Balancing Structural and Temporal Constraints in Multitasking Contexts

Dario D. Salvucci (salvucci@drexel.edu)
Department of Computer Science, Drexel University
3141 Chestnut St., Philadelphia, PA 19104, USA

Tuomo Kujala (tuomo.kujala@jyu.fi)
Department of Computer Science and Information Systems, University of Jyväskylä
P.O. Box 35, FI-40014 Jyväskylä, Finland

Abstract

Recent research has shown that when people multitask, both the subtask structure and the temporal constraints of the component tasks strongly influence people's task-switching behavior. In this paper, we propose an integrated theoretical account and associated computational model that aims to quantify how people balance structural and temporal constraints in everyday multitasking. We validate the theory using data from an empirical study in which drivers performed a visual-search task while navigating a driving environment. Through examination of illustrative protocols from the model and human drivers as well as the overall fit on the aggregate glance data, we explore the implications of the theory and model for time-critical multitasking domains.

Keywords: Multitasking; driving; cognitive architectures

Introduction

Multitasking often occurs in time-critical situations, such as answering a ringing phone while babysitting, cooking over a stove, or driving a vehicle. In these situations, the structure of one or both tasks may impose a sense of urgency to complete a task, perhaps due to the environment (e.g., a pot boiling over) or perhaps due to self-imposed pressures. For instance, consider Janssen, Brumby, and Garnett's (2012) example of texting the message "Running late" while driving: even as driving is clearly the task with the highest priority, a driver who is almost done typing the message ("Running lat...") will feel strongly compelled to finish typing before continuing to drive. In such situations, we continually balance our urgency to complete one task with the urgency imposed by other tasks.

A wealth of recent research has explored how we multitask, both in constrained laboratory studies and in complex real-world environments. One general finding is that task structure—how a task breaks down into smaller subtasks—strongly affects how people perform that task concurrently with other tasks (e.g., Borst, Taatgen, & van Rijn, 2015; Iqbal & Bailey, 2005). Complementary studies have shown that a task's temporal constraints can strongly affect multitasking; arguably the most studied context is that of driving, for which studies have explored the relationship between driving urgency (or uncertainty) and time looking away from the road (e.g., Kujala et al., 2015; Lee, Gibson, & Lee, 2015).

In this paper, we examine the critical relationship between structural and temporal constraints on multitask behavior and performance. Although there are multitasking scenarios in which either the structural or the temporal constraints are dominant, they are both generally present in some form: structural constraints (e.g., the chunking of a telephone number: Brumby, Howes, & Salvucci, 2009; Janssen, Brumby, & Garnett, 2012) still appear in very time-critical domains like driving, and temporal constraints (e.g., the limited time for answering a ringing phone) still appear in primarily structurally-driven domains. Although most studies have focused on only one of these constraints at a time, recent studies have examined how overall task priorities affect task switching (e.g., Brumby, Howes, & Salvucci, 2009; Janssen, Brumby, & Garnett, 2012) and glance durations between tasks (Lee, Gibson, & Lee, 2015). As yet, however, a rigorous cognitive process model that quantifies this relationship has proven elusive. Here, we propose a theoretical framework to help understand and quantify the balance between subtask structure and temporal constraints.

Balancing Structure and Time

The issues that arise in balancing structural and temporal constraints are perhaps best illustrated by behavioral protocols collected from people acting in multitasking contexts. For this purpose, we delve more deeply into a recently collected data set in which drivers (N=12) were asked to perform visual-search tasks while driving (Kujala & Salvucci, 2015). In the experiment, drivers searched through multiple screens of 6, 9, or 12 songs for a particular target song. The screens were laid out in one of two ways: a *Grid* layout, with 2, 3, or 4 rows of 3 songs each; or a *List* layout, with the songs listed vertically top to bottom. If the target song did not appear on a given screen, the driver was asked to press the down-arrow button to advance to the next screen. Here we are concerned only with screens without the target, in which drivers searched through all the songs. Drivers performed the search tasks in a driving simulator, on a display to the right of the steering wheel, and navigated a straight three-lane road at highway speed (80 km/h), occasionally performing the search task when requested.

