

Task Interruption: Resumption Lag and the Role of Cues

Erik M. Altmann (ema@msu.edu)

Department of Psychology
Michigan State University
East Lansing, MI 48824

J. Gregory Trafton (trafton@itd.nrl.navy.mil)

Naval Research Laboratory, Code 5513
4555 Overlook Avenue S. W.
Washington, DC 20375

Abstract

The consequences of interrupting someone in the middle of a complex task are of considerable practical and theoretical interest. We examine one behavioral measure of the disruption caused by task interruption, namely the *resumption lag*, or the time needed to “collect one’s thoughts” and restart a task after an interruption is over. The resumption lag (in our task environment) was double the interval between uninterrupted actions (3.8 s vs. 1.9 s), indicating a substantial disruptive effect. To probe the nature of the disruption, we examined the role of external cues associated with the interrupted task, finding that cues available immediately *before* an interruption facilitate performance immediately *afterwards* (reducing the resumption lag). This *cue-availability* effect suggests that people deploy preparatory perceptual and memory processes, apparently spontaneously, to mitigate the disruptive effects of task interruption.

Introduction

For better or worse, interruptions are part of everyday life. For better, interruptions are an essential part of efficient communication, among people and between people and machines. For worse, interruptions can be annoying, and can seem disruptive. For example, the annoyance of unwanted telephone solicitations drove the recent overwhelming popularity of “do not call” registries in the United States. Similarly, consider the “software assistant” included in Microsoft products in the late 1990s. If Word, for example, detected what it thought was a letter being drafted, it would freeze the keyboard and demand to know if the user needed “help” — a feature that in more recent editions of the software is no longer enabled by default.

There are many parameters to how an interruption is structured — including interruption duration, for example, or whether the interrupted person has control over interruption timing (McFarlane & Latorella, 2002) — and there is also a range of different behavioral measures on which to assess the impact of an interruption. In one classic result, Zeigarnik (1927/1938) found that interrupted tasks were actually remembered better, in terms of recalled detail, than tasks that were allowed to run to completion. In more applied work, however, interruptions have been found to have detrimental effects on situational awareness in dynamic task environments like aviation (Latorella, 1996), where losing one’s place in a checklist during takeoff, for example, can have catastrophic results (NTSB, 1988).

The current study examines the disruptive effects of interruption in terms of the time needed to resume the primary (interrupted) task after the secondary (interrupting) task is complete. In less formal terms, we examine the time needed to collect one’s thoughts, or pick up the thread

again, when an interruption is over and we can return to what we were doing before. There is surprisingly little research focused on this measure, and what there is is distributed across a variety of domains and paints no clear picture of whether interruptions are disruptive or not. For example, in a study of interruption of administrative and clerical workers, disruptive effects of interruption were difficult to detect (Zijlstra, Roe, Leonora, & Krediet, 1999), and interruption in a simple question-answering task can actually improve performance, measured in terms of overall accuracy and time-on-task (Speier, Vessey, & Valacich, 2003). In the mainstream cognitive psychology literature, research on “task switching” (Monsell, 2003) would seem to be relevant, as studies in this domain typically focus on the “switch cost” associated with shifting from one task to another. However, this literature is perhaps not so aptly named; the “tasks” used in task-switching studies take a few hundred milliseconds to complete, with switch cost a small and not particularly relevant fraction of that (Altmann, in press). We are interested in higher-level tasks with greater ecological validity, where switch cost is measured not in tens of milliseconds, but in seconds or longer.

In operational terms, the dependent measure in the current study is the *resumption lag*, illustrated in Figure 1. The resumption lag is the time interval separating the end of the secondary task and the first subsequent action taken by the human operator in the primary task. We report first a comparison of this resumption lag to an estimate of the time interval that usually separates actions in the primary task, to give a sense of the absolute magnitude of the disruptive effect; to preview, the mean inter-action interval, in the highly interactive primary task we are using, is roughly 2 s, and the resumption lag is roughly 4 s, indicating a substantial disruption both in absolute and relative terms.

