Journal of Experimental Psychology:
Learning, Memory, and Cognition
2004, Vol. 30, No. 3, 697–713

0278-7393/04/$12.00 DOI: 10.1037/0278-7393.30.3.697

Task Switching in a Hierarchical Task Structure: Evidence for the Fragility of the Task Repetition Benefit

Mei-Ching Lien and Eric Ruthruff
National Aeronautics and Space Administration Ames Research Center

This study examined how task switching is affected by hierarchical task organization. Traditional task-switching studies, which use a constant temporal and spatial distance between each task *element* (defined as a stimulus requiring a response), promote a *flat* task structure. Using this approach, Experiment 1 revealed a large switch cost of 238 ms. In Experiments 2–5, adjacent task elements were grouped temporally and/or spatially (forming an *ensemble*) to create a hierarchical task organization. Results indicate that the effect of switching at the ensemble level dominated the effect of switching at the element level. Experiments 6 and 7, using an ensemble of 3 task elements, revealed that the element-level switch cost was virtually absent between ensembles but was large within an ensemble. The authors conclude that the element-level task repetition benefit is fragile and can be eliminated in a hierarchical task organization.

In the course of daily activities, people face myriad tasks that could potentially be performed. Consequently, the key to achieving one's goals is the ability to select, prepare, and perform the most relevant task(s), rather than simply responding to the most salient impending stimuli in a reflexive, bottom-up manner. Because executive control of task set plays such an important role in human performance, this topic has become a focal point of recent cognitive psychology research. One widely used approach to studying executive control is the presentation of a series of single-task trials in which each task is either a repetition or a switch from the previous task. Studies using this task-switching paradigm have found that response time (RT) is generally longer for task switches than it is for task repetitions (the *switch cost*), even when participants have ample opportunity to prepare to switch tasks (the *residual switch cost*).

In the typical task-switching paradigm, participants receive an undifferentiated stream of discrete task elements (i.e., a stream of stimuli, each requiring a separate response) with a constant spatial and temporal spacing. Thus, there is no obvious reason for participants to form a strong higher level connection between any particular set of task elements. We refer to this type of task organization as a *flat* task structure. In contrast, the vast majority of real-world tasks have a *hierarchical* task structure. For instance, the sequential tasks of washing, drying, and folding clothes are joined by a higher level goal of "doing laundry." Although hierarchical organization has been found to have profound effects on memory (e.g., chunking) and perception (e.g., gestalt grouping), there has been little effort to examine how hierarchical organization of task elements affects task switching. Our goal in this study, therefore, is to examine this issue and develop a more complete model of executive task control.

Mei-Ching Lien and Eric Ruthruff, National Aeronautics and Space Administration (NASA) Ames Research Center, Moffett Field, California.

This research was supported by funding from the National Research Council and the Psychological and Physiological Stressors and Factors Project of NASA's Airspace Systems Program. We thank Iring Koch, Sander Los, Michael Masson, Nachshon Meiran, Jeff Miller, and Jennifer Stolz for their helpful comments on drafts of this article. We also thank Phil Allen, Jim Johnston, Joel Lachter, and Roger Remington for their insight and valuable discussions.

Correspondence concerning this article should be addressed to Mei-Ching Lien, Mail Stop 262-4, NASA Ames Research Center, Moffett Field, CA 94035. E-mail: mclien@mail.arc.nasa.gov

Background on Task Switching

Before reviewing previous task-switching studies, it is necessary to first clarify a few basic terms. One potential point of confusion is that the term *task* has sometimes been used to refer to a rule used to map stimuli to responses and sometimes to refer to an instance in which the rule is applied. For the sake of clarity, we use *task type* to refer to the stimulus–response (S-R) mapping rule, and we use *task element* to refer to an instance of a task (a stimulus requiring a response). This distinction between task types and task elements will be especially critical in the present study, in which we examine a hierarchical organization of task elements that differ in task type. One exception to this convention is that we refer to the act of switching between different task types simply as *task switching* (rather than *task type switching*).

In Jersild's (1927) influential task-switching study, he asked participants to perform two different types of speeded, choice RT tasks (Task Types A and B for short) in alternating blocks (e.g., ABABAB. . .) and pure blocks (e.g., AAAA. . .). Results showed that responses were often much slower in the alternating blocks than they were in the pure blocks (i.e., a switch cost). Importantly, switch costs were especially large when both task types operated on a common stimulus domain (e.g., odd–even judgments versus greater than–less than judgments on digits).[1] Because each stim-

[1] In fact, Jersild (1927) found small negative switch costs when the tasks were performed on separate sets of stimuli. Spector and Biederman (1976) later replicated this result and argued that it occurs because the switch condition benefits more from stimulus preview than does the repetition condition. Without preview, Spector and Biederman found a small, although nonsignificant, switch cost.

ulus is associated with, or "affords," the performance of both task types in this case (sometimes called *bivalent;* Fagot, 1994), the stimulus might automatically activate both the appropriate and inappropriate task sets. Consequently, more time might be needed to inhibit activations associated with the inappropriate task set in this *dual-affordance* condition than in the *single-affordance condition,* in which each stimulus is associated with only one relevant task type.

Using a variant of Jersild's (1927) method, Allport, Styles, & Hsieh (1994) found larger switch costs when participants shifted from difficult task types (presumably requiring strong task-set control) to easy task types (presumably requiring weak task-set control) than when participants shifted in the opposite direction. To account for this counterintuitive result, they proposed that switch costs are due to *task-set inertia*—the involuntary carryover of task-set activation from one task element to the next. Such persisting activation, hypothesized to be beyond participants' control, should facilitate performance when the task type repeats but interfere when the task type switches (Allport & Wylie, 2000). Because task-set activation is presumed to passively decay over time, the task-set inertia hypothesis implies that switch costs should decrease as the time interval between each response and the subsequent stimulus onset (response–stimulus interval; RSI) increases. Although increasing the RSI does reduce switch costs, Meiran (1996, 2000) argued persuasively that the RSI effect reflects, in part, an active preparation process rather than just a passive decay of task-set activation (for further evidence against task-set inertia theory, see De Jong, 2000, and Hübner, Futterer, & Steinhauser, 2001).

One drawback of the experimental design used by Jersild (1927) and Allport et al. (1994) is that participants must keep two task sets available in alternating blocks, but only one in pure blocks. Consequently, the costs of task switching are confounded with the costs of holding two S-R mappings in mind. To avoid this confound, Rogers and Monsell (1995) presented two visual-manual task types in alternating runs (e.g., AABB. . .) so that task switches and task repetitions were intermixed within the same block. One type of task was to classify a digit as odd or even, and the other was to classify a letter as consonant or vowel. In most of their experiments, a digit and a letter were presented adjacently on the screen (a dual-affordance condition). Therefore, to assist participants in tracking which task type needed to be performed on a given trial, the stimuli were presented in one of four squares, with the top two squares indicating one task type and the bottom two squares indicating the other task type. After each response, the next stimulus was displayed in the square located immediately clockwise from the square used in the previous trial.

Using the alternating-runs paradigm, Rogers and Monsell (1995) confirmed the existence of a substantial switch cost (approximately 200 ms) with dual-affordance stimuli, which remained large even when participants had plenty of time to prepare for the upcoming task switch (e.g., a long RSI of 1,200 ms). This residual switch cost led Rogers and Monsell to propose that online task reconfiguration is needed when the task switches, but not when it repeats. A crucial feature of this hypothesis is that the appearance of the stimulus is needed to trigger completion of the reconfiguration (hence the label *online*). They proposed that "this is the completion of a stagelike process of reconfiguration and that completion can be triggered only exogenously by the arrival of a stimulus suitably associated with the task" (p. 229). Logan and Gordon (2001) proposed a quantitative model along these same lines, in which the set of task parameters must be transmitted from working memory to the subordinate processor when the task switches but not when it repeats (see also Mayr & Kliegl, 2000).

Another piece of evidence supporting the online reconfiguration hypothesis comes from Rogers and Monsell's (1995) Experiment 6, in which a run of four elements of one task type was followed by a run of four elements of the other task type (e.g., AAAABBBB. . .). Mean RT was approximately 230 ms slower, and the error rate was 3% higher for the first element of each run (a task switch) than it was for the second element of each run (a task repetition). Importantly, no differences were found between the second, third, and fourth elements of each run (all of which were task repetitions). Gopher, Armony, and Greenshpan (2000) reported essentially the same result with an even longer run length. These findings strongly suggest that performance depends primarily on whether the task type is the same or different from the one performed on the immediately preceding element, just as predicted by the online reconfiguration hypothesis.

Effects of Hierarchical Task Organization

Unlike most task-switching studies, task elements in the real world tend to be hierarchically organized. A hierarchical task organization is especially likely to occur when multiple task elements serve the same immediate goal. For example, the task elements of washing, drying, and folding clothes can all be organized under the single higher level task of doing laundry. Hierarchical task organization seems especially likely to occur with contingent task elements, in which the output of one task element serves as the input to the next task element (e.g., see Carlson & Lundy, 1992; Wenger & Carlson, 1995). However, a hierarchical task organization can also occur for task elements that are performed together closely in time and/or space, even when those task elements are not contingent on one another and do not serve the same immediate goal. For instance, if the grocery store happens to be located next to the video store, the task of grocery shopping and the task of renting a video might often be performed closely in space and time and hence be linked together at a higher level of mental organization (e.g., a task of doing mall errands).

A hierarchical task organization might also be formed in many laboratory studies. For instance, dual-task experiments using the psychological refractory period (PRP) paradigm often present two independent task elements (Task 1 and Task 2) close together in time within each trial, followed by a long intertrial interval before the next pair of tasks. In addition to this temporal grouping, the task elements within a trial are generally preceded by a common fixation display and followed by a common feedback display. Arguably, this stimulus presentation strongly encourages participants to group Task 1 and Task 2 together at a higher level of mental task organization. There is, in fact, strong evidence that hierarchical task organization occurs in the PRP paradigm (e.g., De Jong, 1995; Lien, Schweickert, & Proctor, 2003; Luria & Meiran, 2003).

Although the general issue of how hierarchical organization affects performance has been examined extensively in the literatures on memory, perception, and (to a lesser extent) dual-task performance, it has not received much attention in the task-switching literature (for an exception, see Kleinsorge & Heuer, 1999). One useful clue, however, can be found in a dual-task experiment reported by De Jong (1995; Experiment 3). In that experiment, a visual-manual task (Type A) and an auditory-vocal task (Type B) were presented on each trial, either in the order AB or the order BA. Within each block, this order was either fixed from trial to trial (e.g., AB-AB-AB-AB. . . , where the hyphen indicates the break between trials) or alternating (e.g., AB-BA-AB-BA. . .). For our purposes, the key result is that RT to Task 1 (defined as the task element presented first within each trial) was significantly slower (albeit by only 11 ms) in the alternating-order condition than it was in the fixed-order condition. In other words, responses to Task 1 were actually faster when it was a task switch from previous Task 2 than when it was a task repetition. This intriguing *switch benefit* at the task-element level is opposite to the usual finding from traditional task-switching studies using single-task paradigms with flat task structures (e.g., Rogers & Monsell, 1995). De Jong, who was interested primarily in how preparatory control of Task 1 and Task 2 would affect Task 2 performance, did not comment on this unusual finding in detail.

If one considers switching only at the level of elementary tasks, as in traditional treatments of task switching, it is difficult to explain why De Jong (1995) found a switch benefit rather than a switch cost on Task 1. However, this result can be explained straightforwardly if one considers the possibility of switch costs at higher levels of task organization (in this case, at the level of the pair of task elements on each trial). For ease of discussion, we henceforth refer to the combination of task elements (e.g., Task 1 and Task 2) presented on each trial as a *task ensemble.*[2] Note that Task 1 in the fixed-order condition (e.g., AB-AB-AB-AB. . .) involved a task type repetition at the ensemble level but a switch at the element level, whereas Task 1 in the alternating-order condition (e.g., AB-BA-AB-BA. . .) involved a task type switch at the ensemble level but a repetition at the element level. Thus, in effect, this experimental design pitted the effects of switching at the element level against the effects of switching at the ensemble level (see Table 1). Viewed from this perspective, Task 1 responses were fast in the fixed-order condition because the ensemble repeated and slow in the alternating-order condition because the ensemble switched. Thus, the results are consistent both with the hypotheses that the dual-task presentation led participants to form a hierarchical task organization (covering the ensemble of Task 1 and Task 2) and that switching at the task-ensemble level influences performance.