The visual-search and driving task includes the types of structural and temporal constraints found in many multitasking contexts. The search task comprises two subtasks that repeat for each screen, namely to ***search*** through the on-screen items and to ***press*** the button when finished. At the same time, the driving task involves increasing urgency over time as the driver looks away from the road, eventually reaching the point where the driver needs to look back. Thus, the balance between completing the search task and driving safely created the key challenge for the driver in managing both tasks concurrently.

The original treatment of these data (Kujala & Salvucci, 2015) focused on aggregate analysis and modeling of human behavior in this task. In focusing on individual behavior protocols in the data, however, we found that the aggregate analysis shrouded the interesting cognitive strategies that arose on individual trials. To this end, we now focus on examining the human protocols from the study, and in the next section develop a much more in-depth model that matches both individual and aggregate behavior.

Switching at Subtask Boundaries. As mentioned earlier, many studies have shown that people tend to switch tasks at subtask boundaries. Participants in the present study were no exception, and often switched from ***search*** to driving and from ***press*** to driving after each of these subtask boundaries. Figure 1 shows a classic example of this type of switching, taken from a person performing the search task on the Grid-6 display. Specifically, the figure shows the timeline of screen glances and button presses as the person searches through 3 consecutive screens for the target. Throughout the timeline, we see an alternation between glances for each subtask: a first glance for ***search*** and a second for ***press***. The emerging pattern is one of switching at subtask boundaries, with the total number of glances equal to the total number of subtasks completed (2 glances for each of 3 screens).

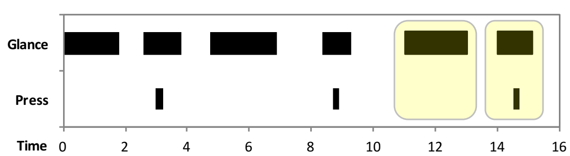

Figure 1: Human protocol showing task switching at subtask boundaries (highlighted: switch after ***search***, switch after ***press***).

Self-Interrupting during Subtasks. Although switching at subtask boundaries is commonly observed in the study protocols, there are at least two other common behaviors. One such behavior involves interrupting oneself (or "giving up": Bogunovich & Salvucci, 2011) during a subtask, when a person decides during a subtask that s/he needs to switch immediately rather than complete the subtask. Figure 2 illustrates this behavior in the timeline of glances and presses for a person doing 3 trials in the Grid-12 condition. Here, for each trial, the person divides the ***search*** subtask into three roughly equal components, and finally makes a fourth shorter glance to press the button.

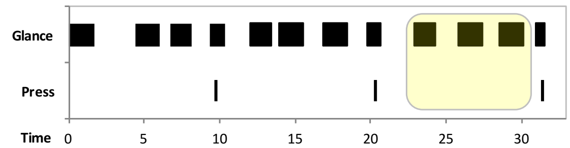

Figure 2: Human protocol showing self-interruption during subtask (highlighted: multiple glances during ***search***).

Continuing beyond Subtask Boundaries. Besides self-interrupting before subtask boundaries, people also exhibit the behavior of continuing beyond subtask boundaries—that is, reaching the subtask boundary, but then continuing to the next subtask rather than switching away from the task. Figure 3 shows the sample protocol for one example in the List-12 condition (see highlighted area): the driver presses the button and then immediately continues to search the next screen. Unfortunately, for some of these cases, continuing to the next subtask results in a long off-road glance lasting over 2 seconds, illustrating a perhaps unintuitive effect: shorter tasks and faster behavior may sometimes actually lead to longer glance times at a task, because it tempts a person into continuing with the next subtask.

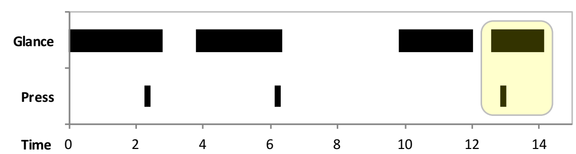

Figure 3: Human protocol showing subtask continuation (highlighted: ***press*** followed by ***search*** in same glance).

A Theory and Computational Model

The main goal of our work here is to better understand the interplay between subtask boundaries and temporal urgency illustrated by the above examples. In particular, we aim to develop a computational model that can run in simulation and thus produce behaviors directly comparable to those of human participants. In this section we describe the details of our theoretical account and computational model, to be validated with human data in the next section.

