When to Switch? Understanding How Performance Tradeoffs Shape Dual-Task Strategy

Duncan P. Brumby (Brumby@cs.ucl.ac.uk)
Nina del Rosario (Nina.Del@gmail.com)
Christian P. Janssen (C.Janssen@ucl.ac.uk)
UCL Interaction Centre
University College London
London WC1E 6BT UK

Abstract

A novel dual-task paradigm was used to investigate how people adapt their task interleaving behavior to meet a specific performance objective. The study required participants to encode and enter a series of route instructions from a secondary display while driving a simulated vehicle. Explicit instructions were given to give greater priority to either safe driving or rapid completion of the secondary navigation task. Results showed that participants met the required task objective by varying the frequency and duration of visits to the secondary task display, and by also varying the amount of time given up to steering control in between visits. We explain these data using a framework for modeling driver distraction effects. The model predicted the observed shift in task performance between the two focus conditions and also the observed change in task interleaving strategy. Taken together these results support the idea that people can strategically control the allocation of attention in multitask settings to meet specific performance criteria.

Keywords: Multitasking, cognitive modeling.

Introduction

Consider for a moment a driver following a set of written directions to reach an unfamiliar destination. As the driver approaches a junction, they might want to consult their directions, and in doing so must consider the risks of taking their eyes off the road ahead. A safe driver, given the opportunity, might pull over to study their directions, or if this is not possible, they might choose to make many brief glances to the instructions. A risky driver, on the other hand, may choose to look away from the road for prolonged periods to study the directions in detail. In this way, the frequency and duration of attention shifts between tasks is determined by the relative importance of each task, and also a judgment of safe and acceptable behavior.

It is well known that in many multitasking situations, such as the one sketched above, constraints on the human cognitive architecture limit the extent to which tasks are performed in parallel (Meyer & Kieras, 1997). How people control the allocation of resources to multiple concurrent tasks is a topic of considerable theoretical and practical interest (e.g., Navon, & Gopher, 1979; Norman & Bobrow, 1975; Salvucci & Taatgen, 2008; Wickens, 2002).

One important application of multitasking theory has been to understand driver distraction. Driving is a safety critical task performed by millions of people on a daily basis, and with the growing ubiquity of mobile and in-car devices there are concerns about the deleterious effects of driver distraction. In this area, many studies have investigated the impact of cell phone dialing on driving performance. Typical results show that drivers tend to dial chunks of digits at a time, returning their attention to driving in between each chunk (Brumby, Salvucci & Howes, 2009; Salvucci, 2005). This pattern of task interleaving might reflect the fact that the dialing task has a strong representational structure that is difficult to disrupt, and this could be used to guide decisions about when to switch attention between tasks (Salvucci, 2005). But how might people decide how to interleave tasks in situations where there are no natural cues to guide this decision?

Salvucci and Taatgen's (2008) threaded cognition theory assumes that relatively complex multitasking behavior can emerge from a simple bottom-up process without the need for any explicit top-down control structures. The theory assumes that the cognitive system processes task threads using a least-recently-processed scheduling heuristic. While this theory offers a parsimonious account of multitasking behavior, it is not clear how this account allows the cognitive system to make strategic decisions to favor one task over another. Indeed, a large body of empirical work demonstrated that people can make explicit decisions about how to allocate attention to different tasks in multitask settings by prioritizing performance on one task over another (e.g., Brumby et al., 2009; Horrey et al., 2006; Gopher et al., 1982; Gopher, 1993; Wang et al., 2007).

One possibility for how people might adapt their dual-task strategy to meet a specific task objective is that they monitor the amount of time that has elapsed since they last checked on the more important task. Kushleyeva, Salvucci, and Lee (2005) found that when participants were required to monitor a safety-critical dynamic task, they adapted their monitoring behavior to changes in the temporal demands of the task. This suggests that the safer driver in the example above might simply set a lower threshold for the amount of time that they are prepared to take their eyes off the road, and in doing so, will interleave attention between tasks more frequently.

