Journal of Experimental Psychology:
Learning, Memory, and Cognition
2006, Vol. 32, No. 5, 1120–1132

Copyright 2006 by the American Psychological Association
0278-7393/06/$12.00 DOI: 10.1037/0278-7393.32.5.1120

Contextual Cues Aid Recovery From Interruption: The Role of Associative Activation

Helen M. Hodgetts and Dylan M. Jones
Cardiff University

A series of experiments introduced interruptions to the execution phase of simple Tower of London problems and found that the opportunity for preparation before the break in task reduced the time cost at resumption. Retrieval of the suspended goal was facilitated when participants were given the opportunity to encode retrieval cues during an "interruption lag" (the brief time before engaging in the interrupting task) but was impeded when these visual cues were subsequently altered following interruption. The results provide useful support for the goal-activation model (E. M. Altmann & G. J. Trafton, 2002), which assumes that context—at the points of both goal suspension and goal retrieval—is critical to efficient interruption recovery.

Keywords: interruption, Tower of London, goals, activation, memory

Goals are central to the organization and direction of human behavior, yet the processes involved in remembering and coordinating these intentions are not well understood. The classical view is that pending goals are stored in a stack and retrieved in a last-in, first-out manner, by which the right goal is always supplied at the right time (e.g., Anderson & Lebiere, 1998; Ernst & Newell, 1969; G. A. Miller, Galanter, & Pribram, 1960). However, this view is in more recent times challenged by a growing body of empirical research which demonstrates that—as with ordinary declarative memory elements—the success of goal retrieval may be dependent on processes of decay or interference in memory (Altmann & Trafton, 2002; Anderson & Douglass, 2001; Byrne & Bovair, 1997; Hodgetts & Jones, 2006). In the current work we seek to further investigate those mechanisms involved in goal-directed cognition using task interruption, a means with which to examine the formulation, retention, and execution of task goals. Unlike mere distractions (e.g., the onset of an unexpected sound), which can temporarily draw attention away from a primary task, interruptions require the distinct cessation and resumption of the ongoing task in order to perform a separate intervening activity. Thus, not only do interruptions cause an unexpected break in the cognitive focus of a task, they additionally require the management of multiple task goals. We use as a theoretical basis for this work the cognitive architecture of *adaptive control of thought—rational* (ACT–R; Anderson & Lebiere, 1998) and assess two models that derive from this framework. The models of Anderson and Douglass (2001) and Altmann and Trafton (2002) make particular predictions regarding the nature of goal memory, and specifically we address the question of whether goal retrieval may be aided by the encoding of retrieval cues prior to that goal's suspension.

Although the study of interruptions has a long history (e.g., Zeigarnik, 1927), much of the research has an applied focus and says little about the involvement of underlying cognitive mechanisms (e.g., Bainbridge, 1984; Kirmeyer, 1988). Even among those studies that have used an experimental approach, there is still a lack of consensus as to what makes interruptions disruptive (see McFarlane & Latorella, 2002, for a review). This may be due to the lack of a common theoretical approach as well as the use of disparate methodologies. The present work follows from our previous research on interrupted task performance (Hodgetts & Jones, 2006), in which we found that brief interruption to the execution phase of five-disk Tower of London (ToL) problems (Ward & Allport, 1997) incurred a time cost when the suspended goal was to be retrieved; furthermore, this time cost was exacerbated by both a longer interruption and one that was more complex. The current study adopts the same theoretical basis and methodology as this previous work and as such aids the continuity and coherence of the literature. The theoretical focus for the research is on two models that derive from the cognitive architecture of ACT–R: The models of Altmann and Trafton (2002) and Anderson and Douglass (2001) were originally developed to simulate behavior in the Tower of Hanoi (ToH) task, but the processes of goal suspension and retrieval during problem solving can easily be applied to those processes operating when a task is interrupted and subsequently resumed. Both models eschew ACT–R's traditional goal-stack construct and instead suggest that goals may be subject to the same limitations as ordinary declarative memory elements. In the current experiments we focus on the critical points of goal suspension and goal retrieval and assess how these models may speak to the issues of advance preparation and the use of cues to aid retrieval.

Models of Goal Memory

In a study examining the costs associated with the storage and retrieval of subgoals in 15-move ToH problems, Anderson and Douglass (2001) found that the cost of storing/suspending a sub-

Helen M. Hodgetts and Dylan M. Jones, School of Psychology, Cardiff University, Cardiff, United Kingdom.

The research reported here was supported by an Economic and Social Research Council (United Kingdom) postdoctoral fellowship awarded to Helen M. Hodgetts.

Correspondence concerning this article should be addressed to Helen M. Hodgetts, School of Psychology, Cardiff University, Cardiff, CF10 3AT United Kingdom. E-mail: hodgettshm@cardiff.ac.uk

goal accounted for very little of the variance in regression analyses. From this, Anderson and Douglass inferred that participants do not choose to engage in any costly preparatory processes at the point of goal suspension, even if this means that the suspended goal may eventually be forgotten. They therefore largely discounted the idea of advance preparation, finding instead that the processes operating at goal retrieval were of more interest: Latency data showed that participants' actions were slower at points in the task when they were to retrieve goals, and slower still the longer ago that these goals were formulated. Furthermore, goals suspended for longer were also more likely to be forgotten, a finding that parallels the decay of items in declarative memory. In ACT–R's declarative memory system, the likelihood of retrieving a given item depends on its level of activation at that time. This activation is determined by two components: base-level activation, which increases the more an item is sampled but decays with disuse, and associative activation, a limited type of attentional activation that spreads among those items that are relevant to the current context (Anderson & Schooler, 1991). Anderson and Douglass (2001) suggested that just as base-level activation decreases over time in declarative memory, activation decay may also be responsible for forgetting in goal memory. They produced an ACT–R model that applies the base-level learning equation to goal memory and found that this was able to account well for the experimental data in their study. Furthermore, our previous research on interrupted task performance fits well with their model's concept of goal activation decay, as those goals suspended for longer were found to be more time consuming to retrieve (Hodgetts & Jones, 2006).

The goal-activation model (Altmann & Trafton, 2002) also proposes that goals decay when no longer actively involved in the governance of behavior, but Altmann and Trafton specified additionally a role for associative activation in goal memory. This role for associative activation creates differences between the two models in terms of the processes operating at goal suspension and resumption. In the Anderson and Douglass model, as with the traditional goal stack, these processes are relatively straightforward: Goals are suspended without any prior preparation, and—although subject to decay—they are retrieved in the order in which they were created without any interference or intrusions from other goals. In the goal-activation model, however, there is no predetermined last-in, first-out order for completion; rather, whichever goal is most active at a given time will be the one to govern behavior. The activation of a goal in this model is dependent on not only the length of time since that goal was last sampled (base-level activation) but also environmental factors (associative activation). Altmann and Trafton (2002) emphasized two cognitive constraints, strengthening and priming. To govern behavior, a new (or interrupting) goal must be repeatedly sampled or *strengthened*, a process that rapidly builds up base-level activation above that of other goals to overcome proactive interference. This base-level activation decreases once the goal is selected until eventually it dips below that of newer, more active goals. Subsequent retrieval of this goal is then aided by a process of *priming*: Cues in the mental or physical context boost associative activation of that goal, making it more active than other competing distractors (Figure 1).

Goal retrieval is not the only point on which the two models differ, as the priming constraint also has implications for the processes operating at goal suspension. That is, for successful

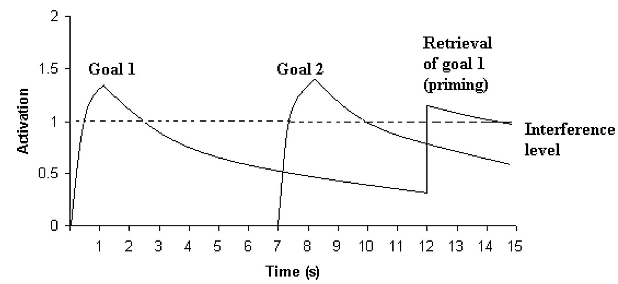

Figure 1. Priming to increase activation of an old goal above that of other competing goals at the interference level. Illustration adapted from Altmann and Trafton (2002).

priming at resumption, associative links must be formed between the cue and the target goal before it is suspended. Effective cues are formed by co-occurrence: They will be to some extent "obvious" so that the system must process them, they must be available in the mental or physical context both when the goal is suspended and when it is to be resumed, and they will prime the target goal and no or few others. On the basis of the goal-activation model, one would predict that the ease of task resumption would be dependent on an opportunity to encode associative cues before goal suspension and also the subsequent availability of these cues at the point of goal retrieval. As well as the prospective encoding of contextual cues, the goal-activation model suggests that another form of preparation may be possible: goal rehearsal. Rehearsal of the current goal before it is suspended will increase base-level activation so that it is later more easily reactivated in memory. Altmann and Trafton (2002) proposed that there is a window of opportunity for this strengthening of goals and encoding of cues just before the current goal is suspended and during the time when the new goal is being strengthened. In terms of an interruption, this opportunity is the interruption lag: the time between the signal for the change in tasks and the actual onset of the secondary activity.