An individual task can generally be thought of as a hierarchy of higher- and lower-level tasks (or goals) and subtasks (or subgoals) (see, e.g., Schraagen, Chipman, & Shalin, 2000). To account for the interaction of structural and temporal constraints, we require that the hierarchical decomposition also specify the timing of the various components. The simplest way to achieve this goal is to assign times to the actions at the leaves of the hierarchy tree; one might assume constant times (e.g., taken from the keystroke-level model: Card, Moran, & Newell, 1980), or variable times based on aspects of the current environment (e.g., mouse movement over different distances, cognitive delays due to recalling information). Taken further, the task hierarchy could be instantiated as a running computational model that adapts continuously to its environment—for example, a model developed using a computational cognitive architecture (e.g., Anderson, 2007; Newell, 1990). For our purposes here, a task hierarchy augmented with action times is sufficient to serve as a model of an individual task.

Given two task models, we would like to express how structural and temporal constraints are balanced to dictate how people switch tasks in a multitasking context. As a first step toward this goal, we define *urgency* as a measure of a person's desire to work on a given task; in essence, each task in a multitasking context has an associated urgency, and generally speaking, people tend to switch to (or continue with) the most urgent task at a given time. Urgency provides a convenient way to formulate the effects of structure and time into a single measure, and a way to evaluate concurrent tasks and decide whether and when to switch between them.

Urgency and Structural Constraints. Earlier we noted how task structure has a strong influence on multitasking; how can we formulate this influence in terms of urgency? Empirical studies have made clear that people value the completion of a subtask, and thus, a person's urgency should increase as s/he approaches a subtask boundary. If we assume a strong association between urgency and time, we could say that a person receives a "reward" upon completion of each subtask (e.g., Fu & Anderson, 2006), and that the reward is equal to the time spent on that subtask.

The value of receiving a reward at the completion of a subtask may also be propagated back to the actions that led to successful completion. One method allows prior actions to receive a boost to their urgency, but with a discount factor γ between 0 and 1 that reduces the reward with increasing temporal distance from the actual subtask completion. Consider a subtask S^i comprised of actions a_k^i with times (durations) t_k^i. The total reward $R(S^i)$ for this subtask is computed as the sum of its action times:

$$R(S^i) = \sum_k t_k^i$$

We then compute the structural urgency $U_s(a_k^i)$ of a particular action a_k^i as a function of the subtask reward, $R(S^i)$, and the remaining time between the end of the action and the completion of the subtask, $T(a_k^i)$:

$$T(a_k^i) = \sum_{j>k} t_j^i$$

$$U_s(a_k^i) = \gamma^{T(a_k^i)} \cdot R(S^i)$$

The final action receives the full reward ($T(a_k^i) = 0$), and each action before the final action receives the reward discounted by γ and the remaining time to completion. For example, Figure 4 graphs the structural urgency for 10 actions of 300 ms with different values of γ. This produces a hook-like function with larger γ values producing a flatter urgency function (earlier actions receiving more reward) and smaller values producing a sharper curve.

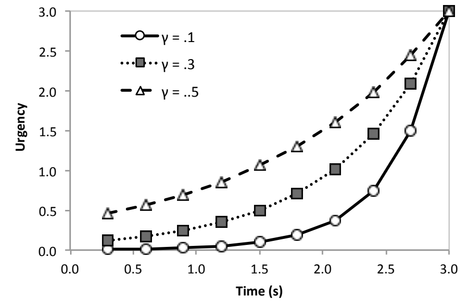

Figure 4: Sample structural urgency profiles for 10 actions each with a duration of 300 ms.

Structural urgency as defined thus far accounts for people's preferences in switching at subtask boundaries. A complementary empirical finding is that as people complete one subtask in a multitasking scenario, they are generally averse to continuing to the next subtask unless they feel they have sufficient time to complete that one as well (Bogunovich & Salvucci, 2011). This finding suggests that people have an awareness of the time needed to complete a full subtask, and that they use this information in deciding whether or not to continue. In terms of urgency, the aversion to continuing to another subtask can be represented as a negative urgency at the start of the subtask. Specifically, we define a continuation penalty when continuing from one subtask to the next, whereby we subtract the full duration of the next subtask from the structural urgency. In the context of a multitasking scenario, continuing to the next subtask will generally have a lower urgency than switching to another task; however, if no other task can proceed and/or other tasks have even lower urgency, the next subtask may then proceed.