The Preparatory Control Structure for an Ensemble

How do participants prepare for an ensemble of multiple task elements in a hierarchical task structure? Consider a paradigm in which two different task types (Task Types A and B) must be performed on each trial (as in the PRP paradigm). One simple hypothesis is that participants prepare the S-R mapping rules for both task types without regard to the order in which these types are to be performed. De Jong (1995; Experiment 1), however, found that responses were slow and error rates were high when participants were cued to expect a certain order of Type A and Type B in the upcoming trial but instead received the opposite order. This result clearly indicates that participants prepare for a particular order of task types rather than the just the task types to be performed.

Another hypothesis for ensemble preparation is that participants prepare exclusively for the task type expected to appear first without explicitly preparing for the task type expected to appear second. This hypothesis was also ruled out by De Jong (1995; Experiment 3), who found that Task 2 RT decreased when the preparation time (the intertrial interval) increased. Luria and Meiran (2003) found similar results using explicit cues that indicated the order of the task types for the upcoming trial in a PRP design. De Jong proposed that there is “a multi-level control structure that prepares the processing system not only for the immediate performance of the first task but also for a timely and rapid switch to the second task” (p. 21). According to this view, participants initially commit fully to Task 1 and, once the critical stages on Task 1 have been completed, they then commit fully to Task 2.

How exactly do participants prepare to carry out Task 2 even before they have finished performing Task 1? It seems unlikely that participants actually load the task set for Task 2 before they have completed the critical stages (e.g., a response-selection stage) on Task 1 (see Lien et al., 2003, for detailed discussion of this point). The reason is that the set for Task 2 might crowd out and/or interfere with the set for Task 1 in working memory; note that, in dual-affordance conditions, the task types are highly incompatible. However, the preparation might involve firmly establishing the intention to switch task set after a certain point in Task 1 processing (e.g., immediately following response selection). This preparation might include (a) the intention to disengage the set for Task 1 and (b) the intention to engage the specific task set needed for Task 2.

Interaction Between Element-Level and Ensemble-Level Effects

Having discussed the nature of ensemble preparation, we now consider how the effects of switching at the ensemble level might interact with effects of switching at the element level. Because De Jong’s (1995) experimental design pitted the task-switching effect at the ensemble level against the effect at the element level, the results are consistent with at least two different interpretations. First, it is possible that both the ensemble level and the element level produced substantial effects, but that the ensemble-level effect was larger than the element-level effect. Second, the hierarchical task organization might have suppressed the effect at the

Table 1
Classification of the Fixed-Order Conditions and the Alternating-Order Conditions in Terms of Element-Level Task Transition (on Task 1) and Ensemble-Level Task Transition

Condition	Element level	Ensemble level
Fixed order (e.g., AB-AB-AB-AB. . .)	Switch	Repetition
Alternating order (e.g., AB-BA-AB-BA. . .)	Repetition	Switch

[2] The term *task ensemble* refers to a set of task elements that are grouped together by some experimental manipulation. It is logically possible in such cases that the individual tasks would be joined as one single *supertask;* that is, the task types defined at the element level might lose all significance. This hypothesis seems unlikely, however, given that the tasks discussed in the present article require an individual response to each individual stimulus, without temporal overlap between task elements.

element level so that there was no cost of switching cognitive operations from one individual task element to the next.

Before considering possible tests of these alternative hypotheses for the interaction between element- and ensemble-level effects (see Experiments 6 and 7), we first need to critically examine whether task ensembles actually have any effect at all. The absence of a switch cost in De Jong's (1995) study might have occurred simply because his experimental design was generally insensitive to element-level task switching. Indeed, De Jong used single-affordance stimuli, which often produce little or no switch cost, combined with a larger than usual amount of practice (3 sessions of 720 trials each). In addition, he used different input modalities (visual and auditory) for the two task types, whereas traditional task-switching studies typically use the same input modality (visual) for both task types (e.g., Allport et al., 1994; Rogers & Monsell, 1995). This separation of input modalities might have minimized the switch cost.

Given these considerations, it is plausible that De Jong's (1995) design would have produced no switch cost on Task 1 even if there were no hierarchical task organization. Because De Jong's study lacked a control condition with a flat task structure, there is no direct way to determine what effect, if any, task ensembles had. Another complicating factor in his study is the presence of temporal overlap between the processing of Task 1 and Task 2 (in a PRP paradigm). Perhaps Task 1 was differentially influenced by temporal overlap in the fixed- and alternating-order conditions. These problems ultimately stem from the fact that De Jong's study was not designed to examine the interaction between ensemble-level and element-level effects. To more clearly pin down the role of hierarchical task organization in task switching, it is necessary to design new experiments for this purpose.

The Present Study

In this study, we examined how task switching is affected by hierarchical task organization. As a starting point, we adopted De Jong's (1995) basic experimental design, in which the two types of tasks (which we call Type A and Type B for ease of discussion) are presented on each trial in either a fixed order (e.g., AB-AB. . .) or an alternating order (e.g., AB-BA. . .) within a block. Consistent with the PRP terminology used by De Jong, we refer to the first task element within each trial as Task 1 and to the second task element as Task 2. However, we made several critical changes to De Jong's dual-task design. To facilitate comparison with traditional task-switching studies, we presented nonoverlapping, single-task trials. To maximize the sensitivity of our design to the effects of task switching, we used dual-affordance stimuli, which generally produce very large element-level switch costs (~200 ms) in the traditional task-switching paradigm. Consequently, if this element-level switch cost can be eliminated when there is a task ensemble, it would be very impressive and would provide strong evidence for an effect of hierarchical task organization.

To evaluate the ensemble-level effect, it was necessary to first measure the baseline switch cost with a flat task structure, in which the task ensemble structure clearly was not present. Thus, in Experiment 1, we used a variant of the traditional task-switching paradigm with a long, constant RSI (1,500 ms) between successive task elements. Although we labeled the task elements Task 1 and Task 2 for continuity with later experiments, nothing in the overt events of Experiment 1 distinguished Task 1 and Task 2. Because there was no stimulus support for a hierarchical organization of tasks, we assume that each task element was typically performed independently (i.e., in a flat task structure).

To validate the existence of ensemble-level effects, we manipulated stimulus factors in Experiments 2–5 that should, on the basis of common sense and gestalt psychology, increase mental ensemble-level task organization. Specifically, we presented stimuli for Task 1 and Task 2 close together in time and/or space. We began with a modest manipulation in Experiment 2 (shortening the RSI between Task 1 and Task 2) and then made further modifications in the subsequent experiments to provide increasingly strong support for an ensemble-level organization of Task 1 and Task 2.

Measurement of Switch Costs and Predictions

The critical predictions in our study concern Task 1 performance. As shown in Table 1, Task 1 was an element-level switch but an ensemble-level repetition in the fixed-order condition (e.g., AB-AB. . .), whereas Task 1 was an element-level repetition but an ensemble-level switch in the alternating-order conditions (e.g., AB-BA. . .). Thus, we can determine the relative effects of element-level task switching and ensemble-level task switching by comparing the performance of Task 1 in these two conditions. As a measure of the *relative switch cost* in Experiments 2–5, we subtracted Task 1 RT in the alternating-order blocks (element repetition, ensemble switch) from that in the fixed-order blocks (element switched, ensemble repetition):

$$\text{Relative switch cost} =$$
$$\text{Task 1 RT (fixed order)} - \text{Task 1 RT (alternating order)}.$$

Positive values for the relative switch cost indicate that the element-level effect was larger than the ensemble-level effect. Negative values (as in De Jong, 1995, Experiment 3) indicate that the ensemble-level effect was larger than the element-level effect.

As in traditional task-switching studies using dual-affordance stimuli, we expected to observe a large positive value (about 200 ms) for the baseline Task 1 switch cost in Experiment 1 (in which there was no stimulus support for task ensembles). In the progression of Experiments 2–5, we provided increasingly strong stimulus support for task ensembles by increasing the temporal and/or spatial grouping of Task 1 and Task 2 on each trial. If there is no effect of ensemble-level task switching, then the relative switch cost on Task 1 in Experiments 2–5 should be similar to that observed in Experiment 1. In contrast, if there is an effect of ensemble-level task switching, it should cause the relative switch cost on Task 1 to decrease in Experiments 2–5 relative to Experiment 1 and, possibly, even reverse and become negative (as in De Jong, 1995). Looking ahead to the results, the relative switch cost on Task 1 did, in fact, become negative in Experiments 4 and 5. In Experiments 6 and 7, therefore, we used a slightly different task paradigm to isolate the contribution of the element-level effect from the ensemble-level effect and determine whether the element-level effect was still present when there was strong stimulus support for a task ensemble.

Note that the RSI (the opportunity for advance preparation) has been shown to have a modest effect on switch costs (e.g., Rogers & Monsell, 1995). Although we manipulated the RSI between Task 1 and Task 2 across experiments, the RSI leading up to Task

1 was always 1,500 ms. Because the critical data used to assess the relative switch cost come only from Task 1, the comparison of the switch cost across experiments was not confounded with the effects of RSI.

Experiment 1

Our purpose in Experiment 1 was to measure the baseline switch cost on Task 1 in a design without task ensembles. A circular frame with four colored quadrants was presented in the center of the screen throughout each block (see Figure 1). The first stimulus of each block was presented in the top quadrant. The subsequent stimuli were presented, with a constant RSI of 1,500 ms, in the quadrant located immediately clockwise from the previous one, similar to the procedure used by Rogers and Monsell (1995). Because there was no stimulus support for a task ensemble, we assumed that participants generally prepared for task elements individually.

As in Rogers and Monsell (1995), we provided two redundant task cues. One cue was the predictable, repeating task sequence, which was constant throughout each block. The other cue was the color of the location in which the stimulus appeared. Black quadrants indicated the magnitude task, and blue quadrants indicated the parity task. The stimuli rotated from quadrant to quadrant in a predictable manner (generally clockwise) after each response.

Method

Participants. Twenty participants from colleges and universities surrounding the National Aeronautics and Space Administration (NASA) Ames Research Center participated in exchange for extra course credit. All participants were required to have normal or corrected-to-normal vision.

Apparatus and stimuli. The stimuli were presented on IBM-compatible microcomputers connected to SONY Trinitron monitors, housed in a dedicated, sound-attenuating booth. Stimulus presentation, timing, and data collection were controlled using E-Prime (2001) software. The circular frame, 8 cm in diameter, was divided into four quadrants (see Figure 1): top, right, bottom, and left quadrants. The stimuli were the digits 1 to 9, excluding 5, presented in a center of the quadrant (2.0 cm from the middle of the circle). The digits were 0.6 cm in width and 1.0 cm in height and subtended 0.63° (width) × 1.04° (height) based on a viewing distance of 55 cm.

Design and procedure. Two different types of numerical judgments, magnitude and parity, were used as tasks. For the magnitude task (Type A), participants determined whether the number was greater or less than 5. They were to press the *Z* key with their left index finger if the number was less than 5 (1, 2, 3, or 4) and the */* key with their right index finger if the number was greater than 5 (6, 7, 8, or 9). For the parity task (Type B), participants judged whether the number was odd or even. They were to press the *Z* key with their left index finger if the number was odd (1, 3, 7, or 9) and the */* key with their right index finger if the number was even (2, 4, 6, or 8).