The Interruption Lag

In general terms, it seems plausible that the time needed to reorient cognitively toward a task following interruption may be reduced if one is able to consolidate one's place in that task before it is suspended. Indeed, research conducted with office-style tasks indicates that workers do often engage in preparatory behaviors when a task is to be suspended (making notes, using markers, etc.) or if possible use the interruption lag to quickly finish off the current goal so that a convenient cognitive breakpoint can be reached (e.g., Burmistrov & Leonova, 2003). The current work examines preparatory processes at a more fine-grained level to test specific predictions rooted in cognitive theory.

We assess whether resumption of an interrupted task may be facilitated by a brief time lag before the participant engages in the interrupting activity. This is the prediction made by Altmann and Trafton (2002), whereas Anderson and Douglass (2001) would expect that an opportunity to prepare task goals will have no effect on retrieval on the grounds that participants will not choose to engage in such a strategy. Existing research shows that performance suffers less when participants have control over the interleaving of tasks, perhaps because of the opportunity to rehearse or consolidate primary task goals before engaging in the interruption (McFarlane, 2002). The benefit of rehearsal is further supported by a study of task alternation in which participants were to change every 5 s between a primary and a secondary task (Monk, Boehm-Davis, & Trafton, 2004). When the secondary task was unfilled rather than a tracking activity, resumption lags on the primary task were smaller by nearly half a second, a difference that the authors attribute to the opportunity to freely rehearse goals during the break in task. However, in some studies that encouraged participants to rehearse goals or to make notes during the interruption lag, performance actually suffered owing to increased interference (e.g., Clifford & Altmann, 2004; S. L. Miller, 2002).

Although a number of interruption studies have shown the benefit of an opportunity for rehearsing or strengthening task goals, few have specifically focused on the role of cues and prospective goal encoding during this time. It is this idea of encoding contextual cues that we focus on in the current work, a feature that is central to the goal-activation model but one that is not incorporated into the Anderson and Douglass (2001) framework. Empirical studies concerning the role of retrieval cues in interruption recovery are inconclusive. Verbal protocols taken during a computer game task indicated that an 8-s interruption lag was used for some retrospective rehearsal of the current state, but mainly participants engaged in prospective goal encoding (Trafton, Altmann, Brock, & Mintz, 2003). A further study using the same task manipulated interruption lag length (either 2, 4, 6, or 8 s in duration) and the availability of primary task information during this time (either visible or replaced with a blank screen). At the 8-s lag only, resumption times were quicker in the cue than in the no-cue condition. However, the data show that performance was not necessarily *improved* by the availability of primary task cues over a longer period, but rather the long lag with no primary task information *reduced* typical performance owing to boredom and perhaps the activation of irrelevant thoughts (Altmann & Trafton, 2004). The specific role of cues in facilitating task resumption is therefore unclear in this study. Given the ambiguity of the interruption literature and the discrepant predictions of the two ACT–R models, the role of preparatory processes in supporting memory for task goals is an issue that warrants further investigation.

The Current Experiments

The current series examines the role of the interruption lag with particular regard to associative priming and the encoding of cues. Altmann and Trafton (2002) proposed that cues linking the target goal to features in the current context may be encoded relatively automatically if time permits, simply by co-occurrence. Anderson and Douglass (2001), on the other hand, believe that if cues were to be encoded it would be a deliberate and effortful process and therefore one that participants would choose rarely to adopt. Unlike previous studies investigating the interruption lag, the following experiments do not use an explicit alert to signal the upcoming interruption on the grounds that this might capture participants' attention too much and perhaps cause a break in the cognitive focus of the task even before the onset of the secondary activity. If participants are attending to the appearance of a visual alert on the screen, then it is unlikely that they will be able to encode associative cues concurrently in an efficient manner.[1] Moreover, an opportunity for preparation afforded simply by the transition in tasks is fairly representative of many computer-initiated interruptions occurring in the office environment. The use of a specific visual warning and/or instructions to rehearse primary task goals would perhaps encourage participants to engage in interruption management strategies that they might not otherwise use in a nonexperimental setting. Instead, the presence of an interruption lag is manipulated by the transition in tasks (i.e., the participant begins the secondary task at his or her own pace, or there is a short pause between tasks). As Altmann and Trafton suggested that the processing of associative cues may occur relatively effortlessly, one might expect the ease of task resumption to be affected simply by the opportunity, and not necessarily the instruction, to engage in preparatory processes before the onset of the secondary task. As well as being of theoretical interest in terms of models of goal memory, the current work may also have some practical applications, for example, in terms of minimizing the cost of computer-initiated interruptions in the workplace through effective interface design.

We use a methodology similar to that of Hodgetts and Jones (2006), wherein brief, undemanding interruptions are incorporated into the execution phase of computerized five-disk ToL problems (Ward & Allport, 1997). Unlike the classic ToL (Shallice, 1982) or the ToH, the five-disk ToL task embodies fewer constraints: All disks are the same size so that any one can be placed on top of any other and each peg holds the same number of disks. Disks on the main display are to be moved one at a time from peg to peg using the mouse, until the configuration of disks on pegs exactly matches a goal state shown at the top of the screen (Figure 2). Participants are required to engage in two separate phases for each trial, planning and execution. Research suggests that participants can efficiently preplan up to two subgoals ahead at the beginning of a trial, and then execution of this plan is supported by a process of online monitoring and updating (Phillips, Wynn, McPherson, & Gilhooly, 2001). We hoped that for the six-move problems used in our experiments, participants would be able to plan a sequence of moves at the start of a trial and then execute this solution relatively continuously if uninterrupted.

On interrupted trials participants were required to complete a mood checklist, a relatively brief and undemanding secondary task but one that would occur unexpectedly during the execution phase. The time taken for participants to resume the ToL following interruption is used as an indication of the accessibility of goals and the ease of retrieval under different conditions. Also recorded are the number of trials in each condition that are not solved in the minimum number of moves. According to the goal-activation model, errors in goal retrieval may be more frequent in those conditions in which insufficient contextual cues render the suspended goal less active than competing distractors.

Experiment 1A

Experiment 1A tested whether allowing a brief interruption lag before the onset of the secondary task can facilitate task resump-

[1] In a short pilot study, the start of the interruption lag was signaled by a change in the background color of the screen. However, this procedure was abandoned when it became apparent that the change to the task environment hindered the processing of task-relevant cues and actually worsened performance.

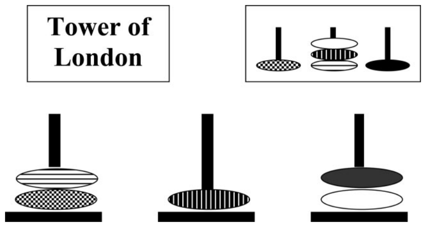

Figure 2. The five-disk Tower of London problem.

tion. The availability of cues was manipulated by having two testing sessions, one in which the display of the interrupting task immediately blocked the view of the ToL and another in which the ToL task remained visible in the background throughout the interruption period. According to Altmann and Trafton (2002), resumption times should be quicker in this latter condition as the less abrupt transition between tasks would allow a window of opportunity to encode retrieval cues and boost the activation of the current goal before fully engaging in the secondary task. Conversely, Anderson and Douglass (2001) would expect no effect of interruption type. They predict instead that participants will not engage in any rehearsal process before switching tasks, and as such the cost of retrieving the suspended goal would be the same in both conditions.