Another possibility is that people select strategies to meet a desired dual-task performance tradeoff objective. Brumby, Salvucci, and Howes (2009) have shown that in the case of manually dialing a standard US telephone number while driving, dialing three or four digits at a time is a particularly efficient strategy because any more interleaving incurs additional time costs without significant improvement in lane keeping, and any less interleaving sacrifices safety. To demonstrate this claim, Brumby et al. derived performance predictions for a range of dual-task strategies using a computational model. This approach of explicitly considering the performance tradeoffs involved for choosing between various dual-task allocation strategies is similar to that of defining a Performance Operating Characteristic (Norman & Bobrow, 1975; Navon & Gopher, 1979). The analysis by Brumby et al. showed that one limitation of the dialing-while-driving paradigm is that interleaving at the natural subtask boundaries of this task often corresponds with the most efficient dual-task interleaving strategy, in terms of completing the secondary dialing task in a relatively safe and timely manner.

In this paper, we investigate multitasking behavior using a novel dual-task paradigm. The paradigm, developed by Del Rosario (2009), requires participants to look at a secondary display to encode and enter a series of route instructions while driving a simulated vehicle. The benefit of this paradigm, over the classic dialing-while-driving paradigm, is that it does not have an external representational structure that can be used to guide decisions about when to interleave. Thus, participants are free to interleave the tasks how they like.

We use this paradigm to investigate how people adapt their dual-task interleaving behavior to meet varying performance objectives. In particular, we manipulate the experimental instructions and feedback given to participants to encourage either safe driving or rapid completion of the secondary navigation task. We consider how this change in task objective affects task performance and also the decision about when to interleave attention between tasks. Finally, we seek to apply Brumby, Salvucci, and Howes' (2009) model of how people interleave cell phone dialing and driving to this novel dual-task paradigm. An important question is whether the model will generalize to this new task setup, and if so, whether it will predict how people choose to interleave in each condition.

Experiment

Method

Participants. Sixteen participants (five female) took part in the study. Participants were unpaid volunteers, aged between 21- and 42-years (M=28.3 years). All had a valid driver's license and at least two years of driving experience.

Materials. The experiment used a dual-task setup in which participants had to complete a secondary navigation task while driving a simulated vehicle. Figure 1 shows how the two task displays were arranged.

For the driving task, participants were required to navigate the center lane of a three-way highway environment. The simulation environment was displayed on a 30-inch monitor and controlled by a Logitech G25 Racing Wheel. Participants were only required to steer the vehicle to maintain a central lane position. The vehicle's speed was held at a constant 55 miles/h (88.5 km/h). To reinforce safe lane keeping, safety cones were placed at either side of the

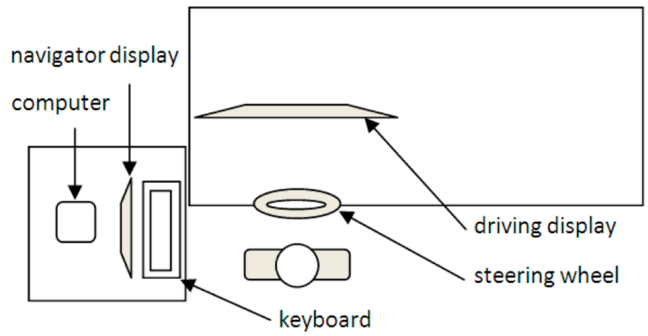

Figure 1. A schematic representation of how the driving and navigation displays were arranged.

driver's central lane. Noise was added to the vehicle dynamics, causing the vehicle to gradually drift about in the lane. This meant that the participant had to actively control and monitor the vehicle's lateral position and heading to maintain a central lane position.

For the navigation task, participants had to look at and enter a sequence of ten directions (lefts or rights). The to-be-entered sequence was randomly generated with the constraint that five left and five right directions were included and that there were no more than three consecutive repeating directions. The sequence of commands was represented either graphically (<=) or textually ("Left"), and was presented as a single vertical list on a 17-inch monitor positioned to the left of the participant (see, Figure 1).

The experiment was designed so that participants would be forced to sequentially interleave their attention between the two tasks. This was achieved by allowing only one of the task displays to be visible at any one time. By default the driving display was visible and the navigator display was blanked out. Participants activated the navigator display by moving their left hand from the steering wheel and using it to hold down the space bar on the keyboard in front of the navigator display. While the space bar was depressed the navigator display was presented and the driving display was blanked out. This meant that participants could not monitor the vehicle's position in the lane while encoding instructions for the navigation task. After viewing the instructions on the navigator display, participants had to return their hand to the steering wheel to use the left and a right paddle controls positioned under the steering wheel to enter the route instructions from memory.

