

Running head: RECOVERY FROM INTERRUPTION

Timecourse of Recovery from Task Interruption: Data and a Model

Erik M. Altmann

Michigan State University

J. Gregory Trafton

Naval Research Laboratory

In press, *Psychonomic Bulletin & Review*

Document statistics:

Abstract (p. 2): 120 words

Body, references, and author note (pp. 3-19): 3,996 words

Figures (pp. 20-21): 2

Correspondence:

Erik M. Altmann
Department of Psychology
Michigan State University
East Lansing, MI 48824

517-353-4406 (voice)
517-353-1652 (fax)
ema@msu.edu

Abstract

Interruption of a complex cognitive task can entail, for the “interruptee”, a sense of having to recover afterwards. We examined this recovery process by measuring the timecourse of responses following an interruption, sampling over 13,000 interruptions to obtain stable data. Response times dropped in a smooth curvilinear pattern for the first 10 responses (15 sec or so) of post-interruption performance. We explain this pattern in terms of the cognitive system retrieving a displaced mental context from memory incrementally, with each retrieved element adding to the set of primes facilitating the next retrieval. The model explains a learning effect in our data in which the timecourse of recovery changes over blocks, and is generally consistent with current representational theories of expertise.

Timecourse of Recovery from Task Interruption: Data and a Model

Interruptions are pervasive in everyday task environments. Working on a manuscript, for example, one might be interrupted by a phone conversation, then, some minutes later, return to the manuscript and need a few moments to regain the train of thought. We examined these first few seconds of post-interruption performance to develop empirical data on the timecourse of recovery after a cognitively complex task is interrupted, and to develop a simple but formal model of the underlying attentional and memory processes.

The prevalence of task interruption in everyday and workplace environments has been documented empirically (e.g., Chisholm, Dornfeld, Nelson, & Cordell, 2001; Czerwinski, Horvitz, & Wilhite, 2004; McFarlane & Latorella, 2002), and catches the eye of the popular media from time to time (e.g., Thomson, 2005), yet psychological studies of the *effects* of task interruption on cognition and performance paint a somewhat inconsistent picture. Specifically, time to perform a task may increase when performance is interrupted (e.g., Gillie & Broadbent, 1989), as one might expect, but may also decrease, possibly as a function of increased arousal (Speier, Vessey, & Valacich, 2003) or strategic compensation between interruptions (Zijlstra, Roe, Leonora, & Krediet, 1999). In response to these conflicting findings, some studies have aimed to develop more sensitive and reliable measures by focusing on performance immediately after each interruption, to factor out effects of noise and other confounding variables on baseline performance between interruptions.

One measure commonly reported in such studies is the *resumption lag*, or time between interruption offset and the first subsequent task-related response (Hodgetts &

Jones, 2006a, 2006b; Iqbal & Bailey, 2005; Monk, Boehm-Davis, & Trafton, 2004; Swets, 2006; Trafton, Altmann, Brock, & Mintz, 2003). In this study, we extend previous work on the resumption lag to develop an empirical and theoretical *timecourse function* characterizing the transition from this slow first response after an interruption to the faster responses that follow. We used a cognitively complex task environment, introduced below and described in more detail in supplemental materials posted on-line (www.msu.edu/~ema/timecourseofrecovery), to attain reasonable ecological validity with respect to real-world tasks that require some cognitive engagement to perform well.

Experiment

Method

Participants. Three hundred seventy five Michigan State University undergraduates participated in exchange for partial credit toward a course requirement; 25 were randomly assigned to each between-participants cell.

Materials. The task that was periodically interrupted is a computer game (Brock & Trafton, 1999; Trafton et al., 2003) that has two key characteristics, for present purposes. First, performance involves a high rate of responding – roughly one response per second, at baseline – allowing a fine-grained sampling of behavior as the timecourse of recovery unfolds. Second, the task is cognitively complex, such that one would expect there often to be a substantial cognitive state for an interruption to disrupt. The objective of the game is to defeat a set of enemy positions by deploying military tanks, but planning to meet this objective is complicated by various interacting constraints. For example, each tank has limited carrying capacity, requiring allocation decisions – more

fuel allows a tank to travel further, but more munitions allow it to do more damage along the way.

