Interruption-recovery training transfers to novel tasks

Winston E. Jones (winston.jones@psllab.org)
Jarrod Moss (jarrod.moss@msstate.edu)
Department of Psychology, PO Box 6161
Mississippi State, MS 39762 USA

Abstract

Interruption interference is a significant decrease in performance that follows task interruption. This interference is often studied using a primary and interrupting task pair. Evidence suggests that interruption interference can be reduced through practice by exposing individuals to many interruptions. However, the evidence that this skill transfers beyond the tasks being trained is less clear. In particular, these practice effects may only occur when the same interrupting/primary task pairs are involved. A transfer paradigm was implemented to assess the transfer of interruption-recovery skill. Participants in separate conditions performed either the Tower of Hanoi or the Tower of London primary task during a training block and a transfer block of trials. Trials were interrupted by another task, and the primary measure was the time taken to resume the interrupted task. Significantly lower resumption times at the beginning of the transfer block lead to the conclusion that interruption recovery skill can be transferred to a novel task.

Keywords: interruptions; skill acquisition; transfer

Introduction

Interruptions involve the suspension of task performance in order to perform a second, interrupting task. Handling interruptions has been found to result in robust costs to both speed and accuracy when compared to non-interrupted performance (e.g., Hodgetts & Jones, 2006; Trafton et al., 2003). Therefore, along with research that has sought to identify the sources of interruption interference, research has also explored possible means for mitigating interruption interference by improving recovery from interruption via training (e.g., Cades et al., 2011).

Research examining interruption interference has often used the Memory for Goals theory to explain interference effects (Altmann & Trafton, 2002). In this theory, the likelihood of successful interruption recovery depends on the level of activation of the interrupted task's goal-state information in memory. This theory is implemented with the ACT-R architecture (Anderson, 2007), which includes a long-term declarative memory in which activation of memories decays with time unless strengthened by rehearsal or retrieval. If the activation of a memory drops below the retrieval threshold, then retrieval will fail unless further activated by a cue in the environment that spreads activation to that memory. Level of activation in memory affects the speed and accuracy of retrieval. Therefore, the resumption lag (i.e., the time taken to resume a primary task following the end of an interruption) is affected by the amount of goal state rehearsal prior to or during an interruption because this rehearsal raises the goal's level of activation of in memory.

Initial evidence that interruption handling could be improved through practice used a three-block paradigm where individuals performed a task while being interrupted multiple times either with a warning prior to the interruption or with no warning (Trafton et al., 2003). Resumption lag decreased from earlier to later blocks, providing evidence that interruption recovery improved with practice. However, this effect only occurred in the no-warning condition, so the practice effect might be an adaptation only to events that did not include warning time for preparation.

The possibility of improving interruption recovery through training was further addressed in a later study using a similar primary/interrupting task paradigm in which resumption lag times were compared across individuals who had varying exposure to interruptions in a three-session AAA AAB ABB design where A blocks received no interruptions and B blocks received interruptions (Cades, Trafton, & Boehm-Davis, 2006). Individuals improved their resumption lag with more exposure to interruptions. These results supported the possibility of a general interruption recovery skill by which interruption handling can be trained.

To examine the possibility that interruption recovery skill was independent of task skill, a similar three-session study design was used that manipulated the onset of a new interrupting task rather than the onset of interruptions (i.e., an AAA, AAB, ABB design where the letters now correspond to different interrupting tasks). Resumption lag did not improve and temporarily got worse during the first block including the new interrupting task (Cades et al., 2011). These results support the conclusion that resumption lag improvements might be confined to a practiced primary/interrupting task pair. However, it is still not clear whether improvements in resumption lag might transfer to a new primary task if the interrupting task were held constant.