Entry errors on the navigation task were associated with a time cost. If an input error occurred (e.g., a left paddle action was performed when a right action was required), the trial was terminated and the participant was instructed that they had to repeat the trial with a new list of instructions.

Design. A 2x2x2 (task-focus x representation x visual cue) mixed design was used, where task-focus was the between-subjects factor. To manipulate task priority, participants were instructed to either focus on completing the secondary navigation task as quickly as possible (the navigation-focus condition) or to focus on keeping the car as close as possible to lane center (the steering-focus condition).

Features of the secondary navigation task were manipulated as within-subjects factors. The route instructions were presented in a graphical or a textual format. In addition, a salient visual cue, indicating the current position in the list, was either present or absent.

The main dependent variables of interest were the time taken to complete the secondary navigation task and the impact that completing this task had on driving performance. The driving simulator logged the lateral distance of the vehicle from the center of the lane at a rate of 200 Hz. Driving performance was indexed by calculating the root mean square error (RMSE) of these lateral deviation samples over the period of time that the participant was working on the secondary navigation task. In addition, we were also interested in how participants chose to interleave the two tasks. To index task interleaving we consider the number and duration of each secondary task visit, as well as the time in between two visits.

Procedure. Participants were randomly assigned to one of the focus conditions, with the exception that effort was made to balance gender across conditions. Participants were given an opportunity to practice both the navigation and driving task separately. Once familiar with each task, participants completed four blocks of dual-task trials, one for each of the route representation and visual cue conditions. Trials were grouped by condition, and the order was counter-balanced across participants. For each block, participants were required to complete 10 error-free trials, up to a maximum of 15 trials per block. This dissuaded participants from making errors on the secondary navigation task.

Experimental instructions were given to encourage participants to prioritize either safe driving (steering-focus) or rapid completion of the navigation task (navigation-focus). To reinforce these instructions participants received feedback at the completion of every trial on their performance on the relevant variable. Specifically, participants in the steering-focus condition received feedback about the vehicle's RMSE lateral deviation, while participants in the navigation-focus condition received feedback on total trial time.

Results and Discussion

Due to space limitations we do not report data on how task performance was affected by manipulating features of the navigation task (see, Del Rosario, 2009, for details). Instead, we focus our analysis on how varying the instructions given to participants to prioritize one task over the other affected performance and decisions of how to interleave tasks. The primary dependent measures of interest were the time taken to complete the secondary navigation task and the lateral deviation of the vehicle from the center of the lane. We consider four separate measures to index task interleaving strategy: the number of visits to the navigator display per trial, the average duration of each visit, the number of navigation task items entered following each visit, and the average time between visits.

Figure 2 shows task time for the navigation task plotted against average RMSE lateral deviation for the driving task. There is a clear effect of task objective on how participants

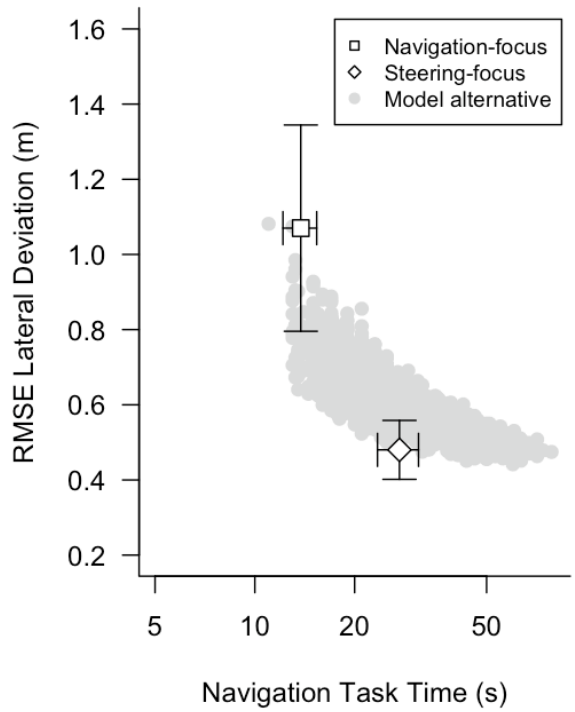

Figure 2. Data plot showing task time and RMSE lateral deviation across for varying task objectives. Error bars on human data points represent 95% confidence intervals. Model data points show performance predictions for different task interleaving strategies.

