, children's responses were systematically biased toward A: children typically reached in between A and B, and these reaches were biased in the direction of A regardless of the layout of A and B in the sandbox. Furthermore, as with 8 to 10-month-old infants in the canonical A-not-B task (see Marcovitch & Zelazo, 1999; Smith et al., 1999), the magnitude of 2-year-olds' biases toward A depended on the number of A trials: biases toward A were stronger after three A trials than after one A trial.

Figure 2 shows a simulation of the DFT that illustrates how 2-year-olds' errors on the B trials in the sandbox task may have come about. This figure shows how activation evolves in a planning field operating in the self-sustaining mode (e.g., a 2-year-old). The inputs to this field are identical to the inputs used in Figure 1 with one exception—there is no task input. This reflects the homogeneous nature of the task space in the sandbox task. As can be seen in Figure 2d, the absence of task input has a dramatic effect on activation in the planning field during the delay: the self-sustaining activation peak at the B location "drifts" toward the A location during the delay. Consequently, the 2-year-old will reach in between A and B after the memory delay.

Why does activation drift during the delay? Self-sustaining peaks like the one shown in Figure 2d are sensitive to inputs: they are attracted toward inputs

provided that the location of the input and the location of current activation in the planning field are relatively close. This is the case in Figure 2: the right edge of the activation peak in the planning field at B (see Figure 2d) overlaps with the left edge of the memory input at A (see Figure 2c). Thus, when the specific input at B goes to zero (i.e., when the toy is hidden), the remaining input to the planning field—the memory input—is larger on the right side of the self-sustaining peak than the left. As a consequence, the peak in the planning field begins to drift toward A (i.e., toward the memory input), because there is slightly stronger local excitation on the A-ward side of the peak. Note that this did not occur in the A-not-B task depicted in the bottom panel of Figure 1, because the task input at B helped keep the peak centered at the B location.

In summary, the simulation in Figure 2 demonstrates how the DFT can account for 2-year-olds' location memory biases in the A-not-B version of the sandbox task (Spencer et al., 2001). These results demonstrate that the processes proposed to explain the performance of 8- to 12-month-old infants in the A-not-B situation can be usefully generalized to explain the performance of older children in other spatial memory tasks. Furthermore, this simulation reveals a new characteristic of the field model—self-sustaining peaks can show continuous spatial drift during delays. In the present study, three specific predictions derived from this characteristic of the DFT were tested.

Specific Predictions

The first prediction tested was that young children's location memory responses will show a continuous spatial drift over short-term delays in tasks with a homogeneous task space. This prediction is consistent with 2-year-olds' reaches in between the A and B locations in Spencer et al. (2001) and with the continuous nature of children's biases in studies by Huttenlocher, Newcombe, and Sandberg (Huttenlocher et al., 1994; Sandberg, Huttenlocher, & Newcombe, 1996). Memory delays, however, were not systematically manipulated in these previous studies. In the present study, the spatial drift hypothesis was tested by measuring children's location memory responses following delays of 0 s, 5 s, and 10 s.

Second, the DFT predicts that in tasks with a homogeneous task space, children's responses will drift toward their longer term memory of previously moved-to locations—the memory input—when this input is close to the location currently activated in working memory. Spencer et al. (2001) demonstrated that 2-year-olds' responses were biased toward an A

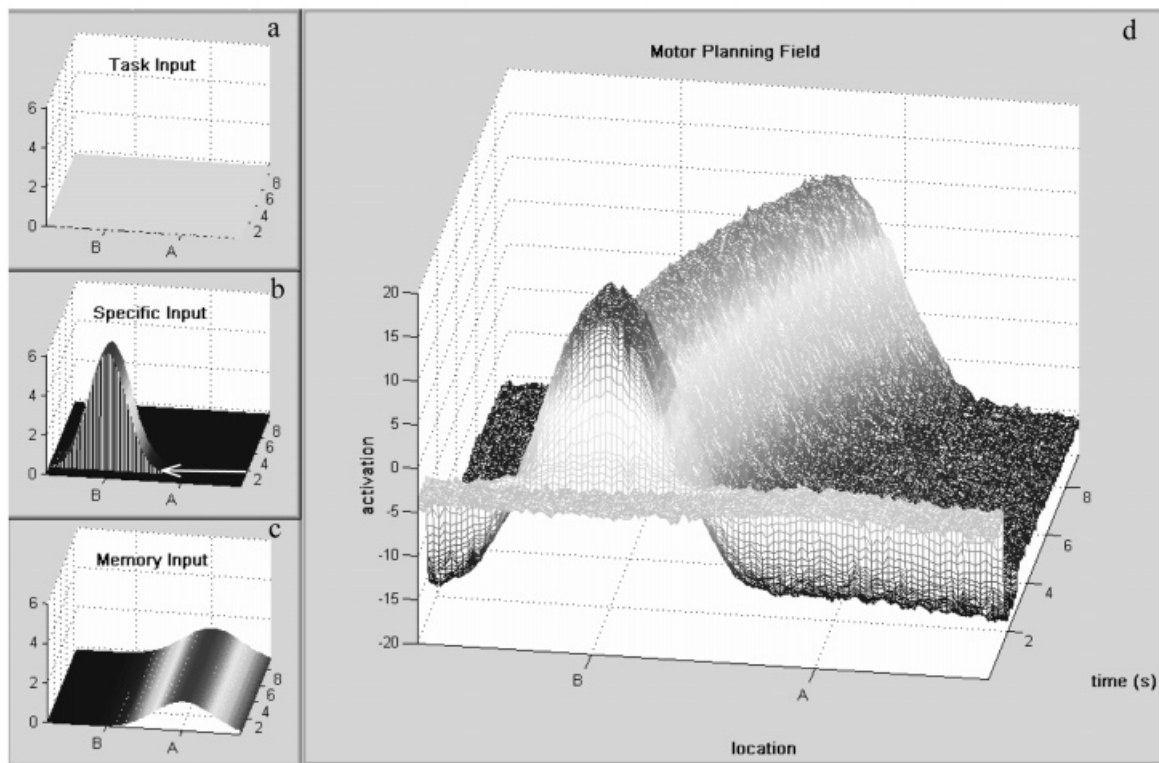


Figure 2 The dynamic field theory with homogeneous task input. Figure details are identical to Figure 1 with one exception—the task input (a) was set to 0. See text for further details.

location in the sandbox task at an 8 to 10-inch (20.3 cm to 25.4 cm) separation. The present study examined whether 3-year-olds also show this characteristic of spatial memory. This is important because several accounts of the A-not-B error predict that biases toward information stored in longer term memory should be less prevalent over development (e.g., Diamond, 1990a; Diamond, 1990b; Munakata, 1998). Thus, 3-year-olds—in contrast to the 2-year-olds in Spencer et al.—may not show response biases toward previously moved-to locations over delays.

Finally, according to the DFT, the direction and magnitude of the delay-dependent drift in children's responses will depend on the strength and spatial structure of the memory input. Specifically, if there is strong memory input on one side of an activation peak in working memory, 3-year-olds will show significant response biases toward this input even at short delays (e.g., 5 s). In contrast, if memory input is far from the locations currently activated in working memory, responses should be accurate even at long delays (e.g., 10 s). The characteristics of the memory input were manipulated in two ways. As in Spencer et al. (2001), the memory of one location was differen-

tially strengthened by increasing the number of trials to a biased or "A" location. In addition, the spatial separation between target locations was increased: if targets are far apart, there is less chance that activation associated with one location will be attracted toward memory input at a second location.

Readers familiar with recent accounts of the A-not-B error may question whether these three predictions are novel. For instance, three other accounts of the A-not-B error predict that infants will make errors in between A and B provided that there are multiple hiding locations in the task space (Diamond et al., 1994; Marcovitch & Zelazo, 1999; Munakata, 1998; Munakata et al., 1997). Might these models also explain why older children make errors in between A and B in a sandbox task? The answer is "not necessarily," because none of these accounts specify a priori how the presence of lids in the canonical A-not-B task and the absence of lids in the sandbox task affect children's ability to remember locations—and these subtle task differences matter. For instance, according to most accounts of the A-not-B error, such errors disappear around 12 months of age. Thus, if 2-year-olds were put in an A-not-B task with multiple hiding lo-

cations, they should not make errors in between A and B. In contrast, Spencer et al. (2001) demonstrated that 2-year-olds do make A-not-B-type errors in the sandbox task. To explain why children of the same age show differences in performance across tasks, an account must specify how the task structure affects young children's ability to remember locations. As can be seen in Figures 1 and 2, the DFT is up to this challenge. In this sense, the predictions of the DFT are novel.

EXPERIMENT 1

In the first experiment, the three predictions of the DFT were tested using a new task—the spaceship task (see also, Hund & Spencer, 2002; Spencer, & Hund, 2002, in press). In this task, children see a spaceship appear and then disappear on a large, homogeneous tabletop. When they hear a “go” signal, children move a marker (a rocket) to the location of the target (the spaceship). Thus, as in the sandbox task, children must remember a target location in the absence of salient location cues. Although somewhat artificial, this aspect of the task is critical to testing the mechanistic predictions of the DFT about how sustained activation in working memory (i.e., in the planning field) drifts during delays. Importantly, the spaceship task has several advantages over the sandbox task. First, because a computer controls the presentation of the target and the “go” signal, it is possible to control the location of the target and the duration of the memory delay with a high degree of precision. Second, a motion analysis system was used to record children's responses (i.e., the position of the rocket). This system can record the location of the marker with millimeter precision, allowing for the detection of even subtle changes in children's responses.

To examine whether children's responses show a continuous spatial drift toward a longer term memory of the target locations, the strength and spatial structure of the memory input across conditions was manipulated. Specifically, the spatial separation between adjacent targets was varied. The predicted effect of this manipulation is illustrated in Figure 3. Recall that the memory input in the DFT is built up by repeatedly moving to the same locations across trials (Thelen et al., 2001). Thus, during an experimental session, the memory input comes to reflect the spatial layout and frequency with which the targets were visited. Figure 3a depicts the memory input that might be constructed by having children move to three locations (-20° , 0° , and 20°) equally often. The dotted lines in this figure capture the activation in longer term memory associated with each individual target

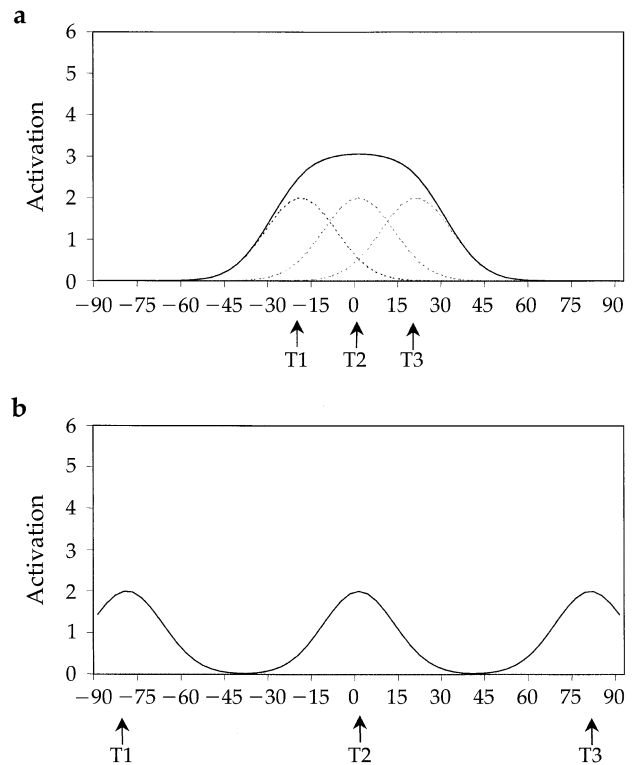


Figure 3 Example of activation in the memory input with three targets that are (a) close together and (b) far apart. Arrows indicate target (T) locations; dashed lines reflect activation associated with each target; solid lines are the sum of the individual target activations.

location. Because these activation distributions overlap, they will sum in memory, resulting in the activation distribution shown by the solid line. Thus, after moving repeatedly to targets at -20° , 0° , and 20° , children's longer term memory will be biased toward an averaged location near 0° . Consequently, when the 20° location, for instance, is the target on a particular trial, there will be stronger memory input to the left of this location and activation in working memory should drift toward 0° during the delay.