We then report on two factors that have the potential to reduce the resumption lag. Both focus on the *interruption lag*, also illustrated in Figure 1. The interruption lag is a brief transitional interval immediately preceding an interruption, during which the operator knows of the pending interruption but is not yet engaged by it. An example is the time between the phone starting to ring and the act of actually taking the call; during this interval, there is a brief opportunity to complete a thought, for example, or negotiate quickly with a conversation partner (physically present) how and when to resume after the call is over. Many real-world interruptions, even more urgent ones, afford a brief interruption lag; even when a fire alarm sounds, one is still likely to take time to save changes to an electronic document, for example, before evacuating.




Figure 1: Time course of an interruption. For example, if the alert is the phone starting to ring, then the secondary task is the ensuing phone conversation.

Two characteristics of the interruption lag are examined here: (1) whether or not the primary task display is perceptually available during this brief period, and (2) the actual duration of this period. The relevance of these two factors is predicated on a memory model of what makes interruptions disruptive (Altmann & Trafton, 2002; Trafton, Altmann, Brock, & Mintz, 2003). The basic premise of the model is that during an interruption, the cognitive representations that support performance of the primary task will decay, in particular relative to the cognitive representations that support performance of the secondary task. Thus, when resuming the primary task, retrieval cues will be necessary to re-activate the relevant representations.

This memory analysis predicts, qualitatively, that the interruption lag — the brief interval *before* an interruption — has a crucial role to play in facilitating resumption *after* the interruption. During the interruption lag, when the operator is aware that he or she will soon be interrupted but can still focus mentally on the primary task, there is an opportunity to “prepare to resume,” for example by prospectively encoding goals to accomplish at resumption (Trafton et al., 2003). To the extent that people do engage in such preparatory processing, it should help to have the primary-task display perceptually available during the interruption lag, to allow retrieval cues to be quickly accessed and accurately encoded. Thus, to build on earlier evidence that people do prepare to resume (Trafton et al., 2003), we asked here whether cue availability is a factor in this process. In the *cue* condition the primary task display was preserved during the interruption lag, whereas in the *no-cue* condition the screen went blank during the interruption lag (starting with onset of an alert signaling the

pending interruption). Moreover, because processes like perceptual search and memory encoding take time, we varied the duration of the interruption lag across experiments, to examine what length of interruption lag would render cue availability effective in reducing the resumption lag.

Experiments

We conducted four experiments, with interruption lags of two, four, six, and eight seconds respectively. These values were based on evidence that an 8-second interruption lag is enough to allow people to (at least partially) prepare to resume (Trafton et al., 2003). The primary task involved planning and resource allocation subject to constraints, and thus involved a substantial amount of state information to be represented cognitively. The secondary task was less complex but nonetheless involved a series of tightly-spaced forced-choice decisions unfolding over a 30- to 45-second period. Interruption timing was under system, rather than operator control, a factor that tends to aggravate the disruptive effects of interruption (McFarlane & Latorella, 2002).

The independent variable within each experiment was whether or not cues were available during the interruption lag. The main dependent variable was the resumption lag (Figure 1), but we also compare the resumption lag to the mean interval between primary-task actions, to estimate the overall disruptive effect of an interruption.

Method

Participants Ninety-six Michigan State University undergraduates participated in exchange for partial credit toward a course requirement. Each of the four experiments involved 24 participants, randomly assigned to the cue and no-cue conditions (described below).

Materials The primary task was a complex resource-allocation task (Trafton et al., 2003) in which participants were asked to defeat a set of simulated destinations using simulated tanks. Participants selected which destinations to attack and in what order, and issued tanks with appropriate amounts of fuel and munitions. Fuel was consumed in reaching a destination, and munitions were consumed in engaging it, but tank payload was limited, as was the total number of tanks and other resources available. Points were awarded for defeating destinations and subtracted for consuming resources.

Figure 2 shows a view of the primary-task display as it normally appears when the participant is doing the task. There is a central window with buttons for allocating tanks to missions, choosing destinations, and displaying a map with distances between destinations. There is also an area for displaying mission outcomes (whether a destination was defeated, whether a tank ran out of fuel, etc.) To the left are windows showing the supply pool (available fuel, munitions, and tanks) and windows for outfitting heavy and light tanks with varying amounts of fuel and munitions. To the right is a window showing the participant’s scoring




Figure 2: Screen shot of the primary task display during normal performance (see text for summary).

history from past simulations, and windows showing the status of selected tanks and destinations.