The four quadrants were colored blue or black to indicate the task type to be performed on stimuli appearing within that quadrant. Black indicated that the magnitude task should be performed, whereas blue indicated that the parity task should be performed. The first stimulus of each block appeared in the top quadrant. After participants responded to that digit, feedback ("wrong" on error trials or a blank message on correct trials) appeared in that quadrant for 300 ms. A new digit appeared 1,200 ms later in the quadrant located immediately clockwise from the previous quadrant (i.e., the right quadrant). Thus, the total RSI was 1,500 ms. Figure 1 shows the sequence of events for stimuli presented in the top and right quadrants,

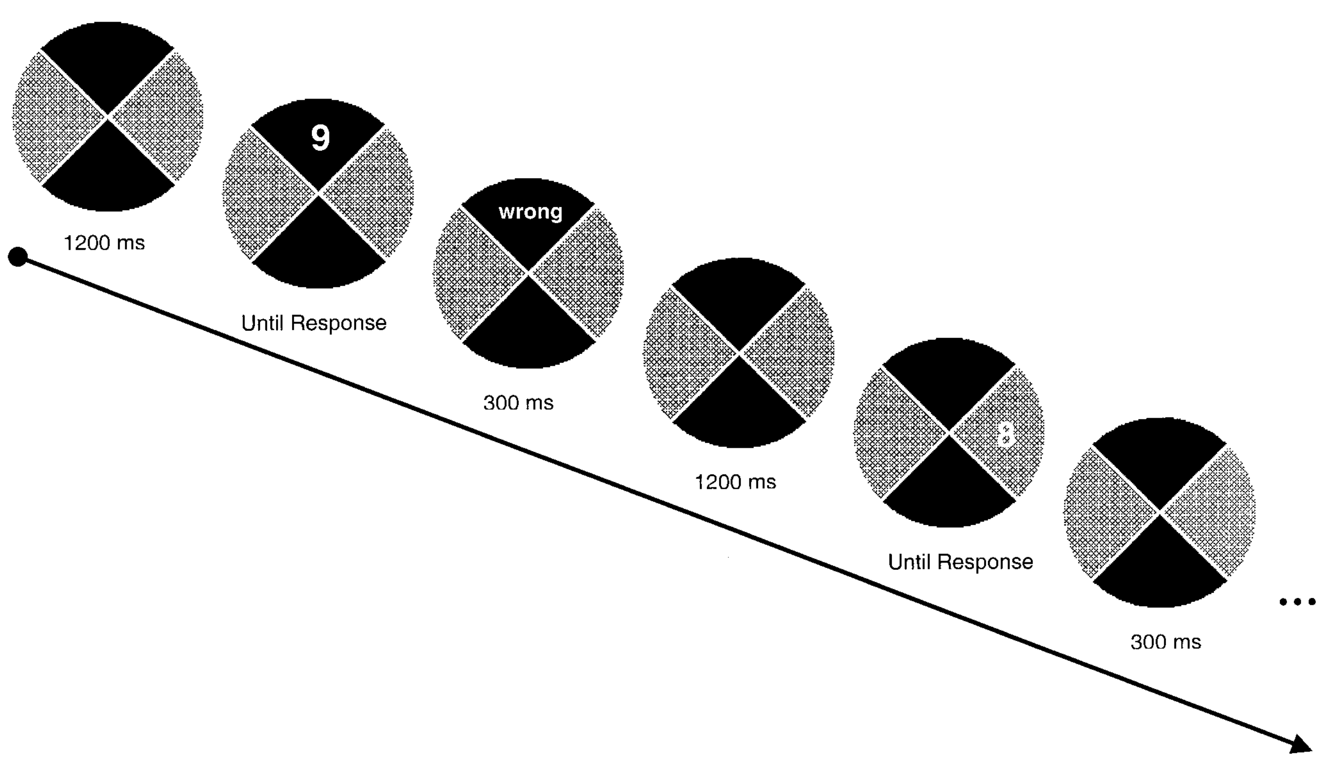

Figure 1. An example of the time course of stimulus presentation used in Experiment 1. The black quadrant corresponds to the magnitude task (Type A), and the gray quadrant (which was blue in the actual experiment) corresponds to the parity task (Type B).

with the ABAB task sequence as an example. The digits continued rotating clockwise around the circle throughout the block (i.e., through the top, right, bottom, and left quadrants, etc.). Therefore, the distance from one stimulus to the next was constant across trials, as in Rogers and Monsell (1995). Another consequence of this procedure is that the color for a quadrant was constant throughout each block but changed across blocks, depending on the specific task sequence to be performed.

Although there was no stimulus support for an ensemble of Task 1 and Task 2 in Experiment 1, for continuity with the subsequent experiments, we refer to task sequences of ABAB and BABA as the *fixed-order* condition and task sequences of ABBA and BAAB as the *alternating-order* condition (see Figure 2 for an example of the stimulus arrangement for each condition). Participants performed each of the four block types once within the session, using one of the following four orders: (a) ABAB, ABBA, BAAB, BABA; (b) ABBA, ABAB, BABA, BAAB; (c) BAAB, BABA, ABAB, ABBA; or (d) BABA, BAAB, ABBA, ABAB. Each order was used equally often across participants.

Participants first performed 64 practice trials, which served to acquaint them with the tasks and the paradigm. For each of the four block types mentioned above, participants performed a practice block of 32 trials (to acquaint them with the new task sequence) followed by four blocks of 64 regular trials each. Participants were told that both speed and accuracy were important. They were also encouraged to take a brief break between blocks.

Data Analyses

The first cycle of four task elements in each experimental block, serving as warm-ups, were omitted from the analyses. Also omitted were task elements in which the stimulus was the same as the previous stimulus, to avoid contamination of task repetition effects with stimulus repetition effects. For RT analyses, data were also omitted if the current or previous response was an error. RTs outside the range of 100–4,000 ms were treated as outliers; this led to the elimination of an additional 0.35% of responses (see Ulrich & Miller, 1994).

Task type, magnitude versus parity, had little effect in these experiments and did not consistently interact with other factors; consequently, this variable was not included as a factor in the final data analyses. Because our primary concern was the relative switch cost on Task 1 between the alternating-order and the fixed-order conditions, data analyses were reported only for Task 1 and included only task-type order (fixed vs. alternating) as a factor. An alpha level of .05 was used to determine statistical significance.

Results and Discussion

The main purpose of Experiment 1 was to provide a baseline for the switch cost on Task 1. As described above, this switch cost was measured by subtracting Task 1 RT (or the proportion of error, PE) in the alternating-order condition from that in the fixed-order condition. Mean RT and PE for the four task elements in each cycle (Task 1, Task 2, Task 1, and Task 2) are shown in Table 2 for each block type. The effect of order condition on RT was significant, $F(1, 19) = 35.91$, $p < .01$, $MSE = 31{,}565$; mean RT

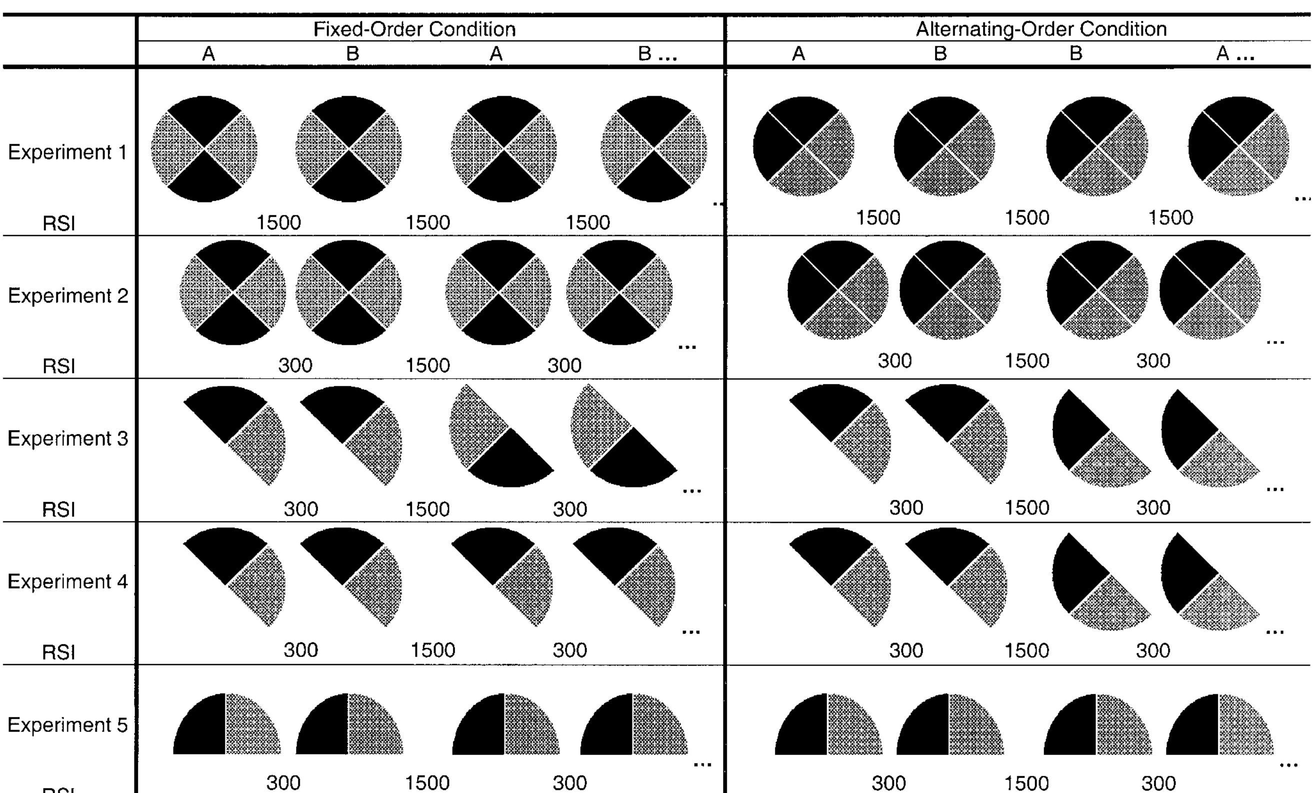

Figure 2. An example of the stimulus arrangement for the fixed-order conditions (ABAB) and the alternating-order conditions (ABBA) of Experiments 1–5. The black quadrant corresponds to the magnitude task (Type A), and the gray quadrant (which was blue in the actual experiment) corresponds to the parity task (Type B). RSI = response–stimulus interval.

Table 2
Mean Response Times (RTs; in Milliseconds) and Proportions of Errors (PEs) for Task 1 and Task 2 of Each Cycle in the Fixed-Order Condition and the Alternating-Order Condition in Experiments 1–5

Condition and sequence	Task 1		Task 2		Task 1		Task 2	
	RT	PE	RT	PE	RT	PE	RT	PE
Experiment 1								
Fixed order								
ABAB	786	.04	856	.05	818	.05	864	.06
BABA	917	.09	853	.05	894	.06	790	.04
Alternating order								
ABBA	597	.04	832	.04	611	.04	746	.05
BAAB	649	.03	736	.05	605	.03	830	.04
Experiment 2								
Fixed order								
AB-AB	985	.06	947	.05	960	.05	933	.05
BA-BA	1,024	.04	974	.07	992	.06	956	.06
Alternating order								
AB-BA	899	.07	919	.04	849	.03	939	.08
BA-AB	843	.02	904	.05	892	.04	934	.07
Experiment 3								
Fixed order								
AB-AB	1,039	.03	861	.04	897	.02	844	.03
BA-BA	984	.04	802	.02	962	.04	808	.04
Alternating order								
AB-BA	938	.05	947	.06	914	.03	873	.04
BA-AB	981	.05	855	.04	837	.03	915	.03
Experiment 4								
Fixed order								
AB-AB	926	.03	811	.05	909	.05	834	.04
BA-BA	907	.03	783	.07	903	.04	791	.06
Alternating order								
AB-BA	1,008	.08	895	.05	956	.05	834	.03
BA-AB	1,003	.03	872	.03	997	.04	856	.02
Experiment 5								
Fixed order								
AB-AB	864	.02	871	.03	853	.03	875	.03
BA-BA	1,045	.03	917	.02	982	.05	894	.04
Alternating order								
AB-BA	1,016	.03	1,016	.04	1,086	.03	1,016	.03
BA-AB	1,011	.04	1,022	.03	1,028	.03	1,040	.02

Note. The relative switch cost is measured on Task 1 by subtracting performance measures in the alternating-order condition from those in the fixed-order condition. A = magnitude task; B = parity task.

for Task 1 was 616 ms in the alternating-order condition but 854 ms in the fixed-order condition. In other words, the switch cost on Task 1 was 238 ms. The effect was in the same direction for the PE data, $F(1, 19) = 8.19$, $p < .05$, $MSE = 0.0017$, with a switch cost of .03.