There are theoretical reasons why one might find transfer in the case of a similar interrupting task being paired with a new primary task. Singley and Anderson (1989) describe transfer as being dependent on the amount of overlapping production rules between two tasks. More recently, Taatgen (2013) has proposed that transfer between two tasks can occur as long as they share common Primitive Information Processing Elements (PRIMs) at a finer level of detail than the production rule. In the case of interruption recovery skill, the PRIMs responsible for interrupting task execution or those responsible for preparing to engage in a common interrupting task might provide an account of transfer between new primary tasks with the same interrupting task. Furthermore, the interrupting or primary tasks may not have to be identical, but instead they may simply have to share enough PRIMs. Because many interruptions involve the encoding of interrupted task-state information and a switch between tasks, this account indicates that it may be possible to see a more task-general interruption recovery process (i.e., task state or goal rehearsal).

Prior to engaging in a detailed task analysis, including understanding the overlapping PRIMs, it seemed prudent to first explore whether transfer could be observed in an interruption paradigm in which transfer of resumption lag improvements could be examined. The current study therefore tests the general interruption recovery skill hypothesis in which transfer of resumption lag improvements is assessed when both the interrupting task and primary task differ between training and transfer. In addition, because previous studies concluded that recovery skill is task-pair specific, the study design included an assessment of whether transfer occurred when the interrupting task stays the same but the primary task changes (i.e., a common-interrupting-task hypothesis).

A counterbalanced experiment design (e.g., A-B B-A) was implemented to evaluate transfer using between-group comparisons (Singley & Anderson, 1989). In these analyses, Task A trial blocks (i.e., yellow cells in Table 1) are compared to each other and Task B trial blocks (i.e., red cells in Table 1) are compared to each other across conditions. Using the example shown in Table 1, transfer from task B to task A is examined by comparing the performance at the beginning of the transfer block of the B-A condition to see if it is better than performance at the beginning of the training block of the A-B condition.

Table 1: Design of experiment to assess transfer

Condition	Training	Transfer
A-B	Task A	Task B
B-A	Task B	Task A

Method

Participants

Participants were 178 Mississippi State University students who received partial course credit. Eight participants were excluded from analyses because they did not perform the interrupting tasks with at least 70% accuracy, and three were removed for non-compliance during the experiment.

Tasks and Materials

Two primary tasks (Tower of London and Tower of Hanoi) and two interrupting tasks were used in this study, all of which were implemented and presented via computer. In the primary tasks, participants had the goal of moving a set of five discs one-at-a-time until the arrangement on the three pegs matched a target image that was displayed at the top of the screen as shown in Figure 1 for the Tower of London.

The Tower of Hanoi was selected as a task in this study because previous research has used it to explore the behavioral effects on goal/subgoal retrieval after various durations of suspension (Anderson & Douglass, 2001), and it has also been used in interruption research (Altmann & Trafton, 2002). It has been used in this prior research due to the subgoaling required to complete the task. Therefore, during an interruption, there is significant task-related information that, if recalled, would speed up resuming the task. In Tower of Hanoi trials, participants had to move each of five different-sized discs into the goal configuration with the constraints that only one disc could be moved at a time and larger discs could not be placed on top of smaller discs.

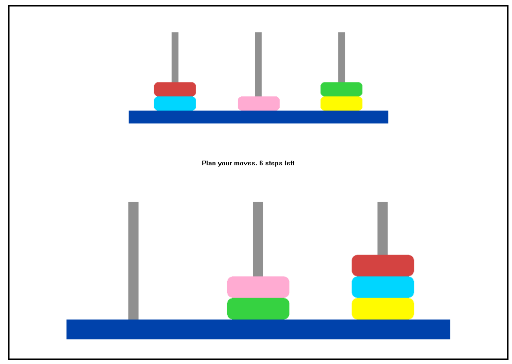

Figure 1: Tower of London task trial screen. The goal configuration is at the top, the current state is presented at the bottom of the screen. The Tower of Hanoi task interface was similar except that the discs were of different sizes.

The Tower of London was selected because of its similarity (i.e., visual features, response method, and task goal) to the Tower of Hanoi as well as its use in previous interruption studies from which this version of the task was adapted (Hodgetts & Jones, 2003, 2006). It also has a significant subgoaling component. In Tower of London trials, all discs were the same size, but no more than three discs were allowed on a single peg at a time.