traded performance between the two tasks, in that, participants that were instructed to prioritize the navigation task completed it relatively quickly (M=13.76s, SD=2.31s), but in doing so had poor lateral control of the vehicle (M=1.07m, SD=0.41m). Conversely, participants that were instructed to prioritize safe driving completed the navigation task relatively slowly (M=27.30s, SD=5.57s) but were better able to maintain lateral control of the vehicle (M=0.48m, SD=0.10m). A 2x2x2 mixed factorial ANOVA found a significant effect of task objective on task time, $F(1,14)=40.26$, $p<.001$, $MSE=72.76$, and RMSE lateral deviation, $F(1,14)=15.87$, $p<.001$, $MSE=.35$.

We were also interested in participants' interleaving strategy, which was indexed by considering when participants choose to access the navigation task display. The data presented in Figure 3 show that the reason why participants in the steering-focus condition were better able to maintain lateral control of the vehicle than participants in the navigation-focus condition was because they made more visits to the navigation display (4.5 visits vs. 3.3 visits), $F(1,14)=3.67$, $p=.07$, $MSE=6.49$, entered fewer items following each visit (2.4 items vs. 3.4 items), $F(1,14)=5.19$, $p=.04$, $MSE=3.23$, and gave up more time to steering control between visits to the secondary display (5.34s vs. 2.57s), $F(1,14)=21.05$, $p<.001$, $MSE=6.25$.

The results of the study show that participants in the steering-focus condition interleaved more frequently and spent more time in between glances to the secondary display stabilizing the vehicle than participants in the navigation-focus condition. However, it is not immediately clear why participants adapted their strategy in the way that they did. Changing the task priority lead to only a single extra item, on average, being encoded and entered following each visit to the secondary display. In contrast, participants spent nearly twice as long in between visits to the navigation display in the steering-focus condition. But why did participants opt to spend more time between visits rather than interleave much more frequently? To explain the observed pattern of task interleaving we apply a modeling framework developed to explain behavior in a dialing-while-driving paradigm (Brumby et al., 2007, 2009).

Model

Our modeling approach focuses on deriving performance predictions for various strategies for completing the navigation task while driving. The model represents basic task operators (i.e., encoding a single instruction from the navigation display, or performing a steering control update) as discrete processing units that are limited by a serial bottleneck. Within this framework, we systematically consider every possible dual-task strategy that could have been adopted. Specifically, given that the navigation task required participants to enter 10 route instructions, we can consider at least 2^9 = 512 different task interleaving strategies (i.e., where strategies differ in terms of whether after encoding an item, another item is encoded or attention is returned to driving). For each of these strategies we also consider varying the amount of time that is given up to steering control in between visits to the secondary display.

We assume that glancing at the navigation display interferes with steering control. We estimate core parameters for the navigation task directly from the data. With these parameters fixed, we derive performance predictions for various dual-task interleaving strategies using a pre-existing model of steering control processes. For each strategy we derive predictions for critical performance metrics, namely, task time and lane keeping performance. The aim of this analysis is to explain the observed shift in dual-task performance between conditions, and also the precise change in low-level task interleaving behavior.

Navigation task. The navigation task is modeled at the granularity of the time taken to encode and enter route instructions. We estimate the time taken to perform these basic activities from the empirical data. Specifically, we estimate the time taken to:

- Shift attention from one task to the other
- Encode an item from the navigation display
- Input an instruction using the paddles

The time to switch attention from the secondary display to the driving task can be approximated by considering the average time between the release of the space bar (signaling the end of a visit) and the first paddle action being performed after the visit. Analysis shows that the average time between these events was approximately 1 second. A limitation of this measure as index of the cost of switching attention between tasks is that it assumes that the participant immediately commenced entering the instructions prior to returning their hand to the steering wheel.