The baseline switch cost on Task 1 obtained in this experiment (238 ms) is consistent with the findings of traditional task-switching studies (e.g., Rogers & Monsell, 1995). Note that the alternating-order condition (e.g., ABBAABBA. . .), by itself, is the same as the alternating-runs paradigm of Rogers and Monsell with a run length of two. Therefore, we can use this condition to measure the switch cost in the exact same way in which Rogers and Monsell did; that is, we can compare the performance of task element repetitions (Task 1) and task element switches (Task 2). Using this measure of the switch cost, we obtained a 170-ms switch cost. This effect is roughly similar in magnitude to the costs obtained in previous studies and to the effect we obtained by comparing Task 1 performance in the fixed-order and alternating-order conditions. Having demonstrated that our basic experimental design (without task ensembles) produced the usual switch cost, we proceeded to examine the effect of task ensembles.

Experiment 2

The purpose of Experiment 2 was to determine whether responses to Task 1 would still be faster in the alternating-order condition (e.g., AB-BA. . .) than they were in the fixed-order condition (e.g., AB-AB. . .) when there was stimulus support for an ensemble of Task 1 and Task 2. In De Jong's (1995) Experiment 3, the stimulus for Task 2 often appeared while participants were still performing Task 1 (i.e., a PRP paradigm). The temporal proximity between Task 1 and Task 2 might have been the critical factor that led to the formation of a hierarchical task organization. Furthermore, studies of response sequence learning have shown that temporal grouping can increase the likelihood that a series of responses will be chunked together (e.g., Koch & Hoffmann, 2000; Stadler, 1993). Therefore, as an initial attempt, we simply shortened the RSI between Task 1 and Task 2 to only 300 ms, leaving the RSI between Task 2 to the next Task 1 at 1,500 ms. This modification brought Experiment 2 one step closer to the PRP paradigm used in De Jong's Experiment 3. However, in our design the stimulus for Task 2 did not appear until after participants had responded to Task 1. Thus, unlike De Jong's study, in our study there was no temporal overlap in the processing of Task 1 and Task 2.

As described in the introduction, we measured relative switch costs by comparing the performance on Task 1 between the fixed-order and alternating-order conditions. If there is no effect of ensemble-level task switching, the relative switch cost on Task 1 should be similar to the 238-ms switch cost obtained in Experiment 1 (note that the RSI leading up to Task 1 was the same in both experiments). In other words, Task 1 RT should still be roughly 238 ms faster in the alternating-order condition than in the fixed-order condition. In contrast, if there is an effect of ensemble-level task switching, it should cause the relative switch cost to decrease relative to Experiment 1, or even to reverse and become negative (as in De Jong, 1995).

Method

Participants. There were 20 participants in this experiment. These participants were recruited from the same population as those in Experiment 1, but none had participated in that experiment.

Apparatus, stimuli, and procedure. The method was the same as in Experiment 1 except that the RSI between Task 1 and Task 2 was reduced from 1,500 ms to 300 ms, leaving just enough time for the presentation of the Task 1 feedback message.

Results and Discussion

Mean RT and PE for the four task elements in each cycle are shown in Table 2. The data analyses were performed in the same way as in Experiment 1. Less than 0.5% of responses were omitted because of RT cutoffs. The effect of order condition on Task 1 was significant for RT, $F(1, 19) = 8.26$, $p < .01$, $MSE = 34{,}643$, but not for PE, $F(1, 19) < 1.0$. Mean RT for Task 1 was 871 ms in the alternating-order condition (e.g., AB-BA. . .) and 991 ms in the fixed-order condition (e.g., AB-AB. . .). In other words, when we shortened the RSI between Task 1 and Task 2, the relative switch cost on Task 1 was reduced from 238 ms in Experiment 1 to 120 ms in Experiment 2 (see Figure 3). This reduction in the switch cost was significant, $F(1, 38) = 4.24$, $p < .05$, $MSE = 33{,}104$, consistent with the hypothesis that there is an effect of switching between different task ensembles.

Experiment 3

In Experiment 2, increasing the temporal contiguity between Task 1 and Task 2 reduced the switch cost relative to Experiment 1. However, unlike De Jong's (1995) Experiment 3, in our experiment, a switch cost was still obtained. One potential reason for the difference in results is that temporal contiguity failed to reliably produce an ensemble of Task 1 and Task 2. Perhaps the 300-ms RSI could not easily be distinguished from the 1,500-ms RSI, or perhaps participants had difficulty tracking which task elements had a short preparation time (short RSI) and which ones had a long preparation time (long RSI). Consequently, participants might not have formed an ensemble of Task 1 and Task 2. Moreover, if participants did not know in advance that they would have a sufficiently long preparation time, they might have failed to engage in the appropriate preparation for the upcoming task element (see Rogers & Monsell, 1995, for a similar argument). One line of evidence for this possibility is that responses on Task 1 and Task 2 were, overall, much slower in Experiment 2 than they were in Experiment 1 (see Table 2).

Experiment 3 was designed to further increase the stimulus support for a task ensemble. Besides the temporal dimension, spatial location is one of the most salient physical dimensions in perception. Research on gestalt organizational principles has shown that objects are optimally grouped as a unit when they are connected and located within the same explicit boundary on the display; these methods of grouping are known as connectedness and common region, respectively (Rock & Palmer, 1990). Thus, Experiment 3 examined the switch cost with both temporal grouping and spatial grouping between Task 1 and Task 2. Specifically, instead of presenting all four quadrants of the circular frame at all times (as in Experiments 1 and 2), only the semi-circular frame containing the two quadrants for the current ensemble was presented (see Figure 2). The RSI between Task 1 and Task 2 was still 300 ms, as in Experiment 2.

Method

Participants. Twenty college students, from the same participant pool as the participants in Experiments 1 and 2, were in this experiment. None had participated in those experiments.

Apparatus, stimuli, and procedure. The apparatus, stimuli, and procedure were the same as in Experiment 2, except as noted. Only two quadrants of the semicircular frame appeared on the screen in each trial (i.e., the top and right quadrants appeared first, then the bottom and left quadrants, then the top and right quadrants again; see Figure 2). After

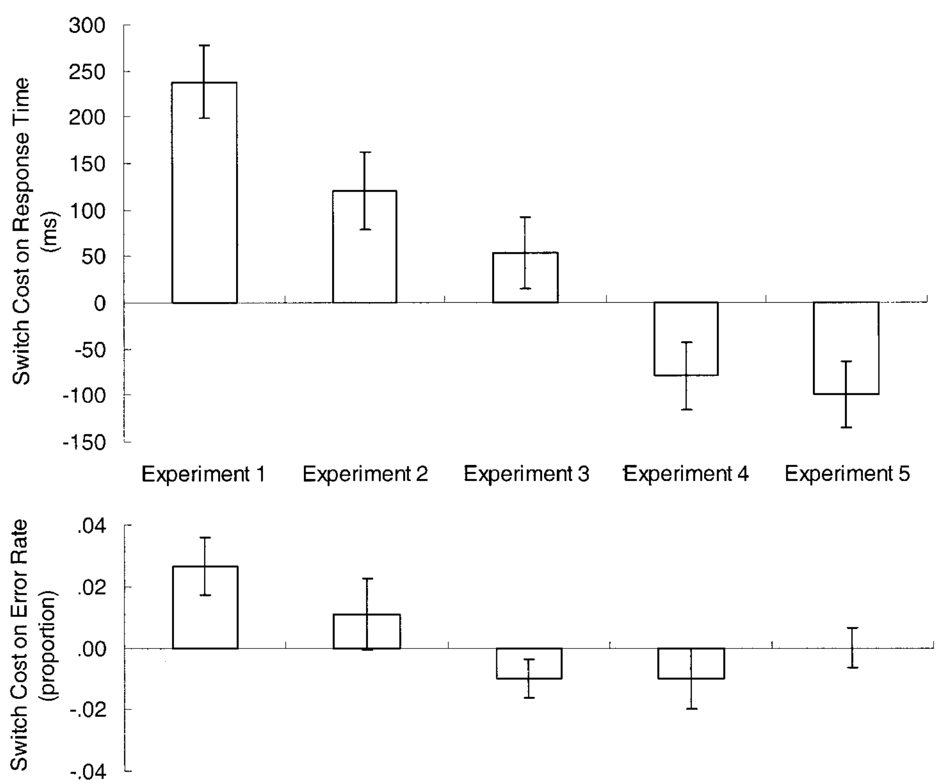

Figure 3. The relative switch cost for mean response time and proportion of errors on Task 1 in Experiments 1–5. Error bars represent the standard error of the mean.

participants responded to Task 2, feedback for Task 2 was displayed inside the quadrants for 300 ms. The feedback and the semicircular frame then disappeared, and the semicircular frame for the next two quadrants appeared. After 1,200 ms, the stimulus for the next Task 1 appeared. Therefore, as in Experiments 1 and 2, the total RSI leading up to Task 1 was 1,500 ms.

Results and Discussion

Mean RT and PE for the four task elements in each cycle are shown in Table 2. The data analyses were performed in the same way as in Experiments 1 and 2. Less than 0.4% of responses were omitted because of the RT cutoffs. In contrast to Experiments 1 and 2, the effect of order condition on Task 1 was not significant for RT, $F(1, 19) = 1.90$, $p = .1843$, $MSE = 29{,}677$, or PE, $F(1, 19) = 2.50$, $p = .13$, $MSE = 0.0008$. Mean RT for Task 1 was 917 ms in the alternating-order condition and 970 ms in the fixed-order condition. Thus, when Task 1 and Task 2 were spatially and temporally grouped, the relative switch cost declined to less than one fourth of the cost found in Experiment 1 (53 ms vs. 238 ms). This decline was statistically significant, $F(1, 38) = 11.18$, $p < .01$, $MSE = 30{,}621$, consistent with the hypothesis that there is an effect of ensemble-level task switching.

Experiment 4

Even when the ensemble repeated in Experiment 3, its position always changed (see Figure 2). This change in position may have reduced the stimulus support for ensemble-level task repetition. To prevent this from occurring, we always presented a particular ensemble of Task 1 and Task 2 (e.g., AB) in the same position (e.g., on the top and right quadrants) in Experiment 4. If the consistent stimulus presentation from trial to trial can enhance stimulus support for ensemble-level task repetition, an even stronger ensemble-level effect should be obtained. In other words, the relative switch cost on Task 1 should decrease further or even reverse and become negative (as in De Jong, 1995).

Method

Participants. Twenty participants from the same participant pool as the participants in Experiments 1–3 were in this experiment. None had participated in those experiments.

Apparatus, stimuli, and procedure. The apparatus, stimuli, and procedure were the same as in Experiment 3, except as noted. The top and right quadrants were always colored black and blue, respectively, to indicate the sequence of Type A (the magnitude task) followed by Type B (the parity task). The bottom and left quadrants were always colored blue and black, respectively, to indicate the sequence of Type B followed by Type A.

Consequently, the presentation locations in each task sequence were as follows. In the AB-AB sequence of the fixed-order condition, the presentation locations were top, right, top, right, and so on. In the BA-BA sequence, the presentation locations were bottom, left, bottom, left, and so on. In the alternating-order condition, the presentation locations in the AB-BA sequence were top, right, bottom, left, and so on, and the presentation locations in the BA-AB sequence were bottom, left, top, right, and so on.