Method

Participants. Twenty-four students at Cardiff University received £5 (approximately U.S.$3.00) for their participation.

Apparatus and materials. The ToL computer program was written in Visual Basic 6.0 and run on a personal computer. The main display showed a ToL configuration with five different colored, equal-sized disks arranged on three pegs (Figure 2). In the top right-hand corner of the display was a 9 × 5-cm box containing a diagram of the goal state for that particular trial. Only disks on the main display could be moved, and this was done by clicking on buttons positioned below each of the three pegs. Clicking first on the peg holding the chosen disk displayed the words *FROM HERE* below it, and then clicking on the destination peg displayed the words *TO HERE*. The top disk on the chosen peg would then move across to its allotted position. No more than three disks could be held on a peg at one time, and attempting to do so would activate a *beep* sound to warn the participant that another move should be selected. It was hoped that this added constraint would further emphasize the need for thorough planning at the start of each trial. Also, this restricted the movement of disks such that for the particular six-move problems selected for this study, only one optimal solution path was possible when this constraint was in place, thus allowing a direct comparison between moves on perfect trials. A pop-up box appeared upon completion of each trial notifying participants how many moves had been taken. Clicking a button labeled *OK* would initiate the next trial and display a new start and goal configuration each time.

The interrupted problems each required six moves to solution, and interruption occurred upon completion of the third move. These trials all followed a similar structure: The first and fourth moves were the beginning of two-move sequences to place disks in their target location, whereas the remaining moves all placed disks directly in their goal positions (see Appendix A for ToL problems). Each interrupted problem was matched to a control trial, which was essentially the same problem with the same solution path but with the colors of the disks changed. Equivalent problems were always located at least eight trials apart. The rest of the trials were fillers requiring four, five, or six moves to completion.

Participants were interrupted on 6 out of 25 trials. The interruption was triggered by clicking on the *TO HERE* button that completed the third move. On interruption trials in Session A, the whole screen turned white and a 12 × 8-cm white box appeared in the center of the screen containing a series of six mood statements, positioned one below the other, describing a mood continuum (e.g., *extremely happy*, *fairly happy*, *slightly happy*, *slightly sad*, *fairly sad*, *extremely sad*). The participant was required to click with the mouse on the statement that best applied to his or her mood at that point and then click a button labeled *Continue*, positioned directly below the mood statements, to return to the ToL display. In Session B, interruption trials were different in that instead of covering the whole screen, the checklist appeared in the same-sized box in the top left-hand corner of the screen, therefore leaving the main ToL display and the goal state still clearly in view. In neither condition could any disks be moved during the time that the checklist was displayed; participants were to click the *Continue* button in order to reactivate the interrupted ToL problem and remove the mood assessment task from the screen. The checklists were different for each of the six interruptions and involved the following mood continuums: *happy–sad*, *tense–relaxed*, *bored–interested*, *angry–calm*, *tired–alert*, and *confident–unsure*. Resumption time was measured from clicking the *Continue* button on the checklist until completion of the next move (clicking the *TO HERE* button) on the ToL display.

Design. A repeated measures design was used so that each participant took part in two testing sessions, A and B, the order of which was counterbalanced. In Session A, the interrupting mood checklist would always cover the whole screen, and in Session B it would always appear in just the top left-hand corner. The two sessions were completed at least a day apart so that the manipulation was not so obvious to participants. Each session included six interrupted trials (Trials 4, 7, 12, 15, 19, and 25) and six matched no-interruption control trials (Trials 6, 9, 13, 16, 20, and 23).

Procedure. Participants were told that they would be taking part in a study to investigate problem solving and mood, and that this was to be assessed on two different days. At the beginning of the first session, a standardized instruction sheet was given that explained the aim of the ToL task and how the disks were to be moved on the screen. It was emphasized that the solution to each problem (up to six moves) should be planned thoroughly at the start of each trial. Participants were told they should aim to execute the solution continuously without pausing midtask to plan their next move. The instructions also warned participants that they would be asked to assess their mood at certain points in the experiment—they were to do this as quickly and accurately as possible and then continue with the main task as normal. Participants completed two practice trials in order to gain familiarity with the task, and more were given if the participant thought them necessary. The whole experimental session typically lasted half an hour. The participant returned on a later day to complete the second testing session. This began with the participant rereading the instruction sheet and completing two practice trials before beginning the session of experimental trials as before.

Results and Discussion

The main focus of these experiments was the speed with which participants execute the planned ToL sequence. Accuracy data were recorded (Appendix B), but in line with previous findings (Hodgetts & Jones, 2006) interruption of the ToL task did not appear to affect the accuracy of problem solving, and so this is not discussed further in this article. One participant who did not complete all of the 25 trials was removed from analyses. Move time data were recorded for each condition (Figure 3) and subjected to a 2 (interruption vs. control) × 2 (session: full-screen vs. corner-screen interruption) repeated measures ANOVA. Partial eta squared (η_p^2) was used as a measure of effect size. There was a significant main effect of interruption, with participants taking

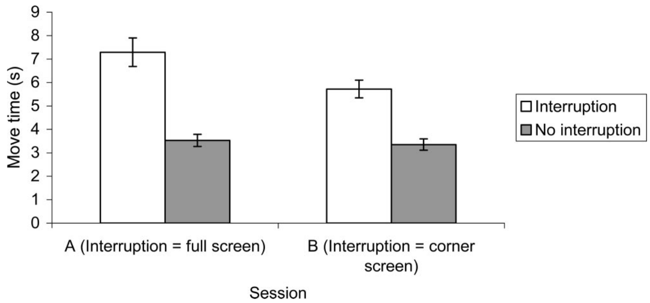

Figure 3. Mean time taken (s) to make the fourth move for both interruption and control trials in each session. Error bars show standard error.

significantly longer in both sessions to make their fourth move following interruption relative to when plan execution was continuous, $F(1, 22) = 60.97$, $MSE = 3.54$, $\eta_p^2 = .74$, $p < .001$. There was a main effect of interruption type, $F(1, 22) = 5.81$, $MSE = 3.07$, $\eta_p^2 = .21$, $p < .05$, such that move times were slower in Session A, when the mood checklist covered the whole screen, than in Session B, when it appeared just in the corner. There was also a significant interaction, $F(1, 22) = 4.86$, $MSE = 2.28$, $\eta_p^2 = .18$, $p < .05$, which shows, as would be expected, that there was no difference between sessions for move times in the control condition but that in the interruption condition participants resumed the primary task significantly faster when the mood checklist appeared in the corner of the screen rather than covering the whole display.

Resumption time data were analyzed separately for "perfect" trials, that is, those that were completed in the minimum of six moves. It is possible that resuming a task in which six moves were ultimately taken was different from resuming when the problem was eventually completed in seven or more moves. To test this possibility, the data were pooled across participants, and paired trials selected according to whether the participant had completed both the full-screen and corner-screen experimental trials in exactly six moves. With 23 participants each completing six interruption trials in each session, there were 138 possible paired data points. It was common for participants to take seven rather than six moves to solution, so only 58 of the 138 were perfect trial pairings. Paired *t* tests conducted between these two experimental conditions showed an effect of interruption type for perfect trials, $t(57) = 2.09$, $p < .05$; imperfect trials, $t(79) = 2.58$, $p < .05$; and the whole data set, $t(137) = 3.09$, $p < .05$. Regardless of whether the trial was completed in the minimum of six moves, participants were still quicker to resume the task when the interruption appeared in just the corner of the screen rather than when it obscured the whole screen.

These findings seem to fit well with the predictions of the goal-activation model, suggesting that the opportunity to prepare task goals before engaging in the secondary activity can reduce the subsequent time needed to resume the primary task. In the full-screen condition, the immediate onset of the interruption would have meant that there was little time to consolidate the to-be-suspended goal or to encode associative cues; as such, reactivation of that goal to a level above the interference threshold would have been a time-consuming process. When both tasks were visible on the screen at the same time, the transition between them was less abrupt. It could be assumed that this allowed time to engage in preparatory processes before turning attention wholly to the interrupting task, for example, by boosting the activation of the next planned move or by processing associative cues that would prime later retrieval of the suspended goal. Anderson and Douglass (2001) would have difficulty in accommodating these findings, because they believe that participants lack the motivation to rehearse and so no difference should be evident between conditions.