Each interruption is triggered by a mouse click chosen pseudo-randomly by the software to spread interruptions out across a block of performance. The goal is 12 interruptions per 20-minute block, but there may be fewer when there are long pauses between responses. A selected mouse click triggers onset of a visual alert, which is a 2-inch-square window containing a pair of line-drawn “eyeballs” appearing in an unused corner of the screen and remaining visible through the warning interval. At onset the alert window flashes three times to simulate the eyeballs “blinking”, providing a series of visual transients spread over 600 msec.

After the warning interval, the interrupted-task display is replaced by the interrupting-task display, a simulated radar screen on which icons appear successively and have to be classified according to simple rules (e.g., Brock, Stroup, & Ballas, 2002). Interruption duration varied between 30 and 45 sec, a range in which variance seems to have little effect on resumption lag (C. A. Monk, personal communication, June 19, 2006) or overall performance (Gillie & Broadbent, 1989).

After the interruption, the interrupted-task display is restored without warning, and without the alert window. The state of the task at this point is exactly as it was after the participant’s last response before the interruption.

Procedure. Participants were tested individually. Each was trained for about 20 minutes by walkthroughs of the two tasks separately, then was left alone to work through three 20-minute blocks.

Design. The two independent variables of primary interest, both within-participants, were the serial *position* of responses following the interruption (levels 1-10) and *block* within session (levels 1-3).

The dependent variable was response time (RT). (Error measures were difficult to define, and overall game scores were too variable to register any effects.) Participants used a mouse to make responses, and we coded one click as one response, except for repeated consecutive clicks on the same object, which we aggregated into one response (affecting mostly fast, repeated clicks on a given scrolling widget). RT for position 1 was timed from interruption offset, and RT for later positions was timed from the previous response. Thus, each participant contributed 30 data points, with each point the median RT for that position for that block.

We manipulated two other variables, both between participants, to link our timecourse results to previous work. The first was *duration* of the warning interval between alert onset and interruption onset (Hodgetts & Jones, 2006b; Trafton et al., 2003); this had five levels: 0.3, 1.3, 2.3, 3.3, and 4.3 sec. (The integer components 0 through 4 were timed, and the fraction .3 was due to system overhead in launching the interrupting task.) The second was the *mode* of participant-task interaction during the warning interval, which had three levels. In *perception-enabled* mode, the display remained perceptually available through the warning interval (though frozen in its state at onset of the visual alert). In *perception-disabled* mode the display was replaced with a dark screen during the warning interval (except for the visual alert). These two modes allowed us to test for effects of perceptual cues being available during the warning interval (Hodgetts & Jones, 2006a). Responses were disabled during the warning interval

for both modes, because they would have been difficult to make in the perception-disabled condition. In *fully-enabled* mode, the display was visible and responses were enabled during the warning interval, so that participants could continue to work. We viewed this condition as having the highest ecological validity – one could imagine typing out a final thought as the phone rings before actually taking the call, for example – and expected the lowest recovery costs here, anticipating that participants would benefit from having some control over when to suspend their performance (Iqbal & Bailey, 2005; McFarlane & Latorella, 2002).

The full design, then, was 10 (position) x 3 (block) x 5 (duration) x 3 (mode). The 375 participants yielded 13,377 total interruptions.

Results

Aggregate data appear in Figure 1 (the full data set is posted with the supplemental materials). In each block, RT dropped smoothly from position 1 roughly to asymptote by position 10, with effects measured in seconds. Across blocks, RT decreased overall, but the shape of the recovery function also changed, with asymptote reached earlier in block 1 than in blocks 2 and 3.