We can approximate the time needed to encode a single route instruction by assuming that the number of items entered after a visit corresponds to the number of items that were encoded during that visit. Taking the average visit duration, we can calculate the average encoding rate to be approximately 500ms per item (i.e., in the navigation-focus condition, visits were on average 1.67s long and 3.4 items were entered after each visit). This assumes that participants never encoded items that were later forgotten or simply not entered. We shall revisit the implications of this assumption in the general discussion.

Finally, to estimate the time taken to input an instruction using the paddle, we consider the average time between two consecutive paddle entries. This yields an average time interval of 250ms between each paddle event. We assume that participants were able to perform steering updates while using the paddle to enter the route instructions, and that all instructions were entered before the next visit occurred. With these basic parameters set we can consider how this task might have interfered with driving performance.

Driving task. We use a simple mathematical model, taken from Brumby, Salvucci, and Howes (2009), which describes how people tend to adjust the heading of a vehicle based on its position in the lane. The model captures the basic idea that as the vehicle strays closer to the lane boundary, drivers react by making sharper corrective steering movements, which in turn, increase the lateral velocity of the vehicle, returning it to a central lane position more rapidly. The model assumes that discrete steering control updates are performed once every 250ms, which adjust the lateral velocity of the vehicle as follows:

$$\textit{Velocity} = 0.2617 \times LD^2 + 0.0233 \times LD - 0.022 \qquad (1)$$

where, *LD* represents lateral deviation from lane center, and there is an upper bound on velocity of 1.7m/s. In between steering updates, external factors can influence the vehicle's heading. To model this, we permute the vehicle's heading every 50 milliseconds with a value drawn from a Gaussian distribution with a mean of zero and a standard deviation of 0.09. Next we describe how this model is used to derive predictions of changes in a simulated vehicle's lateral deviation over time given discrete periods of driver attention and inattention.

For each of the 512 different strategies, we consider alternatives that give more or less time up to steering control in between visits to the navigation display. Specifically, we consider steering periods of between 250ms and 5000ms, at intervals of 250ms. This combined with the number of basic task interleaving strategies considered yields a fairly large set of 6,644 alternatives. For each, 50 simulations were run and performance averaged.