If participants are rehearsing to-be-suspended task goals in the corner-screen condition, then what form might this preparation take? At the start of each trial, participants are instructed to plan all moves, which they then execute in sequence with support from online monitoring processes. At the point of interruption, participants must unexpectedly suspend their planned next move and resume this intention after completion of the secondary task. We see the suspended "goal" in this task as the intention to move a particular disk to a particular peg, rather than the higher order goal of completing the problem. It is this association of a particular disk to a new target peg—their planned but thwarted move—that participants must retain over the course of the interruption and reactivate upon returning to the ToL display. Informal discussion with some participants after the experiment indicated that they were not using verbal information to remember their planned solution sequence (e.g., "move red disk from Peg 1 to Peg 3") but rather tried to remember the moves as a series of "spatial shifts" between pegs. It is therefore possible that consolidating a to-be-suspended goal may simply involve an eye movement to reaffirm the association of disk to peg rather than involving more intensive rehearsal strategies of the type that would be required in a serial recall task, for example.

The time spent on the interrupting task was recorded and analyzed. If participants were taking longer to complete the task in the corner-screen condition, then this may indicate that they were engaging in preparatory processes in addition to completing the checklist task. Mean interruption duration was 6.03 s ($SD = 2.47$) in the full-screen condition and 6.10 s ($SD = 2.00$) in the corner-screen condition, but a paired t test showed that the difference was nonsignificant, $t(22) = 0.16$. Although this measure does not seem to support the idea that participants take time to prepare, it does not necessarily follow that this proposal can be discounted. It is possible that participants were taking time to consolidate their place in the ToL task before switching in the corner-screen condition, but they may have then compensated for this time by working slightly more quickly on the checklist task. This would be in keeping with research that has shown workers to sometimes overcompensate for potential performance decline when interrupted (Zijlstra, Roe, Leonora, & Krediet, 1999). Another possibility is that the time difference between preparing and not preparing is very small in magnitude. If preparatory processes take only milliseconds (e.g., an eye movement), then perhaps a significant difference between conditions would be unlikely to emerge with the current sample size.

Previous studies have noted that participants become better at dealing with interruptions with practice (Trafton et al., 2003), so the resumption time data were analyzed according to order throughout the experiment. However, a repeated measures ANOVA showed that there was no effect of order of interruption either in Session A, $F(5, 110) = 1.22$, $MSE = 26.81$, $p = .31$, or in Session B, $F(5, 110) = 1.16$, $MSE = 7.28$, $p = .36$. Participants did not become more adept at coping with the interruptions, perhaps because the occurrence of the mood checklist was still relatively infrequent and unexpected (only 6 out of 25 trials).

Times taken to make the first move at the start of each trial (planning time) were recorded. If planning times were very similar to resumption times, then this would suggest that participants were replanning their moves following interruption. Mean planning times for the critical six-move trials were as follows: Session A (interrupted and matched control trials) = 19.13 s (*SD* = 10.16) and Session B (interrupted and matched control trials) = 20.02 s (*SD* = 12.81). Task resumption times were quite a lot quicker than the initial planning times, suggesting that some residual knowledge survives the interruption and that participants are actually retrieving old goals rather than simply planning anew.

Although it might be tempting to conclude at this point that the observed benefit of the corner-screen interruption was due to preparatory processes operating at the point of goal suspension (in line with the goal-activation model's predictions), the ToL task was of course visible throughout the whole of the interruption and not just during the transition in tasks. It is possible, then, that interruption recovery was quicker in the corner-screen condition not because of the availability of primary task cues at goal suspension but because of the availability of the task just before resumption. Or, perhaps participants' attention during the mood checklist was not truly taken away from the ToL while it was still in view, meaning that primary task goals were maintained at a higher level of activation throughout the interruption and were therefore quicker to retrieve.

It is difficult to speculate on this point from the existing literature, as the presence of the primary task during the interruption is not a factor that has been explicitly manipulated before. Indeed, in most computer-based interruption studies, the primary task is not visible during the secondary task, as the computer screen is cleared and replaced with the interrupting activity (e.g., Detweiler, Hess, & Phelps, 1994; Gillie & Broadbent, 1989; Hess & Detweiler, 1994). However, two experiments involving instant messenger interruptions, Cutrell, Czerwinski, and Horvitz (2000) and Czerwinski, Cutrell, and Horvitz (2000), retained the same "pop-up box" style, with the secondary task area positioned so that the primary task (searching through a list of book titles) was still largely in view. Both studies manipulated the influence of explicit retrieval aids—either a highlighter function (Cutrell et al., 2000) or a reminder of the search item throughout the task (Czerwinski et al., 2000)—but the benefits were small in magnitude. Perhaps task resumption was already aided by the fact that the primary task was available in the background throughout the interruption, so these cues to reorient the participant provided no additional benefit to task resumption. Using a data entry task, one study found that on-screen interruptions were more disruptive than telephone calls or walk-in visitors (Storch, 1992), a result that was attributed to task similarity, as both activities were computer based. However, another critical feature was that the on-screen pop-up interruption locked participants out of the main task, whereas this was still available during the other interruption scenarios; perhaps, therefore, performance suffers less when the participant is allowed more control over the integrating of the two tasks. From the current experiment we cannot be sure whether the benefit of the corner-screen interruption is due to the involvement of preparatory processes at the point of goal suspension or whether the presence of primary task cues throughout the secondary task was the critical feature. This is assessed in the following experiment.

Experiment 1B

In Experiment 1A, task resumption times were quicker when the mood checklist was displayed in the corner of the screen rather than when it covered the whole screen. Experiment 1B investigates this finding further by testing whether—as the goal-activation model might suggest—the first few seconds of the interruption are the critical ones from the point of view of having the ToL task in view. Experiment 1B included a third condition in which a full-screen interruption was preceded by a 2-s pause, during which time the main display froze and no disks could be moved. Although no explicit alert was given, this pause provided a brief time lag when participants would be attending to the main display but would not be able to advance the current problem state. This brief time period should be beneficial to participants by providing a critical opportunity to consolidate current task goals before the interruption ensues. It would therefore be predicted that task resumption times in the pause condition would be reduced relative to the immediate condition and would benefit performance in a similar way to that of the corner-screen interruption.

Method

Participants. Twenty-four Cardiff University students received £3 (approximately U.S.$3.00) for their participation.

Apparatus and materials. The same ToL task was used as in Experiment 1A, but the six interruptions received by the participant took the following form: two interruptions in which the mood checklist appeared in the top left-hand corner of the screen; two in which the checklist immediately covered the whole screen; and two in which the whole screen was covered but following a pause of 2 s. In this last condition, the main display froze for 2 s during which time none of the disks could be moved. The mood checklist then appeared, covering the whole display as it did in the second condition.

For this and the subsequent experiments, each new trial began with an enforced planning time of 8 s during which the word *PLAN* was displayed in the top left corner of the screen and no disks could be moved. After this time, the word *SOLVE* was displayed to signal to the participant that the program was now active. Participants were able to spend more than the 8 s to plan if they so wished, but an obligatory planning time was thought to discourage participants further from moving disks early in the trial before having formulated a full solution sequence.

Design. A repeated measures design was used in which participants completed two of each of the three types of interruption: full-screen interruption, corner-screen interruption, and full-screen interruption preceded by a 2-s pause. Interruptions always occurred on the same trials for all participants (4 and 12; 7 and 19; 15 and 25), but the type of interruption was counterbalanced across three groups in a partial Latin square design. There were also six matched control trials.

Procedure. There was one testing session of 25 trials, including six interruptions. The same basic procedure was used. Participants were told that they would be required to complete mood checklists at various points, but they were told nothing about the types of interruption that they would encounter.