Figure 3b shows the memory input built up by having children move repeatedly to three locations that are far apart: at -80° , 0° , and 80° . In this case, the activation associated with each individual location is identical to the individual activation profiles shown in Figure 3a. The only difference is that the individual targets are far apart. This has a dramatic effect on the summed activation: the summed activation is distributed symmetrically above each target location (see solid line). Thus, when one of the locations, for instance 80° , is selected as the target, activation in working memory should not drift

during the delay due to symmetric memory input around 80° . The memory input will help children's responses remain accurate.

Note that the target locations in the present study varied only in their angular separation from the midline of the table; they were always the same distance from a starting location. Only target direction was manipulated because previous findings indicate that children younger than 9 years show less systematic response biases when locations vary in both direction and distance than when locations vary in direction only (Sandberg et al., 1996). Nevertheless, both aspects of children's responses were examined.

Method

Participants

A total of thirty-three 36- to 40-month-olds ($M = 39.1$ months, $SD = 1.28$ months) participated in Experiment 1. Participants were recruited from a database at the University of Iowa, and were tested in two sessions that were scheduled within a week of each other. After each session, the participants received a \$5 gift certificate to a local toy store, and following the second session they received a certificate of participation. The parents of all participants gave informed consent. Twenty-two of the participants were female and 11 were male. Data were collected from 13 additional participants; however, these data were not analyzed for the following reasons: 9 participants did not complete the task (5 from the condition in which the targets were separated by 80°), and 4 participants did not have at least one trial of useable data to each target and delay following data processing (2 from the 80° condition). Thus, clearly there was difficulty in getting children to complete the 80° condition. We return to this issue in the Discussion section for Experiment 1.

Apparatus

The participants stood at a $4\text{ ft} \times 4\text{ ft}$ ($1.22\text{ m} \times 1.22\text{ m}$) horizontal table (see Figure 4). The top of the table was a uniform piece of Plexiglas™ with an arc removed from one side. The participants stood within the arc with a parent or guardian seated in a chair directly behind them. The top of the table was at approximately belly height. The Plexiglas was covered with black tinting on one side to keep the participants from seeing the light fixtures placed underneath. The room lights were dimmed and black curtains were hung along the walls to the front and sides of the table and across the ceiling. This prevented reflections from appearing on the tabletop that could be used as reference points. A yellow sticker was placed along the midline axis of the table, 15 cm from the front edge of the table. This was the starting point for each trial. A rocket ship 8.5 cm high and 3 cm in diameter sat on this sticker. Children used the rocket to indicate where "spaceships" were hiding.

Light-emitting diodes (LEDs) were placed below the tabletop every 10° from -90° to 90° . The LEDs were positioned in a semicircle with a radius of 15 cm centered at the starting point. An X-ray film was placed on top of the LED board. This film was black except for spaceship shapes that were aligned with the LEDs. A thin piece of white paper on top of the X-ray film diffused the LED light. The LED voltage was adjusted to a level at which there were no visual afterimages. A computer equipped with an I/O board connected to an external relay board controlled the lights. The board was able to trigger the target lights with better than 10-ms precision. The computer controlled the type and timing of all stimuli presented in the experiment using customized software. Prerecorded messages were played through two speakers on either side of the table. These messages led the participants through the game and gave feedback after each trial.




Figure 4 Diagram (b) of experimental apparatus. Inset (a) shows the possible target locations.

Participants' movements of the rocket were recorded at 150 Hz using an optical–electronic motion analysis system (Optotrak, Northern Digital, Inc., Waterloo, ON, Canada). This system tracks small (radius = 3.5 mm), individually pulsed infrared emitting diodes (IREDs) within a specified three-dimensional volume with better than 1-mm precision. One IRED was attached to the tip of the rocket to track children's responses as they moved the rocket from the starting position to the remembered location.

Task and Procedure

When each participant arrived at the laboratory, the study was described to the parent while the child played with toys and became comfortable with the laboratory environment. After the parent signed consent forms, the experimenter played a warm-up game on the floor with the child to teach the child the basics of the task. The child was told that he/she was going to play a game. The goal of the game was to help "Buzz Lightyear™" find his lost spaceships. The experimenter gave the child the toy rocket and then turned over two flashcards, one with a spaceship on it and one with a star. The experimenter pointed out the spaceship card and placed both cards face down on the floor. When the experimenter said "Go," the child was encouraged to place the rocket on top of the spaceship card. If the child "found" a spaceship, a star sticker was placed on a certificate of participation as a reward. The warm-up game was played until the child successfully found at least two spaceships in two different locations.

Next, the child and parent moved over to the experimental table to start the spaceship task. The session began with demo trials to help the child learn the game. These trials were identical to the test trials except that the experimenter performed the task. The experimenter controlled the number of demo trials. Generally, children required only one demo; however, demo trials were repeated if a child did not attend to a complete demo trial or was not willing to participate.

Each demo or test trial began when the computer said, "Let's look for a spaceship." Following a random pretrial delay, a spaceship was illuminated for 2 s on the table in front of the child. The child's task was to move the rocket from the starting location to where the spaceship was hiding when the computer said "Go, go, go." If the participant moved the rocket before the "Go" signal, the computer gave a verbal warning such as, "Don't forget to wait for the go." After each trial, the target was reilluminated for 1.5 s so that the child could compare the actual location with the loca-

tion of the rocket (i.e., the remembered location). The child received verbal feedback from the computer based on whether he/she found the spaceship (the distance between the response and the target was <5 cm), was close to the spaceship (the response–target distance was >5 cm and <8 cm), or was not so close (the response–target distance was >8 cm; as a basis for comparison, note that targets separated by 20° were 5.21 cm apart, and targets separated by 40° were 10.26 cm apart). For each spaceship the child found, a star was placed on the participation certificate. The parent or guardian was instructed not to talk during a trial or give any signal that would help the child find the spaceship, but was asked to give positive feedback after each trial.

Experimental Design

Children were randomly assigned to one of four conditions. In each condition, children moved to three target locations: one to the left of the midline axis of the table, one aligned with the midline axis (0°), and one to the right of midline (see Figure 4). Children moved to these locations after three delays: 0 s, 5 s, and 10 s. Across conditions, the angular distance separating the three targets was varied. The separation between adjacent targets in the four conditions was 20°, 40°, 60°, and 80°. For example, in the 20° condition, targets were located at –20°, 0°, and 20° (see Figure 4).

Children participated in two sessions that were each approximately 20 minutes long. The two sessions were identical except that the warm-up game was not played before the second session. Each session consisted of 6 trials at each delay (2 to each target), for a total of 18 trials per session. Thus, across the two sessions, children were asked to complete 36 trials, 4 to each target at each delay. The target presentation and order of the delays were randomized such that no more than two trials to the same location occurred in a row. Participants were encouraged to complete all 18 trials during each session; however, during some sessions, children stopped playing the game early. On average, children who were included in the final analyses completed 29.48 trials across the two sessions ($SD = 5.61$).

Method of Analysis

The Optotrak data were used to identify a starting and ending location for each trial. This was difficult for some trials. For example, sometimes children moved the rocket to a location, touched the table at that location, and then moved the rocket back to the

start. When this happened, the location at the end of the trial was near the starting position. To overcome such difficulties, computer-automated software was used to identify the start and end of the movement. The start of the movement was defined as the first data frame in a trial with a tangential velocity >2 cm/s. This "resting level" criterion was used by Spencer and Hund (in press) to distinguish low-level noise from the movement of the hand. Several trials were checked via visual inspection to verify that this criterion worked effectively with 3-year-olds' movements. The end of the movement was identified by searching backward from the end of the trial to the last frame that had a velocity less than 2 cm/s and a z-coordinate value (vertical dimension) less than 11 cm and greater than 6 cm (recall that the rocket was 8.5-cm high). A z-coordinate greater than 11 cm meant that the rocket was still in the air above the table. A z-coordinate less than 6 cm meant that the rocket had fallen over.

After the start and end locations were selected, the computer calculated the directional and distance error for each trial. Specifically, the computer calculated the angle between the line connecting the start location and the target location and a line connecting the start location and the ending location. Negative directional errors indicate counterclockwise errors relative to the target direction. The computer also calculated the distance along the start–end line between the ending location and the target distance. Negative distance errors indicate that children undershot the target.

All trials from the computer-automated analyses were checked for computer selection mistakes using an interactive version of the automated analysis software before final analysis. The interactive software allowed for manual editing of the starting and ending locations; however, all manual selections were required to meet the starting and ending criteria outlined above. After data processing, four children did not have at least one trial to each target at each delay. Therefore, data from these children were removed from the final analyses. For the remaining children, the median directional and distance error to each target at each delay were analyzed. Unlike the mean response, children's median responses were less affected by the variability in performance from trial to trial.

Results

Directional Error

The directional errors to each target at each delay are shown in Figure 5 for the four separation conditions. Negative errors indicate counterclockwise errors; positive errors indicate clockwise errors. As can

be seen in the figure, responses to the left and right targets generally drifted toward the center target as the delay increased. Specifically, errors to the left target drifted clockwise (positive errors) over delays, whereas errors to the right target drifted counterclockwise (negative errors) over delays. The drift toward the center target was greatest in the 20° and 60° conditions, and smallest in the 40° condition. Errors to the center target remained generally accurate across the 10-s delay in all conditions.

Children's median directional errors were analyzed in a three-way repeated-measures ANOVA with separation as a between-subjects factor and target and delay as within-subjects factors. There was a significant main effect of target, Wilks' $\Lambda = .62$, $F(2, 28) = 8.67$, $p = .001$, and a significant Delay \times Target interaction, Wilks' $\Lambda = .29$, $F(4, 26) = 15.79$, $p < .001$. No other effects in the ANOVA reached significance. Thus, in contrast to our prediction that children's biases toward the memory input would decrease at large target separations, children's responses to the left and right targets were biased toward the center target, even in the 80° condition.

The Delay \times Target interaction is shown in Figure 6. Tests of simple effects revealed that directional error to the right target, $F(2, 64) = 10.90$, $p < .001$, and left target, $F(2, 64) = 9.69$, $p < .001$, increased significantly over delay, whereas directional error to the center target, $F(2, 64) = .62$, *ns*, did not. Follow-up *t*-tests indicated that directional errors to the right target increased significantly between 0 s and 5 s, $t(32) = 2.80$, $p < .005$, and between 5 s and 10 s, $t(32) = 2.08$, $p < .025$. Similarly, directional errors to the left target increased significantly between 0 s and 5 s, $t(32) = -2.34$, $p < .025$. Directional errors to the left target also increased between 5 s and 10 s, but this increase was only a marginal effect, $t(32) = -1.60$, $p < .07$. These results are consistent with the predictions of the DFT: children's responses to the left and right targets showed continuous spatial drift toward the center target during the memory delay.