There were two possible interrupting tasks: addition and matching. The addition task required solving randomly-generated two-digit addition problems. Participants typed responses for these trials into a text box using the keyboard as shown in Figure 2. The matching task presented visual matching trials from Thurstone's Perceptual Speed test (Thurstone & Jeffrey, 1984). Participants solved these problems by selecting a figure that matched the target figure using the keyboard. Each trial had a target figure on the left and a set of five potential matches as shown in Figure 3.

Design and Procedure

Each participant was randomly assigned into one of six experimental conditions. The six conditions in this study were designed in order to address the primary research question regarding the ability to transfer interruption recovery skill across both primary and interrupting task changes. The order of tasks in each condition is shown in Table 2 such that there are three separate A-B pairs that can be assessed as explained earlier in Table 1. Participants performed two blocks of 30 puzzle trials over the course of the experiment session (i.e., training and transfer blocks) with the primary puzzle task switching between blocks. The order of puzzle presentation was counterbalanced across groups so that half received Tower of Hanoi trials in the first block and Tower of London trials in the second block while the other half received Tower of London trials in the first block and Tower of Hanoi trials in the second block. This design sets up the necessary between-subject transfer comparison between conditions 1 and 2 because the transfer block in condition 2 is the same primary task as the training block in condition 1 and vice versa. Conditions 1 and 2 serve as a replication of a prior unpublished study in which the interrupting task was not manipulated (Jones & Moss, 2013). The method of the prior study was identical to that of conditions 1 and 2 but did not include conditions 3-6. Conditions 3-6 have a different interrupting task between their first and second blocks. The order of the interrupting tasks was also counterbalanced across these four conditions.

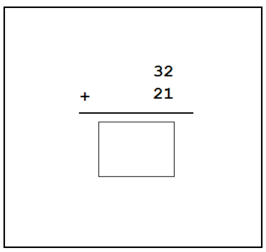

Figure 2: Example two-digit addition trial.

Figure 3: Example visual matching trial.

Instructions to participants described the Tower of London and Tower of Hanoi tasks as “puzzles”, and participants were told that all puzzle trials could be completed in exactly six moves. A six-second delay at the beginning of each puzzle trial, along with a written prompt, instructed participants to plan their moves in order to minimize task difficulty. This planning prompt was used to encourage use of a planned sequence of subgoals. Move attempts that did not follow the six-move solution path resulted in a written error message, and the attempt did not move the selected disc. Thus, participants were not allowed to deviate from the minimum solution path.

Interruptions occurred without warning in predetermined puzzle trials throughout the experiment. Only one interruption could occur during any given puzzle trial. Interruptions occurred following critical moves within the six-move sequence of each trial. Critical moves were identified as those that involved a 2- or 3-move sequence to move blocking disc(s) out of the way in order to move a disc to its goal location. The critical move was the final move in such a sequence. All interruptions occurred at these moves. In other words, participants were interrupted just after they had moved the blocking disc(s) out of the way but before they could move the blocked disc. Recalling this goal information should speed up resumption of the task relative to having to reconstruct the next move from the interface. Critical moves always occurred after either the second, third, or fourth move in a six-move solution sequence. All interruptions lasted approximately 20 s, during which time participants completed as many trials of the interrupting task as possible. The interrupting task was designed to terminate when a component trial was completed and the duration of the interruption had been at least 20 s.

All participants had 31 total interruptions over the course of the session. Tower of Hanoi blocks had 17 interruptions and Tower of London blocks 14 interruptions. Originally, the intent was to have 15 interruptions across both tasks, but a programming error led to more interruptions being presented in the Tower of Hanoi task.

Written instructions for the primary and interrupting task were provided just prior to each block of trials. Following the written instructions, two practice trials were given to familiarize participants with the response interface. No interruptions occurred during practice trials. Tower of London and Tower of Hanoi practice trials were simpler one- and three-move trials meant to demonstrate both how to select and move discs and how match the target disc arrangement displayed at the top of the screen. Addition and matching practice trials were identical in complexity to those seen during interruptions.