Figure 3 shows (at a more reduced scale) the state of the display during the interruption lag in the no-cue condition. To signal the pending interruption (and thus mark the start the interruption lag), an “eyeball alert” appeared in the top-right corner of the display. In the no-cue condition, the primary-task display was blanked out simultaneously with alert onset, whereas in the cue condition the primary-task display was preserved. In both conditions, with the start of the secondary task (and thus the end of the interruption lag), the primary-task display (whatever its state) was completely erased and replaced with the secondary-task display.

During the interruption lag, the cursor was hidden and disabled, so that all physical interaction with the primary task ceased. After an interruption, the primary-task display was reinstated in the same form it was in at the moment the eyeball alert appeared, with the following exceptions: The window that was active then — that is, the window that the participant was working in at the moment the alert appeared — was de-activated at task resumption, and the cursor was moved to the top-left corner of the screen. The effect was to eliminate the active window and the mouse cursor as




Figure 3: Screen shot of the primary task display during the interruption lag in the no-cue condition. The “eyeball alert” in the top right corner onsets at the start of the interruption lag and remains until start of the secondary task. In the cue condition the eyeball alert is identical but the primary task display (as in Figure 2) is preserved during the interruption lag.

retrieval cues that participants could deploy strategically to remind themselves of what they had been doing before the interruption, and therefore make resumption lag more sensitive to our experimental manipulations.

The secondary task involved evaluating “tracks,” or targets on a radar screen, as friendly or hostile, according to attributes like speed and shape of icon (Trafton et al., 2003). A screenshot of the secondary-task display appears in Figure 4. Each instance of the secondary task lasted 30 to 45 seconds; afterwards, the participant was returned directly back to the primary task, as described above.

Design Interruption lag varied between experiments, as described above. Within each experiment, the independent variable was cue availability during the interruption lag (cue, no-cue), which was manipulated between subjects. The main dependent variable was resumption lag, the time from the end of the secondary task to the first subsequent action (mouse click) in the primary task. The other measure of interest, for comparison with resumption lag, was the inter-action interval, the mean time between actions in the primary task. Reported values for these measures are means of participant medians. For the inter-action interval, values below 1 s were discarded first, to eliminate anticipation errors, as well as ballistic components of motor plans, such as the second click of a double-click action.

Procedure Participants were tested individually, in sessions lasting roughly 90 minutes. A session began with a training period, in which participants learned to perform both tasks separately and were shown an example of how the computer would switch them from one task to the other and back again. After training, there were three 20-minute blocks of actual task performance. Within each block there were 10 interruptions, each triggered by a mouse click selected randomly to occur within a time window with quasi-random boundaries, to make interruption timing difficult for participants to predict.

At no point was the hypothetical function of the




Figure 4: Screen shot of the secondary task display. Participants classified objects moving across a simulated radar display, according to color, shape, and speed.




Figure 5: Resumption lag compared to inter-action interval, across experiments. The difference is an estimate of the disruption caused by task interruption.

interruption lag in facilitating resumption mentioned to participants; thus, any effects of cue availability can be attributed to spontaneous use of preparatory strategies.

Results We conducted two analyses of variance (ANOVAs) for each experiment. The first ANOVA, for which the data appear in Figure 5, compared mean resumption lag to mean inter-action interval; for all four experiments the difference was highly reliable, $ps < .0001$. The second ANOVA, for which the data appear in Figure 6, compared cue resumption lags to no-cue resumption lags. For 2- and 4-second interruption lags, there was no effect of cue availability, $F_s < 1$. For the 6-second interruption lag, the cue-availability effect was marginal, $F(1,22)=4.1, p=.056$. For the 8-second interruption lag, the cue-availability effect was reliable, $F(1,22)=5.7, p<.03$.

We also conducted an omnibus ANOVA to compare across experiments, with cue availability and interruption lag as factors. The cue-availability effect was marginal, $F(1,88)=3.6, p=.060$, the interruption-lag effect was not reliable, $F(3,88)=1.2, p>.30$, and the two factors did not interact, $F(3,88)=1.4, p>.25$.