Distance Error

Figure 7 shows the mean distance error to each target at each delay for the four target separation conditions. Positive errors indicate that children overshot the target, and negative errors indicate that children undershot the target. Generally, 3-year-olds undershot the targets slightly at the shortest delay (0 s), and overshot the targets following the 5 and 10 s delays. These data were analyzed in a three-way repeated-measures ANOVA with separation as a between-subjects factor and target and delay as within-subjects

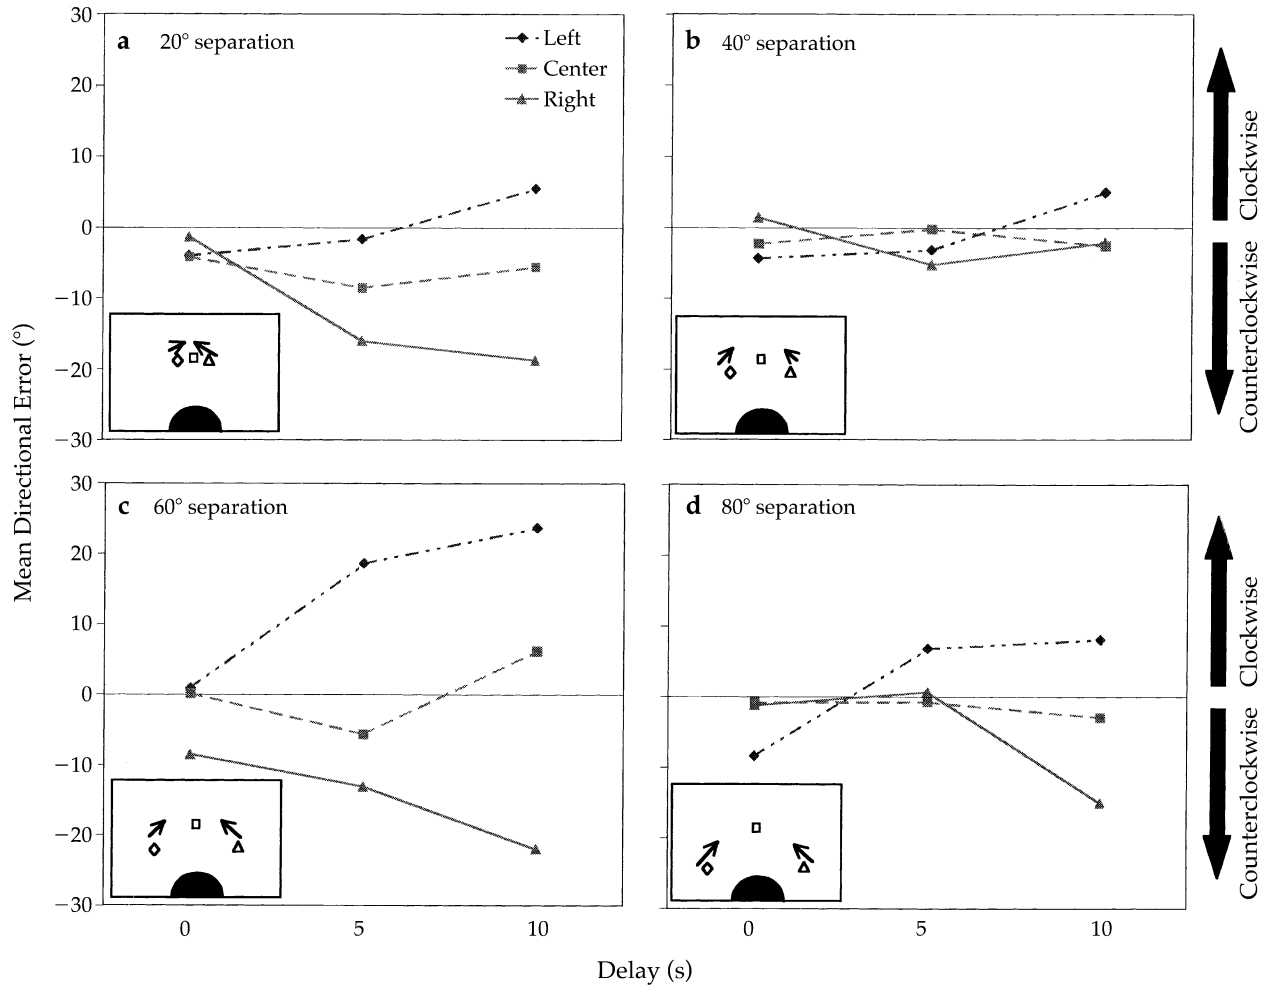


Figure 5 Mean directional error in the (a) 20°, (b) 40°, (c) 60°, and (d) 80° target separation conditions for movements to the left (dot-dashed line, diamond symbol), center (dashed line, square symbol), and right (solid line, triangle symbol) targets over delays. Diagrams of the table top (insets) with arrows denoting the direction of errors are included for reference. The symbols marking the locations of the spaceships on the table correspond to the symbols used for the different targets on the graph.

factors. There was a significant delay main effect, Wilks' $\Lambda = .43$, $F(2, 27) = 17.82$, $p < .001$, but no other significant effects. Thus, the increase in distance error over delays in Figure 7 was statistically reliable. In contrast to directional error, however, which increased systematically across delays, follow-up t tests indicated that distance error increased significantly between 0 s and 5 s, $t(31) = 4.47$, $p < .001$, but not between 5 s and 10 s, $t(33) = -1.83$, *ns*. No other effects in the ANOVA reached significance.

Discussion

As predicted by the DFT, 3-year-olds' directional responses in the spaceship task showed continuous

spatial drift, even at short delays. Children's directional responses to the left and right targets were biased toward the center target, and this bias increased from 0 s to 5 s and from 5 s to 10 s. The drift in directional responses toward the center target was consistent with the prediction that children's responses would be biased toward a longer term memory of the target locations. Specifically, we predicted that children's longer term memory of the left, center, and right targets would blend together, producing strong memory input at a central location—the average of the three targets (see Figure 3a). Responses to the center target were also consistent with this proposal. Responses to the center target were accurate across delays, suggesting that the memory input helped stabilize working memory at this location.




Figure 6 Mean directional error collapsed across target separation conditions for movements to the left (dot-dashed line, diamond symbol), center (dashed line, square symbol), and right (solid line, triangle symbol) targets over delays. A diagram of the tabletop (inset) with symbols marking the left, center, and right targets and arrows denoting the direction of errors is included for reference.

Nevertheless, data were not consistent with the prediction of the DFT that children would show less bias toward the center target at larger target separations. Although there was a small reduction in directional bias toward the center target at the 40° and 80° separations, there were no statistically significant differences in directional biases across separations. It is possible that we did not sample a large enough spatial range to probe the characteristics of the memory input. In the 80° condition, the targets were separated by 19.2 cm. Three-year-olds' longer-term memory may be broadly distributed such that there was still considerable overlap in longer-term memory at this separation. This interpretation is consistent with data from Spencer et al. (2001). These researchers found that 2-year-olds' responses on the B trial of an A-not-B sandbox task were strongly biased toward an A location, even when the absolute location of "A" varied across a 20.3-cm (8 in) spatial range. Spencer et al. proposed that 2-year-olds' memory of the different A locations blended together, creating a strong memory input around an average A location.

Although it may be possible to test this spatial range proposal by increasing the target separation in the spaceship task, it is important to note that it was difficult to get the 3-year-old children to complete the 80° condition. This condition may be particularly dif-

ficult because children turned their heads to one side or the other when the left and right targets were illuminated, and then reoriented their attention to the center of the table during the delay. Such recurrent reorientation of gaze in the absence of salient visual cues may be quite difficult for 3-year-olds. If this is the case, the data reported here may have been influenced by a selection bias in the 80° condition: only very skilled children may have been able to complete this condition, perhaps leading to slightly more accurate responses. Thus, although it would be interesting to increase the target separation in future experiments, we suspect the challenges present in the 80° condition would be more prevalent at even larger target separations and would preclude strong tests of the model predictions.

In addition to the significant biases in directional error, children tended to overshoot the target at long delays. Given that children had to remember a single target distance, it is somewhat surprising that their errors increased over delays. It is possible that this increase reflects a dependency between directional and distance errors, that is, 3-year-olds may not represent the direction and distance of a target along two independent dimensions. Instead, biases in the representation of information along one dimension may produce biases along the second dimension. This would

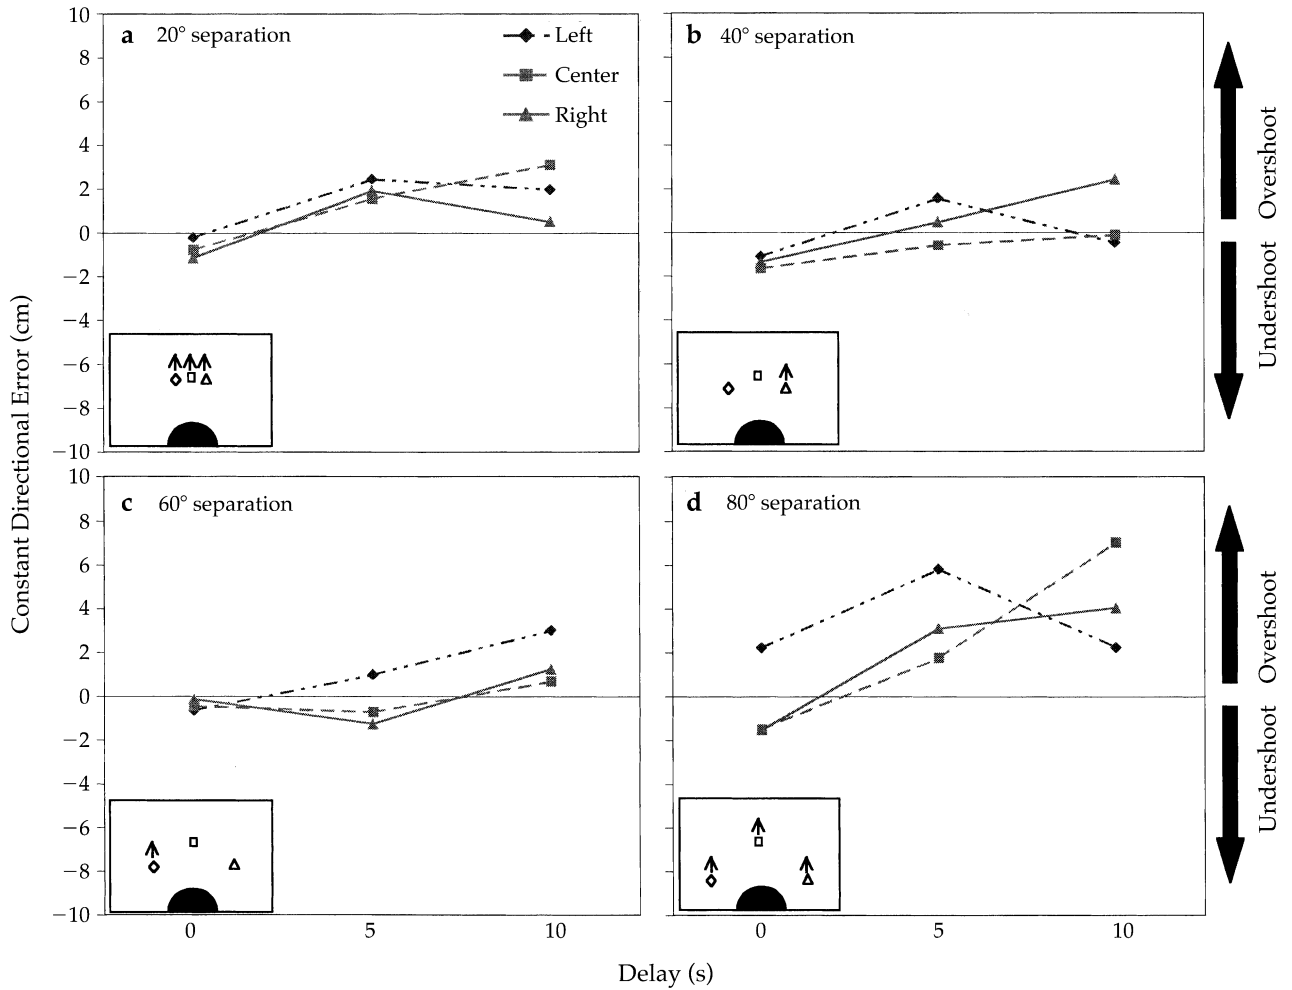


Figure 7 Mean distance error for the (a) 20°, (b) 40°, (c) 60°, and (d) 80° target separation conditions for movements to the left (dot-dashed line, diamond symbol), center (dashed line, square symbol), and right (solid line, triangle symbol) targets over delays. Diagrams of the tabletop (inset) with symbols marking the left, center, and right targets and arrows denoting children's distance errors are included for reference.

contrast with data showing that 9-year-old children and adults represent direction and distance independently (Ghez et al., 1997; Gordon, Ghilardi, & Ghez, 1994; Huttenlocher et al., 1991; Messier & Kalaska, 1997; Rosenbaum, 1980; Sandberg et al., 1996).