Analysis of variance revealed main effects of position, $F(9,3240)=672.1$, $p<.001$, $\eta^2=.651$, and block, $F(2,720)=123.1$, $p<.001$, $\eta^2=.255$, but none of duration, $p=.108$, or mode, $p=.253$. There were three reliable interactions. The first was block x position, $F(18,6480)=3.33$, $p<.001$, $\eta^2=.009$, reflecting the change in shape of the recovery timecourse across blocks, which we interpret later in terms of our model. The second was block x mode, $F(4,720)=2.54$, $p=.039$, $\eta^2=.014$, which we do not try to interpret. The third was position x duration, $F(36,3240)=2.59$, $p<.001$, $\eta^2=.028$, which we probed

by separating position 1 (resumption lag) from positions 2-10. Position 1 RTs appear in Figure 2, averaged over block. These showed a duration effect, $F(4,360)=4.66$, $p=.001$, $\eta^2=.049$, as in previous studies (Hodgetts & Jones, 2006b; Trafton et al., 2003), but no mode effect, $F<1$, and no interaction, $F<1$. Neither variable interacted with block, $F_s<1$. Positions 2-10 showed no position x duration interaction, $p=.306$; thus, the position x duration interaction over positions 1-10 reflected the duration effect on position 1.

Planned comparisons on position 1 RTs (Figure 2) revealed no benefit for the perception-enabled condition over the perception-disabled condition, $p=.275$, unlike the finding by Hodgetts and Jones (2006a); their task environment was quite different, such that the functional significance of perceptual cues could conceivably have been greater. Fully-enabled differed neither from perception-enabled nor from perception-disabled, $F_s<1$, yielding no support for our hypothesis that control over when to stop working during the warning interval might facilitate resumption. It may be that control over actual interruption onset is what facilitated resumption in previous studies (Iqbal & Bailey, 2005; McFarlane & Latorella, 2002).

In sum, the main effect of position explained substantial variance in our data (65.1%) whereas warning-interval manipulations explained much less (2.8% for the position x duration interaction, with other effects not significant), hinting at structural constraints on recovery that preparatory processing cannot overcome; indeed, four additional seconds of warning interval recouped only about a second of resumption lag (Figure 2), which is hardly a net gain. Nonetheless, that we could replicate the duration effect found in previous studies suggests that our design tapped similar processes, and thus that our timecourse function, and the model we develop next, may generalize.

Modeling the Timecourse of Recovery

The empirical timecourse function in Figure 1 is smoothly curvilinear, suggesting a single recovery process playing out over time. To explain this process, we start with the premise that the mental representations supporting performance of a cognitively complex task are themselves complex, including elements like goals and subgoals, plans to achieve them, previous outcomes relevant now, etc. Crucially, we assume that these elements are associatively linked, reflecting some mix of semantic relationships (as in a semantic network), procedural relationships (as among steps of a plan or script), and episodic relationships (as in retrieval structures, Ericsson & Kintsch, 1995; see also Edwards & Gronlund, 1998, Experiment 1). We also assume that the strength of these associative links follows a gradient, falling off with increasing psychological distance between elements. Thus, in a hypothetical procedure with three steps, the first and second steps would be strongly linked, and the first and third would be more weakly linked. These assumptions imply that when the first step of such a plan is retrieved, it will prime retrieval of the second step strongly and retrieval of the third step weakly.

These assumptions explain the timecourse of recovery in terms of a cumulative priming effect as task-relevant representational elements are retrieved after interruption offset. If the hypothetical three-element procedure above happens to be what is performed at task resumption, step 1 would be retrieved first, leading to response 1, and also priming step 2; step 2 would be retrieved next, leading to response 2, and also, together with step 1, priming step 3. In general, as the number of retrieved steps grows, priming for successive retrievals will approach ceiling, and non-priming processes will come to be the limiting factors on response time.

To formalize the model tractably, we take the fall-off in link strength with psychological distance to be geometric, and assume that one representational element is retrieved per response. Then, for the response in position p after an interruption, total priming $A(p)$ flowing from already-retrieved elements to the to-be-retrieved element is

$$A(p) = -1 + \sum_{i=1}^p \text{assoc}^{i-1}, \text{assoc} < 1 \quad (1)$$

where assoc is the strength of the link between immediately neighboring elements in the to-be-retrieved representation. For example, with $\text{assoc}=0.5$, the response 1 element would be primed with 0 units of associative activation, the response 2 element with 0.5 units, the response 3 element with 0.75 units, etc.