Results and Discussion

Three outliers with resumption times exceeding three standard deviations from the mean were excluded from analyses. Times taken to make the fourth move were recorded for each condition (Table 1). A repeated measures analysis of variance (ANOVA) was conducted on the data, and owing to a violation of sphericity a Huynh–Feldt correction was used. A significant effect of interruption condition was obtained, $F(1.83, 36.58) = 10.76$, $MSE = 5.35$, $\eta_p^2 = .35$, $p < .0001$, and Fisher's post hoc comparisons assessed differences between conditions at the $p < .05$ level. Move times in the control condition were significantly quicker than move times in any of the interruption conditions, a finding that is in line with the previous experiment. Also, the finding of Experiment 1A was replicated: Participants were significantly quicker to resume the task when the interruption appeared in the corner of the screen than when it immediately covered the whole screen. In the critical interruption condition, in which the checklist covered the whole screen after a 2-s pause, task resumption times were found to be no different from the corner-screen condition but were significantly quicker than the immediate full-screen condition. For this and subsequent experiments, we do not provide a comparison with "perfect" data sets. Too few participants completed problems in all conditions without error on any trial; however, given that no differences were apparent between perfect and imperfect trial pairings in Experiment 1A, we do not feel that our findings are compromised by using the whole data set. In general, the solution path for imperfect trials did not deviate too greatly from the optimal solution sequence, and therefore the data for perfect and imperfect trials were relatively comparable.

The data were analyzed for order effects to test whether task resumption times reduced with practice. In particular, we wanted to test whether participants improved between the first and second occurrence of the pause or corner-screen interruption conditions, as it may be that later in the experiment they were more prepared to use the time between tasks to consolidate task goals. A 3 (interruption type) × 2 (order of interruption) repeated measures ANOVA confirmed the effect of interruption type, $F(2, 40) = 5.67$, $MSE = 7.95$, $\eta_p^2 = .22$, $p < .01$, but the effect of order did not reach significance, despite a trend for quicker resumption times on the second occurrence of each interruption, $F(1, 20) = 1.88$, $MSE = 9.36$, $p = .19$. There was no interaction between order and interruption type, $F(2, 40) = 0.06$. Interruption duration was recorded, but as in Experiment 1A, there was no difference between conditions in terms of the time spent on the interruption, $F(2, 40) = 1.04$, $MSE = 1.72$, $p = .36$.

It seems that the key factor in determining the ease of goal retrieval is the availability of task-related cues at the specific point of goal suspension and not necessarily the availability of the ToL task for the duration of the interruption. This finding provides support for the goal-activation model and demonstrates the importance of the interruption lag as a critical time for preparing to resume an interrupted task. In previous studies showing a benefit of the interruption lag (e.g., Altmann & Trafton, 2004; Trafton et al., 2003), preparatory periods of up to 8 s in duration have been used, but the current results show that even the briefest opportunity to encode associative cues and/or boost the activation of the target goal is beneficial. These experiments have therefore demonstrated the benefit of an interruption lag at the point of goal suspension, but in order to further investigate the role of associative cues, we turn now to look at the point of goal retrieval.

Table 1
Mean Time Taken (s) to Make the Fourth Move in Experiment 1B

Interruption condition	*M*	*SE*
No interruption	3.50	0.20
Corner screen	4.91	0.33
Full screen, pause	4.81	0.24
Full screen	6.65	0.62

Experiment 2A

The previous experiments show that there is a clear advantage of time to encode cues *before* the interruption, but how is performance affected if changes are made to the visual display *after* the interruption? In Altmann and Trafton's (2002) model, for a suspended goal to overcome retroactive interference and to govern behavior once more, it must receive priming from contextual cues in order to boost its activation level above that of other competing goals. Associative priming is dependent on links being formed between the target goal and cues in the current environment before suspension, and it is also dependent on the availability of these cues when the goal is to be retrieved. Thus, one would expect task resumption to be impeded if these cues are no longer available, perhaps forcing the system to use less efficient cues to reactivate the goal or to attempt reconstruction from the environment.

An investigation of the role of associative priming of course begs the question as to what exactly the cues are that prime goal retrieval. In the simulation of the ToH task, Altmann and Trafton (2002) suggested that the disks themselves act as cues to retrieve the suspended higher order goal. Because subgoals are created to free up a larger disk, a suspended parent goal is easily retrieved on completion of the subgoal, as the last action will uncover the next disk to be moved. The disks are processed automatically during goal encoding, so that the goal and the cue will be associatively linked during planning. In the current study, goals are suspended in order to attend to an unrelated task rather than to complete a problem-critical subgoal. As the sequence of moves is broken by an interruption, the last action will not automatically cue the next disk as it does in the ToH simulation. However, even with an unrelated intervening task, it would still be reasonable to assume that the disks themselves cue retrieval of the suspended goal. The disks are processed during the initial planning stage and are strengthened further during online planning and monitoring. When a move is unexpectedly interrupted, the next disk will already be linked associatively with a peg because of the look-ahead nature of problem solving in this task. To retrieve this suspended goal, it seems reasonable to assume that the ToL display that the participant returns to will cue retrieval of the last planned move, as the target disk and its associated peg will have been strengthened just before the interruption.

Experiment 2A therefore manipulated whether the display that participants returned to was the one they had left at the point of interruption or whether it was different and would therefore not prime retrieval. This was achieved by changing around the colors of the disks on the returning display. The actual configuration of disks on the pegs was left exactly the same (i.e., the same number of disks on each peg) on both the main display and goal state diagram; only the colors of the disks were changed. This meant that exactly the same solution path was required to complete the problem, but the color cues to retrieve this path were removed.

Method

Participants. Participants were 26 undergraduate psychology students of Cardiff University who each received course credit for their participation.

Materials and apparatus. The same ToL task was used as in the previous experiment, but the mood checklist would immediately cover the whole screen for every interruption. On half of the interruption trials, the colors of the disks were changed around following the secondary task, although the positioning of actual disks on pegs stayed exactly the same. The colors of the disks in the goal state box were also changed so that they were congruent with those on the main display, thus retaining the same solution path.

Design. A repeated measures design was used in which all participants completed six interruption trials. For half of participants, the display was changed following interruption on Trials 4, 7, and 15, and the display remained the same on Trials 12, 19, and 25. This was reversed for the other half of participants. There were also six control problems.

Procedure. The procedure was the same as in the previous experiments, with participants completing at least two obligatory practice trials. No warning was given to participants that the display might change.

Results and Discussion

Time taken to make the fourth move in a solution sequence was recorded (Table 2), and one outlier was excluded from analyses. A repeated measures ANOVA revealed a significant difference between conditions, $F(2, 48) = 24.03$, $MSE = 2.21$, $\eta_p^2 = .50$, $p < .001$. Fisher's post hoc comparisons showed significant differences between each of the three conditions at the $p < .05$ level. Times taken to make the fourth move were quickest when plan execution was continuous, slower when the chain of goal retrieval was broken by an interruption, and slower still if the disk colors on the ToL problem had been changed around after this break in task. There was no difference between conditions in terms of the duration of the interruption, $t(24) = 0.08$.

The current results show that changes to contextual cues at the point of task resumption cause disruption to goal retrieval, a finding that is compatible with the assumptions of the Altmann and Trafton (2002) model. During the initial planning stage and subsequent online monitoring of the problem, each successive move will be strengthened in advance of actual execution. Therefore, because of the look-ahead nature of these problems, at the point of interruption the next disk to be moved will already have been linked associatively with the target peg (linked more strongly, of course, if an interruption lag has allowed further strengthening of this goal). When participants return from the interruption, they must then attempt to retrieve this planned move. It seems plausible that when the screen once again displays the ToL task showing the same disk configuration, this arrangement will cue participants to attend to the location of the last strengthened disk. Research on visual attention suggests that selection of information by location is stronger and more efficient than selection by other attributes such as color or shape (Von Wright, 1968). Attention is drawn to a particular spatial location, which, together with the target-colored disk, will reactivate the previously intended move. Only brief verification is needed with the goal state in this case, as the move to be executed is exactly the same as that previously strengthened. In the color-change condition, however, although exactly the same move is required (the original solution path is retained), this move will not be reactivated so easily, as the disk in the target location is not the color expected. Retrieval of the target move based on disk location alone will require more time and effort than if the color cues were congruent. Also, such retrieval will perhaps require more verification from the goal state diagram before the move is confidently executed, therefore increasing task resumption times.

Given the significant increase in task resumption times when the disk colors were changed, it seems plausible to argue that the disks themselves are acting as the cues to prime retrieval of previously planned intentions. However, it is also possible that simply any change to the visual display would disrupt performance and that the disks are not necessarily critical cues to goal retrieval. This explanation was assessed in the following experiment.