An alternative possibility is that children's directional and distance errors reflect a bias toward the center of the table. Huttenlocher et al. (1991) proposed that children and adults use the centers of spatial regions to help them remember locations. According to their CA model, young children treat large, homogeneous spaces such as the tabletop in the spaceship task as a single category with a spatial "prototype" at the center (see Huttenlocher et al., 1994). When a target is presented, children encode the "fine-grained" or metric location of the target

and the location of the prototype. During recall, these two types of information are combined, producing a bias toward the prototype because all target locations within a category are "weighted" with the same prototype. In the present experiment, 3-year-olds' directional biases toward the center target may, in fact, reflect a bias toward the center of the table. The overshoot in children's distance responses may reflect the same bias.

In the present experiment, the three target locations were positioned symmetrically around the midline axis of the table. Consequently, the memory input and the prototypical location in the task space overlapped. In the next experiment, these two types of information were disassociated by rotating the targets to one side of the table.

EXPERIMENT 2

Results from Experiment 1 clearly demonstrated that 3-year-olds' directional responses drift systematically during 0- to 10-s delays. The cause of these directional biases is not clear, however. According to the DFT, these biases reflect attraction toward an averaged memory input at the center target. According to the CA model, however, these biases reflect attraction toward a spatial prototype at the center of the table.

These alternative accounts were tested by rotating targets to the left or right of the midline of the table. In a -40° rotation condition, targets were located at -60° , -40° , and -20° , whereas in a 40° rotation condition, targets were located at 20° , 40° , and 60° (see Figure 4). If the biases toward the center target in Experiment 1 were caused by a pull toward the memory input, as predicted by the DFT, then directional responses to the left and right targets in each condition (e.g., 20° , 60°) should be biased toward the center target (e.g., 40°) over delays. Furthermore, responses to the center target (e.g., 40°) should be accurate over delays. Indeed, according to the DFT, directional errors in both rotation conditions should be identical to children's errors in the 20° separation condition from Experiment 1 (see Figure 8a). In all of these conditions, the three targets were separated by 20° . Consequently, the spatial structure of the memory input should be identical, and working memory should drift comparably over delays.

It is important to emphasize that according to these predictions, children should make different delay-dependent errors in the rotation conditions relative to the same absolute locations used in Experiment 1. For instance, in Experiment 1, there was a significant increase in directional error toward the center target over delays when targets were located at $\pm 20^\circ$. The DFT predicts that directional responses to the $\pm 20^\circ$ locations in Experiment 2 will change directions and drift away from the center of the table. Thus, how memory drifts over delays (or fails to drift) is experience dependent: responses depend on the trial-by-trial construction of the memory input.

In contrast to the DFT, the CA model predicts that children's responses to all three targets in each rotation condition should be biased toward the center of the table—the spatial prototype. Thus, children's biases in the rotation conditions should be identical to responses to the same absolute locations in Experiment 1. Figure 8b shows data from the left targets in the 20° , 40° , and 60° conditions of Experiment 1. Because these same absolute target locations were used in the -40° rotation condition, biases toward the spatial prototype in this condition should be identical to the biases

shown in Figure 8b. Similarly, Figure 8c shows data from the right targets in the 20° , 40° , and 60° conditions of Experiment 1. Again, because these same absolute locations were used in the 40° rotation condition, biases toward the spatial prototype in this condition should be identical to the biases shown in Figure 8c.

Method

Participants

A total of nineteen 36- to 40-month-olds ($M = 38.3$ months, $SD = 1.36$ months) participated in Experiment 2. Five of the participants were female and 14 were male. Data were collected from 3 additional participants; however, these data were not analyzed for the following reasons: 1 participant did not complete both sessions, data from 1 participant had to be excluded due to an experimenter error, and 1 participant did not have at least one trial of useable data per target and delay after data processing. Participants were recruited from a participant database at the University of Iowa, and were tested in two sessions that were scheduled within a week of each other. After each session, the participants received two small toys, and after the second session, they received a certificate of participation. The parents of all participants gave informed consent.

Apparatus, Task, and Procedure

The apparatus for the -40° rotation condition was the same as in Experiment 1. The apparatus for the 40° rotation condition was identical from the children's perspective; however, two modifications were made to the apparatus to enhance its functioning. First, the top of the table in the 40° rotation condition was a rear projection screen rather than a piece of Plexiglas. Second, a video projector below and to the rear of the table projected images onto the table's surface. The projected spaceship images and the image that marked the starting location were the same size as in Experiment 1 and had a similar luminance. All other apparatus details were identical to Experiment 1. The task and procedure were the same as in Experiment 1.

Experimental Design

Participants were randomly assigned to the -40° rotation and 40° rotation conditions. In the -40° rotation condition, targets were located at -60° , -40° , and -20° . In the 40° rotation condition, targets were located at 20° , 40° , and 60° . As in Experiment 1, partic-




Figure 8 (a) Mean directional error predicted by the DFT in both ($\pm 40^\circ$) rotation conditions for movements to the left (dot-dashed line, diamond symbol), center (dashed line, square symbol), and right (solid line, triangle symbol) targets over delays. (b, c) Mean directional error predicted by the CA model for the (b) -40° rotation and (c) 40° rotation conditions. (d, e) Mean directional error for movements to the left, center, and right targets over delays for the (d) 40° rotation and (e) -40° rotation conditions of Experiment 2. Diagrams of the tabletop (insets) with symbols marking the left, center, and right targets and arrows denoting the direction of errors are included for reference.

ipants moved to these locations after three delays: 0 s, 5 s, and 10 s. All other details of the experimental design were identical to Experiment 1 with one exception: the total number of trials was reduced because

several participants in the previous experiment had difficulty completing all of the trials. There were 3 trials to each target at each delay for a total of 27 trials across the two sessions—14 in session one and 13 in

session two. On average, children completed 21.53 trials ($SD = 3.73$) across both sessions.

Method of Analysis

The method of analysis was the same as in Experiment 1.

Results

Directional Error

Children's directional errors are shown in Figure 8 along with the predictions generated from the DFT and CA models. As can be seen in Figure 8d, children's directional errors in the 40° rotation condition were similar to the predictions of the DFT: responses to the left and right targets were biased toward the center target over delays, although biases to the left or "inner" target (20°) were relatively large even at the shortest delay. In contrast, data from the -40° rotation condition (see Figure 8e) generally reflected the predictions of the CA model: responses to all three targets at the 5-s delay were biased clockwise or toward the center of the table.

Median directional errors were analyzed in a three-way repeated-measures ANOVA with rotation as a between-subjects factor and target and delay as within-subjects factors. There was a significant main effect of target, Wilks' $\Lambda = .56$, $F(2, 16) = 6.43$, $p < .01$. Responses to the left target were rotated clockwise or toward the center target ($M = 8.62$), and responses to the right target were rotated counterclockwise or toward the center target ($M = -7.68$). In contrast, responses to the center target were relatively accurate ($M = -1.59$).

The ANOVA also revealed a significant Delay \times Rotation interaction, Wilks' $\Lambda = .45$, $F(2, 16) = 9.76$, $p < .01$. This interaction is shown in Figure 9. For comparison, data from the 20° separation condition from Experiment 1 are also shown. As can be seen in this figure, children's directional error was more strongly biased toward the center of the table over delays in the -40° rotation condition than in the 40° rotation condition. Tests of simple effects revealed a significant increase in directional error over delays in the -40° rotation condition, $F(2, 34) = 6.82$, $p < .01$, but not in the 40° rotation condition, $F(2, 34) = 2.96$, *ns*. As discussed previously, the biases toward the center of the table in the -40° rotation condition are consistent with the CA model. Responses in the 40° rotation condition, however, should have been more strongly biased toward the center of the table according to this model. In contrast, the DFT predicted that data from




Figure 9 Mean directional error over delays in the 40° rotation (solid line, diamond symbol) and -40° rotation (dashed line, square symbol) conditions from Experiment 2 and the 20° separation condition (dot-dashed line, triangle symbol) from Experiment 1. A diagram of the tabletop (inset) with stars marking the targets in the 40° and -40° rotation conditions and arrows denoting the direction of the overall error is included for reference.

both rotation conditions would be comparable with the data from the 20° separation condition of Experiment 1. It is clear from Figure 9 that this was only the case in the 40° rotation condition.

Tests of the DFT predictions. To test the predictions of the DFT directly, a set of planned comparisons was conducted. According to the DFT, there should not be a significant difference between directional errors in the rotation conditions and directional errors in the 20° separation condition from Experiment 1. Thus, two three-way repeated-measures ANOVAs were conducted comparing data from each rotation condition with data from Experiment 1. Experiment (1, 2) was a between-subjects factor in each ANOVA, whereas target and delay were within-subjects factors. Given that a subset of data from Experiment 1 were included in each analysis, only significant experiment effects are reported.

In the first analysis, data from the -40° rotation condition and data from Experiment 1 were compared. There was a significant experiment effect, $F(1, 16) = 8.72$, $p < .05$, and a significant Delay \times Experiment interaction, Wilks' $\Lambda = .55$, $F(2, 15) = 6.14$, $p < .025$. This interaction is shown in Figure 9; children's responses in the -40° rotation condition were biased toward the center of the table over delays and not toward the center target as predicted by the DFT. In contrast, results from the second analysis supported the predictions of the DFT: there were no significant experiment effects

when data from the 40° rotation condition were compared to data from Experiment 1. Across these two analyses, therefore, there was mixed support for the DFT.

Tests of the CA model predictions. To test the predictions of the CA model directly, a set of planned comparisons was conducted in which data from each absolute location used in Experiments 1 and 2 were directly compared. Six 2-way ANOVAs were conducted, one at each of the following locations: $\pm 20^\circ$, $\pm 40^\circ$, and $\pm 60^\circ$. In each ANOVA, Experiment (1, 2) was a between-subjects factor and delay was a within-subjects factor. As above, only significant experiment effects are reported. According to the CA model, there should be no significant experiment effects in these analyses.

There were significant experiment effects at two spatial locations: -40° and 20° . There was a significant Delay \times Experiment interaction at -40° , $F(2, 30) = 3.49$, $p < .05$. Tests of simple effects indicated that directional errors to this location increased significantly over delays in Experiment 1 (see the dashed line in Figure 8b), $F(2, 14) = 4.28$, $p < .05$, but not in Experiment 2 (see the dashed line in Figure 8e), $F(2, 16) = 1.44$, *ns*. The lack of significant delay effects in Experiment 2 is not consistent with the CA model, although the direction of bias at -40° (i.e., clockwise) was generally in the predicted direction. In addition to the effects at -40° , there was an experiment main effect at 20° , $F(1, 17) = 22.41$, $p < .001$. In Experiment 1, children showed significant directional biases toward the center target and toward the center of the table (see the dot-dashed line in Figure 8c), whereas in Experiment 2, directional biases at 20° were toward the center target and away from the center of the table (see dot-dashed line in Figure 8d). Thus, the experiment effects at 20° were not consistent with the CA model.