Results and Discussion

Mean RT and PE for the four task elements in each cycle are shown in Table 2. The data analyses were performed in the same way as in the previous experiments. Less than 0.2% of responses were omitted because of the RT cutoffs. The effect of order condition on Task 1 was significant for RT, $F(1, 19) = 4.80$, $p < .05$, $MSE = 26{,}127$, but not for PE, $F(1, 19) = 1.06$, $p = .32$, $MSE = 0.0019$. Mean RT for Task 1 was 991 ms in the alternating-order condition but only 912 ms in the fixed-order condition. Thus, in contrast to Experiments 1–3, mean RT of Task 1 was actually *slower* in the alternating-order than it was in the fixed-order conditions (see Figure 3). That is, the relative switch cost on Task 1 was −79 ms, indicating that the ensemble-level effect was larger than the element-level effect. This large decline in the relative switch cost from Experiment 1 (238 ms) to Experiment 4 (−79 ms) was significant, $F(1, 38) = 34.89$, $p < .01$, $MSE = 28{,}846$. Thus, Experiment 4 provides further evidence that there is an effect of ensemble-level task switching when the stimulus support for an ensemble is strong.

Experiment 5

In Experiments 3 and 4, the display contained only the two quadrants for the current Task 1 and Task 2. After participants responded to both Task 1 and Task 2, these quadrants disappeared, and the quadrants for the next Task 1 and Task 2 appeared. It is possible that the appearance of the frame prompted participants to interpret the color of the quadrants for the next Task 1 and Task 2. Perhaps this cue interpretation (which would itself be a type of "task") intervened between the task repetitions and somehow eliminated the element-level repetition benefit. We refer to this possibility as the cue-interpretation hypothesis. Because the task sequence was given in advance and repeated throughout each experimental block (plus the immediately preceding practice block), it seems unlikely that participants needed to interpret the color cue on each trial. Nevertheless, it is conceivable that the sudden appearance of the colored frame prompts cue interpretation even when it is not necessary. Experiment 5 was designed to directly address this possibility.

Experiment 5 also addressed a potential confound in Experiment 4. In that experiment, each task type in the fixed-order condition was always presented in the same location, whereas each task type in the alternating-order condition was presented in two different locations across trials. For example, the Task Types A and B in the sequence AB-AB were always displayed in the top and right quadrants, but the Task Types A and B in the sequence AB-BA were presented in the top and right quadrants and then the bottom and left quadrants. It is conceivable that the consistent display location of Types A and B in the fixed-order condition (e.g., AB-AB. . .) led to a reduction in RT and hence can account for some of the negative switch cost observed in Experiment 4.

To address the cue-interpretation hypothesis and the potential confound between task type order and the number of display locations, the critical change we made in Experiment 5 was to display the same semicircular frame on the screen throughout the entire experiment. As shown in Figure 2, this frame contained a black quadrant (Type A) on the left and a blue quadrant (Type B) on the right. Because the frame never changed (just as in Experiment 1), it is unlikely that cue interpretation was necessary. Furthermore, the location for a particular task type was always the same, in both the fixed-order condition and the alternating-order condition. Thus, this experiment deconfounds task type order with the number of display locations per task type.

Method

Participants. Twenty participants from the same participant pool as that used in Experiments 1–4 were used in this experiment. None had participated in those experiments.

Apparatus, stimuli, and procedure. The apparatus, stimuli, and procedure were the same as in Experiment 4, except as noted. As shown in Figure 2, there were only two quadrants (left and right), located adjacent to each other. The left and right quadrants were always colored black and blue, respectively, to indicate that the magnitude task would always be on the left side and that the parity task would always be on the right side. Immediately after the feedback for the previous trial, a fixation cross was presented for 300 ms in the quadrant for the next Task 1 (similar to Experiment 5 of Rogers & Monsell, 1995). The purpose of the fixation point was to ensure that participants would know the stimulus location for the upcoming Task 1. The stimulus for Task 1 appeared 900 ms later. Thus, the RSI leading up to Task 1 was still 1,500 ms, as in previous experiments.

In the fixed-order condition, the presentation locations in the AB-AB sequence were left, right, left, right, and so on, and the presentation locations in the BA-BA sequence were right, left, right, left, and so on. In the alternating-order condition, the presentation locations in the AB-BA sequence were left, right, right, left, and so on, and the presentation locations in the BA-AB sequence were right, left, left, right, and so on. Note that the location of Task 1 was always the same as the previous Task 2 in the alternating-order condition, which should highlight the fact that it was an element-level task repetition.

Results and Discussion

Mean RT and PE for the four task elements in each cycle are shown in Table 2. The data analyses were performed in the same way as in the previous experiments. Less than 0.3% of responses were omitted because of the RT cutoffs. As in Experiment 4, the effect of order condition on Task 1 was significant for RT, $F(1, 19) = 7.73$, $p < .05$, $MSE = 25{,}415$, but not for PE, $F(1, 19) < 1.0$. Mean RT for Task 1 was 1,035 ms in the alternating-order condition (e.g., AB-BA. . .) but 936 ms in the fixed-order condition (e.g., AB-AB. . .). Thus, similar to Experiment 4, the relative switch cost was −99 ms, indicating that the ensemble-level effect was once again larger than the element-level effect. Comparing Experiment 1 with Experiment 5, the relative switch cost on Task 1 was reduced from 238 ms to −99 ms on RT and from .03 to .00 on PE (see Figure 3). This large decline in the relative switch cost was significant for both RT, $F(1, 38) = 39.91$, $p < .01$, $MSE = 28{,}490$, and PE, $F(1, 38) = 4.84$, $p < .05$, $MSE = 0.0013$. Thus, these findings are consistent with the hypothesis that there is a large effect of ensemble-level task switching when the stimulus support for an ensemble is strong.

In Experiment 5, the same two-quadrant frame (see Figure 2) was presented on the screen at all times. Because the frame did not disappear and reappear (as in Experiment 4), there is no reason for the frame to prompt cue interpretation on every trial. In addition, this experiment eliminated the potential confound in Experiment 4 between task type order and the number of display locations. Nevertheless, a negative switch cost, similar in magnitude to the effect in Experiment 4, was still observed.

Although Experiments 4 and 5 revealed evidence for ensemble-level effects, they do not indicate whether the alternating-order condition (e.g., AB-BA. . .) was slower than the fixed-order condition (e.g., AB-AB. . .) because of the need to switch ensembles or because of the need to maintain readiness for two different ensembles at the same time (see Los, 1996; Luria & Meiran, 2003; Meiran, 2000; Rogers & Monsell, 1995). Further research would be needed to determine the relative contributions of these two components. Rather than pursue this issue in the subsequent experiments, however, we instead focused on the critical issue of whether there is any element-level switching effect in a hierarchical task organization.

Experiment 6

Experiments 1–5 showed that when there was stimulus support for a task ensemble (due to temporal and/or spatial grouping), the relative switch cost on Task 1 decreased relative to the baseline switch cost found in Experiment 1 and even reversed (see Figure 3). These results indicate that effect of switching at the ensemble level dominated the effect of switching at the element level. As described in the introduction, there are at least two different ways to interpret this finding. One is that both the ensemble-level effect and the element-level effect existed and contributed roughly additively to task performance, but the ensemble-level effect was stronger than the element-level effect. A second interpretation is that the presence of a task ensemble somehow eliminated the element-level effect; that is, there was no longer a cost of switching between cognitive operations from one task element to the next. Because the design in Experiments 2–5 pitted the element-level and ensemble-level variables against each other (see Table 1), we cannot use the data to distinguish between these two interpretations.

In Experiment 6, we tested these two possible interpretations by eliminating the confounding of the element-level and ensemble-level variables. Specifically, we directly measured the contribution of the element-level switch effect, holding the ensemble-level task condition constant (by using ensemble repetitions only). To achieve this goal, we used the same task types—magnitude (Type A) and parity (Type B)—as in Experiments 1–5, but simply added a third element (Task 3) to the ensemble. The sequence of task types within an ensemble were repeated throughout each block. Consider the task sequence of AAB-AAB-AAB, and so on, versus the task sequence of ABA-ABA-ABA, and so on. Task 1 was always an element-level task switch from Task 3 of the previous ensemble in the former task sequence (AAB) but was always an element-level task repetition in the latter task sequence (ABA). Note that the task sequences underlying the element-level switch condition (e.g., AAB-AAB-AAB. . .) are essentially identical to those of the element-level repetition condition (e.g., ABA-ABA-ABA. . .)—the only difference is the location of the hyphens indicating the ensemble boundary.

We noted earlier that the ensemble-level switch cost in Experiments 2–5 could be due to two different components: the need to switch between ensembles or the need to maintain readiness for two different ensembles at the same time. Because Experiment 6 used the ensemble repetition condition only, neither of these two components should have any effect.

To provide stimulus support for a task ensemble, we adopted a similar approach to that of the fixed-order condition of Experiment 5: The three task elements of each trial were presented within a single semicircular frame divided into three sections (which we call *slices*), separated in time by a short RSI (0 ms in this experiment). As in Experiments 1–5, the RSI leading up to Task 1 was 1,500 ms.

If the element-level effect is still present, then a substantial switch cost, perhaps as large as that in Experiment 1 (238 ms), should be observed. However, if the element-level effect is eliminated by the ensemble of Tasks 1–3, then no switch cost on Task 1 should be observed.

Method

Participants. Twenty participants from the same participant pool as the participants in Experiments 1–5 were in this experiment. None had participated in Experiments 1–5.

Apparatus, stimuli, and procedure. The apparatus, stimuli, and procedure were the same as in Experiment 5, except as noted. A semicircular frame with three slices (left, middle, and right) was displayed in the center and remained on the screen throughout the whole block. The three elements for each task sequence (Task 1, Task 2, and Task 3) were presented in the left, middle, and right slices, respectively, separated by a 0-ms RSI. The feedback for the three task elements appeared immediately after the response to Task 3 was made and lasted for 300 ms. Task 1 for next ensemble appeared in the left slice 1,200 ms after the offset of the feedback message. Thus, the total RSI between trials was 1,500 ms, as in Experiments 1–5.

We used two different task sequences (AAB-AAB. . . and BBA-BBA. . .) for the element-level switch condition and two different task sequences (ABA-ABA. . . and BAB-BAB. . .) for the element-level repetition condition. To minimize the number of times participants changed ensembles, half of the participants received only the task sequences of AAB and ABA (in separate blocks), whereas the other half of the participants received only the task sequences of BBA and BAB. The order of these task sequences within a session was counterbalanced across participants. Participants first performed one practice block of 96 trials and four experimental blocks of 48 trials of one task sequence, followed by one practice block of 48 trials and four experimental blocks of 48 trials of the other task sequence.

Results and Discussion

Mean RT and PE for the three elements in each task sequence are shown in Table 3. The data analyses were performed in a similar way as in Experiments 1–5. Less than 0.2% of responses were omitted because they fell outside the RT cutoffs. The main purpose of Experiment 6 was to determine whether the element-level switch cost on Task 1 still occurs when there is strong stimulus support for a task ensemble. The element-level switch cost on Task 1 was measured by comparing the element-level switch condition (e.g., AAB-AAB. . .) to the element-level repetition condition (e.g., ABA-ABA. . .). This effect was significant on RT, $F(1, 19) = 21.32$, $p < .01$, $MSE = 9{,}197$; mean RT for Task 1 was 952 ms in the repetition condition but only 812 ms in the switch condition (a −140 ms switch cost). Although the effect was in the opposite direction on PE (the switch cost was 0.01), it was small and not statistically significant, $F(1, 19) = 3.86$, $p > .05$, $MSE = 0.0007$. Even in follow-up data analysis excluding 3 participants who showed a relatively large switch cost on PE, results still showed a significant element-level effect on Task 1 RT (a switch benefit of 153 ms), $F(1, 16) = 19.79$, $p < .01$, $MSE = 10{,}020$, but not on Task 1 PE (a switch cost of 0), $F(1, 16) < 1.0$. Thus, the switch benefit on Task 1 RT observed in Experiment 6 was not due to a speed–accuracy trade-off. Furthermore, the switch benefit of 140 ms was significantly different from the switch cost of 238 ms observed in Experiment 1 (the control experiment), $F(1, 38) = 58.25$, $p < .01$, $MSE = 11{,}860$. The switch cost on PE (.01) was also numerically smaller than that observed in Experiment 1 (.03), although the difference was not statistically significant, $F(1, 38) < 1.0$. Thus, the lack of an element-level switch cost on Task 1 is consistent with the hypothesis that the presence of a task ensemble virtually eliminates the element-level effect.