Results

Resumption Lag Improvements

Consistent with prior research, resumption lag was the primary dependent measure to assess the impact of interruptions. Resumption lag was measured as the amount of time between the end of the final trial of the interrupting task and the first mouse click in the primary task.

The first set of analyses was performed on conditions 1 and 2. This analysis assesses the common-interrupting-task

Table 2: Tasks performed in each condition and resumption lag results. Standard errors are in parentheses.

Condition	N	Training Block			Transfer Block		
		Primary Task	Interrupting Task	Resumption Lag (s)	Primary Task	Interrupting Task	Resumption Lag (s)
1	31	Tower of Hanoi	Addition	5.65 (0.30)	Tower of London	Addition	2.40 (0.18)
2	26	Tower of London	Addition	3.47 (0.19)	Tower of Hanoi	Addition	3.83 (0.33)
3	28	Tower of Hanoi	Addition	5.67 (0.37)	Tower of London	Matching	2.80 (0.23)
4	28	Tower of London	Addition	3.34 (0.22)	Tower of Hanoi	Matching	4.00 (0.44)
5	28	Tower of Hanoi	Matching	5.45 (0.44)	Tower of London	Addition	2.80 (0.22)
6	26	Tower of London	Matching	3.46 (0.24)	Tower of Hanoi	Addition	3.66 (0.39)

transfer hypothesis because in these conditions the primary task changed but the interrupting task did not. Because it was anticipated that the largest effect of transfer would be seen early in the second block due to skill acquisition, only the first four interruption events for the training and transfer tasks were used. Separate two-way ANOVAs for each primary task with a within-subject factor of interruption serial order (first through fourth interruption) and a between-subject factor of block (training, transfer) were run. In other words, for the Tower of Hanoi ANOVA, the resumption lag for the first four interruption events in condition 1 (training block) were compared to the first four interruption events in condition 2 (transfer block). Resumption lags were smaller during transfer in both the Tower of Hanoi, $F(1, 55) = 16.17$, $p < 0.001$, and the Tower of London, $F(1, 55) = 16.62$, $p < 0.001$. These results replicate prior unpublished results using the same method for these two conditions (Jones & Moss, 2013) showing that this cross-task transfer effect is robust.

The second set of analyses involving conditions 3-6 was performed across condition pairs that manipulated both primary and interrupting tasks. Within a condition, neither the primary task nor the interrupting task was the same between training and transfer blocks. Conditions 3 and 6 were compared and conditions 4 and 5 were compared in the same manner as condition 1 and 2 above. For the Tower of Hanoi, transfer resumption lags in condition 6 were smaller than training in condition 3, $F(1, 52) = 14.02$, $p < 0.001$. Likewise, the comparison of resumption lag times between the Tower of Hanoi blocks of conditions 4 and 5 were significantly lower in the transfer block, $F(1, 54) = 5.30$, $p = 0.025$. This result is consistent with the hypothesis that interruption recovery skill transfers even when neither the primary nor interrupting task is identical. The comparison of the Tower of London resumption lags in conditions 4 and 5, however, only showed a significant difference at the $\alpha = .10$ level, $F(1, 54) = 3.12$, $p = 0.084$. Similarly the Tower of London resumption lags between conditions 3 and 6 were also only significant at the $\alpha = .10$ level, $F(1, 52) = 3.92$, $p = 0.053$, indicating that recovery skill transfer might still be occurring between these trial blocks but with a smaller effect size.