To map the activation of to-be-retrieved element p to RT for the response that depends on retrieving p , we adopt the relationship

$$RT(p) = F \exp[-A(p)] \quad (2)$$

where F is a scale parameter that absorbs time for non-retrieval processing. This is the relevant formalism from the ACT-R cognitive theory (Anderson et al., 2004), on which our earlier theoretical work (Altmann & Trafton, 2002; Trafton et al., 2003) is also grounded.

The primed-retrieval model is then simply

$$RT(p) = F \exp\left[1 - \sum_{i=1}^p \text{assoc}^{i-1}\right], \quad (3)$$

with two parameters, F and assoc , to be estimated from data.

To fit the model, we estimated F and assoc separately for each participant for each block by minimizing root mean squared deviation (RMSD) from the data. The fits, averaged over participants within block, appear in Figure 1 (dashed lines). Descriptive measures of fit, also averaged over participants within block, were $r^2=.571$, $.659$, and

.646, and RMSD=.731 sec, .543 sec, and .531 sec, each for blocks 1, 2, and 3, respectively.

To test goodness of fit inferentially, we performed runs tests (see Bradley, 1968) on the signs of model-data deviations. If the model's timecourse function were flatter, for example, than whatever function was actually reflected in our data, deviations for early and late positions would tend toward one sign and deviations for middle positions would tend toward the other, producing fewer runs of same-sign deviations than if the deviations were randomly positive or negative. The runs statistic is asymptotically normal, so to obtain enough deviations per test we concatenated blocks within participant, producing 375 30-deviation series. The mean of the runs statistic across participants (-0.003) did not differ from 0, $t(374)=-0.071$, $p=.94$, suggesting that deviations were in fact random and that the model captures the systematic variance due to position. This test also suggests that the curvilinear pattern in Figure 1 is not an artifact of averaging; a curvilinear pattern is a structural prediction of the model, and any systematic deviations from this pattern at the participant level, beyond those that the F and $assoc$ parameters could accommodate, would have manifested in runs of same-sign deviations.

To test the model at a more conceptual level, we asked whether it could help interpret practice effects in our data. The model tracks the overall improvement with practice with a decrease in the scale parameter F across blocks (4.31, 3.97, 3.84), $F(2,720)=24.61$, $p<.001$, $\eta^2=.064$. More interesting is the block x position interaction reflecting a change in shape of the timecourse function particularly between blocks 1 and 2 (Figure 1), which the model tracks with an increase in the $assoc$ parameter (.464, .504, .504), $F(2,720)=26.79$, $p<.001$, $\eta^2=.069$. $Assoc$ represents connection strength between

the elements hypothetically being retrieved after interruption, so we would say that the block x position interaction reflects these connections being strengthened as participants settle into the use of particular procedures. Thus, the model offers a plausible account of a practice effect that may otherwise have been difficult to interpret.

Discussion

Our data show a recovery process operating over the first 15 sec or so after interruption offset (Figure 1), with a cumulative cost across positions of 4 to 5 sec. We characterize this process as reconstructing the rich episodic mental context necessary to perform a cognitively complex task, when this context is disrupted by an interruption. As each element of the context is reactivated, it contributes to the total amount by which the *next* retrieval target is primed. Changes in connection strength between representational elements, represented by changes in the *assoc* parameter, affect the shape of the recovery timecourse, explaining the block x position interaction in our data as a strengthening of links between related elements as participants improve their understanding of the task.

The shape of the recovery timecourse in Figure 1 is not complicated and may not be surprising, but related work suggests it could have been otherwise. In one particular kind of *task-switching* procedure, the participant performs one simple task, like judging the parity of a stimulus digit, repeatedly for a run of trials, after which a new cue is presented to signal the task for the next run. There are commonly two tasks – the other might be to judge digit magnitude relative to 5 – and the task cued for a given run is randomly selected. Presentation of a *repeat* cue, namely the same cue presented for the previous run, has the flavor of an interruption in which the participant must suspend performance briefly (just to process the cue) then resume that same task again afterwards.