Table 3
Mean Response Times (RTs; in Milliseconds) and Proportions of Errors (PEs) for Task 1, Task 2, and Task 3 as a Function of Task 1 Condition (Repetition or Switch From the Previous Task 3) in Experiment 6

Condition and sequence	Task 1		Task 2		Task 3	
	RT	PE	RT	PE	RT	PE
Task 1 repetition						
ABA	882	.02	922	.05	797	.06
BAB	1,021	.03	1,040	.07	1,001	.05
Task 1 switch						
AAB	758	.03	546	.01	769	.05
BBA	866	.05	629	.03	829	.08
Switch cost	−140	.01	393	.04		

Note. The switch cost on Task 1 is measured by subtracting performance measures in the Task 1 repetition condition from those of the Task 1 switch condition. The switch cost on Task 2 is measured by subtracting Task 2 repetition performance (Task 1 switch condition) from Task 2 switch performance (Task 1 repetition condition). A = magnitude task; B = parity task.

One unanticipated finding from this experiment was that RTs for all three task elements (Task 1, Task 2, and Task 3) were slower in the element-level repetition conditions (e.g., ABA) than they were in the element-level switch conditions (e.g., AAB). It is especially notable that 100 ms of slowing occurred for Task 3, which was an element-level switch in both conditions (see Table 3); this slowing was significant, $F(1, 19) = 8.08$, $p < .05$, $MSE = 12{,}443$. Thus, these data suggest that the ABA (and BAB) ensembles are somehow more difficult to perform than the AAB (and BBA) ensembles. Why should one type of ensemble be more difficult than the other, even though the ensemble type always repeated? Following De Jong (1995), we proposed that the preparation for an ensemble involves preparation for Task 1 as well for the subsequent switch(es). Consequently, the preparatory control structure prior to an AAB ensemble might contain preparation for performing Type A (twice) and the subsequent switch to Type B. In contrast, the preparatory control structure for the ABA ensemble might contain the preparation for Type A, the subsequent switch to Type B, followed by another switch back to Type A. In other words, the AAB ensemble requires preparation for just one switch, whereas the ABA ensemble requires preparation for two switches. Thus, consistent with De Jong's model of ensemble preparation, the data from Experiment 6 suggest that a critical factor determining performance within an ensemble is the number of switches that must be prepared.

The existence of an ensemble-level difficulty effect in this experiment suggests that caution is required in interpreting the −140 ms switch cost on Task 1. Assuming an ensemble difficulty effect of about 100 ms (as suggested by Task 3 performance) on Task 1, the adjusted estimate for the element-level switch cost on Task 1 will be −40 ms. Thus, even after making this correction, there is still no evidence for an element-level switch effect on Task 1. Nevertheless, it would obviously be desirable to measure the switch cost without contamination from the ensemble-level difficulty effect. We designed Experiment 7 with this purpose in mind.

Although there was no evidence for an element-level switch cost between ensembles (from Task 3 to the next Task 1), it is important to note that Task 2 RT was in fact 393 ms slower when it was an element-level switch from Task 1 than when it was an element-level repetition (see Table 3), $F(1, 19) = 91.70$, $p < .01$, $MSE = 16{,}856$. Even adjusting for an estimated ensemble-level difficulty effect of 100 ms, the switch cost on Task 2 would still be about 293 ms. This effect size is similar to that observed in traditional task-switching studies, with a flat task structure, when the RSI is short (see Rogers & Monsell, 1995). Thus, it appears that the hierarchical task organization had no effect on the element-level switch cost within an ensemble. We will discuss the implications of these findings in the General Discussion section.

Experiment 7

Experiment 7 was designed to replicate the results of Experiment 6 without contamination from the ensemble-level difficulty effect. In this experiment, half of the participants received a task sequence of "AB?" throughout the entire experiment, where the "?" indicates that Task 3 (Type A or Type B) was not revealed until participants had responded to Task 2. For the purpose of counterbalancing, the other half of the participants received the repeating task sequence "BA?" throughout the whole experiment. As in the Experiment 6, this method allows us to measure the element-level effect from Task 3 to the next Task 1. Consider, as an example, the task sequence ABB-ABA-ABA. . . and so on. Following the ensemble ABB, the next Task 1 (Type A) is an element-level switch, whereas following the ensemble ABA, the next Task 1 (Type A) is an element-level repetition. Regardless of whether Task 1 was an element-level repetition or an element-level switch, the complexity of the ensemble preparatory control structure for the upcoming trial was always the same (e.g., prepare for Type A and the subsequent switch to Type B). Thus, this design allowed us to measure the element-level effect without contamination from the ensemble-level difficulty effect.

Method

Participants. Twenty participants from the same participant pool as the participants in Experiments 1–6 were in this experiment. None had participated in Experiments 1–6.

Apparatus, stimuli, and procedure. The apparatus, stimuli, and procedure were the same as in Experiment 6, except as noted. Half of the participants received only the task sequence of AB?, whereas the other half of the participants received only the task sequence of BA?. A consequence of this design is that, for a given participant, Task 1 was the same type (either Type A or Type B) throughout the experiment, as was Task 2. In the AB? sequence, for example, Task 1 was always Type A (the magnitude task), and Task 2 was always Type B (the parity task).

As in Experiment 6, the left and middle slices of the semicircular object were always colored blue or black, depending on the task sequence. The right slice (Task 3), however, contained a white cross-hatched pattern. Immediately after participants had responded to Task 2, this cross-hatched pattern was removed to reveal the color of the slice, which indicated what task type needed to be performed. The feedback for the three task elements appeared immediately after the response to Task 3 was made and lasted for 300 ms. The white cross-hatched pattern then immediately reappeared in the right slice. After 1,200 ms, the stimulus for Task 1 of the next ensemble appeared in the left slice. Thus, the total RSI leading up to Task 1 was 1,500 ms, just as in Experiments 1–6.

Results and Discussion

Mean RT and PE for the three elements within an ensemble are shown in Table 4 as a function of Task 1 condition (repetition versus switch from the previous Task 3) and Task 3 condition (repetition versus switch from Task 2). The data analyses were performed in a similar way as in Experiments 1–6. Less than 0.3% of responses were omitted because they fell outside the RT cutoffs. The element-level switch effect on Task 1 was not significant on RT, $F(1, 19) = 2.72$, $p = .12$, $MSE = 11{,}515$, or PE, $F(1, 19) < 1.0$. Mean RT on Task 1 was 1,002 ms in the Task 1 repetition condition and 962 ms in the Task 1 switch condition (a nonsignificant −40 ms switch cost). PE on Task 1 was .03 in the repetition condition and .04 in the switch condition (a nonsignificant .01 switch cost). These results suggest that there was little or no element-level switch cost, replicating Experiment 6. The lack of an element-level switch cost is consistent with the hypothesis that the presence of a task ensemble somehow eliminates the element-level effect.

Although there was no element-level switch cost between ensembles, there was once again a large switch cost between elements within an ensemble. Task 3 RT was in fact 308 ms slower when it was a switch from Task 2 than when it was a repetition, $F(1, 19) = 82.84$, $p < .01$, $MSE = 22{,}868$. The switch cost on Task 3 was also evident on PE, $F(1, 19) = 17.61$, $p < .01$, $MSE = 0.0014$; PE on Task 3 was .06 higher when it was a switch from Task 2 than when it was a repetition (see Table 4). The switch costs within an ensemble were about as large as those observed in traditional task-switching studies with a flat task structure. Thus, these data further support the conclusion that the hierarchical task organization has no effect on the element-level switch cost within an ensemble. We will consider the implications of this dissociation between the switch costs observed between and within ensembles in the General Discussion section.

General Discussion

Task switching is typically associated with a substantial time cost. Residual switch costs have been observed even with a long

Table 4
Mean Response Times (RTs; in Milliseconds) and Proportions of Errors (PEs) for Task 1, Task 2, and Task 3 as a Function of Task 1 Condition (Repetition and Switch From the Previous Task 3) and Task 3 Condition (Repetition and Switch From Task 2) in Experiment 7

Condition and sequence	Task 1		Task 2		Task 3	
	RT	PE	RT	PE	RT	PE
Task 1 repetition						
Task 3 repetition	1,022	.04	1,003	.07	843	.06
Task 3 switch	982	.02	1,002	.08	1,092	.11
Task 1 switch						
Task 3 repetition	968	.04	953	.07	714	.04
Task 3 switch	956	.04	968	.07	1,081	.11
Switch cost	−40	.01			308	.06

Note. The switch cost on Task 1 is measured by subtracting performance measures in the Task 1 repetition condition from those of the Task 1 switch condition. The switch cost on Task 3 is measured by subtracting Task 3 repetition performance from Task 3 switch performance.

RSI (e.g., Goschke, 2000; Meiran, 1996; Rogers & Monsell, 1995), predictable task-switch sequences (e.g., Tornay & Milán, 2001), task cuing (e.g., Koch, 2001; Logan & Bundesen, 2003), and practice (e.g., Meiran, 1996). Thus, even when task type is known and ample preparation time is provided, reconfiguration for a task switch appears to be incomplete. Rogers and Monsell (1995) argued that this residual switch cost occurs because task-set reconfiguration can be completed only after the stimulus has been presented. Logan and Gordon (2001) proposed a quantitative model along the same lines, in which switch costs reflect the time required to transmit task-set parameters from working memory to the subordinate processor. There are differences between these and other task-switching theories, however; as Gopher et al. (2000) summarized, “what all authors seem to agree on is that whatever factors are involved, their influence is not amenable to voluntary, advanced preparation” (p. 311).

In this study, we examined how the switch cost is affected by hierarchical task organization. Experiment 1 provided a baseline against which to compare the results of subsequent experiments. Because we used a constant 1,500-ms RSI between task elements, there was no stimulus support for an ensemble (i.e., the task structure was flat). We obtained a substantial switch cost of 238 ms, replicating traditional task-switching studies with dual-affordance stimuli.

To evaluate how hierarchical task organization affects switch costs, we manipulated stimulus factors in Experiments 2–5 that should, according to common sense, encourage participants to form a hierarchical task organization. In Experiment 2, we increased the temporal grouping of Task 1 and Task 2 by shortening the RSI between them to only 300 ms (keeping the RSI leading up to Task 1 at 1,500 ms). The relative switch cost on Task 1 RT was reduced from 238 ms in Experiment 1 to only 120 ms in Experiment 2. In addition to the RSI reduction, we spatially grouped Task 1 and Task 2 in Experiment 3 by presenting them within a common frame. The relative switch cost on Task 1 RT was further reduced to a nonsignificant 53 ms. In Experiment 4, we enhanced the stimulus support for ensemble repetition (in the fixed-order condition) by presenting the ensemble (e.g., AB) in the same location on every trial. Interestingly, the relative switch cost declined to −79 ms, indicating that the ensemble-level effect dominated the element-level effect. In Experiment 5, we presented the same two colored quadrants for both the fixed- and alternating-order conditions on the screen throughout each block to eliminate potential confounds in Experiment 4 (see above). We observed a −99 ms switch cost on Task 1, replicating the finding of Experiment 4. In sum, we observed a dramatic decline in the relative switch cost on Task 1 (from 238 ms to −99 ms) as the stimulus support for an ensemble of Task 1 and Task 2 increased (see Figure 3).