Distance Errors

Median distance errors were analyzed in a three-way repeated-measures ANOVA with rotation as a between-subjects factor and target and delay as within-subjects factors. There was a significant main effect of delay, Wilks' $\Lambda = .63$, $F(2, 16) = 4.72$, $p < .025$. No other effects reached significance. As in Experiment 1, children undershot the targets slightly at the 0-s delay ($M = -0.53$), and overshot the targets at the other delays (5 s: $M = 2.36$; 10 s: $M = 2.11$). Follow-up *t* tests revealed that there was a significant increase in distance error between the 0-s and 5-s delays, $t(18) = -3.18$, $p < .01$, but not between the 5-s and 10-s delays, $t(18) = -0.27$, *ns*.

Discussion

The goal of Experiment 2 was to determine if the delay-dependent response biases reported in Experiment 1 were due to an attraction toward memory input (DFT) or attraction toward a spatial prototype at the center of the table (CA model). These alternatives were tested by asking children to remember target locations rotated to the left or right of the midline axis of the table. This dissociated the center of the three targets (i.e., the center of the memory input) from the center of the table.

Results from the rotation conditions provided mixed support for each model. As in Experiment 1, children's responses to the $\pm 60^\circ$ locations showed a large spatial drift over delays; however, the direction of these response errors (i.e., inward) was consistent with both models. Thus, data from these targets provided an important replication of the time-dependent effects from Experiment 1, but did not help to falsify one model or the other. Responses to the $\pm 40^\circ$ locations had more potential to tease apart predictions of each model. The CA model predicted that responses to these locations would be biased toward the center of the table, whereas the DFT predicted that responses would be accurate over delays. In general, responses to $\pm 40^\circ$ were biased toward the center of the table, thereby supporting the CA model. Comparisons to data from Experiment 1, however, indicated that children's memory of the -40° location in the present experiment did not change significantly over delays, thereby supporting the DFT model. Thus, results to the $\pm 40^\circ$ locations supported aspects of both models, but were not completely consistent with either model.

Responses to the $\pm 20^\circ$ locations provided the strongest test of each model's predictions, because the predicted effects were in opposite directions: the CA model predicted that responses would be biased toward the center of the table; the DFT predicted that responses would be biased away from the center of the table and toward the center target. Responses to the 20° location were consistent with the DFT—children's responses were biased toward the center target and away from the center of the table. Responses to the -20° location generally supported the CA model, however, responses were biased toward the center of the table at the 5-s delay. There was a bias toward the center target at the 0- and 10-s delays, but these effects were small. Thus, once again, data were not completely consistent with either model.

How can these mixed results be explained? One possibility is that the spatial drift in 3-year-olds' location memory responses is influenced by both factors—memory input and the midline axis of the table. This

might explain why responses to the $\pm 20^\circ$ locations showed variations in the direction of bias over delays: memory for these locations was pulled in opposite directions, both toward and away from the midline of the table. The combination of both effects might explain the mixed results at the $\pm 40^\circ$ locations as well. In this case, responses were biased toward the center of the table, but the memory input may have prevented the memory of the -40° location from drifting significantly over delays.

The proposal that young children's spatial memory is biased toward two factors is consistent with data from Spencer et al. (2001). These researchers found that 2-year-olds' responses were biased toward an A location in a sandbox task. The magnitude of this bias, however, varied depending on the locations of A and B in the sandbox: responses showed less bias toward A when A was in a direction away from the center of the sandbox relative to B, than when A was toward the center of the sandbox relative to B.

If children's spatial memory is influenced by both the geometric properties of the task space—the midline of a table, the center of a sandbox—and a longer-term memory of the possible target locations, then neither model tested here provides a complete account of 3-year-olds' spatial memory abilities. The DFT does a good job of explaining the continuous spatial drift in children's responses, and the memory input effects reported here and in Spencer et al. (2001). There is no clear mechanism in this model to explain the bias toward the center of the table/sandbox, however. In contrast, the CA model explains how spatial memory is affected by the geometric properties of the task space. Nevertheless, this model does not provide a formal account of delay and memory input effects. It is important to note that these limitations of each model do not preclude a complete account of spatial memory. Rather, they point out how each model must be modified in the future. We return to these issues in the General Discussion.

EXPERIMENT 3

The final experiment examined two central issues raised by the results of Experiment 2. First, Experiment 2 provided the first clear evidence of a bias toward memory input with 3-year-old children. This evidence, however, was mixed because children's responses to the -20° location were not consistently biased away from midline and toward the center target. This contrasts with the findings of Spencer et al. (2001) that showed that 2-year-olds' responses were always strongly biased toward memory input at an A location, even when A was in a direction away from the

center of the sandbox. Thus, one goal of Experiment 3 was to provide evidence that 3-year-olds' responses to targets on both sides of midline can be pulled away from midline in the spaceship task with strong memory input.

The second goal of the present experiment was to investigate the construction of the memory input from trial to trial. According to the DFT, response biases should change from trial to trial as children construct a longer term memory of the possible target locations (see Thelen et al., 2001). In contrast, the CA model predicts that response biases should remain relatively similar across trials, because these biases are largely determined by the geometric properties of the task space. In Experiments 1 and 2, the trial order was randomized; thus, it was not possible to tease apart memory input and spatial prototype effects on a trial-by-trial basis. In Experiment 3, this issue was investigated directly by using an A-not-B-type trial order—repeated trials to an A location followed by random trials to A and B.

Three-year-olds moved to two locations used in Experiment 2: $\pm 40^\circ$ and $\pm 60^\circ$. The A location was always the outer target (i.e., $\pm 60^\circ$), whereas the B location was the inner target ($\pm 40^\circ$). Children moved to the A location four times in a row. This was followed by eight randomly ordered trials—four to A and four to B. According to the DFT, biased experience moving repeatedly to A should build up a strong memory input at the A location, which should have three effects: (1) children's responses to A should be more accurate relative to performance to the same locations in Experiment 2, (2) children's responses to B should be pulled toward A (i.e., away from midline) over delays because of strong memory input at the A location, (3) there should be evidence of a bias toward A on the first B trial. Note also that the bias toward A on the B trials might diminish across repeated trials to the B location as the memory input at B becomes stronger.

Although the CA model does not predict trial-by-trial effects, it is useful to consider what this model predicts for the direction of effects on the A and B trials: on both the A and B trials, children's responses should be biased toward the center of the table. Therefore, as with the $\pm 20^\circ$ locations in Experiment 2, the CA and DFT models make opposite predictions in the present experiment.

Method

Participants

A total of eighteen 36- to 41-month-olds ($M = 39.7$ months, $SD = 1.07$ months) participated in Experi-

ment 3. Nine of the participants were female and 9 were male. Data were collected from 3 additional participants; however, these data were not analyzed for the following reasons: 2 participants did not complete the task and data from 1 participant had to be excluded due to experimenter error. All other participant details were identical to Experiment 2.

Apparatus, Task, and Procedure

The apparatus for this experiment was the same as the apparatus used in the 40° rotation condition of Experiment 2. The task and procedure were the same as in Experiments 1 and 2.

Experimental Design

Participants were randomly assigned to one of two conditions. In the -40° rotation condition, the targets were located at -40° and -60°; and in the 40° rotation condition, the targets were located at 40° and 60°. Children moved to the ±60° target—the A target—eight times per session: twice after a 0-s delay, three times after a 5-s delay, and three times after a 10-s delay. Children moved to the ±40° target—the B target—four times per session: twice after a 5-s delay, and twice after a 10-s delay. Only two delays were used on the B trials to keep the total number of trials comparable with the totals used in Experiment 2 (27 trials). The first 4 trials of each session were always to the A location (±60°) at a randomly selected delay. The remaining trials were randomized such that no more than two trials to the same target location occurred in a row. Participants were encouraged to complete all trials during each session; however, during some sessions, the children stopped playing the game early. On average, children completed 22.33 trials ($SD = 1.90$) across both sessions.

Method of Analysis

The method of analysis was the same as in Experiments 1 and 2.

Results

Directional Error

The goal of Experiment 3 was to determine if 3-year-olds' memory for locations is strongly affected by the memory input. Thus, rather than examining whether children's directional responses drifted in a clockwise or counterclockwise manner over delays, this experiment examined whether directional re-

sponses drifted toward A over delays. Errors toward A were positive; errors away from A were negative. Note that, because the A location (±60°) was always away from midline with respect to the B location (±40°), this measure also captured errors away from (positive) or toward (negative) the center of the table. Note also that there were no trials to B at the 0-s delay. Consequently, errors at 0 s were not analyzed.

According to the DFT, biasing how often children moved to A should decrease directional errors to the A location. As can be seen in Figures 10c and 10d, directional errors to A (the outer target) in both rotation conditions were near zero at 5 s and 10 s. This stands in sharp contrast to the large biases toward the center of the table when children moved to the same targets at the same delays in Experiment 2 (see Figures 10a and b). The DFT also predicted that 3-year-olds' responses to B (the inner target) should be pulled toward A. This was indeed the case. Errors to the B location were positive, or toward A, in both rotation conditions (see Figures 10c, d). These biases were generally in the opposite direction (i.e., away from the center of the table) relative to errors to the same locations (±40°) in Experiment 2 (see Figures 10a and b).

Children's median directional errors were analyzed in a three-way repeated-measures ANOVA with rotation (-40°, 40°) as a between-subjects factor and target (A, B) and delay (5 s, 10 s) as within-subjects factors. There was a main effect of target, Wilks' $\Lambda = .24$, $F(1, 16) = 51.92$, $p < .001$. Follow-up t tests indicated that responses to the B locations differed significantly from 0 error, $t(17) = 3.99$, $p = .001$, whereas responses to A did not, $t(17) = -.98$, ns . These results are consistent with the predictions of the DFT.

Comparisons to data from Experiment 2. In a second directional analysis, children's median directional errors from the present experiment were compared to directional errors to the ±40° and ±60° targets from Experiment 2. Specifically, a repeated-measures ANOVA was conducted with Experiment (2, 3) and rotation (-40°, 40°) as between-subjects factors and target (inner, outer) and delay (5 s, 10 s) as within-subjects factors. Given that a subset of data from Experiment 2 were included, only significant experiment effects are reported.

There was a significant experiment main effect, $F(1, 33) = 14.20$, $p = .001$. Children's directional errors in Experiment 2 were generally negative or toward the midline of the table (see Figure 10). In contrast, directional errors in Experiment 3 were generally positive or away from midline (i.e., toward A). Thus, the large difference in directional error across experiments shown in Figure 10 was a statistically reliable effect.




Figure 10 Mean directional error for movements to the outer ($\pm 60^\circ$; dashed line, diamond symbol) and inner ($\pm 40^\circ$; solid line, square symbol) targets at the 5- and 10-s delays in the (a) -40° and (b) 40° rotation conditions of Experiment 2 and the (c) -40° and (d) 40° rotation conditions of Experiment 3. Positive scores indicate errors toward the outer target (the A location); negative errors indicate errors away from the outer target. Diagrams of the tabletop (insets) with symbols marking the inner and outer targets and arrows denoting the direction of errors are included for reference.

Trial-by-trial analyses. In a final set of directional analyses, the trial-by-trial predictions of the DFT (for the number of children contributing data to each trial-by-trial analysis below, see the caption of Figure 11) were examined. First, children's errors on the first four A trials in each session were examined. As can be seen in Figure 11a, children's responses were generally quite accurate on the first four A Trials, with directional error near 0 by Trial A4. There was, however, a small bias toward midline, particularly during Session 2. Data from each session were analyzed in a repeated-measures ANOVA with rotation (-40° , 40°) as a between-subjects factor and trial (A1, A2, A3, A4) as a within-subjects factor. There were no significant effects for either session. Furthermore, *t* tests indicated that the directional errors in Session 1, $t(13) = -.08$, *ns*, and Session 2, $t(17) = -1.04$, *ns*, did not dif-

fer significantly from 0 error. Thus, as predicted by the DFT, children's responses were accurate following repeated trials to an A location.