Urgency and Temporal Constraints. In addition to the urgency contributed by its structure, a task will often have an associated temporal urgency—a feeling that compels a person to complete the task as soon as possible. Temporal urgency is influenced by the amount of time passed since switching away from the task, with urgency (typically) increasing with the passage of time. We define temporal urgency $U_t(\Delta t)$ as a function of the time since switching away Δt. The specific form of this function depends heavily on the task domain: highly time-critical domains will have a steep function with urgency rapidly increasing over time, whereas less time-critical domains will have flatter functions. In the next section we will see a concrete example of such a function for a time-critical task domain.

Deciding when to Switch Tasks. In the case of multiple concurrent tasks, we use a decision mechanism similar to the conflict resolution mechanism in ACT-R (Anderson, 2007) to determine which task will progress at a given time. First, for each task, the total urgency U is computed as the sum of its structural urgency, temporal urgency, and noise factor ϵ:

$$U(a_k^i, \Delta t) = U_s(a_k^i) + U_t(\Delta t) \ + \epsilon$$

As in ACT-R, the noise ϵ is sampled from a logistic distribution, with mean $\mu = 0$ and scale s, a free parameter to be estimated (described shortly). Then, the task with the larger urgency is allowed to proceed.

Computational Simulation. We implemented the proposed framework as a Java simulation system to enable rapid testing of models and parameters settings. The system takes as input a model as described above (with subtasks and associated action times), and generates sequential behavioral protocols as output. The protocols can then be analyzed for more aggregate measures, such as the common measures of glance times to be used shortly in the forthcoming study.

Study: Visual Search and Driving

To test the proposed approach, we return to the domain discussed earlier, namely visual search and driving. The prior model in Kujala & Salvucci (2015) focused mainly on the temporal constraints of the driving task, dynamically adjusting an in-car glance duration threshold according to the stability of the vehicle after switching back to the driving task. While this prior model provides a good account of the aggregate data, it does not conform well to the individual protocols shown earlier—notably, because it does not account for the structural constraints of the visual-search task (e.g., it is equally likely to switch early or late in the visual search, whereas people show a tendency to avoid switching late in the task). Here, we use the theoretical framework above to model the two tasks of visual search and driving, and then illustrate how the theory and simulations produce behavior that better resembles that of human participants.

Models of Visual Search and Driving

As discussed earlier, the visual-search task breaks down into a repeated iteration of two basic subtasks: a ***search*** of each of the on-screen items; and (assuming the target is not found, which is always the case for the screens analyzed here) a ***press*** of the down-arrow button to advance to the next screen. The ***search*** subtask includes an encoding action for each of the on-screen items—that is, 6, 9, or 12 actions to match the items in that particular condition. The ***press*** subtask contains a single action to press the button. The time required for each action was derived from simulations of the earlier ACT-R model of this task (Kujala & Salvucci, 2015): 368 ms per item for ***search*** in the Grid layout; 291 ms per item for ***search*** in the List layout; and 564 ms for ***press*** in all cases.

The model of driving used here is derived from the ACT-R model of driver behavior (Salvucci, 2006). The core subtask is a cycle that visually encodes the near and far points of the road and updates the vehicle controls accordingly. These actions require a total of 200 ms to complete the cycle, and thus this is also the duration of the driving subtask, which is simply repeated while the model is actively driving. However, because the focus of our analysis here is on behavior in the visual-search task, we simply assume that the model drives for 1 second (5 cycles) and then switches back to the visual-search task.