If one considers switching only at the level of elementary tasks, as in traditional treatments of task switching, it is difficult to explain the present results. One important assertion of many traditional task-switching theories is that advance preparation is inherently incomplete. Consequently, the obvious expectation would have been much slower performance when Task 1 was an element-level switch (e.g., AB-AB. . .) than when it was an element-level repetition (e.g., AB-BA. . .) across Experiments 1–5. Experiments 4 and 5, however, showed the opposite pattern. Hence, the hypothesis that a stimulus-triggered reconfiguration is always needed for element-level task switches but not for element-level task repetitions, by itself, cannot explain the present findings.

Evidence for Hierarchical Task Organization

The switch costs showed a clear and strong trend across Experiments 1–5: The more steps that were taken to increase the stimulus support for a task ensemble, the smaller the relative switch cost became (see Figure 3). In fact, the relative switch cost even reversed (from 238 ms in Experiment 1 to −99 ms in Experiment 5). The net change in the switch cost was 337 ms, suggesting that (a) the temporal and/or spatial grouping led participants to form a higher level mental task organization (i.e., at the ensemble level) and that (b) the ensemble level had a profound impact on performance.

Although our main results suggest that participants formed a higher level task organization corresponding to the ensemble level, it is worthwhile to ask whether there is any converging evidence for this hypothesis. As noted above, De Jong (1995) and Luria and Meiran (2003) have provided evidence for a higher level task organization in a paradigm similar to ours (except that their PRP design involved temporal overlap between the processing of Task 1 and Task 2). They concluded that participants prepared in advance for Task 1 and the subsequent switch to Task 2 prior to each trial (i.e., an ensemble). One might therefore expect Task 1 RT to increase when there was an ensemble of Task 1 and Task 2, because instead of preparing for just Task 1, participants also prepared for a subsequent task switch. Cross-experiment comparisons showed that Task 1 RT was in fact significantly slower in Experiments 2–5 than in Experiment 1 (see Figure 4), $Fs(1, 38) \geq 12.18$, $ps < .01$, $MSEs \leq 131{,}173$.

Furthermore, the hypothesis that participants formed an ensemble-level task organization predicts that ensemble repetition should not only benefit Task 1 but that it should also benefit Task 2. Thus, Task 2 RT should be faster in the fixed-order than in the alternating-order conditions. In contrast, if task elements are prepared individually, then performance on Task 2 should be roughly equivalent in both conditions. From Figure 4, it can be seen that Task 2 was in fact faster in the fixed-order than in the alternating-order conditions in Experiments 3–5, $Fs(1, 19) \geq 7.11$, $ps < .05$, $MSEs \leq 93{,}544$. This ensemble repetition benefit on Task 2 provides further evidence for an ensemble-level mental task organization.

Another interesting finding, consistent with the hypothesis that participants formed an ensemble-level task organization, concerns the effects of RSI on Task 2 (which was always an element-level task switch in Experiments 1–5). Even though the RSI leading up to Task 2 was only 300 ms in Experiments 4 and 5, mean Task 2 RT in the fixed-order condition was comparable to that in Experiment 1, in which the RSI was 1,500 ms. In other words, when the ensemble repeated in Experiments 4 and 5, there was no apparent cost of dramatically decreasing the RSI. This finding suggests that ensemble preparation prior to Task 1 can largely compensate for the lack of preparation time immediately prior to Task 2.

Did the Element-Level Effect Vanish?

The results of Experiments 2–5 suggest that the ensemble-level effect can dominate the element-level effect. Because the experimental design pitted the ensemble-level effect against the element-level effect (see Table 1), there are two obvious candidate explanations for the results. One is that the element level still has a substantial effect on Task 1 (possibly as large as the 238-ms effect

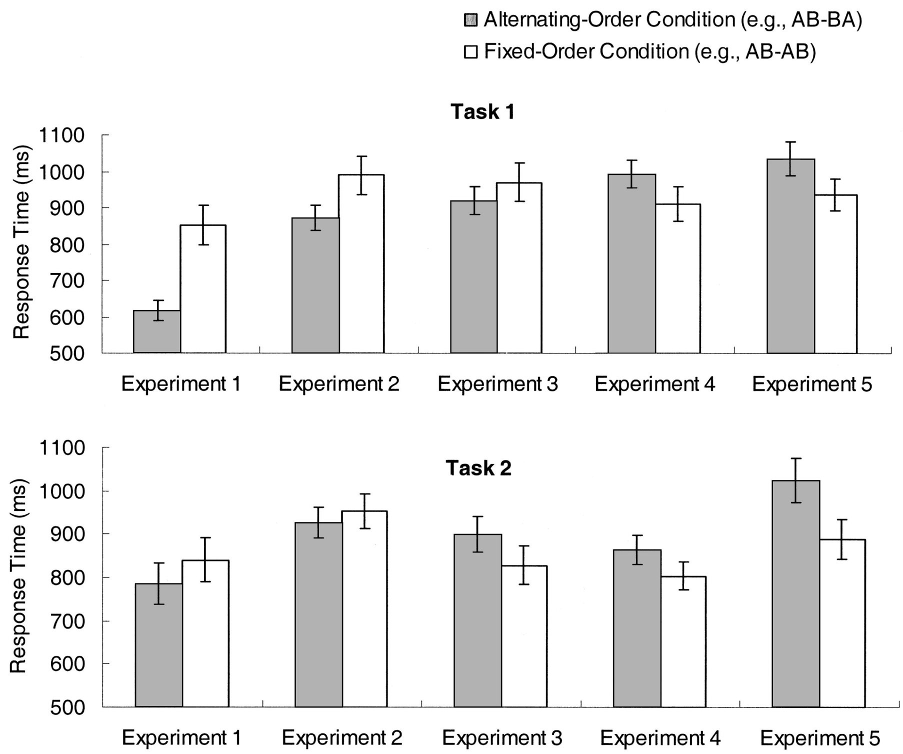

Figure 4. Response times for Task 1 and Task 2 as a function of order condition (fixed order and alternating order) in Experiments 1–5. Error bars represent the standard error of the mean.

observed in Experiment 1), but this effect was counteracted by even larger ensemble effects working in the opposite direction. Another explanation is that the hierarchical task organization strongly suppressed the element-level effect, so that there was no cost of switching cognitive operations from one individual task element to the next.

Experiment 6 was designed to test these two alternative hypotheses by isolating the contribution of the element-level effect from that of the ensemble-level effect. To accomplish this goal, we presented three temporally and spatially grouped task elements on each trial. The ensemble of three task elements repeated throughout every block (e.g., ABA-ABA. . .), so that all conditions involved ensemble repetition. Depending on the task sequence within an ensemble, Task 1 could be either an element-level repetition (e.g., ABA-ABA. . .) or an element-level switch (e.g., AAB-AAB. . .). This experiment revealed a −140-ms element-level switch cost on Task 1. We argued that this result might be due in large part to an ensemble-level difficulty effect (i.e., the ensemble ABA involved more switches than the ensemble AAB, and thus required a more complicated preparatory control structure). Adjusting for this ensemble-level difficulty effect (estimated to be about 100 ms), the actual element-level switch cost on Task 1 was estimated to be −40 ms. Thus, little or no element-level effect was evident when there was strong stimulus support for task ensembles.

The purpose of Experiment 7 was to replicate Experiment 6 without contamination from the ensemble-level difficult effect. Participants received a repeating task sequence such as “AB?,” where the “?” indicates that the type of Task 3 (A or B) was not revealed until they had responded to Task 2. Depending on the identity of Task 3, the subsequent Task 1 could be either a repetition or a switch at the element level. Because the same ensemble repeated throughout the whole experiment (e.g., AB?-AB?-AB?-. . .), ensemble-level difficulty was presumably constant across the Task 1 repetition and switch conditions. Results showed a nonsignificant −40 ms element-level switch cost on Task 1. The absence of the element-level switch cost on Task 1 RT in both Experiments 6 and 7 is consistent with the hypothesis that hierarchical task organization strongly suppresses or even eliminates the element-level effect.

Although there was no evidence for an element-level effect between ensembles in Experiments 6 and 7, there was still a large element-level effect within an ensemble. For instance, there was a roughly 300-ms element-level switch cost between Task 1 and Task 2 in Experiment 6 and between Task 2 and Task 3 in Experiment 7. Thus, a strong ensemble-level task organization appears to eliminate the element-level switch effect between but not within ensembles. Before discussing how to account for this dissociation, we first consider possible explanations for the lack of a switch cost between elements belonging to different ensembles.

What Happened to the Element-Level Effect?

There are two obvious ways in which the element-level effect between ensembles could be eliminated: (a) The repetition condition could lose its benefit, or (b) the switch condition could lose its cost. In other words, when there is an ensemble of multiple task elements, element-level repetitions could become much slower, or element-level switches could become much faster. When we compared Experiment 1 (flat task structure) to Experiments 4 and 5 (hierarchical task structure), we found little change in Task 1 RT for element-level switches, $Fs(1, 38) \leq 1.41$, $ps > .05$, $MSEs \leq 104{,}185$. In contrast, there was a sharp increase in Task 1 RT for element-level repetitions (see Figure 4), $Fs(1, 38) \geq 57.87$, $ps < .01$, $MSEs \leq 60{,}876$. These data therefore suggest that the benefit of repeating the cognitive operation from one task element to the next is fragile and can be eliminated when there is a hierarchical task organization.

Further evidence for the fragility of the element-level task repetition benefit comes from a recent study by Logan and Bundesen (2003; Experiment 3). As in the present study, their participants performed parity (odd or even) and magnitude (high or low) tasks on digit stimuli. A valid task cue was presented in the center of the screen, followed some time later (0–900 ms) by the digit. The task cue could be either a name cue (e.g., "Parity" or "Magnitude") or a mapping cue (e.g., "Odd–Even" or "High–Low"). Thus, it was possible for the task type to repeat even when the cues were different. Although task repetition with identical cues (e.g., "Parity," then "Parity") produced a large repetition benefit (about 150 ms), task repetition with different cues (e.g., "Parity," then "Odd–Even") produced only a very small benefit. Thus, their experiment provides further evidence that the benefit of repeating cognitive operations between task elements is fragile.

Implications for Theories of Executive Task Control

The present findings lead to several conclusions regarding executive task control. First, hierarchical mental task organization does occur, even with very simple noncontingent, nonoverlapping tasks. Thus, it is reasonable to believe that hierarchical task organizations arise frequently in the real world as well as in the laboratory. Second, hierarchical task organization has strong effects on performance: (a) Switching between task ensembles results in a cost, and (b) the transition from one task ensemble to the next (regardless of whether it is an ensemble switch or repetition) virtually eliminates the element-level switch cost. This surprising elimination of the element-level switch cost appears to be due to the loss of the benefit of repeating cognitive operations from one task element to the next. Traditional task-switching models, which focus on the time cost associated with changing the cognitive operations (or S-R mapping rules) from one task element to the next, provide no obvious explanation for the present findings.

To succeed, any candidate account must consider the interaction among different levels of the task hierarchy. The only existing candidate account we know of is the hierarchical switching model proposed by Kleinsorge and colleagues (e.g., Kleinsorge & Heuer, 1999; Kleinsorge, Heuer, & Schmidtke, 2001). Their model was based on a study of switch costs among several hierarchically related task types. They used a spatial judgment and a numerical judgment, each of which had a version with a compatible S-R mapping and a version with an incompatible S-R mapping. They examined whether the dimensional organization among these task types (with a judgment level, a mapping level, and a response level) would be reflected in the switch costs. They found interactions between switch effects at these different levels. For instance, switch costs were actually smaller when both the judgment and mapping switched compared to when only the judgment switched. To account for these results, they proposed that (a) participants represent tasks as several hierarchical levels (with judgment at the top of the hierarchy, the mapping in the middle, and the response at the bottom); (b) there are switch costs at all levels of the hierarchy; and (c) when participants switch one level, they automatically switch all lower levels as well. For example, if the judgment level (the highest level) switches, then the mapping and response levels will also be switched. If those lower levels did not need to be switched, then they will need to be switched back, resulting in an additional time cost.