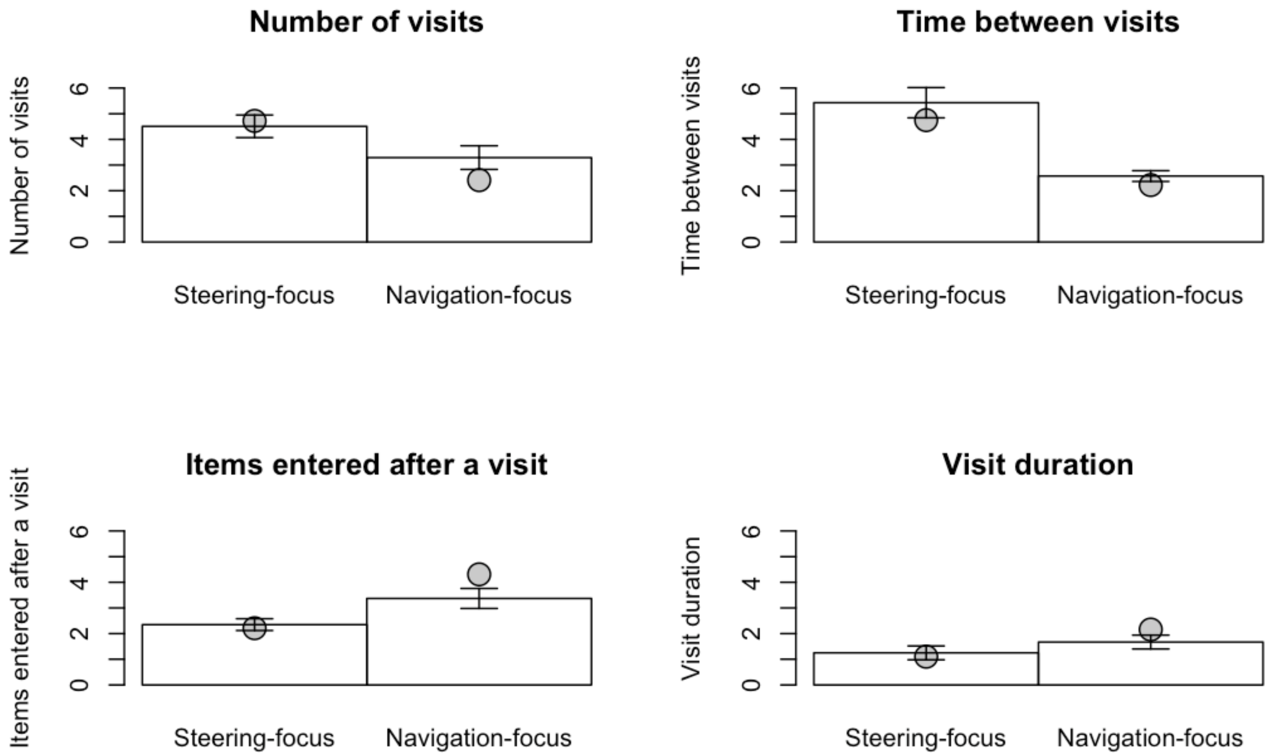

Figure 3. Data and model predictions for various navigation task measures. Bar charts show human data, with error bars representing standard errors of the mean. Circular data points represent model predictions for each priority condition. The values are the means for model alternatives that fall within the Confidence Interval in Figure 2 (see text for details).

Model Results

Figure 2 shows the predicted RMSE lateral deviation and task time for each of the 6,644 strategies that were evaluated along with the human data for each priority condition. The model predicts a clear dual-task performance tradeoff between strategies that complete the navigation task quickly and have relatively poor lane keeping performance, and those that complete the navigation task more slowly giving relatively better driving performance.

The shape of the tradeoff curve predicted by the model is noteworthy. There is a clear tipping point where improvements in lane keeping performance become smaller with increased task time. The human data for the steering-focus condition lie at this tipping point in the tradeoff curve, suggesting that participant adapted their strategy to meet the performance objective of minimizing lateral deviation while completing the secondary task in a reasonable amount of time (note that time is represented on a logarithmic scale). In contrast, data from the navigation-focus condition lie at close to the leftmost extreme of the strategy space, where faster performance is associated with poor lane keeping.

Figure 2 shows that there are many different strategies that fall within the predicted performance bounds of the human data for each condition. To get a better sense of how this performance tradeoff was achieved, we consider how these strategies allocated attention between the tasks. Specifically, we consider for each condition the subset of strategies that fall within the 95% confidence interval (CI) of the human data for each condition.

For the navigation-focus condition there were 34 strategies that fell within the CIs of the human data, while for the steering-focus condition there were 307 strategies that fell within the CIs of the human data. For each of these best-fitting strategies we define the same four measures of task interleaving behavior used in the analysis of the human data (i.e., the number of visits to the navigator display per trial, the average duration of each visit, the number of navigation task items entered following each visit, and the average time between visits). For each measure, we calculate the mean across the subset of best fitting strategies for each condition. In doing so, we get a better sense of how the best fitting strategies for each condition differed, and can compare these indexes of behavior to the human data.

Figure 3 shows these mean model predictions along with the corresponding human data for each condition. The fit of the model to these low-level task interleaving measures is remarkable, in that the model explains why participants in the steering-focus condition would have chosen to double the time between visits and encode one extra item per visit in order to reach the tipping point in the tradeoff curve.

General Discussion

A novel dual-task paradigm was used to investigate how people adapt their behavior to meet a specific performance objective. In the study, participants were required to encode and enter a series of route instructions while driving a simulated vehicle. Explicit instructions were given to participants to give greater priority to either safe driving or rapid completion of the navigation task. Results showed that participants met the required task objective by varying the number and duration of visits to the navigation display, and by also varying the amount of time given up to steering control between visits. These findings support the idea that people can strategically allocate attention in multitask settings (e.g., Brumby et al., 2009; Horrey et al., 2006; Gopher et al, 1982; Gopher, 1993; Wang et al., 2007).