Discussion

The first empirical finding was that resumption lag is substantially longer than the mean interval between actions (Figure 5). This affords one measure of the disruptive effect of interruptions, at least in this highly interactive task in which, without interruption, actions occur at a rapid pace: The first action after an interruption took longer to execute than the first action after another primary-task action. In absolute terms, the resumption lag was 3.8 s – double the 1.9 s inter-action interval, which was measured rather conservatively by excluding all inter-action intervals under 1 s. This difference would appear to be of considerable practical interest in dynamic task environments, for example involving real airplanes or even automobiles traveling at highway speed, in which the world can look substantially different after an additional few seconds have elapsed.

The disruptive effect of interruption, as illustrated in Figure 5, was large and robust, which may agree with our intuitions about interruptions but doesn't necessarily agree




Figure 6: Resumption lag for cue and no-cue conditions, across experiments. The difference is marginally reliable with a 6-second interruption lag and reliable with an 8-second interruption lag (see text).

with other controlled studies (e.g., Speier et al., 2003; Zijlstra et al., 1999). Three factors may have contributed to the robust effect measured here in terms of resumption lag. First, resumption lag is a local measure, taken immediately after every interruption; in contrast, other studies have reported global measures, such as overall time on task (Gillie & Broadbent, 1989; Speier et al., 2003). In our study, one global measure is a participants' total score over a session, but this was highly variable and showed no interpretable trends (so we did not report it). One specific problem with global measures is that they allow for compensatory strategies to work against the disruptive effects of interruption. Zijlstra et al. (1999), for example, speculate that their administrative workers compensated for interruption by using the time in between interruptions more efficiently.

A second factor may have been the relatively substantial cognitive state required to perform our primary task. In many scenarios in this task, beyond the resource-allocation tradeoffs involved, it was a challenge simply to piece together missions that would actually succeed in defeating destinations. In other studies, the primary task was simpler (Speier et al., 2003) and may have been more automated (Zijlstra et al., 1999), and therefore placed a smaller premium on maintaining complex representations in working memory.

Finally, a third factor may have been our implementation decision to trigger interruptions using mouse clicks, rather than strictly on the basis of time passage. Our rationale was that motor actions are often selected and programmed with the intention of achieving specific goals, so we reasoned that action-triggered interruptions would be more likely to disrupt these goals, which are one critical element of cognitive state; even in mundane tasks it's not uncommon to have "Now what was I doing?" moments, and it may be that these are more effectively induced by linking interruptions to actions rather than leaving interruption timing entirely to chance. Indeed, in Zeigarnik's (1927/1938) classic study, the experimenter was charged with judging when the

participant was engrossed, in order to interrupt with the greatest impact.

The second empirical finding was that cue availability during the interruption lag (*before* the interruption) affected performance at task resumption (*after* the interruption, 30 to 45 s later), at least for longer interruption lags (Figure 6). One interpretation of this result, consistent with our memory analysis earlier, is that the various cognitive operations required to locate and encode retrieval cues during the interruption lag take somewhere between 6 and 8 seconds to complete (in our task environment). In other words, longer interruption lags afford enough time to link cognitive representations to external cues to facilitate retrieval later. However, this interpretation would also seem to predict that resumption lag in the cue condition should decrease at longer interruption lags, because cues are facilitating resumption. Instead, though, Figure 6 suggests that longer interruption lags drove an *increase* in resumption lag in the no-cue condition. Cue availability and interruption lag did not interact in the cross-experiment ANOVA, so the increase in no-cue resumption lags could be spurious, but given the exploratory nature of this work it seems useful to consider alternative accounts of why the cue-availability effect was limited to longer interruption lags.

One possible explanation of the increase in no-cue resumption lags might implicate changes in alertness or arousal – participants might simply have gotten bored, staring at a blank screen for 6 or 8 seconds. Some studies suggest that task interruption serves to increase arousal and stress, and thus improve overall (globally-measured) performance, at least on simple tasks (Speier et al., 2003); perhaps a long interruption lag, without visual information to focus on, moderates this effect. However, if arousal were to play a role in the cue-availability effect, it would remain to explain how a change in arousal *before* the interruption could affect performance *after* the interruption, tens of seconds later. Perhaps a drop in arousal caused participants' minds to wander in a way that activated irrelevant thoughts that in turn interfered with relevant cognitive representations. In such an account, however, memory would again play a central role in mediating the effect of pre-interruption variables on post-interruption performance.