Empirically, position 1 RT after a repeat cue is substantially slower than position 2 RT, but position 2 RT itself is typically at or below baseline (e.g., Altmann, 2004); in other words, recovery is complete by position 2, rather than playing out over multiple positions as it did here. Thus, empirical evidence from the task-switching domain suggests that a curvilinear speedup following task interruption was not a foregone conclusion. With our model in hand, however, we can explain the difference in recovery patterns in terms of a difference in task complexity; because the procedure for judging a digit's parity or magnitude spans only one response, there are no further retrievals on positions 2 and later to be facilitated by a cumulative priming effect.

Our model is consistent with old assumptions about the associative nature of thought (see, e.g., Anderson & Bower, 1973), and also with modern conceptions of expertise. The construct of long-term working memory, in particular, posits that an important dimension of expert performance is the use of rich associative scaffolding to store intermediate products (Ericsson & Kintsch, 1995; see also Oulasvirta & Saariluoma, *in press*). In some sense we build on this construct by adding processing assumptions about incremental retrieval during task resumption and by assuming a representational gradient in which neighboring elements of the target representation are better retrieval cues than elements separated by greater psychological distance.

Our model also lends a mechanistic interpretation to intuitive attentional constructs like flow (Csikszentmihalyi, 2005) and situational awareness (Endsley, 1995). Flow refers to a generalized sense of being in the moment, whereas situational awareness emphasizes the human operator's awareness of the current state of a highly dynamic task environment. In terms of our model, we would say that one condition for flow or

situational awareness – probably not a sufficient condition, but quite possibly a necessary one – is that the mental context is sufficiently well stocked with primes that the associative activation reaching the next retrieval target is at or near ceiling.

There are of course likely to be other explanations for our data. These do not seem to include standard practice effects, which play out over much longer timescales (e.g., Newell & Rosenbloom, 1981). They may, however, include a *warm-up effect* (e.g., Anderson, Fincham, & Douglass, 1999), the finding from the memory domain that successive recalls of a given item get faster. This speedup plays out over seconds, like our recovery timecourse, and is reset by a retention interval, like our timecourse is reset by an interruption. There is also intuitive appeal to the notion of “warming up” to a task again after an interruption. Nonetheless, a strict mapping of the warm-up effect to our data would seem to require that a single item is recalled successively on each position after every interruption, such that improved access to this one item is what drives recovery. This one item could be a short-term goal representation of some kind, as we have previously assumed (Trafton et al., 2003). However, we prefer our account, at the moment, because it offers a ready explanation of the block x position interaction in our data, and because the premise about reconstructing a rich, interconnected representation after an interruption seems quite close to first principles of performance in a cognitively complex task. Future work will have to examine these and other competing models, perhaps even ones that start from the same premise but make different processing assumptions that lead to alternative formal expressions.

References

- Altmann, E. M. (2004). Advance preparation in task switching: What work is being done? *Psychological Science, 15*, 616-622.
- Altmann, E. M., & Trafton, J. G. (2002). Memory for goals: An activation-based model. *Cognitive Science, 26*, 39-83.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111*, 1036-1060.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- Anderson, J. R., Fincham, J. M., & Douglass, S. (1999). Practice and retention: A unifying analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*, 1120-1136.
- Bradley, J. M. (1968). *Distribution-free statistical tests*. Englewood Cliffs, NJ: Prentice-Hall.
- Brock, D., Stroup, J. L., & Ballas, J. A. (2002). Effects of 3D auditory display on dual task performance in a simulated multiscreen watchstation environment. In *Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society*. Baltimore, MD.
- Brock, D. P., & Trafton, J. G. (1999). Cognitive representation of common ground in user interfaces. In J. Kay (Ed.), *User modeling: Proceedings of the 7th international conference*. New York: Springer-Wien.