Table 2
Time to Make Fourth Move (s) in Experiment 2A

Interruption condition	*M*	*SE*
No interruption	2.85	0.18
No change	4.83	0.32
Change	5.69	0.47

Experiment 2B

Experiment 2B investigated whether the decrease in performance observed in Experiment 2A was specific to the changing of the disks or whether any change in context would hinder the retrieval of previous task goals. Given the critical nature of the disk color for solving the ToL task and achieving the goal state, it was predicted that an alteration to this feature of the task environment would have a greater impact on goal retrieval than some other change in context. To test this possibility, the current experiment included trials in which the colors of either the disks or the pegs were changed. In addition, it was expected that the results of Experiment 2A would be replicated.

Method

Participants. Twenty-six undergraduate psychology students from Cardiff University received course credit for their participation in the study.

Apparatus and materials. The same ToL program was used as in the preceding experiments, but the interruptions were as follows: two in which the returning display was the same as when it was left; two in which the colors of the pegs had changed from black to pale grey; and two in which the disk colors were changed, as in Experiment 2A.

Design. A repeated measures design was used in which each participant completed two interruptions in each of the three conditions, the order of which was counterbalanced, as well as six matched control trials.

Procedure. The same procedure was used as in Experiment 2A. Participants read through standardized instructions, although in this experiment, as the peg-color change would be quite obvious to participants, they were given the following warning: "During execution of the task there may be some changes to the visual display. These changes will not affect the solution path that you are executing and you should try to ignore them."

Results and Discussion

Move times were recorded and are shown in Table 3. A repeated measures ANOVA demonstrated a significant effect of interruption condition, $F(3, 75) = 43.90$, $MSE = 1.69$, $\eta_p^2 = .64$, $p < .001$, and differences between conditions were analyzed further using Fisher's post hoc tests. In keeping with all previous experiments, the time taken to make the fourth move in control trials was significantly quicker than in any of the interruption trials. There was no difference in resumption times between the condition in which the color of the pegs changed following interruption and the condition in which no change was made to the visual display. When the colors of the disks were changed following interruption, however, resumption times were significantly slower compared with each of the other conditions ($p < .01$). There was no difference between conditions in terms of length of interruption, $F(2, 50) = 0.43$. It is perhaps interesting that in this experiment participants were warned of potential changes to the visual display, but still there was an effect of the color change manipulation.

The current findings provide further support for the idea that the disks themselves act as cues to retrieve the previously planned move. The disks are encoded during planning, and intended moves are strengthened during the online monitoring stage of problem execution. At the point of interruption, a link will already have been forged between the target disk and its goal location, so any change to the color of that disk will impede this association and subsequently delay retrieval of the suspended goal. The ToL pegs, however, cause no greater disruption to performance when changed in color than when no change occurs. This is probably because although the pegs are integral to the task and the main display, the actual color of the pegs is not an important factor for successful problem completion. As such, peg color is unlikely to be processed as a potential cue with an associative link to a target move, even if they are sampled at the same time.

It is, however, possible that changes to the main visual display other than those critical to problem completion may also affect interruption recovery. For example, the background color of the screen is not directly relevant to the solution path, but following interruption, a change to this general feature of the task environment may be particularly attention grabbing and impede retrieval of suspended goals. Equally, it is possible that changes to the disks themselves could be made without impairment to performance, even though the current work has shown them to act as important retrieval cues. For example, changing the shape of the disks at resumption, from ellipses to rectangles, would involve alteration to elements of the display known to influence goal retrieval, yet the shape of the disks is not as problem critical as the color. Along a similar vein, it would be interesting to test the effect of disk-color change following interruption in the ToH task: The same change would be made to the visual display as in the current experiment, but disruption may be less, as the ToH is solved by disk size and not color.

Table 3
Time Taken (s) to Make the Fourth Move in Experiment 2B

Condition	*M*	*SE*
No interruption	2.74	0.19
No change	5.72	0.32
Peg change	5.68	0.30
Disk change	6.62	0.46

Experiment 3

In Experiment 3 we aimed to bring together the findings from the article by assessing both the effects of a time lag at the point of goal suspension and the availability of retrieval cues at the point of task resumption, together in a 2 × 2 design. It was expected that the same main effects of interruption lag (Experiment 1) and disk change (Experiment 2) would be replicated but also that there would be a significant interaction: The presence of an interruption lag to encode task cues before goal suspension should be beneficial only if these cues (disk colors) are still available and unchanged at task resumption. This should strengthen support for the function of the interruption lag as a time specifically to encode primary task retrieval cues. Experiment 3 used a lag of 3 s before the onset of interruption, an increase on the 2-s pause used in Experiment 1B. Although this alteration was not intended for statistical comparison with the earlier experiment, it may provide an additional point of interest, as it seems plausible that a longer lag may facilitate task resumption further.

Method

Participants. Participants were 40 undergraduate psychology students from Cardiff University who received course credit for their time. None had participated in any of the previous experiments.

Materials and apparatus. The standard ToL program was used, but eight interruptions were now included among the 25 trials. Following completion of the third move, either interruptions would occur immediately or the screen would freeze for 3 s before the mood checklist appeared. The second dimension was that interruption trials would return the participant to either the same display seen before the mood checklist or one in which the colors of the disks were changed. In all four conditions, the mood checklist covered the whole screen.

Design. A repeated measures design was used. There were 25 trials in total, of which 8 included an interruption (comprising four interruption conditions). Two interruptions were immediate with the display remaining the same (I-S); two were immediate with the display changed (I-C); two occurred after a pause but the display remained the same (P-S); and two occurred after a pause and the display changed (P-C). Interruptions occurred on Trials 4 and 12, 7 and 21, 10 and 19, and 15 and 25, with the position of I-S, I-C, P-S, and P-C interruptions counterbalanced in a partial Latin square design.

Procedure. The same procedure was used as in the preceding experiments. So as not to influence participants' behavior by alerting them to the experimental manipulations, the instructions gave no warning of any pauses in the program or any potential changes to the visual display.

Results and Discussion

Resumption time data for 2 participants exceeded three standard deviations from the mean and were excluded from analyses. Task resumption times were recorded for each condition (Figure 4) and subjected to a 2 (display change) × 2 (interruption lag) repeated measures ANOVA. There was a significant effect of display change, with participants being significantly slower to resume the task when the colors of the disks had been rearranged, $F(1, 37) = 16.44$, $MSE = 3.94$, $\eta_p^2 = .31$, $p < .001$. The main effect of interruption lag did not quite reach statistical significance, $F(1,$

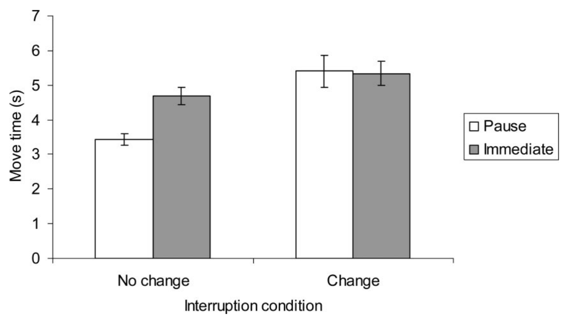

Figure 4. Mean time taken (s) to make the fourth move in conditions either with or without an interruption lag and with or without disk-color change. Error bars show standard error.

37) = 3.64, MSE = 3.65, $\eta_p^2 = .09$, $p = .06$, although there was a significant interaction between interruption lag and display change, $F(1, 37) = 4.97$, $MSE = 3.31$, $\eta_p^2 = .12$, $p < .05$. The interaction, as expected, demonstrated that participants were significantly quicker to resume the task when there was a time lag before changing tasks, but only when the display was not subsequently changed. When the colors of the disks were changed at resumption, the benefit of a time lag at goal suspension was removed. This finding is in line with the predictions of the goal-activation model: The opportunity to encode retrieval cues before the onset of an interruption will be useful only if these cues are still available after the interruption to prime goal retrieval. An analysis of time spent on the interruption revealed no significant differences between conditions, $F(3, 111) = 0.65$.