Beyond the above structural details, we also require some formalization of the temporal constraints of the driving task in particular. There have been several attempts to quantify a driver's cognitive state while looking away from the road, most notably in terms of uncertainty (see Kujala et al., 2015): as time passes, the driver's uncertainty about the external environment (lane position, speed, other vehicles, etc.) gradually increases until it reaches a point at which the driver feels compelled to look back to the road. We translate these ideas into the temporal urgency of the driving task as follows. When the driver has stabilized the driving task, the urgency of further driving updates is rather low. We define a value U_{stable} to denote the low urgency of driving in this condition, a value analogous to the uncertainty threshold in prior work (e.g., Kujala et al., 2015). This value is presumed to be negative to indicate a lack of urgency—that is, it indicates that time might better be spent on some other task. In all, we define the temporal urgency of driving as:

$$U_t(\Delta t) = U_{stable} + \Delta t$$

When the driver looks away from the road ($\Delta t = 0$), the temporal urgency is equal to U_{stable}; however, as time passes and the driver continues the secondary task, the urgency of driving climbs steadily, eventually passing zero and becoming positive if the driver does not switch back to driving within U_{stable} seconds.

When we combine these models of visual search and driving, we can visualize their competing urgencies as a function of time, as illustrated in Figure 5. The urgency of the search task builds gradually because of the increasing urgency to finish the task, ramping up quickly toward the end of the subtask. Meanwhile, the urgency of driving starts low (at the assumed U_{stable} level) but increases over time due to increasing levels of uncertainty. At each point, the two urgencies are compared using the noisy conflict resolution process described earlier, forcing a switch to driving if the urgency of driving exceeds that of the search task. The graph on the right shows the probability of switching to driving at the various times: highest in the middle of the search subtask, and lowest early in the process (because driving still has a very low urgency) and late in the process (because there is high urgency to complete the search subtask). The resulting probability distribution is thus an emergent property of the theoretical mechanisms.

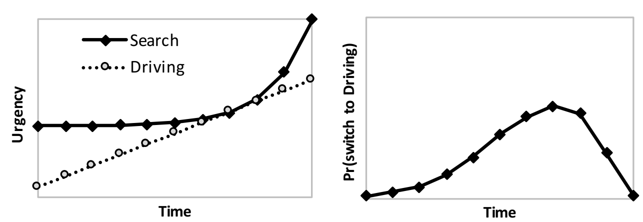

Figure 5: Sample urgency graph for visual search and driving, with associated probability of switching to driving.

For these models, we estimated the three free parameters (the urgency value U_{stable}, the scaling factor γ, and the noise scale s) by running 1000 simulations per parameter-value combination and finding the values that produced the best fit on the aggregate data described later. The estimated values were $U_{stable} = -1.7$, $\gamma = 0.1$, and $s = 0.3$.

Model Behavior and Results

Kujala & Salvucci (2015) examined five separate measures (30 data points total) of aggregate behavior by the human participants. Because of space constraints and because our focus lies primarily in the individual protocols, we avoid a detailed comparison of these aggregate measures here. However, it should be noted that for these five measures, the overall fit of the current model was very much on par with that of the previous model. Table 1 includes the correlation (*R*) and normalized root-mean-squared-error (*RMSE*/*mean*) for both models for all measures.

Table 1: Model-human correlations and errors for prior model (Kujala & Salvucci, 2015) and current model.

	Prior Model		*Current Model*	
	R	*NRMSE*	*R*	*NRMSE*
Number of in-car glances	.99	.32	.96	.20
Total in-car glance duration	.97	.08	.99	.05
Mean in-car glance duration	.81	.13	.62	.15
Maximum in car glance duration	.94	.30	.83	.05
Percent glances over 2 seconds	.65	.31	.62	.31

While the current model matches aggregate behavior as well as the prior model, the current model importantly provides a much better account of the behavior of individual participants and trials by accounting for both temporal and structural constraints. Figure 6 shows one such behavior for the model in the Grid-6 condition, namely the classic strategy of switching at subtask boundaries. The upper portion of the graph shows a timeline of the model's glances and button presses—again, analogous to our earlier analysis of human data. The lower portion shows the competing urgency between search and driving over time. For the first three screens, including behavior for the section screen at the point labeled **A1**, the model begins the ***search*** subtask; the urgency of driving steadily grows from its starting U_{stable} value, but the urgency of completing the ***search*** subtask grows as well. When ***search*** is done, the urgency to continue with the next subtask (***press***) includes the continuation penalty defined earlier, namely subtracting the duration of the next subtask; in essence, the model is checking whether there is sufficient time to complete the next subtask, and if not, it switches back to driving. At the next opportunity, though, the model completes the pressing subtask and switches back. At the point labeled **A2**, the model switches slightly earlier but then completes the search as well as the button press on the next glance.

