

Dialing while Driving? A Bounded Rational Analysis of Concurrent Multi-task Behavior

Duncan P. Brumby (Brumby@cs.drexel.edu)

Dario D. Salvucci (Salvucci@cs.drexel.edu)

Department of Computer Science, Drexel University
Philadelphia, PA 19104 USA

Andrew Howes (HowesA@manchester.ac.uk)

Manchester Business School, University of Manchester,
Manchester, M15 6PB UK

Abstract

When people conduct multiple tasks in tandem, such as dialing a cell phone while driving a car, they often interleave the two tasks, for instance by returning attention to the primary driving task after entering bursts of three or four digits at a time. In order to explain why people tend to interleave these tasks at this particular interval, a control model of steering behavior is described that focuses on understanding how environmental and psychological constraints interact to determine driver performance. We use this model to predict the amount of time that people are prepared to stray from the driving task while engaging in a secondary in-car task and, by consequence, the degree of task interleaving. In particular, a modeling experiment was conducted to determine the consequences of systematically varying the time interval between consecutive steering updates for driving performance. The results of this analysis were then used to demonstrate why returning attention to driving after entering bursts of three or four digits at a time is a particularly efficient strategy: It does not allow driving performance to become too egregious, while at the same time keeping the additional time costs that are incurred as a result of interleaving tasks minimal.

Introduction

While you are driving in your car it is not too difficult to sometimes direct your attention away from the road in order to engage in a secondary task, such as dialing a number on a cell phone. In this complex real-world multitasking scenario, people tend to interleave the two tasks by returning attention to driving after entering bursts of three or four digits at a time (e.g., Salvucci, 2005). One potential explanation for why people choose to interleave these tasks at this particular interval is that the representational structure of the telephone number (e.g., for a 7-digit telephone number this might follow a xxx-xxxx structure) provides a series of natural break points at which to return attention to driving. Although this account has intuitive appeal, it is not entirely obvious that people necessarily have to return attention to steering control after dialing three or four digits. Why not more or less digits at a time? Alternatively, if someone were engaged in some other secondary in-car task that, for instance, demanded longer interaction episodes than dialing (e.g., scrolling through a long list of media content on an Apple iPod), would they still make glances back to the road with the same regularity?

In this paper, we present a bounded rational analysis (Howes, Vera, & Lewis, 2007) of concurrent multi-task

behavior, in order to better understand how long people should be prepared to look away from the road when engaging in a secondary dialing task while driving. In this analysis we focus on understanding how functional-level features of the task environment (Gray, Neth, & Schoelles, in press) and the constraints imposed by the cognitive architecture (Anderson et al., 2004) interact to make some multitasking behaviors more preferable than others. For instance, it seems rather obvious that deprived of regular attention, driving performance will rapidly fall below criterion, with potentially disastrous consequences. At the same time though the benefits of frequently interleaving play against the costs of switching between tasks (e.g., Allport, Styles, & Hsieh, 1994). In particular, switching between tasks often carries costs associated with the physical realignment of the body relative to external resources and the mental recovery of state information associated with each task. Given this trade-off between the potential costs and benefits of frequently interleaving, how do people decide when to switch back and forth between tasks?

One possible factor that might determine when task interleaving is desirable is the shape of the payoff function for the primary task (Payne, Duggan, & Neth, in press; Son & Sethi, 2006). Payne et al. conducted a series of experiments designed to investigate how people allocate time between two Scrabble tasks. Each task required participants to generate as many words as possible from a fixed set of letters in a given amount of time. Importantly, the tasks differed in the number of words that could be readily generated from their respective letter sets. This meant that the tasks had different payoff functions because the rate at which a participant could find a novel word from a particular set of letters differed between the two tasks. This difference between tasks' payoff functions were not known in advance of the study to the participants, and Payne et al. were interested in how participants would allocate their time between the two tasks. As one might expect, Payne et al. found that participants eventually learned to allocate more time to the more productive task (i.e. the task with the greater payoff function) but most still chose to switch between tasks rather frequently. Payne et al. also found that participant's "giving-up time" (i.e., the time between finding the most recent word and the decision to switch tasks) was longer in the less productive task. Taken together these two effects appear to work against each other (i.e., longer visits to the easier task, but shorter giving-up times for this task), but Payne et al. demonstrate that a stochastic model, based on

Green's (1984) assessment rule of optimal foraging theory, account for these data.

Son and Sethi (2006) present a formal analysis that shows that optimal time allocation between tasks is dependent on characteristics of the environment. Son and Sethi give the example of a learning environment where a learner's time must be allocated between multiple tasks (e.g., consider a student studying for a set of final exams). Son and Sethi demonstrate how time pressure as well as the nature of a task's learning curve can lead to different allocations of time between tasks. Moreover, the work of both Payne et al. and Son and Sethi is of interest here because it points to the potential role of a task's payoff function in determining precisely when people are likely to switch from one task to another.