As in Experiment 1, we analyzed resumption times for practice effects to test whether participants were able to develop strategies in the pause condition that may help them to recover more quickly from the second interruption of that type that occurred. The data were subjected to a 4 (interruption type) × 2 (order) repeated measures ANOVA, which showed a significant effect of both interruption condition, $F(3, 111) = 8.42$, $MSE = 7.23$, $\eta_p^2 = .19$, $p < .01$, and the two blocks of interruptions, $F(1, 37) = 23.77$, $MSE = 7.11$, $\eta_p^2 = .39$, $p < .01$. There was no significant interaction, $F(3, 111) = 1.72$, $MSE = 6.84$, $p = .17$. For all four types of interruption, resumption times were quicker in the second block, indicating that participants became generally more efficient at dealing with the unexpected break in task, and this was not necessarily related to the development of strategies in the pause condition. It is surprising that we found an effect of practice in this experiment but not in Experiment 1; however, this could perhaps be accounted for by differences in sample size ($n = 38$ in Experiment 3 compared with $n = 23$ in Experiment 1A and $n = 21$ in Experiment 1B).

At first glance, it would appear that using a 3-s pause in this experiment reduced task resumption times to a greater extent than the 2-s pause used earlier (mean resumption time of 4.78 s in Experiment 1B, and mean resumption time of 3.43 s in the equivalent P-N condition in Experiment 3). However, because of the general variation in move times between experiments, one must be cautious in concluding that the longer interruption lag reduced resumption times further. Average times across conditions in Experiment 1 are slightly slower than those in Experiments 2 and 3, even in control conditions. This variation is not accounted for by practice effects, as no participant took part in more than one of the experiments. It is possible that slight differences in the experimental procedures may have affected between-experiment comparisons: In Experiment 3 there were 8 interruptions (owing to an increased number of experimental conditions), but only 6 out of 25 trials were interrupted in each of the preceding experiments. However, the difference is most likely accounted for by the slightly different participant samples. Experiments 1A and 1B used paid participants from a sample of Cardiff University undergraduate students, whereas Experiments 2 and 3 used only psychology undergraduates who were participating for course credit and may have been more used to completing computer-based cognitive tasks. Further analysis specifically manipulating interruption lag length within participants would therefore be required in order to draw any firm conclusions on this point.

General Discussion

The current series of experiments assessed the processes involved in the suspension and resumption of goals and specifically addressed the question of whether preparation can aid recovery from interruption. A brief time lag before the onset of the secondary activity was found to reduce task resumption times by allowing an opportunity for participants to prepare goals prior to suspension (Experiments 1 and 3). Moreover, support was found for the proposal that this preparatory time is indeed used for the encoding of retrieval cues, as goal retrieval was impeded when changes were made to the returning visual display (Experiment 2), to the extent that there was no benefit of the interruption lag if the potential retrieval cues had been altered (Experiment 3).

The current findings are difficult to accommodate within the Anderson and Douglass (2001) model of goal memory. Although Anderson and Douglass believe that the strengthening/rehearsing of task goals before interruption may occur, they judge it to be an effortful process and therefore one that participants would choose rarely. They argue that in most cases a forgotten goal can be reconstructed from the environment and so unless circumstances preclude this, there is generally no need to engage in effortful preparatory processes. In the ToL, the cost of forgetting is low because the task environment supports goal reconstruction, and therefore evidence of preparation would not have been expected in this paradigm. However, relative to interruptions with immediate onset, those that were preceded by a brief time lag led to faster task resumption times, a finding that suggests the involvement of preparatory processes during the transition between tasks. Although the Anderson and Douglass (2001) model would have difficulty explaining this effect of the interruption lag, it may be better equipped to explain the findings of the disk-change manipulation. It is possible that goals may be forgotten and must be reconstructed from the environment; in this case, changes to the visual display would perhaps affect the ease of goal reconstruction.

Altmann and Trafton's (2002) goal-activation model is able to account well for the current findings: In keeping with the model's predictions, it would seem that context—at the points of both goal suspension and goal retrieval—plays an important role in efficient interruption recovery. Participants were quicker to retrieve the suspended goal when there was the opportunity to consolidate that goal before the onset of interruption, perhaps by encoding retrieval cues or by strengthening the association of a target disk to its target location. Moreover, task resumption times were increased when the colors of disks were changed following interruption, indicating that changes to the task context impaired priming and retrieval of the previously planned action. However, we should consider that other factors could also account for the increased resumption times in this condition. For example, following interruption it is possible that participants take time to verify with the goal state that the required solution path is still the same, and this verification process may take longer in the color-change condition than in the unchanged condition. In this regard, eye-tracking data would be useful to unpack the specific processes involved during task resumption as well as assess those processes operating at goal suspension. It would be interesting also to assess whether there is an optimal length for the interruption lag. Using explicit warning to signal interruption lags of 2, 4, 6, and 8 s, Trafton et al. (2003) found that the longer lags were more beneficial. However, it seems reasonable to suggest that eventually too long an interval (or an interval that varies unpredictably) might constitute an additional task load and therefore impede encoding.

The current findings, while of interest in terms of cognitive theory, may also have practical applications in terms of the design of efficient work environments. The brief and undemanding interruptions used in the current study can be compared to those of computer-initiated messages, such as e-mail alerts, “save” prompts, pop-ups, printing notifications, or instant messages: Although distracting a worker for just a matter of seconds, these experiments show, even such short breaks in task can incur a time cost when one resumes the ongoing activity. Experiment 1A manipulated whether the interrupting pop-up box covered the whole screen or appeared just in the corner, a factor that parallels the differences between the alerts of some e-mail software. It was found that the less intrusive the pop-up box was, the quicker was the resumption of the primary task. When both tasks are available on the screen at the same time, an opportunity is available to boost the activation of the to-be-suspended goal, or to encode retrieval cues before dealing with the interruption, so that this goal is more easily retrieved later on. Of course, these findings are only exploratory, using a structured cognitive task, and further study using more office-based activities would provide more weight to these suggestions within an applied domain.

One point of concern at a practical level is the cost–benefit ratio of the interruption lag. Although the presence of a time lag reduced task resumption times, overall processing time is not reduced when the duration of the preceding pause is also taken into account. A 2-s time lag before engaging in the interruption saved approximately 2 s at the point of resumption (Experiment 1B). It seems, therefore, that processing is simply taking place before the interruption rather than after, achieving no overall reduction in task time. This problem was also encountered by Trafton et al. (2003), who found that an 8-s interruption lag saved only 3 s at resumption in their complex resource allocation task. Given that performance in these ToL experiments is assessed at the level of individual move times, there is not much scope for large differences in task performance. Participants are taking on average 6 s to resume the task with no interruption lag and 2 or 3 s to make the equivalent move in a no-interruption control trial (reflecting mainly program and motor constraints). Mean resumption time in the interruption lag condition is about 4 or 5 s, which falls between these two. Although the difference is small, it demonstrates a fairly consistent benefit, and there is little opportunity for performance times to vary greatly. Another primary task, for example, one involving deeper semantic processing, might demonstrate a greater advantage of time to encode retrieval cues before engaging in the interruption.

Interruption effects are often considered a manifestation of other cognitive phenomena, such as task switching (e.g., Allport & Wylie, 2000; Monsell, 2003) or prospective memory (e.g., Brandimonte, Einstein, & McDaniel, 1996), and as such the topic has not received commensurate attention in the cognitive literature. However, task interruption—the unexpected suspension and subsequent resumption of task goals—is a separate area worthy of investigation and one that is becoming of increased interest in both applied and cognitive domains. The existing literature is inchoate and has yet to yield an overarching theoretical framework to detail the specific cognitive processes involved. The experiments reported here contribute to the establishment of basic interruption effects using the ToL task and with reference to the well-established theoretical foundation of ACT–R. The ToL allows for a fine-grained analysis of performance at the level of individual moves and retains a controlled task environment similar to that in which the ACT–R models were originally developed. Moreover, the work builds on previous research that used the same methodology, a continuity that adds to the coherence of the literature (Hodgetts & Jones, 2006).

The current findings were found to be compatible with Altmann and Trafton’s (2002) goal-activation model, which, rather than relying on the goal stack’s last-in, first-out retrieval mechanism, uses the constructs of activation and associative priming to determine the selection of pending goals. This process of priming seemed to be facilitated by a brief time lag before the onset of the secondary task and impeded when task-related cues were altered at retrieval. The Anderson and Douglass (2001) model also provides a useful basis for the study of task interruption, but it would need further elaboration to accommodate all of the current findings. Further investigation is required to address the cost–benefit ratio of the interruption lag, so that related findings can be of practical benefit as well as theoretical interest.