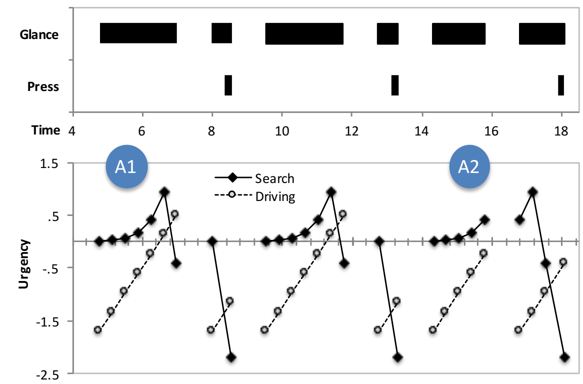

Figure 6: Model protocol showing task switching at subtask boundaries (after ***search*** and ***press***).

Not surprisingly, self-interruption during subtasks becomes more common as the number of on-screen items increases. Figure 7 shows an example in the Grid-9 condition. At point **B1**, we see a typical behavior in which the model reads several items during one glance, several more in a second glance, then finishes reading and finally makes the button press on the third glance. At point **B2**, the model splits up the item reading differently, but the end result is still a total of three glances to complete the ***search*** and ***press***, instead of two glances in the canonical behavior in Figure 6. The behaviors for the other screens show similar patterns; note that because of the noisy conflict resolution mechanism, a lower urgency can sometimes "win" over a higher one, producing a similarly large variety of protocols as for human drivers.

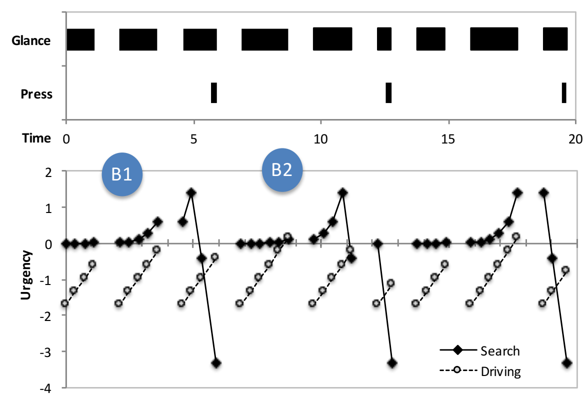

Figure 7: Model protocol showing self-interruption during subtask (multiple glances during ***search***).

As the number of on-screen items decreases, or the duration of individual actions decreases (from Grid to List), the model adapts by occasionally continuing beyond subtask boundaries, as shown in Figure 8 in the List-6 condition. At point **C1**, the model finishes searching the 6 on-screen items so quickly that the urgency for the next subtask, ***press***, is very close to that of driving; in this case (with noise), the model continues and presses the button before switching back to driving. Point **C2** illustrates a different form of continuation: after pressing the button, the urgency of driving is still very low, and again the model decides to continue and begin searching the on-screen items. The end result for this segment of behavior is a total of 4 glances to complete 3 screens, instead of the 6 glances (2 per screen) that would result from switching at subtask boundaries. While continuation after pressing the button was observed as a relatively infrequent behavior for human and model alike, the presence of this behavior at all (again, in both human and model) emphasizes the flexible nature of the balance between structural and temporal constraints.

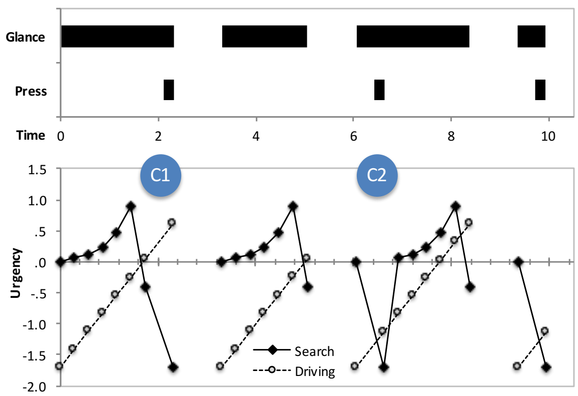

Figure 8: Model protocol showing subtask continuation (**C1**: ***press → search***; **C2**: ***search → press***).