In order to assess whether these marginally significant results in the Tower of London task were due to lack of power, data from the Tower of London training blocks in conditions 4 and 6 were combined (i.e., collapsing across differences in the interrupting task). Data from the Tower of London transfer blocks in conditions 3 and 5 were also combined. As a prerequisite for this collapse across interrupting task, the data were examined using a set of t-tests to see whether the interrupting task did indeed have an effect on the resumption lag. For example, the resumption lags from condition 3 training were compared to the resumption lags from condition 5 training (differing in interrupting task). None of the t-tests revealed a significant difference in resumption lag due to interrupting task, so the data were collapsed across the interrupting task differences. Resumption lag data from the combined conditions were analyzed using the same ANOVA design as the earlier analyses. Resumption lag times in the combined Tower of London transfer blocks were significantly lower than those for the combined training block, $F(1, 108) = 7.11$, $p = 0.009$. This finding is consistent with the marginally significant results being due to lack of power, and it supports the hypothesis that resumption lag improvements transfer to new interrupting/primary task pairs.

Primary and Interrupting Task Results

Table 3 shows that interruption task accuracy was high across conditions and did not differ significantly between conditions. Solution times for the primary tasks were analyzed to evaluate whether significant skill in performing the disc-moving tasks was transferred between tasks. Such transfer could potentially be related to interruption skill transfer. Solution time data for the Tower of Hanoi and Tower of London were compared using the same block-condition pairings used to analyze the resumption lag data. All solution time analyses were performed using only non-interrupted trial data, because interrupted-trial solution time improvements could reflect improvements in interruption recovery skill. Because solution time was more reflective of individual problem difficulty than a serial-order skill acquisition effect, all non-interrupted trials within a block were averaged and *t*-tests were used.

In conditions 1 and 2, training on one primary task did not lead to lower times in the transfer block for either the Tower of Hanoi, $t(55) = 1.21$, or the Tower of London, $t(55) = 0.35$. For conditions 3 and 6, Tower of Hanoi times did differ, $t(52) = 3.18$, $p = 0.003$, but Tower of London times did not, $t(52) = -0.61$. For conditions 4 and 5, Tower of Hanoi times were lower in the transfer block, $t(54) = -2.35$, $p = 0.02$, but Tower of London times were not, $t(54) = 1.25$. In summary, primary task transfer was limited and not as consistent as the resumption lag transfer effects.

Table 3: Non-interrupted primary task performance and interruption task accuracy. Standard error in parentheses.

Condition	Training Block				Transfer Block			
	Primary Task	Primary Task Sol. Time (s)	Int. Task	Interruption Accuracy	Primary Task	Primary Task Sol. Time (s)	Int. Task	Interruption Accuracy
1	Hanoi	19.02 (1.05)	Addition	89.67% (0.84)	London	13.51 (0.87)	Addition	88.92% (0.91)
2	London	14.74 (0.99)	Addition	89.32% (1.01)	Hanoi	17.25 (0.99)	Addition	89.36% (0.88)
3	Hanoi	20.19 (1.25)	Addition	90.85% (1.12)	London	14.19 (0.90)	Matching	91.17% (1.31)
4	London	15.12 (0.79)	Addition	91.77% (1.22)	Hanoi	16.10 (0.82)	Matching	92.98% (0.71)
5	Hanoi	19.47 (1.19)	Matching	93.52% (0.62)	London	13.66 (0.86)	Addition	91.29% (1.28)
6	London	14.97 (0.90)	Matching	92.40% (1.20)	Hanoi	15.52 (0.72)	Addition	90.71% (1.16)

Discussion

Cross-task transfer of interruption recovery skill was explored in the current study. Previous evidence had suggested that resumption lag times decreased with practice dealing with interruptions, but there was no evidence that interruption recovery skill could transfer to new tasks. The current study used a paradigm designed to maximize detection of transfer. Two hypotheses were explored. First, the common-interrupting-task hypothesis predicted transfer between primary tasks if the interrupting task was the same in training and transfer blocks. This hypothesis was supported by the results showing transfer in conditions 1 and 2 of the current study. One explanation of these results could be that it was skill on the interrupting task that was transferring, but this explanation is not consistent with the results testing the second hypothesis.