Even though the hierarchy discussed by Kleinsorge and colleagues (e.g., Kleinsorge & Heuer, 1999; Kleinsorge et al., 2001) concerns task types (or the task space), whereas the hierarchies discussed in the present study concern task elements, it is nevertheless possible to apply the principles underlying their model to the present study. Specifically, one could propose that switching at the ensemble level automatically produces a switch at the element level (but not vice versa). On Task 1 of the fixed-order condition (e.g., AB-AB), only a switch at the element level is required. On Task 1 of the alternating-order condition (AB-BA), however, a switch at the ensemble level is required. Because the ensemble switch (by hypothesis) automatically causes an unnecessary element-level switch, a further element-level switch will be required. Thus, the fixed-order condition involves only an element-level switch, but the alternating condition involves both an ensemble-level switch and a compensatory element-level switch. Thus, this hierarchical model can explain why Task 1 was slow in the alternating-order condition of Experiments 4 and 5 (i.e., the negative relative switch cost). On the other hand, this model also predicts that when the higher level does not switch, as in our Experiments 6 and 7, there should be a substantial cost of switching at the lower level (e.g., Kleinsorge & Heuer, 1999, found a switch cost of roughly 200 ms when only the mapping level switched). Contrary to this prediction, we found little or no effect of element-level switching on Task 1 in both Experiments 6 and 7. Thus, this version of the hierarchical switching model cannot easily account for the present results. In the next section, we propose an alternative account.

A Dual-Route Model of Executive Task Control

In the following dual-route model of executive task control, we propose that there are two different processing routes that can be used to select a response: conditional and unconditional.[3] Element-level task repetitions can, under ideal circumstances, use the unconditional processing route, which involves a simple reapplication of the mental state that carried over from the previous trial. Element-level task switches cannot use this unconditional processing route because the previous mental state was for a different task type. Consequently, they must instead rely on the conditional processing route, which requires supervisory control of task set (e.g., loading a new S-R mapping rule) before or during task

[3] The terms conditional and unconditional routes should not be confused with the terms *automatic* and *controlled processing,* which have often been used in the literature on attention (especially visual search). According to Schneider and Shiffrin (1977), automatic processing occurs "without subject control, without stressing the capacity limitations of the system, and without necessarily demanding attention," whereas controlled processing "requires attention, is capacity-limited," and "is controlled by the subject" (p. 1). Although our conditional route may involve controlled processing, we do not necessarily imply that our unconditional route has all the properties ascribed to automatic processing.

execution (Hübner et al., 2001; Logan & Gordon, 2001; Norman & Shallice, 1986; Rogers & Monsell, 1995).

We assume that the conditional route is slower than the unconditional route for one or more of the following reasons: (a) Participants occasionally neglect to complete the reconfiguration for the new task type (De Jong, 2000); (b) the configuration cannot be completed in advance of the stimulus onset, so an online reconfiguration operation is needed (Rogers & Monsell, 1995); or (c) even if the task-set reconfiguration is completed, the conditional route still cannot reach the level of efficiency that the unconditional route achieves as a result of task repetition (Ruthruff, Remington, & Johnston, 2001). In traditional task-switching studies with flat task structures, element-level repetitions use the fast, unconditional route, whereas element-level switches use the slow, conditional route; thus, a switch cost is obtained.

By definition, the unconditional route involves the reapplication of the previous mental state, and thus is incompatible with any attempts at deliberate task-set configuration (advance preparation). If a task-element repetition condition does, for some reason, rely on the deliberate, conditional processing route, then the repetition benefit will be lost. Consider the condition of Logan and Bundesen's (2003) explicit task-cuing paradigm, in which the task type repeated but the task cue changed. Because it is not immediately clear to participants that the task type will repeat, they might rely on the deliberate, conditional route. Thus, the dual-route model can easily explain why little repetition benefit was observed in that condition. In our paradigm, in which the task organization is hierarchical, people may deliberately prepare for all of the elements belonging to the upcoming ensemble. Following De Jong (1995) and Luria and Meiran (2003), we argue that this advance preparation involves preparation for Task 1 and any subsequent switch(es) within the ensemble. Because this advance preparation involves multiple task elements, it would be impossible to simply reapply the mental state from the previous task element (even when the task type repeats). Consequently, the conditional route would be used, and the benefit of task repetition would be lost.

Although both element-level Task 1 repetitions and switches use the conditional route, they will not necessarily be equally fast. The reason is that the speed of the conditional route is likely to be sensitive to the degree of advance preparation. When the ensemble repeats (as in our fixed-order condition; e.g., AB-AB. . .), the preparatory control structure repeats, and thus the degree of advance preparation for Task 1 should be relatively high. In contrast, when the ensemble switches (as in our alternating-order condition; e.g., AB-BA. . .), the preparatory control structure switches, and thus the degree of advance preparation for Task 1 should be relatively low. Accordingly, Task 1 RT should be somewhat faster in the fixed-order condition than in the alternating-order condition, as observed in Experiments 4 and 5 (see Figure 4). In Experiment 6 the ensemble always repeated, but the ensemble preparation might have been more complicated in the Task 1 repetition condition (ABA), which required preparation for two task switches, than in the Task 1 switch condition (AAB), which required preparation for just one task switch. Assuming that greater ensemble complexity results in a lower degree of advance preparation, the dual-route model can easily explain why mean RT was slower in the Task 1 repetition condition than in the Task 1 switch conditions (e.g., 140 ms slower for Task 1 and 100 ms slower for Task 3).

Although the dual-route model predicts no benefit of repeating elementary cognitive operations between ensembles, it does predict a benefit of repeating elementary cognitive operations within an ensemble. The reason is that advance preparation is presumed to occur prior to an ensemble but not within an ensemble. Thus, task element repetitions within an ensemble can still use the unconditional route, producing a large repetition benefit (similar to that observed in a flat task structure), as obtained in the present Experiments 6 and 7. To conclude, this dual-route model of task switching is straightforward, yet it provides a satisfactory account for our findings as well as the main findings from traditional task-switching studies. Nevertheless, this model is post hoc and needs to be tested further.

Related Findings

The present results suggest that people sometimes form control structures covering multiple task elements (i.e., an ensemble) at the same time. Evidence for a higher-level control structure also comes from studies of sequence learning (e.g., Cohen, Ivry, & Keele, 1990; Keele, Cohen, & Ivry, 1990; Koch & Hoffmann, 2000). These studies typically involve only one task type, but the sequence of responses across trials contains repeated patterns. Results suggest that the response sequence is likely to be coded as a series of chunks that are easy to carry out (e.g., Cohen et al., 1990), and this hierarchical response representation can be affected by manipulations of temporal and spatial grouping (e.g., Koch & Hoffmann, 2000; Stadler, 1993). Our study also contained repeated patterns (e.g., AB-AB or AB-BA), but the pattern repeated at the task type level, not at the response level (in fact, the response pattern was completely random). In addition, the repeated task sequence was not the major factor leading to the hierarchical task organization in Experiments 2–7. Indeed, we hypothesized that the mental task organization was not hierarchical in Experiment 1, even though it contained exactly the same repeating task sequence as in Experiments 2–5. Rather, the key factor leading to the hierarchical task organization in our study appears to have been the temporal and/or spatial grouping of adjacent task elements.

Summary

In the present study, we extended previous task-switching studies by examining situations where there was a strong hierarchical organization of task elements. When we temporally and/or spatially grouped task elements, participants appeared to form a mental representation covering the entire set of elements (i.e., the ensemble). In this hierarchical task structure, we found virtually no cost of task switching between elements belonging to different ensembles (or, more precisely, there was no benefit of task repetition). However, we found the usual, large cost of task switching between elements belonging to the same ensemble. We argue that the benefit of repeating cognitive operations between task elements is fragile and can be eliminated by advance preparation for an upcoming task ensemble. To provide a more detailed explanation of the present findings, we proposed a dual-route model (including a conditional route and a unconditional route). In sum, this study provides an important step toward an understanding of executive control in both flat and hierarchical task structures.

References

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch

(Eds.), *Attention and performance XV* (pp. 421–452). Cambridge, MA: MIT Press.

Allport, A., & Wylie, G. (2000). Task switching, stimulus-response bindings, and negative priming. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 35–70). Cambridge, MA: MIT Press.

Carlson, R. A., & Lundy, D. H. (1992). Consistency and restructuring in learning cognitive procedural sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 127–141.

Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 17–30.

De Jong, R. (1995). The role of preparation in overlapping-task performance. *Quarterly Journal of Experimental Psychology, 48,* 2–25.

De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 357–376). Cambridge, MA: MIT Press.

E-Prime [Computer software]. (2001). Pittsburgh, PA: Psychology Software Tools.

Fagot, C. (1994). *Chronometric investigations of task switching.* Unpublished doctoral dissertation, University of California, San Diego. Retrieved April 2, 2002, from http://www.pashler.com/manuscripts.html

Gopher, D., Armony, L., & Greenshpan, Y. (2000). Switching tasks and attention policies. *Journal of Experimental Psychology: General, 129,* 308–339.

Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task set switching. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 331–355). Cambridge, MA: MIT Press.

Hübner, R., Futterer, T., & Steinhauser, M. (2001). On attentional control as source of residual shift costs: Evidence from two-component task shifts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 640–653.

Jersild, A. T. (1927). Mental set and shift. *Archives of Psychology, 14*(81, Whole No. 89).

Keele, S. W., Cohen, A., & Ivry, R. I. (1990). Motor programs: Concepts and issues. In M. Jeannerod & S. Kornblum (Eds.), *Attention and performance XIII* (pp. 77–110). Hillsdale, NJ: Erlbaum.

Kleinsorge, T., & Heuer, H. (1999). Hierarchical switching in a multidimensional task space. *Psychological Research, 62,* 300–312.

Kleinsorge, T., Heuer, H., & Schmidtke, V. (2001). Task-set reconfiguration with binary and three-valued task dimensions. *Psychological Research, 65,* 192–201.

Koch, I. (2001). Automatic and intentional activation of task sets. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 1474–1486.

Koch, I., & Hoffmann, J. (2000). Patterns, chunks, and hierarchies in serial reaction-time tasks. *Psychological Research, 63,* 22–35.

Lien, M.-C., Schweickert, R., & Proctor, R. W. (2003). Task-switching and response correspondence in the psychological refractory period paradigm. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 692–712.

Logan, G. D., & Bundesen, C. (2003). Clever Homunculus: Is there an endogenous act of control in the explicit task cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance, 29,* 575–599.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review, 108,* 393–434.

Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed blocks of trials. *Acta Psychologica, 94,* 145–188.

Luria, R., & Meiran, N. (2003). Online order control in the psychological refractory period paradigm. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 556–574.

Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1124–1140.

Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1423–1442.

Meiran, N. (2000). Reconfiguration of stimulus task sets and response task sets during task switching. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 377–399). Cambridge, MA: MIT Press.

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

Rock, I., & Palmer, S. (1990). The legacy of Gestalt psychology. *Scientific American, 263,* 84–90.

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General, 124,* 207–231.

Ruthruff, E., Remington, R. W., & Johnston, J. C. (2001). Switching between simple cognitive tasks: The interaction of top-down and bottom-up factors. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 1404–1419.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84,* 1–66.

Spector, A., & Biederman, I. (1976). Mental set and mental shift revisited. *American Journal of Psychology, 89,* 669–679.

Stadler, M. A. (1993). Implicit learning: Questions inspired by Hebb (1961). *Memory & Cognition, 21,* 819–827.

Tornay, F. J., & Milán, E. G. (2001). A more complete task-set reconfiguration in random than in predictable task switch. *Quarterly Journal of Experimental Psychology, 54,* 785–803.

Ulrich, R., & Miller, J. (1994). Effects of truncation on reaction time analysis. *Journal of Experimental Psychology: General, 123,* 34–80.

Wenger, J. L., & Carlson, R. A. (1995). Learning and the coordination of sequential information. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 170–182.

Received March 11, 2003
Revision received October 29, 2003
Accepted November 3, 2003 ■