General Discussion

The complex relationship between structural and temporal constraints presents a fascinating challenge when examining everyday multitasking behaviors, especially those in time-critical contexts. The concept of urgency developed here offers a way to unify these two important factors on multitasking, both in understanding the human behaviors that emerge, and in formalizing rigorous computational models to predict behavior in novel situations. One might consider urgency as related to task priorities that influence behavior through rational adaptation (e.g., Howes, Lewis, Vera, 2009). Empirical work along these lines have focused on manipulating the overall priority of each task (e.g., instructing participants to focus on one task or the other: Janssen, Brumby, & Garnett, 2012). Our treatment here is complementary in focusing on the rise and fall of urgency at a second-by-second level, being closely dependent on the lower-level conditions of each task. Urgency thus helps to formalize how people get "hooked on" subtasks, and how they balance structural urgency of subtasks with the temporal urgency of time-critical task domains.

As a step in this direction, the formulation of urgency has potential for incorporation into larger theories of cognition. For example, the ACT-R cognitive architecture (Anderson, 2007) posits that behavioral rules have an associated utility that can be learned and adapted using reinforcement mechanisms similar to those here (Fu & Anderson, 2006). However, whereas each rule has only one utility, a particular instantiation of a rule can have different urgency values depending on its place in the subtask structure. Urgency is more akin to threaded cognition's (Salvucci & Taatgen, 2011) "least-recently-used" heuristic in choosing the next cognitive thread to run; the heuristic might be subsumed by an improved formulation of urgency as a dynamic property of a complex task.

References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York: OUP.

Bogunovich, P., & Salvucci, D. D. (2011). The effects of time constraints on user behavior for deferrable interruptions. In *Proc. CHI '11*, 3123-3126.

Borst, J. P, Taatgen, N. A., & van Rijn, H. (2015). What makes interruptions disruptive? A process-model account of the effects of the problem state bottleneck on task interruption and resumption. In *Proc. CHI '15*, 2971-2980.

Brumby, D. P., Howes, A., & Salvucci, D. D. (2009). Focus on driving: How cognitive constraints shape the adaptation of strategy when dialing while driving. In *Proc. CHI '09*, 1629-1638.

Card, S. K., Moran, T. P., & Newell, A. (1980). The keystroke-level model for user performance time with interactive systems. *Comm. of the ACM*, *23*, 396-410.

Fu, W.-T., & Anderson, J. R. (2006). From recurrent choice to skill learning: A reinforcement-learning model. *Journal of Experimental Psychology: General*, *135*, 184-206.

Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological Review*, *116*, 717-751.

Iqbal, S. T., & Bailey, B. P. (2005). Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In *Proc. CHI '05*, 1489-1492.

Janssen, C. P., Brumby, D. P., & Garnett, R. (2012). Natural break points: The influence of priorities and cognitive and motor cues on dual-task interleaving. *Journal of Cognitive Engineering & Decision Making*, *6*, 5-29.

Kujala, T., Mäkelä, J., Kotilainen, I., & Tokkonen, T. (2016). The attentional demand of automobile driving revisited: Occlusion distance as a function of task-relevant event density in realistic driving scenarios. *Human Factors*, *58*, 163-180.

Kujala, T., & Salvucci, D. D. (2015). Modeling visual sampling on in-car displays: The challenge of predicting safety-critical lapses of control. *International Journal of Human-Computer Studies*, *79*, 66-78.

Lee, J. Y., Gibson, M., & Lee, J. D. (2015). Secondary task boundaries influence drivers' glance durations. In *Proc. AutomotiveUI '15*, 273-280.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors*, *48*, 362-380.

Salvucci, D. D., & Taatgen, N. A. (2011). *The Multitasking Mind*. New York: OUP.

Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (2000). *Cognitive Task Analysis*. New York: Psychology Press.