- Chisholm, C. D., Dornfeld, A. M., Nelson, D. R., & Cordell, W. H. (2001). Work interrupted: A comparison of workplace interruptions in emergency departments and primary care offices. *Annals of Emergency Medicine, 38*, 146-151.
- Csikszentmihalyi, M. (2005). Flow. In A. J. Elliot & C. S. Dweck (Eds.), *Handbook of competence and motivation* (pp. 598-608). New York: Guilford Publications.
- Czerwinski, M., Horvitz, E., & Wilhite, S. (2004). A diary study of task switching and interruptions. In *CHI 2004 Proceedings* (pp. 175-182). New York: ACM Press.
- Edwards, M. B., & Gronlund, S. D. (1998). Task interruption and its effects on memory. *Memory, 6*, 665-687.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors, 37*, 32-64.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102*, 211-245.
- Gillie, T., & Broadbent, D. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research, 50*, 243-250.
- Hodgetts, H. M., & Jones, D. M. (2006a). Contextual cues aid recovery from interruption: The role of associative activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 1120-1132.
- Hodgetts, H. M., & Jones, D. M. (2006b). Interruption of the Tower of London task: Support for a goal activation approach. *Journal of Experimental Psychology: General, 135*, 103-115.

- Iqbal, S., & Bailey, B. P. (2005). Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *CHI 2005 Proceedings* (pp. 1489-1492). New York: ACM Press.
- 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.
- Monk, C. A., Boehm-Davis, D. A., & Trafton, J. G. (2004). Recovering from interruptions: Implications for driver distraction research. *Human Factors, 46*, 650-663.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1-55). Hillsdale, NJ: Erlbaum.
- Oulasvirta, A., & Saariluoma, P. (in press). Surviving task interruptions: Investigating the implications of long-term working memory theory. *International Journal of Human-Computer Studies*.
- 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.
- Swets, B. (2006). *"Where was I?": A psycholinguistic investigation of conversational interruptions*. Unpublished PhD, Michigan State University, East Lansing.
- Thomson, C. (2005, October 16). Meet the life hackers. *New York Times Magazine*, 40-45.

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.

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.

Author Note

This work was supported by the US Office of Naval Research (EMA: N000140610077, N000140310063; JGT: N0001405WX20011, N0001405WX30020). The authors thank Helen Hodgetts and an anonymous reviewer for their comments and suggestions for improvement.