One could explain the cue-availability effect without reference to memory processes if changes in arousal during the interruption lag persisted across the entire length of the interruption, to influence performance directly at task resumption. Secondary task performance was basically at ceiling for all subjects, so offers little evidence on this possibility. However, if arousal effects were to persist for the entire 30 to 45 seconds of the interruption, one might expect them to persist somewhat beyond as well, and then only gradually dissipate. This would predict that the time between the first and second action after the interruption would also reflect the cue availability effect. Revisiting our data, however, we found no difference, as a function either of interruption lag or cue availability, in the duration of the interval between the first and second actions after an interruption; this measure appears in Figure 7. It seems most likely, then, that even if cue availability and interruption lag interact to affect arousal before an interruption, memory




Figure 7: Inter-action interval between the first and second actions after an interruption, for cue and no-cue conditions, across experiments.

and/or perceptual processes mediate the delayed effect on task resumption.

At least two avenues of future work seem indicated to clarify the effects of our interruption-lag manipulations on speed of task resumption. First, it will be important to repeat these manipulations in context of a factorial design in which interruption lag and cue availability are fully crossed; here, one objection is that the cross-experiment comparison is potentially confounded by changes in the subject population.

Second, although the model that motivated these experiments emphasizes memory processes (Altmann & Trafton, 2002), there are alternative characterizations of why task resumption is time consuming. In particular, one account of automation deficit (Ballas, Kieras, Meyer, Brock, & Stroup, 1999) — like resumption lag, but measured in terms of accuracy — is that it reflects encoding of perceptual information (Kieras & Meyer, 1997) rather than memory retrieval. Thus, in our task environment it could be that the difference between the resumption lag and the baseline inter-action interval (Figure 5) simply reflects the cost of re-encoding the display, and that this re-encoding is what is facilitated by cue availability during the interruption lag. To distinguish between these accounts, one could vary the extent of the cognitive representations required to perform the primary task on one hand, and the perceptual complexity of the display on the other. Under a memory-retrieval model, the cue-availability effect should be linked to complex cognitive states, whereas under a perceptual-encoding model the effect should be linked to complex external displays; in our task environment, these two factors are confounded.

Whatever the ultimate explanation, the cue-availability effect shows an interesting link between what happens before an interruption and what happens later, after tens of seconds of intervening behavior. In practical terms, the effect suggests that interface designs, and possibly training interventions, could exploit cue availability in some way to facilitate resumption in task environments in which interruptions are frequent and seconds matter. In theoretical terms, probing this effect should help us develop constraints on models of memory, perception, and cognitive control as

these functions are deployed in complex dynamic task environments.

Acknowledgements

This research was supported by grants N00014-03-1-0063 (to EMA) and N0001400WX21058 (to JGT) from the US Office of Naval Research. Thanks to the conference reviewers for helpful suggestions for improvement.

References

- Altmann, E. M. (in press). Advance preparation in task switching: What work is being done? *Psychological Science*.
- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science*, *26*, 39-83.
- Ballas, J. A., Kieras, D. E., Meyer, D. E., Brock, D. P., & Stroup, J. (1999). *Cueing of display objects by 3-D audio to reduce automation deficit*. Paper presented at the Fourth annual symposium on situational awareness in the tactical environment, Patuxent River, MD.
- Gillie, T., & Broadbent, D. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research*, *50*, 243-250.
- Kieras, D. E., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, *12*, 391-438.
- Latorella, K. A. (1996). *Investigating interruptions: Implications for flightdeck performance*. PhD thesis, State University of New York, Buffalo, NY.
- 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.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*, 134-140.
- NTSB. (1988). *Aircraft accident report: NWA DC-9-82 N312RC, Detroit Metro, 16 August 1987* (No. NTSB/AAR-88/05). Washington, DC: National Transportation Safety Board.
- Speier, C., Vessey, I., & Valacich, J. S. (2003). The effects of interruptions, task complexity, and information presentation on computer-supported decision-making performance. *Decision Sciences*, *34*, 771-797.
- 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.
- Zeigarnik, B. (1927/1938). On finished and unfinished tasks. In W. D. Ellis (Ed.), *A source book of Gestalt psychology* (pp. 300-314). New York: Harcourt Brace.
- 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 Organizational Psychology*, *72*, 164-185.