In this paper, we present a bounded rational analysis (Howes, Vera, & Lewis, 2007) of possible strategic variability in how people might dial a cell phone while driving. Extending earlier work (Brumby, Howes, & Salvucci, 2007; Brumby, Salvucci, Mankowski, & Howes, 2007), a control model of steering behavior is described that focuses on understanding how environmental constraints (e.g., perturbation of the vehicle's heading over time) and psychological constraints (e.g., people's sensitivity to the lateral position of the vehicle in relation to the center of the lane) interact to determine driver performance. A modeling experiment is conducted to determine the consequences of systematically increasing the time interval between consecutive steering updates for the average lateral deviation of the vehicle from the lane center over time. We show that the particular rate that people tend to make glances back to the road while engaging in a dialing task can be understood in the context of the rate of decline in driver performance over time and the costs of switching back and forth between tasks.

Model of Steering Control

A control model of steering behavior is developed that gives predictions of changes in a simulated vehicle's lateral deviation (i.e., distance from the lane center) over time. The model focuses on understanding how environmental and psychological constraints interact to determine driver performance. The model simulates a vehicle moving at a constant velocity down a straight road. The model performs a series of discrete steering updates that alter the heading (or lateral velocity) of the vehicle dependent on its lateral position in the lane at the time that the steering update is performed. The approach taken is similar to control theoretic accounts of lane keeping (e.g., model 1 in Hildreth et al. 2000), which assume that adjustments to the heading of a vehicle are motivated by the goal of minimizing perceptual input quantities that represent the lateral position and heading of the vehicle.

In order to parameterize the model, driver performance data from two experiments that investigated the effect of cell phone use on driving (Salvucci, 2001; Salvucci & Macuga, 2002) were analyzed to formally characterize how drivers typically adjusted the heading of the vehicle given its lateral

position in the roadway. An underlying assumption of this analysis was that adjustments to the heading of the vehicle were motivated by the driver attempting to maintain a central lane position over time. In particular, the experimental software logged, at a rate of once every 30 ms, the normalized steering wheel angle of the simulated car and its divergence from the center of the lane (in meters). This steering data was then segmented into a series of *steering episodes*, which were defined as periods in which the angle of the steering wheel did not alter over time. For each of these steering episodes, a tuple was defined that represented the duration of the episode (*time*), the change in the lateral position of the vehicle (*distance*), and the average lateral velocity of the vehicle (where *lateral velocity* = *distance* / *time*). Data from all steering episodes across participants from the two studies were pooled, and the lateral velocities of all steering episodes that had a common starting lateral deviation (i.e., originated from the same lateral position in the roadway) were averaged. We report an analysis of these average data.

Figure 1 shows a scatter plot of the relationship between the lateral deviation of the vehicle at the start of a steering episode and its average lateral velocity throughout the episode. It can be seen in the figure that as the car strayed closer to the lane boundary, drivers tended to react by making sharper corrective steering movements, which in turn, increased the lateral velocity of the vehicle, returning it to a central lane position more rapidly. Furthermore, it can be seen that for many steering episodes lateral velocity was negative; indicating that the car was heading farther away from the center of the lane.

Regression analysis was conducted to estimate a best fitting curve to predict the average lateral velocity of a steering episode given the lateral deviation (*LD*) of the vehicle at the start of an episode. It was found that a quadratic function¹,

$$Velocity = 0.2617 \times LD^2 + 0.0233 \times LD - 0.022 \quad (1)$$

provided a high degree of correspondence with the human data ($r^2 = 0.61$), $F(1,80) = 62.61$, $p < .001$. This quadratic model of steering control predicts that as lateral deviation from the lane center increases, there is an increase in the lateral velocity of the vehicle, brought about at discrete steering updates, in order to return the vehicle to a central lane position more rapidly.

Furthermore, the intercept of the curve given by the model (shown in Figure 1) gives some suggestion of the driver's threshold for judging the vehicle's deviation from the lane center. In particular, when the car is near the lane center (i.e., lateral deviation < 0.30 m), predicted lateral velocity is close to zero. This means that the position of the car in the roadway remains more or less constant over time. This implies that the driver was possibly satisfied with the vehicle's position in the roadway if the lateral deviation of the vehicle was less than 0.3 m from the lane center.