Next, errors on the last successive A trial were compared to errors on the first B trial. (Note that for some children, the last successive A Trial occurred on Trial 5 [Session 1: $n = 5$; Session 2: $n = 4$], trial 6 [Session 1: $N = 0$; Session 2: $n = 3$], or trial 7 [Session 1: $N = 2$; Session 2: $n = 1$]). As can be seen in Figure 11b, errors on the last A trial were generally small, whereas errors on the first B trial were positive—toward the A location and away from midline. These data for each session were analyzed in a repeated-measures ANOVA with trial (last A, first B) as a within-subjects factor and rotation (-40° , 40°) as a between-subjects factor. There was a significant effect of trial for Session 1, Wilks' $\Lambda = .66$, $F(1, 14) = 7.19$,




Figure 11 Mean directional error for (a) the first four A trials, (b) the last successive A trial and first B trial, and (c) the first three B trials for Session 1 (dashed line, diamond symbol) and Session 2 (solid line, square symbol). Positive errors indicate errors toward the A location (away from midline); negative errors indicate errors away from the A location (toward midline). Note that data were not available for all participants for all trials in a session. Thus, the number of participants contributing data to each set of means varied: for the first four A trials, $n = 14$ for Session 1 and $n = 18$ for Session 2; for the last A and first B trials, $n = 17$ for Session 1 and $n = 18$ for Session 2; for the first three B trials, $n = 14$ for Session 1 and $n = 17$ for Session 2.

$p < .025$. There was also a significant effect of trial for Session 2, Wilks' $\Lambda = .65$, $F(1, 16) = 8.75$, $p < .01$. Thus, as predicted by the DFT, 3-year-olds showed a significant bias toward the A location on the first B trial. Follow-up t tests confirmed that the directional errors on the first B trial differed significantly from 0 error for Session 1, $t(16) = 2.63$, $p < .025$, and Session 2, $t(17) = 2.52$, $p < .025$.

The third set of analyses examined the first three trials to the B location in each session. According to the DFT, the bias toward the A location might weaken as the number of trials to the B location increases. As can be seen in Figure 11c, the bias toward A generally became smaller across the first, second, and third B trials in each session. A repeated-measures ANOVA for each session with trial (B1, B2, B3) as a within-subjects factor and rotation (-40° , 40°) as a between-subjects factor revealed no significant effects, however. The lack of significant trial effects may have been due to the randomization of A and B trials after the fourth A trial—the A trials that intervened between B1, B2, and B3 may have “refreshed” the bias toward A (for related effects with infants, see Smith et al., 1999; Thelen et al., 2001). Despite the absence of a trial effect, t tests confirmed that the bias toward A on the B trials differed significantly from 0 error during

session 1, $t(13) = 3.73$, $p < .01$, and Session 2, $t(16) = 1.75$, $p < .05$.

Distance Errors

As in the previous two experiments, children overshoot the targets at the 5- and 10-s delays. The average median distance error across delays was 2.05 cm ($SE = .64$ cm). The magnitude of this overshoot was comparable with the average median distance error to the $\pm 40^\circ$ and $\pm 60^\circ$ locations from Experiment 2 ($M = 2.25$ cm; $SE = .91$ cm). Children's median distance errors were analyzed in a three-way repeated-measures ANOVA with rotation (-40° , 40°) as a between-subjects factor and target (A, B) and delay (5 s, 10 s) as within-subjects factors. There was a significant Delay \times Rotation interaction, Wilks' $\Lambda = .76$, $F(1, 16) = 4.94$, $p < .05$. Examination of the data suggested that this effect was largely driven by the median response from one participant in the 40° rotation condition who overshoot the target by 17.13 cm at the 5-s delay (for comparison, the next largest distance error was 10.17 cm). When this participant's data were removed from the analysis, there were no longer any significant distance error effects. Thus, with one exception, children's distance errors to each target at each delay were comparable across conditions.

Discussion

Results from Experiment 3 demonstrate that 3-year-olds' location memory responses are strongly affected by memory input. As predicted by the DFT, children were able to accurately remember the A location over delays after repeatedly moving to this location. In contrast, children's responses on the B trials were biased toward A. Importantly, this effect was most dramatic on the first B trial in each session. Thus, moving repeatedly to the same location created strong memory input at A. This, in turn, caused memory to drift toward A on the first B trial. These data contrast with the predictions of the CA model and with children's directional responses toward the center of the table in the -40° rotation condition of Experiment 2. Thus, 3-year-olds' responses on both sides of the spaceship table can be biased away from midline with strong enough memory input. These data are consistent with data reported by Spencer et al. (2001) with 2-year-old children. As in the present experiment, Spencer et al. reported biases toward an A location, even when A was in a direction away from the center of the sandbox relative to B.

GENERAL DISCUSSION

The goal of the current set of experiments was to test three predictions of the DFT about the nature of 3-year-olds' delay-dependent location memory errors. The first prediction was derived from a novel characteristic of the DFT previously reported in this article—activation within self-sustaining dynamic fields can drift systematically over delays in the absence of salient location cues. Based on this characteristic of the DFT, we predicted that 3-year-olds' responses would show a continuous spatial drift over short-term delays in a task with a homogeneous task space. This was indeed the case. Three-year-olds' directional responses to the left and right targets in Experiment 1 and to the outer targets in Experiment 2 drifted continuously during 0- to 10-s delays. This drift increased significantly from 0 s to 5 s and from 5 to 10 s. Several previous studies have shown that young children's location memory responses are biased in particular directions after a delay (Huttenlocher et al., 1994; Spencer et al., 2001); however, this is the first study with young children to demonstrate that spatial drift increases systematically from second to second. We contend that this systematic spatial drift is an important signature of the processes that underlie the maintenance of location information in working memory—processes formalized by the DFT. It is important to emphasize that continuous spatial drift is not a re-

quirement of the working memory system. Location memory could decay in many different ways. For instance, children's memory for a target location could simply become less accurate over delays. In this case, responses would be more variable after long delays, but the mean of the responses would be centered near the true location.

The second prediction tested was that 3-year-olds' memory responses would be biased toward their longer term memory of previously moved-to locations; that is, toward the memory input. This prediction was also supported. Some evidence of biases toward the memory input was seen in Experiment 2. Experiment 3, however, provided the clearest demonstration of the predicted effect: when we biased how often children moved to an outer target, responses to this "A" location were accurate over delays whereas responses to an inner "B" location were biased toward A. Importantly, there was a significant bias toward A on the first B trial in both experimental sessions. These results are consistent with 2-year-olds' biases in an A-not-B sandbox task (Spencer et al., 2001). Taken together, these studies demonstrate that 2- and 3-year-olds' short-term memory for locations is sensitive to input from longer term memory. The existence of memory input effects at 3 years supports the proposal by Thelen, Smith, and colleagues (Smith et al., 1999; Thelen et al., 2001) that the processes formalized by the DFT are general processes not unique to infancy.

The third prediction tested was that the magnitude and direction of delay-dependent spatial drift would depend on the strength and spatial structure of the memory input. In Experiment 1, we attempted to manipulate the strength of memory input effects by changing the separation between adjacent targets. In contrast to the predictions of the DFT, a significant change in the bias toward the center target across separation conditions was not seen. When the three targets in Experiment 2 were rotated, however, a significant change in directional errors relative to associated changes in the spatial structure of the memory input was seen. The significant delay-dependent drift toward midline at the -40° location in Experiment 1 was no longer significant, and the bias toward midline at 20° reversed direction and drifted away from midline. Finally, as discussed above, manipulations of the strength of the memory input at A and B locations in Experiment 3 had the predicted effects. In summary, changes in the memory input in two of three experiments altered how children's memory drifted over delays in a manner consistent with the DFT.

Taken together, data from the present study pro-

vide support for the proposal by Thelen et al. (2001) that the DFT—originally proposed to explain infants' errors in the classic Piagetian A-not-B task—is a general processes account. Support was found for all three predictions derived from the DFT using a novel task (the spaceship task) with 3-year-old children. Thus, the principles of the DFT are not specific to a particular period in development, nor are they specific to behavior in a particular task.

Although the results of this study demonstrate that the DFT can link the performance of 12- and 36-month-olds, it is also clear from the results that the DFT does not provide a complete account of younger children's spatial memory abilities. Of particular note, 3-year-olds showed biases toward the midline of the task space—biases consistent with the CA model proposed by Huttenlocher and colleagues (1991). In the next two sections, we evaluate how effectively the DFT brings together spatial memory processes in infancy and early childhood.

Implications of 3-Year-Olds' Location Memory Biases for Errors in Infancy

Memory Input Effects

The present study is the first to show systematic biases toward previously moved-to locations with 3-year-old children. These results contradict proposals by Diamond and Doar (1989) and Munakata (Munakata, 1998; Munakata et al., 1997) that longer term memory biases weaken in early development. The data also demonstrate that the structure of the task space plays a pivotal role in the presence or absence of memory input effects. Consequently, it is important for accounts of spatial memory to specify how task cues are used when children are asked to remember locations. We contend that the DFT is the only account of the A-not-B error that does this effectively enough to make a priori predictions about children's behavior in tasks with marked hiding locations and tasks with a homogeneous space.

Although 3-year-olds show biases toward memory input, we agree with Diamond and Munakata that there may be changes in the strength and/or spatial range of such biases in early development. Indeed, results from Experiment 2 that showed weaker memory input effects at -20° may indicate that by 3 years, children are better able to selectively suppress some location cues (for related ideas, see Newcombe & Huttenlocher, 2000). It is important to emphasize, however, that statements about changes in the strength of longer term memory biases must be grounded in the details of the spatial memory task.

Delay-Dependent Biases

As with memory input effects, delay-dependent results from the present study may have important implications for accounts of developmental changes in spatial memory around 12 months. Data from several studies show that infants can tolerate longer delays in the A-not-B task between 8 and 16 months (e.g. Diamond & Doar, 1989). This suggests that infants' ability to sustain information in memory improves during this time. Through the lens of the DFT, however, this statement is overly simplistic. According to this model, the delay during which children can accurately remember a location depends on the strength and spatial characteristics of the memory, specific, and task inputs; and developmental changes in children's ability to sustain activation in working memory. Results from the present study highlight the interplay among these factors. We were able to systematically shift 3-year-olds' ability to accurately remember the same absolute locations across conditions, even though the perceptual cues in the task space and the delays were the same. Thus, memory decay over short-term delays is experience dependent—even subtle changes in children's trial-to-trial experiences can affect their ability to tolerate delays.

The task- and experience-dependent nature of location memory may explain why there is so much variability in the delays that infants can tolerate across tasks. For instance, Brody (1981) noted that estimates of the delays that 8- to 16-month-old infants can tolerate in recall tasks vary between 3 s and 15 s (Millar & Schaffer, 1972; Ramey & Ourth, 1971; Watson, 1967). She suggested that these variations are a function of the type of task an infant is required to perform. The DFT takes this proposal one step farther by formalizing how the task matters. Consequently, the DFT can be used to make predictions about the types of developmental profiles one might see both within a task over development and across different tasks.

Implications for Models of the A-not-B Error

In the present study, the DFT was used to generate hypotheses about the performance of 3-year-olds in a task analogous to tasks used by Huttenlocher and colleagues (e.g., Huttenlocher et al., 1994; Sandberg et al., 1996). Although no other models of the A-not-B error have been applied to such tasks, it is useful to ask whether other models of the A-not-B error might be able to account for the present study's results.