The second hypothesis examined was that cross-task transfer would occur when the interrupting task was different from that trained. The results of the current study's comparison of conditions 3-6 support this hypothesis. Resumption lags were found to be significantly lower at the beginning of the transfer block when participants had already practiced interruptions during an earlier trial block even if the primary task and interrupting task were different. The fact that this second hypothesis was supported in the current study means that there is an apparent discrepancy with those of Cades et al. (2011) who did not find continued improvements in resumption lag when the interrupting task was switched (keeping the same primary task). There are a number of differences between the current study and theirs that might explain the difference, including the complexity of the primary and interrupting tasks. However, a close examination of their Experiment 3 shows that they did not find that resumption lags get worse when the interrupting task was changed, but instead they found that resumption lags did not continue to show a practice-related decrease when the interrupting task was changed. The fact that resumption lags did not show an increase when the interrupting task was changed may be indicative of transfer occurring. Therefore, there may not be a discrepancy at all.

The current study used primary tasks with similar features to control for differences in complexity. However the data indicate that interruption recovery was easier in Tower of London trials than in Tower of Hanoi trials, resulting in generally lower resumption lags and lower effect sizes for resumption lag improvements in Tower of London trials. This effect size difference showed up in the transfer analyses for Tower of London in which a larger sample size was needed to find a statistically significant transfer effect (by collapsing across different interrupting tasks).

In terms of a theoretical understanding of the current results, theories of skill transfer (Singley & Anderson, 1989; Taatgen, 2013), provide mechanisms that might explain the results. There may be common PRIMs involved in recovering Tower of Hanoi and Tower of London task-related information. The number of common PRIMs in the recovery process would then be correlated with the amount of transfer. Based on the Memory for Goals theory, resumption from interruptions always involves the encoding and retrieval of goal-state information when suspending a task (Altmann & Trafton, 2002). Because the primary goal of this recovery procedure (i.e., retrieving an incomplete goal state) is consistent across any interruption event, then interruption recovery skill can be expected to be based on improvements in the ability to encode, maintain, and retrieve goal-state information at certain times before, during, and following an interruption. Because most interruptions involve this common recovery procedure, some degree of skill transfer might occur across interruption events.

At some abstract level, the basic strategy of encoding and retrieving goal-state information is common across all types of resumptions from an interruption. Therefore, the results of the current study, which show transfer of interruption recovery skill, support the concept of a general procedure for recovering from interruptions. However, even if there exist some task-general processes (i.e., productions) for recovering task-related information, these general processes would become specialized to specific tasks as skill is acquired (Anderson, 2007). Therefore it is likely that aspects of the primary and interrupting task become incorporated into such productions, and transfer then becomes partially a function of the similarity between tasks (similarity among primary tasks, interrupting tasks, or both).

It is also possible that characteristics of how the interruption occurs could affect transfer independently of the similarity of different interrupting tasks. For example, differences in resumption lag improvement between warned and unwarned interruptions (e.g., Trafton et al., 2003) can be explained by the opportunity to strengthen goal information in memory prior to the interruption. This warning time provides an opportunity to make primary task resumption easier. In this example, common PRIMs might exist with other interruptions that provide a warning period. Conversely, common PRIMs that improve the success rate of primary task resumption might be present in the recovery process for unwarned interruptions.

Individuals may also develop interruption-handling strategies that take advantage of characteristics of certain interruption types. One example of such a strategy has been suggested in prior research that found that participants who could choose when to switch to an interrupting task tended to delay switching until they reached states of lower processing demand on the primary task (Salvucci & Bogunovich, 2010). In a hypothetical model of this strategy, a production rule would evaluate primary task information in working memory at time of interruption onset to determine whether to take immediate action on the interrupting task or to defer it. This production rule might then transfer to other similar situations.

Similarly, interruptions with inherent structures that facilitate the recovery procedure (e.g., warned interruptions) would only provide minimal skill transfer to interruptions lacking such structure (e.g., unwarned interruptions). Future research is needed to assess how well transfer can occur between different types of interruptions and the differences in processing that separate interruptions into meaningful skill types. Computational models of interruption recovery procedures would be particularly useful for evaluating the skills that individuals acquire to handle different interruption types as well as how this skill might transfer. In particular, it should be fairly straightforward to implement models of the Tower of Hanoi and Tower of London within ACT-R because such models already exist for some versions of the architecture. Then, the PRIM theory could be used to examine predicted transfer in our tasks as well as how transfer would be affected by different experimental or training manipulations. This modeling endeavor is part of planned future research.