References

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

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

Altmann, E. M., & Trafton, G. J. (2004). Task interruption: Resumption lag and the role of cues. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 42–47). Hillsdale, NJ: Erlbaum.

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,* 1331–1346.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396–408.

Bainbridge, L. (1984). Analysis of verbal protocols from a process control task. In E. Edwards & F. P. Lees (Eds.), *The human operator in process control* (pp. 146–158). London: Taylor & Francis.

Brandimonte, M., Einstein, G. O., & McDaniel, M. A. (Eds.). (1996). *Prospective memory: Theory and applications.* Mahwah, NJ: Erlbaum.

Burmistrov, I., & Leonova, A. (2003). Do interrupted users work faster or slower? The micro-analysis of computerized text editing task. In J. Jacko & C. Stephanidis (Eds.), *Human–computer interaction: Theory and practice (Part I)—Proceedings of HCI International* (Vol. 1, pp. 621–625). Mahwah, NJ: Erlbaum.

Byrne, M. D., & Bovair, S. (1997). A working memory model of a common procedural error. *Cognitive Science, 21,* 31–61.

Clifford, J. D., & Altmann, E. M. (2004). Managing multiple tasks: Reducing the resumption time of the primary task. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Cutrell, E., Czerwinski, M., & Horvitz, E. (2000). Effects of instant messaging interruptions on computing tasks. In *Extended abstracts of the CHI 2000 Conference on Human Factors in Computing Systems* (pp. 99–100). New York: ACM Press.

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000). Instant messaging and interruption: Influence of task type on performance. In C. Paris, N. Ozkan, S. Howard, & S. Lu (Eds.), *OZCHI 2000 Conference Proceedings* (pp. 356–361). Sydney, Australia: CSIRO, Mathematical and Information Sciences.

Detweiler, M. C., Hess, S. M., & Phelps, M. P. (1994). *Interruptions and working memory.* Unpublished technical report, Pennsylvania State University.

Ernst, G. W., & Newell, A. (1969). *GPS: A case study in generality in problem solving.* New York: Academic Press.

Gillie, T., & Broadbent, D. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research, 50,* 243–250.

Hess, S. M., & Detweiler, M. C. (1994). Training to reduce the disruptive effects of interruptions. In *Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting* (pp. 1173–1177). Santa Monica, CA: Human Factors and Ergonomics Society.

Hodgetts, H. M., & Jones, D. M. (2006). Interruption of the Tower of London task: Support for a goal activation approach. *Journal of Experimental Psychology: General, 135,* 103–115.

Kirmeyer, S. L. (1988). Coping with competing demands: Interruption and the Type A pattern. *Journal of Applied Psychology, 73,* 621–629.

McFarlane, D. C. (2002). Comparison of four primary methods for coordinating the interruption of people in human–computer interaction. *Human–Computer Interaction, 17,* 63–139.

McFarlane, D. C., & Latorella, K. A. (2002). The scope and importance of human interruption in human–computer interaction design. *Human–Computer Interaction, 17,* 1–61.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior.* New York: Holt, Rinehart & Winston.

Miller, S. L. (2002). Window of opportunity: Using the interruption lag to manage disruption in complex tasks. In *Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society*. Santa Monica, CA: Human Factors and Ergonomics Society.

Monk, C. A., Boehm-Davis, D. A., & Trafton, J. G. (2004). Recovering from interruptions: Implications for driver distraction research. *Human Factors, 46,* 650–663.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences, 7,* 134–140.

Phillips, L. H., Wynn, V. E., McPherson, S., & Gilhooly, K. J. (2001). Mental planning and the Tower of London task. *Quarterly Journal of Experimental Psychology, 54A,* 579–597.

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London B, 298,* 199–209.

Storch, N. A. (1992). *Does the user interface make interruptions disruptive? A study of interface style and form of interruption* (Report UCRL-JC-108993). Springfield CA: Lawrence Livermore National Laboratory.

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,* 583–603.

Von Wright, J. M. (1968). Selection in visual immediate memory. *Quarterly Journal of Experimental Psychology, 20,* 62–68.

Ward, G., & Allport, A. (1997). Planning and problem-solving using the five-disc Tower of London task. *Quarterly Journal of Experimental Psychology, 50A,* 49–78.

Zeigarnik, B. (1927). Das Behalten erledigter und unerledigter Handlungen [The retention of completed and uncompleted actions]. *Psychologische Forschung, 9,* 1–85.

Zijlstra, F. R. H., Roe, R. A., Leonora, A. B., & Krediet, I. (1999). Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organisational Psychology, 72,* 163–185.

(*Appendixes follow*)

Appendix A

Tower of London Problems

Problem	Start	Goal
1	⟨R, G⟩ ⟨P⟩ ⟨B, Y⟩	⟨G⟩ ⟨Y, P, R⟩ ⟨B⟩
2	⟨R⟩ ⟨P, B, Y⟩ ⟨G⟩	⟨R⟩ ⟨B⟩ ⟨Y, G, P⟩
3	⟨B, Y⟩ ⟨⟩ ⟨R, P, G⟩	⟨Y⟩ ⟨G, P⟩ ⟨R, B⟩
4	⟨Y, B⟩ ⟨G, P⟩ ⟨R⟩	⟨B⟩ ⟨G⟩ ⟨P, R, Y⟩
5	⟨R⟩ ⟨Y⟩ ⟨P, G, B⟩	⟨Y⟩ ⟨R, G, P⟩ ⟨B⟩
6	⟨G⟩ ⟨R, P⟩ ⟨Y, B⟩	⟨P, G, Y⟩ ⟨R⟩ ⟨B⟩
7	⟨R, P, Y⟩ ⟨G⟩ ⟨B⟩	⟨G, Y⟩ ⟨P⟩ ⟨B, R⟩
8	⟨Y⟩ ⟨G, B, R⟩ ⟨P⟩	⟨Y, G⟩ ⟨B⟩ ⟨R, P⟩

Note. Pegs are indicated by angular brackets, and disks within pegs are ordered so that the leftmost is on top. R = red; Y = yellow; G = green; B = blue; P = pink.

Appendix B

Accuracy Data

ToL problems not solved in the minimum of six moves were classed as errors and were recorded with respect to condition across all participants. In Experiment 1A, 23 participants each completed 6 interruption and 6 control trials in each session. Out of a total of 138 trials per condition, the number not completed in the minimum number of moves was as follows: Session A (full screen), interruption = 58, control = 52; Session B (corner screen), interruption = 58, control = 49. These differences were nonsignificant according to a chi-square analysis, $\chi^2(1, N = 217) = 0.04$. In Experiment 1B, 21 participants each completed two of each type of interruption, making a total of 42 trials in each condition. Of these, the number not solved in the minimum of six moves was as follows: corner screen = 15; full screen, pause = 19; full screen = 26. These differences did not reach statistical significance at the .05 level according to a chi-square goodness-of-fit test, $\chi^2(2, N = 60) = 3.10$. Control trials produced a similar proportion of errors (57 out of 126 trials). In Experiment 2A, out of a possible 78 trials in each condition, 39 trials in the standard interruption condition and 46 trials in the color-change condition were classed as errors. These differences were not significant at the .05 level according to a chi-square goodness-of-fit test, $\chi^2(1, N = 85) = 0.58$. Also, 73 out of 156 control trials were not completed in the minimum number of moves. In Experiment 2B, out of a total of 52 trials in each condition, the number of error trials was as follows: no change = 12, peg change = 14, disk change = 15. These differences were not statistically significant according to a chi-square goodness-of-fit test, $\chi^2(2, N = 41) = 0.34$. For control trials, 27 out of 156 were completed in more than six moves. In Experiment 3, out of 76 trials in each interruption condition, the number of error trials was as follows: no change, pause = 36; no change, immediate = 30; change, pause = 29; change, immediate = 24. These differences were not significant according to a chi-square analysis, $\chi^2(1, N = 119) < 0.001$.

Received September 8, 2005
Revision received February 2, 2006
Accepted March 11, 2006 ■

Instructions to Authors

For Instructions to Authors, please consult the January 2006 issue of the volume or visit www.apa.org/journals/xlm and click on the "Instructions to authors" link in the Journal Info box on the right.