Although the quadratic model gave a high degree of correspondence with the data, there was also considerable variability with respect to the observed lateral velocities given a particular lateral deviation at the start of an episode. In particular, the standard deviation of the data from the mean

¹ Because of non-positive lateral velocities exponential or power functions could not be applied.

was 0.10 m/s. This suggests that people’s adjustments to the heading of the vehicle were stochastic. In order to develop a stochastic model of steering control, random values were sampled from a Gaussian distribution and added to the value of the updated lateral velocity. Based on an estimate of the average standard deviation observed in the human data, the Gaussian distribution had a mean of 0.00 m/s and standard deviation of 0.10 m/s.

An important functional-level feature of the driving environment is that if left unattended, the heading of the vehicle will also be influenced by external factors, such as bumps in the road, wind, the camber of the road, etc. In order to simulate this feature of the driving environment, the heading of the vehicle was perturbed every 50 ms by a random value sampled from a Gaussian noise distribution. Following estimates from a previous model in the literature (Hildreth et al., 2000), the Gaussian noise distribution had a mean 0.00 m/s and standard deviation 0.10 m/s.

In summary, the model provides a computationally efficient formalism for predicting how drivers typically adjust the heading (or lateral velocity) of a vehicle given its lateral position in the roadway. The model focuses on how functional-level features of the task environment (e.g., perturbation of the vehicle’s heading over time) and psychological constraints (e.g., people’s sensitivity to the lateral position of the vehicle in relation to the center of the lane) interact to determine driver performance. Moreover, it is worth pointing out at this stage that the model does not make any theoretical commitment to the duration of a typical steering update; the model is solely dependent on parameters derived from an analysis of steering performance data and assumptions about the environment. In the next section the model is used to understand how driving performance might decline with increasing periods of driver inattention.




Figure 1: Relationship between lateral deviation at the start of a steering episode and lateral velocity.

Modeling Experiment

We conducted a modeling experiment that symmetrically varied the time interval between steering updates in order to make quantitative predictions of the consequences for lateral deviation over a period of simulated driving. Specifically, we explored steering strategies that updated lateral velocity at an interval of between 50 ms and 6,000 ms, exploring performance at increasing increments of 50 ms. That is, we

evaluated 120 different steering strategies that differed in terms of the duration of time between each steering update over a period of simulated driving. Each steering strategy was run for 1,000 trials, and performance averaged. The vehicle’s lateral deviation at the start of each trial was 0.33 m from the lane center. This initial lateral deviation reflects the average value at the beginning of each trial in the empirical data (taken from Salvucci, 2001). Furthermore, each steering strategy was evaluated over both a single steering episode (single-event) and also over a longer measurement interval of 60-seconds simulated driving (long-term). For each steering strategy we report the lateral deviation — the root-mean-square error (RMSE) of the vehicle’s lateral distance from lane center — over both a single-event and the long-term measurement interval.

Results: The effect of interval between steering updates on lateral deviation

In order to illustrate how the movement of the vehicle is affected by the interval between steering updates, Figure 2 shows a data plot representing changes in lateral deviation over time for different illustrative strategies. Performance is for a single trial. Data points represent periods where the lateral velocity of the vehicle was altered owing to a steering update. Changes in the heading of the vehicle between steering updates are due to environmental noise.

Figure 2 offers a comparison between the performance of steering strategies that conducted relatively frequent updates to the lateral velocity of the vehicle (once every 50 ms) to strategies that updated lateral velocity less frequently (once every 600 ms). It is clear from the figure that there was very little difference in performance between these two strategies; in both cases the vehicle maintained a more or less straight heading (i.e., lateral velocity ≈ 0 m/s) and as a result kept to a consistent lateral position in the lane over time. In contrast, as the interval between consecutive steering updates increased even further (to once every 1800 ms), the vehicle tended to drift more erratically about the lane. This was partial because without frequent steering updates, the heading of the vehicle was perturbed by environmental noise. In order to compensate for this general increase in lateral deviation, the model tended to set a heading when a steering update was eventually performed that gave a large lateral velocity. As can be seen in the figure, these aggressive changes in heading lead the vehicle to move rather erratically about the lane.

We next focus on quantifying the rate at which lateral deviation increases with increasing time between updates of steering control. Figure 3 shows the performance of each strategy over a single steering update (single-event) and also over a longer measurement interval of 60-seconds simulated driving (long-term). The x-axis in the figure represents the interval between steering updates and the y-axis represents mean lateral deviation over 1,000 trials. It is clear that as the time between steering updates increases, lateral deviation generally increases, except, that is, across relatively short intervals between steering updates (< 1 sec). At these shorter intervals, the duration of time in between steering updates did not affect lateral deviation.