This study also has implications for the design of future interruption studies. Previous research has used primary and interrupting tasks with differing procedural and processing elements. These designs have probably stemmed from a common conceptualization that any task can be interrupted by any other task in a real-world setting. Although these designs are useful for evaluating naturalistic performance, studies attempting to examine interruption recovery transfer effects and the mechanisms that underlie recovery skill will likely require more detailed task comparisons in order to accurately evaluate how and when transfer occurs.

Some limitations in the design of the current study should be noted. The programming error that resulted in differing numbers of interruption events between the Tower of London and Tower of Hanoi tasks blocks meant that participants received varying amounts of practice depending on which primary task they received first, which could lead to differences in the amount of acquired skill. Given that Tower of London blocks received fewer interruptions and Tower of Hanoi transfer effects were stronger overall, however, this difference in number of interruptions might not have been great enough to seriously impact the data. Another limitation, however, was that different Tower of London and Tower of Hanoi trials were presented in the same order for all participants meaning that any problem-specific difficulty effects were confounded with skill acquisition effects across time. Randomizing the problem order in future studies would allow problem difficulty effects to be examined separately from skill acquisition within a block of problems.

Also, of course, the two primary tasks in this study were selected to be similar (i.e., both disc moving problems). Transfer in resumption might also be explained as simply due to the similarities in task structure or even similarities in the appearance of the task user interface. Additional studies are planned to assess these possibilities.

Overall, the findings of this study provide new insight into the nature of interruption handling skills. The finding of cross-task transfer in resumption lags following an interruption is novel. Further research is required to document how well, or poorly, interruption recovery skill can transfer to different interruption types as well as delineate the cognitive mechanisms that differentiate interruption events from one another in order to better understand the occurrence of interruption-handling skill transfer.

Acknowledgments

This research was supported by Office of Naval Research grant N000141310251 awarded to the second author. We would also like to acknowledge Adrian Lewis, Skylar Swindle, Caitlin Spencer, Philip Liu, Riley Gordon, and Chelsea Howd for their assistance with data collection.

References

Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39–83.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford Press.

Anderson, J. R., & Douglass, S. (2001). Tower of Hanoi: Evidence for the cost of goal retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(6), 1331–1346.

Cades, D. M., Boehm-Davis, D. A., Trafton, J. G., & Monk, C. A. (2011). Mitigating disruptive effects of interruptions through training: What needs to be practiced? *Journal of Experimental Psychology: Applied*, *17*(2), 97–109.

Cades, D. M., Trafton, J. G., & Boehm-Davis, D. A. (2006). Mitigating Disruptions: Can Resuming an Interrupted Task Be Trained? In *Proceedings of the Human Factors and Ergonomics Society* (Vol. 50, pp. 368–371).

Hodgetts, H. M., & Jones, D. M. (2003). Interruptions in the Tower of London task: Can Preparation Minimise Disruption? In *Proceedings of the Human Factors and Ergonomics Society* (Vol. 47, pp. 1000–1004).

Hodgetts, H. M., & Jones, D. M. (2006). Contextual Cues Aid Recovery from Interruption: The Role of Associative Activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(5), 1120 – 1132.

Jones, W. E., & Moss, J. (2013). *Decreasing the costs of interruptions: Interruption recovery as a trainable and transferable skill.* Poster presented at the meeting of the Psychonomics Society, Toronto, Canada.

Salvucci, D. D., & Bogunovich, P. (2010). Multitasking and Monotasking: The Effects of Mental Workload on Deferred Task Interruptions. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems* (pp. 85–88).

Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill*. Harvard University Press.

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, *120*(3), 439–471.

Thurstone, L. L., & Jeffrey, T. E. (1984). *Space Thinking (Flags)*. Rosemont, IL: London House.

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies*, *58*(5), 83–603.