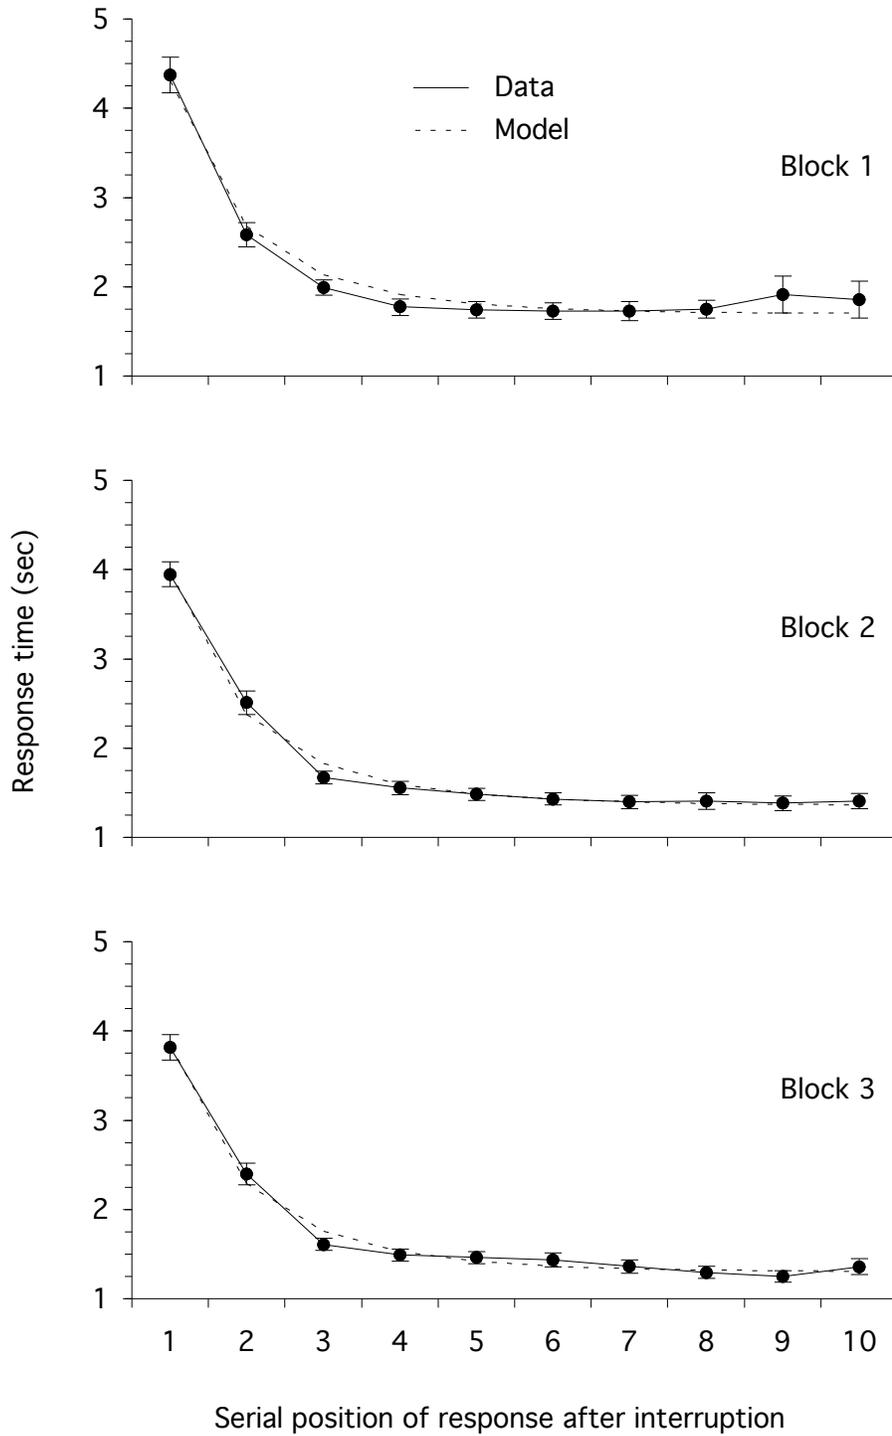


Figure 1: Average response times from the experiment (solid lines) and average model fits (dashed lines), plotted by block (1-3) and serial position after interruption offset (1-10). Error bars are 95% confidence intervals.

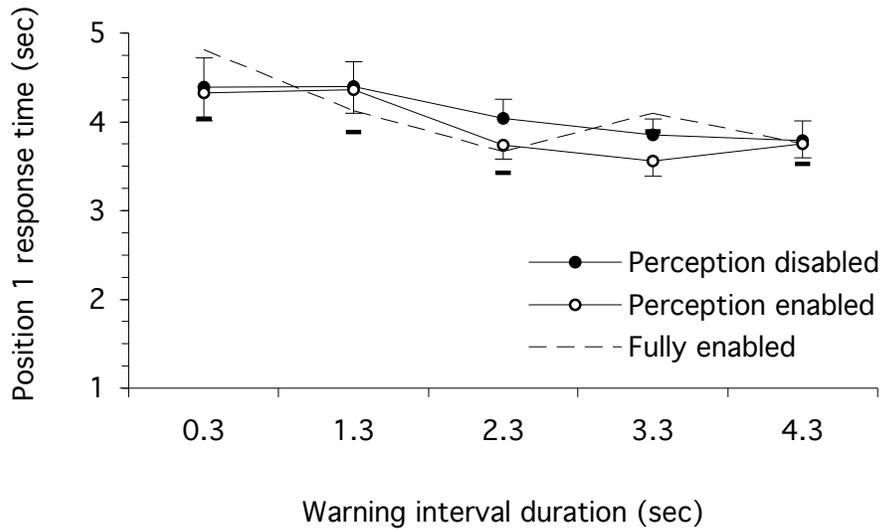


Figure 2: Average response times from the experiment for serial position 1 after interruption offset (i.e., resumption lags), plotted by duration of the warning interval (abscissa) and interaction mode during the warning interval (lines). Error bars are 95% confidence intervals, up for perception-disabled, down for perception-enabled, and thick down for fully-enabled.