We explain participants' decisions about how to allocate attention using an existing framework for modeling driver distraction effects (Brumby et al., 2007, 2009). The model represents basic task operators as discrete processing units that are limited by a serial bottleneck. To apply the model to this new dual-task context, a handful of parameters for the navigation task had to be estimated from the data (i.e., the time taken to encode a single instruction from the navigation display, shift attention back to road, and enter that instruction). With these basic timing estimates fixed, we model the effects of various allocation policies for attending to the secondary navigation display for critical task performance metrics.

The modeling results help explain the observed shift in task performance between the two focus conditions. The model predicts a classic dual-task performance tradeoff between safer driving and shorter task time. Interestingly, the tradeoff curve has a clear tipping point, after which improvements in lane keeping performance become relatively small with increased time investment. Human performance data from the steering-focus condition lie close to this tipping point, and remarkably the modeled strategies in this region of the strategy space corresponded with those adopted by participants.

However, the model did not explain data from the navigation-focus condition as well. Specifically, it under-predicted the number of visits made to the secondary display and over-predicted the number of items entered after each visit (see, Figure 3). The likely explanation for this departure is that the model assumes a perfect and limitless memory, which could enter all ten of the route instructions after a single visit. This is clearly an implausible assumption given the known limits on memory. This aspect of the model could be informed by considering how many items participants would copy over in a single-task setting. Alternatively, we could build on existing work that has modeled memory retrieval processes in similar tasks. For instance, Gray et al.'s (2006) work on modeling the impact of memory constraints in the Blocks World paradigm.

Moreover, because of space limits we could not present an analysis of how features of the navigation task affected performance. Del Rosario (2009) reports that participants could encode textual information faster than graphical information. Future work should point out how the model might explain any shift in strategy based upon changes in time take to encode an item from the display.

In summary, we have used a novel dual-task paradigm to demonstrate that people can strategically allocate attention in multitask settings. A model was used to explain why particular strategies might have been favored in terms of the shape of the performance tradeoff between safer driving and shorter task time.

Acknowledgments

This work was supported by EPSRC grant EP/G043507/1. The study reported here was undertaken as part of the second author's Master's dissertation work. We thank Dario Salvucci for providing the code for the driving simulator used here. We are grateful to three anonymous reviewers for their valuable comments on this work.

References

Brumby, D.P., Howes, A., & Salvucci, D.D. (2007). A cognitive constraint model of dual-task trade-offs in a highly dynamic driving task. In *Proc. SIGCHI Conference on Human Factors in Computing Systems, CHI'07* (pp. 233-242). New York, NY: ACM Press.

Brumby, D.P., Salvucci, D.D., & Howes, A. (2009). Focus on driving: How cognitive constraints shape the adaptation of strategy when dialing while driving. In *Proc. SIGCHI Conference on Human Factors in Computing Systems, CHI'09* (pp. 1629-1638). New York, NY: ACM Press.

Del Rosario, N. (2009). *Reasons to Switch: The Effects of Priority and Information Presentation on Dual-Task Interleaving Strategies.* Master's dissertation, University College London, London, UK.

Gopher, D., Brickner, M., & Navon, D. (1982). Different difficulty manipulations interact differently with task emphasis: Evidence for multiple resources. *Journal of Experimental Psychology: Human, Perception, & Performance, 8,* 146–157.

Gopher, D. (1993). The skill of attention control: Acquisition and execution of attention strategies. In D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience* (pp. 299–322). Cambridge: MA: MIT Press.

Gray, W.D., Sims, C.R., Fu, W.-T., & Schoelles, M.J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review, 113*, 461-482.

Horrey, W.J., Wickens, C.D., & Consalus, K.P. (2006). Modeling drivers' visual attention allocation while interacting with in-vehicle technologies. *Journal of Experimental Psychology: Applied, 12,* 67-78.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review, 86,* 214-255.

Norman, D.A., & Bobrow, D.G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology, 7,* 44-64.

Salvucci, D.D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science, 29*, 457-492.

Salvucci, D.D., & Taatgen, N.A. (2008). Threaded cognition: An integrated theory of concurrent multitasking. *Psychological Review, 115*, 101-130.

Wickens, C. D. (2002). Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science, 3,* 159-177.

Wang, D.D., Proctor, R.W., & Pick, D.F. (2007). Acquisition and transfer of attention allocation strategies in a multiple-task world environment. *Human Factors, 49*, 995–1004.