To account for data from this present study, A-not-B models would have to capture the performance of children in tasks with a homogeneous task space. Al-

though this seems like a straightforward extension of some existing models, this may not be the case. Consider, for instance, one of the most thoroughly specified models of the error—Munakata's Parallel Distributed Processing (PDP) model (Munakata, 1998; Munakata et al., 1997). In its current form, this model cannot produce changes in the magnitude of responses in between three locations because there are only three location nodes in the network (consequently, the network can only make a response at one of three discrete locations). Thus, one would need to add more location nodes, a relatively simple operation. The addition of location nodes, would not necessarily lead to the type of systematic spatial drift observed in this study, however. For this, activation in the network must travel "laterally" across nodes during delays. It is not apparent that the Munakata network will show this type of coherent, traveling activation. Indeed, we suspect that the model would require new constraints on how the nodes are connected to one another. Therefore, although homogeneous tasks seem very similar to tasks with marked locations, it may not be easy for the Munakata network to formally reproduce the types of behaviors reported in this article.

Two other accounts might be able to capture the present study's data because they explain infants' responses in between A and B in a multilid A-not-B task—Diamond et al.'s (1994) inhibitory account and Marcovitch and Zelazo's (1999) response versus representation account. The primary limitation of these accounts—in contrast to the Munakata model—is that they are not formalized process models. That is, these accounts fail to specify the processes that underlie performance in A-not-B-type tasks with enough detail to allow for a priori generalizations to new situations such as the spaceship task used in the present study.

According to Marcovitch and Zelazo's (1999) account, for example, A-not-B errors disappear around 12 months when infants begin to rely more on a representational system that encodes space allocentrically rather than a response system that encodes space relative to action. (Note that our focus is on the Marcovitch and Zelazo account because the Diamond et al., 1994 account has not been mathematically formalized [although see Marcovitch & Zelazo, 1999]). The representation and response systems are formally represented by two functions—one that specifies the resultant contribution of the representational system to behavior, and one that specifies the resultant contribution of the response system to behavior. Children's performance in a given task can be captured by changing the parameters of each function such that the sum of these functions fits the relevant data. Concretely,

the parameters can be set such that the response system dominates at 8 to 10 months when infants perseverate in the A-not-B task, and the representational system dominates at 12 months when infants stop perseverating and correctly reach to B. Therein lies the fundamental limitation of this model—there are no constraints on how the parameters are set. Consequently, it is not clear what the model would predict about 3-year-olds' performance in the spaceship task used in this study. One might contend that if 12-month-olds can rely on a representational system in the A-not-B task with the hiding locations marked, then surely 3-year-olds should be equally reliant on this system when the hiding locations are not marked in the spaceship task. Alternatively, one might claim that the homogeneous nature of the task space requires that children rely more on the response system, in which case 3-year-olds might err in between A and B. Although there may be kernels of truth in these statements, the point is this: the model does not help to decide which of these statements is true a priori. Thus, the model can only be fit to results post hoc, leaving it with little predictive value (although for an interesting application of this model to the perseverative responses of 2-year-olds, see Zelazo, Reznick, & Spinazzola, 1998).

Beyond Infancy: Evaluating the DFT and CA Models

As stated previously, the DFT and CA models can account for different aspects of the present results. The DFT does a good job in capturing the continuous spatial drift noted in this study, as well as many of the memory input effects. In contrast, the CA model can capture children's biases toward the midline of the task space. Although a formal comparison of these models is beyond the scope of the current article, in this section the relative merits of these models are outlined and possible future elaboration of these accounts is discussed.

The Dynamic Field Model

The major weakness of the DFT is that it fails to explain why children showed biases toward the midline of the task space, even when the memory input was shifted to the left or right of midline (see Experiment 2). A simple post hoc account of the midline bias is possible, however. Data from several studies demonstrate that people do not perceive "empty" geometrical spaces as empty per se. Instead, people perceive these figures as being subdivided by symmetry axes such as the vertical, horizontal, and diagonal axes

that cut through the center of a square or rectangle (Beh, Wenderoth, & Purcell, 1971; Blakemore, Carpenter, & Georgeson, 1970; Bouma & Andriessen, 1970; Carpenter & Blakemore, 1973; Hartley, 1982; Wenderoth, 1997; Wenderoth, Johnstone, & van der Zwan, 1989). If we assume that 3-year-olds perceived the midline symmetry axis of the spaceship table, a midline “task input” can be added to the dynamic field model to produce the combination of memory input and midline biases we observed in this study.

Figure 12 shows a set of simulations of the DFT that demonstrate that this model can qualitatively capture delay-dependent biases toward a memory input and biases toward midline with the addition of a midline task input. Each simulation figure is structured in the same manner as the simulations shown in Figures 1 and 2. The left column in each figure shows the task input, specific input, and memory input, with the planning field to the right. The left panels of Figure 12 show simulations of the DFT when targets are presented at -60° (see the specific input in Figure 12a), -40° (see the specific input in Figure 12b), and -20° (see the specific input in Figure 12c) without task input but with memory input. Specifically, there is no task input in Figures 12a, 12b, and 12c, and the memory input is distributed across the three possible hiding locations (-60° , -40° , -20° ; the memory input in these simulations reflects activation built up in longer term memory following repeated trials to the same three locations). The right panels of Figure 12 show simulations of the model when targets are presented at the same three locations (-60° , -40° , -20°) in the presence of both a midline task input and memory input. Note that the memory inputs in Figures 12d, 12e, and 12f are identical to the memory inputs in the simulations in the left panels.

As can be seen in Figures 12a through 12c, the dynamic field model shows biases toward the center target (-40°) when memory input is present, but task input is not. When the target is presented at -60° (Figure 12a), the model builds a peak of activation centered at the target location and, during the delay (4–10 s), activation drifts toward -40° (i.e., inward). Similarly, when the target is presented at -20° (Figure 12c), the peak of activation at the target location drifts toward -40° (i.e., outward) during the delay. In contrast, when the target is presented at -40° , activation remains centered at the target location during the delay. These effects are modulated when a midline task input at 0° is added. With the target at -60° (Figure 12d), activation still shows a strong inward bias. When the target is presented at -40° (Figure 12e) and -20° (Figure 12f), however, there is a slight bias toward midline (i.e., toward 0°). Thus, the presence of a

midline task input effectively attenuates memory input effects when the model is asked to “remember” the center and right targets. This captures the qualitative pattern of results from the -40° rotation condition of Experiment 2. In this condition, there was strong bias toward midline at -60° , a small bias toward midline at -40° that did not increase significantly over delays, and a slight bias toward midline after the 5-s delay at -20° .

The simulations in Figure 12 demonstrate that a modified version of the DFT can capture midline biases. Therefore, the primary weakness of the DFT in this study can, in principle, be overcome. It is worth noting, however, that it is not clear why children in Experiment 2 showed stronger midline biases to the -20° target in the -40° rotation condition than to the 20° target in the 40° rotation condition. Further investigation is needed to clarify this result. It is also important to reiterate that the midline task input was a post hoc addition to the model. Thus, this aspect of the model will require thorough investigation in the future. With this larger goal in mind, it is worth noting that the new version of the model makes some specific predictions. In particular, it should be possible to modify children’s memory responses by changing the strength of the midline task input. For example, if we were to add perceptual cues to the task space—perhaps by projecting a collection of “landmark” dots along midline—children should show stronger midline biases.

The CA Model

The CA model was proposed to explain children’s use of spatial category information in simple recall tasks. Huttenlocher and colleagues (1994) proposed that young children treat homogeneous task spaces as a single category with a prototype at the center. At recall, children combine their fine-grained memory of the target location with categorical information (the location of the prototype). This produces biases toward the spatial prototype; that is, toward the center of the task space. Thus, this model can account for the midline biases found in this study. Indeed, this model provides the dominant account of such biases in the literature.

Can the CA model capture the delay-dependent and memory input effects that were reported? Given that fine-grained and prototypical information are represented along continuous spatial dimensions, the CA model can produce the type of “in between” responses noted in this study. To produce continuous spatial drift, however, one must assume that fine-grained information becomes continuously less certain over delays. As this occurs, children should “weight” prototypical information more heavily, leading to a

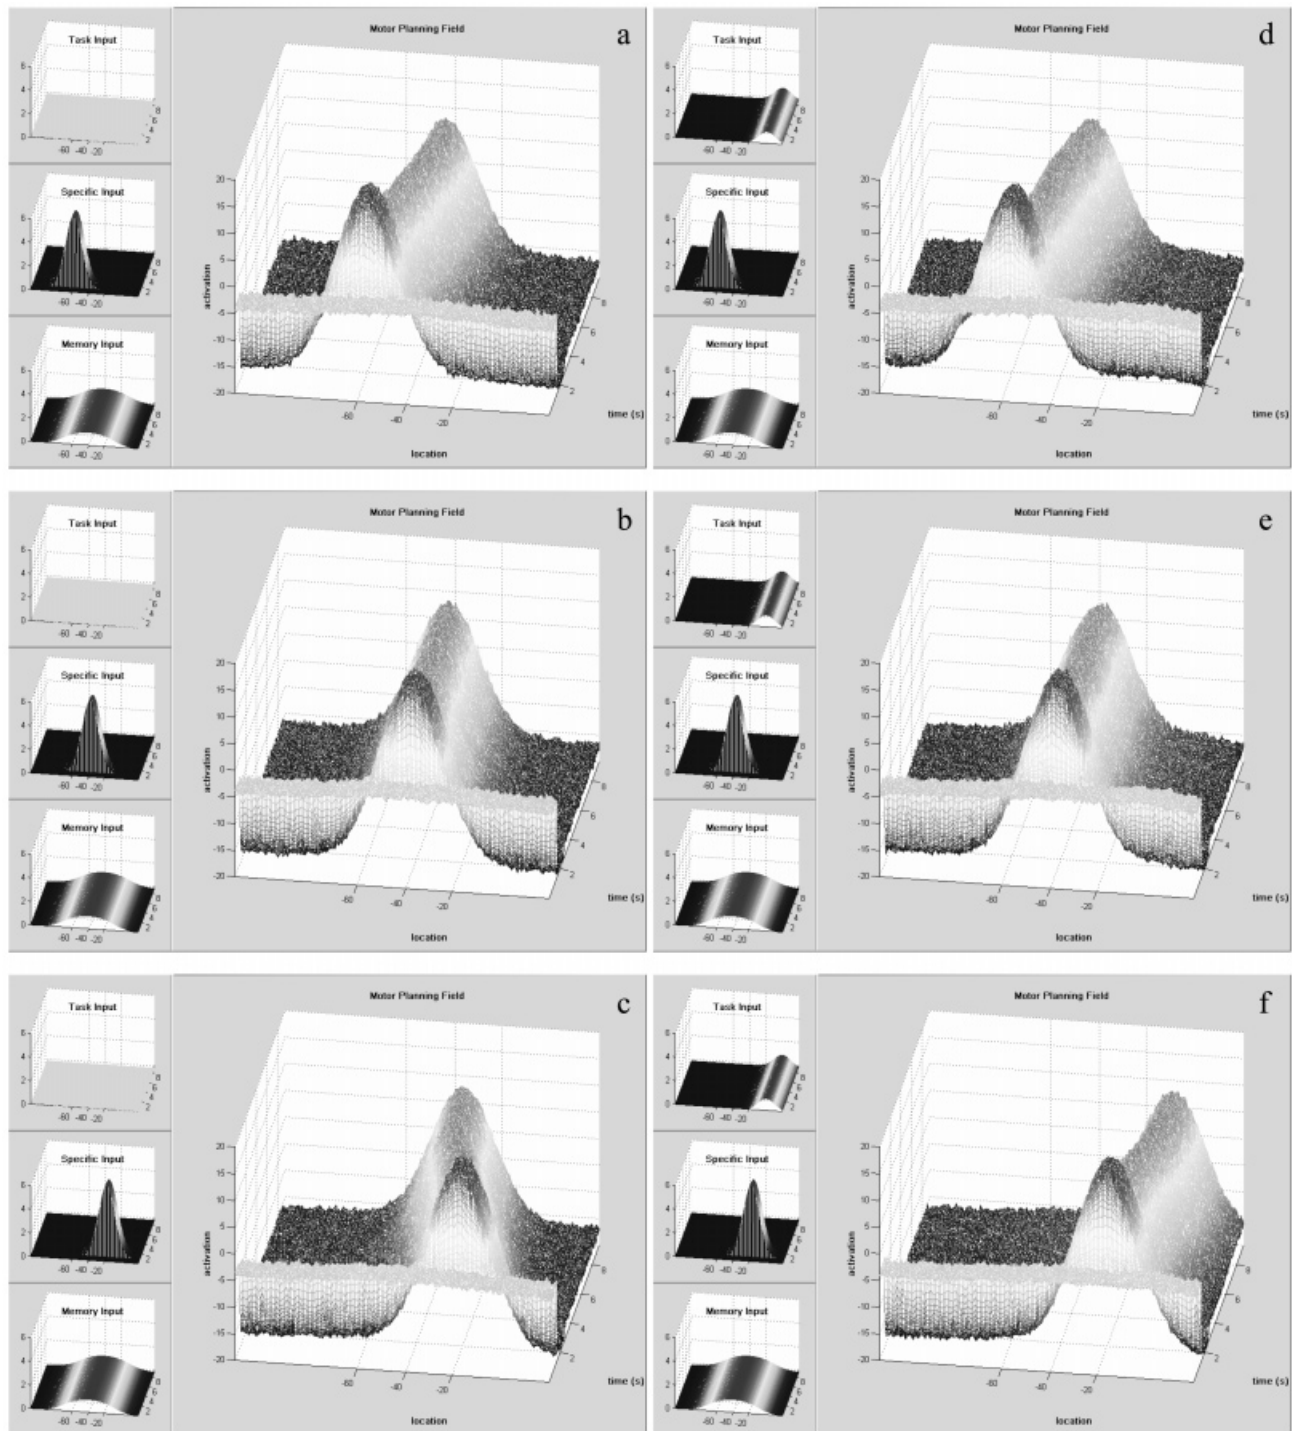


Figure 12 Simulation of the dynamic field theory with targets presented at (a, d) -60° , (b, e) -40° , and (c, f) -20° without a midline task input (left panels) and with a midline task input (right panels). The parameters for all of the simulations were identical except for the presence or absence of task input. See text for further details.

continuous increase in response bias over delays. In principle, then, the CA model can produce delay-dependent spatial drift.

Nevertheless, we contend that the CA model needs greater specification to fully account for the present study's data. In its current form, there are no detailed

constraints on how fine-grained and categorical information are weighted. Thus, to simulate delay-dependent effects, one can modify the weighting function or the represented information “by hand” to produce a more biased outcome. This lack of constraint is especially problematic given that responses to the same locations in the present study showed spatial drift in some conditions, but not in others. It is not clear how these results could be accounted for by uniform decay of fine-grained information over delays.

The memory input effects in the present study present a more serious challenge to the CA model: in its current form, this model has no mechanism for such effects. Huttenlocher et al. (1991) proposed that prototypical information might be influenced by the distribution of target locations within a spatial category, but this proposal has not been formalized (however, see Huttenlocher, Hedges, & Vevea, 2000). Instead, these researchers have placed special emphasis on the centers of geometric regions. We suspect that a rather substantial change is needed if the CA model is to capture memory input effects.

Conclusion

In the present study, three specific predictions derived from the DFT of the A-not-B error were tested. Support was found for all three predictions, demonstrating that the general processes captured by the DFT provide a bridge between phenomena in infancy and phenomena in later development. Results, however, also showed that biases toward the midline of the task space originally reported by Huttenlocher et al. (1994) affect children’s responses. Thus, aspects of the present study’s data were effectively captured by the DFT, whereas other aspects of the data were consistent with the CA model. We hope that the existence of two competing models of spatial memory development will encourage new research in the spatial memory domain that sheds light on truly classic developmental issues.

ACKNOWLEDGMENTS

The authors would like to thank the parents and children who participated in this study. Gabriel Casciato, Michael Chen, Alycia Hund, and Sandy Seeman helped with data collection and analysis. Lloyd Frei, Keith Miller, Dale Parker, and MechDyne, Inc. provided valuable assistance constructing and programming the experimental apparatus. Alycia Hund, Larissa Samuelson, and Esther Thelen provided helpful comments on an earlier version of this manuscript.

ADDRESSES AND AFFILIATIONS

Corresponding author: Anne R. Schutte, Department of Psychology, University of Iowa, E11 Seashore Hall, Iowa City, IA 52242; e-mail: anne-schutte@uiowa.edu. John P. Spencer is also at the University of Iowa.

REFERENCES

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*, 77–87.
- Amari, S., & Arbib, M. A. (1977). Competition and cooperation in neural nets. In J. Metzler (Ed.), *Systems neuroscience* (pp. 119–165). New York: Academic.
- Beh, H., Wenderoth, P., & Purcell, A. (1971). The angular function of a rod-and-frame illusion. *Perception and Psychophysics*, *9*, 353–355.
- Blakemore, C., Carpenter, R., & Georgeson, M. (1970). Lateral inhibition between orientation detectors in the human visual system. *Nature*, *228*, 37–39.
- Bouma, H., & Andriessen, J. J. (1970). Induced changes in the perceived orientation of line segments. *Vision Research*, *10*, 333–349.
- Brody, L. R. (1981). Visual short-term cued recall memory in infancy. *Child Development*, *52*, 242–250.
- Carpenter, R. H., & Blakemore, C. (1973). Interactions between orientations in human vision. *Experimental Brain Research*, *18*, 287–303.
- Diamond, A. (1990a). Development and neural bases of AB and DR. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 267–317). New York: National Academy of Sciences.
- Diamond, A. (1990b). Developmental time course in human infants and infant monkeys, and the neural bases of inhibitory control in reaching. In A. Diamond (Ed.), *The development and neural bases of higher cognitive functions* (pp. 637–676). New York: National Academy of Sciences.
- Diamond, A., Cruttenden, L., & Neiderman, D. (1994). AB with multiple wells: 1. Why are multiple wells sometimes easier than two wells? 2. Memory or memory + inhibition? *Developmental Psychology*, *30*, 192–205.
- Diamond, A., & Doar, B. (1989). The performance of human infants on a measure of frontal cortex function, the delayed response task. *Developmental Psychobiology*, *22*, 271–294.
- Ghez, C., Favilla, M., Ghilardi, M. F., Gordon, J., Bermejo, R., & Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research*, *115*, 217–233.
- Gordon, J., Ghilardi, M. F., & Ghez, C. (1994). Accuracy of planar reaching movements. I. Independence of direction and extent variability. *Experimental Brain Research*, *99*, 97–111.
- Hartelman, P., van der Maas, H. L. J., & Molenaar, P. C. M. (1998). Detecting and modelling developmental transitions. *British Journal of Developmental Psychology*, *16*, 97–122.

- Hartley, A. A. (1982). The role of axes of symmetry in orientation illusions. *Perception and Psychophysics*, *31*, 367–375.
- Hund, A. M., & Spencer, J. P. (2002). *Developmental changes in the relative weighting of geometric and experience-dependent location cues*. Manuscript submitted for publication.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376.
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgement? *Journal of Experimental Psychology: General*, *129*, 220–241.
- Huttenlocher, J., Newcombe, N., & Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, *27*, 115–147.
- Marcovitch, S., & Zelazo, P. D. (1999). The A-not-B error: Results from a logistic meta-analysis. *Child Development*, *70*, 1297–1313.
- Messier, J., & Kalaska, J. F. (1997). Differential effect of task conditions on errors of direction and extent of reaching movements. *Experimental Brain Research*, *115*, 469–478.
- Millar, S. W., & Schaffer, H. (1972). Visual manipulative strategies in infant operant conditioning with spatially displaced feedback. *British Journal of Psychology*, *64*, 545–552.
- 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., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, *104*, 686–719.
- Newcombe, N., & Huttenlocher, J. (2000). *Making space: The development of spatial representation and reasoning*. Cambridge, MA: MIT Press.
- Newcombe, N., Huttenlocher, J., Drumme, A., & Wiley, J. (1998). The development of spatial location coding: Place learning and dead reckoning in the second and third years. *Cognitive Development*, *13*, 185–200.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic.
- Piaget, J., & Inhelder, B. (1956). *The child's conception of space*. London: Routledge and Kegan Paul.
- Ramey, C., & Ourth, R. (1971). Delayed reinforcement and vocalization rates of infants. *Child Development*, *42*, 291–297.
- Rosenbaum, D. A. (1980). Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology*, *109*, 444–474.
- Sandberg, E. H., Huttenlocher, J., & Newcombe, N. (1996). The development of hierarchical representation of two-dimensional space. *Child Development*, *67*, 721–739.
- Schöner, G., Kopecz, K., & Erlhagen, W. (1997). The dynamic neural field theory of motor programming: Arm and eye movements. In P. G. Morasso & V. Sanguinetti (Eds.), *Self-organization, computational maps and motor control* (Vol. 119, pp. 271–310). Amsterdam, The Netherlands: Elsevier-North Holland.
- Smith, L. B., Thelen, E., Titzer, R., & McLin, D. (1999). Knowing in the context of acting: the task dynamics of the A-not-B error. *Psychological Review*, *106*, 235–260.
- Spencer, J. P., & Hund, A. M. (in press). Prototypes and particulars: Geometric and experience-dependent spatial categories. *Journal of Experimental Psychology: General*.
- Spencer, J. P., & Hund, A. M. (2002). *Are we there yet? Spatial memory models fail to capture both geometric and experience-dependent biases in children's memory for locations*. Manuscript submitted for publication.
- Spencer, J. P., Smith, L. B., & Thelen, E. (2001). Tests of a dynamic systems account of the A-not-B error: The influence of prior experience on the spatial memory abilities of 2-year-olds. *Child Development*, *72*, 1327–1346.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A dynamic field theory of infant perseverative reaching errors. *Behavioral and Brain Sciences*, *24*, 1–86.
- van der Maas, H. L. J., & Molenaar, P. C. M. (1992). Stage-wise cognitive development: An application of catastrophe theory. *Psychological Review*, *99*, 395–417.
- Van Geert, P. (1998). A dynamic systems model of basic developmental mechanisms: Piaget, Vygotsky, and beyond. *Psychological Review*, *105*, 634–677.
- Watson, J. S. (1967). Memory and contingency analysis in infant learning. *Merrill-Palmer Quarterly*, *13*, 55–76.
- Wenderoth, P. (1997). The role of implicit axes of bilateral symmetry in orientation processing. *Australian Journal of Psychology*, *49*, 176–181.
- Wenderoth, P., Johnstone, S., & van der Zwan, J. (1989). Two-dimensional tilt illusions induced by orthogonal plaid patterns: Effects of plaid motion, orientation, spatial separation, and spatial frequency. *Perception*, *18*, 25–38.
- Zelazo, P. D., Reznick, J. S., & Spinazzola, J. (1998). Representational flexibility and response control in a multi-step multilocation search task. *Developmental Psychology*, *34*, 203–214.
- Zelazo, P. R., & Zelazo, P. D. (1998). The emergence of consciousness. In H. H. Jasper, L. Descarries, V. F. Castellucci, & S. Rossignol (Eds.), *Consciousness: At the frontiers of neuroscience, advances in neurology* (Vol. 77, pp. 149–165). Philadelphia: Lippincott-Raven.