It is also apparent from Figure 3 that the way in which lateral deviation increased with increasing time between




Figure 2: Data plot representing movement of the car in relation to the center of lane for illustrative steering strategies. Data points represent steering updates. Connecting lines represent movement of the car in between steering updates.

consecutive steering updates was dependent on the total period of simulated driving (i.e., single-event vs. long-term measurement interval). In particular, the rate of decline in steering performance was less when there was only a single steering update than when a strategy was maintained for a longer period of time. This is because when there is only a single steering update event, the vehicle travels at a fairly constant lateral velocity; therefore the distance traveled from the lane center will be dependent on the time until the next steering update. However, when a strategy is maintained for a longer period of time (i.e., 60 sec), the average deviation can grow quite large because, in some sense, the problems start building on each other. That is, as we described earlier, the car not only drifts farther from the lane center with increasing time away from driving, but as a consequence, it is also placed into sharper corrective headings to compensate for being farther from lane center. This interaction between increased lateral velocity and longer intervals between steering updates makes the car move erratically about the lane.

Regression analysis was conducted to estimate the best-fitting curve to account for the relationship between the interval between steering updates and lateral deviation. For performance based on a single-event, it was that an exponential function fit the data very well ($r^2 = 0.98$), $F(1,118) = 5828$, $p < .001$, where

$$\text{Lateral Deviation} = 0.2755^{0.2177 \times \text{Update Interval}} \quad (2)$$

The exponent in this function increased, however, when the strategy was maintained for a longer period of simulated driving, giving

$$\text{Lateral Deviation} = 0.2807^{0.3453 \times \text{Update Interval}} \quad (3)$$

($r^2 = 0.99$), $F(1,118) = 35747$, $p < .001$. This meant that the rate of decline in steering performance increased more dramatically with increasing interval between steering updates. In the next section we derive predictions for driving performance under dual-task conditions by considering possible strategic variability in how people might dial a cell phone while driving.




Figure 3: Data plot showing the relationship between the duration of time between each steering update and lateral deviation. Each steering strategy was run over both a single steering event and also a longer 60-second period.

Predicting Multi-task Performance

We model data from an earlier study that investigated in-car multitasking (Salvucci, 2001). In Salvucci's experiment participants were required to dial 7-digit numbers on a cellular phone that was positioned on a hands-free device while driving. It was assumed that one "power-on" key-press preceded the 7-digit number and that one "send" key-press followed it — giving 9 key-presses in all. Salvucci reports average baseline (or single-task) dial-time for the participant's to enter the 9-keypresses of 5.21 seconds ($S.D. = 1.09$ sec). We use this empirical estimate of dial-time to calibrate the model.

We assume that in normal conditions drivers typically adjust the heading of the vehicle once every 150 ms. This 150 ms estimate is consistent with assumptions adopted in previous computational cognitive models in the literature (e.g., Salvucci, 2005). Moreover, at this baseline interval between steering updates, lateral deviation predictions given by the model ($M = 0.33$ m, $S.D. = 0.02$ m, see Fig. 3) are comparable with reported baseline lateral deviation in Salvucci's (2001) experiment ($M = 0.35$ m, $S.D. = 0.08$ m).

We assume that engaging in a secondary task while driving disrupts the normal pattern of checking and adjusting the heading of the vehicle. In particular, we assume that steering updates cannot occur while the driver's attention is directed towards a secondary in-car task, such as when they are entering keypresses for the dialing task. This assumption is based on the idea that peripheral resources, such as the eyes, will limit the degree of parallel processing between tasks. Moreover there are numerous demonstrations in the literature of central interference affecting driver performance in dual-task conditions (e.g., Brumby, Salvucci, & Howes, submitted; Levy, Pashler, & Boer, 2006).

Furthermore, we assume that switching between tasks carries a cost overhead (or switch cost), which reflects the time required to move visual attention between the outside of the car (i.e., to focus on the road) and the inside of the car (i.e., to focus on the phone). Instead of developing a detailed model of the perceptual/motor processes involved, we use a simple timing estimate of 185 ms to move visual attention

between the phone and the road, or vice versa. This timing estimate was taken from the ACT-R cognitive architecture (Anderson et al., 2004).

Given the above set of assumptions and also the estimates of single-task performance, we derive predictions for lateral deviation and task time in dual-task conditions. Brumby, Howes, and Salvucci (2007) have previously demonstrated that there are at least $2^8 = 256$ possible strategy variants for completing the dial task with more or less interleaving of steering control. Here, we attempt to abstract over this strategy space by conducting an analysis that varies the number of equal length episodes into which the dial task could be divided and explore the consequences for the interval between steering updates. It should be made clear that this level of analysis abstracts over the actual units of the dial task (i.e., entering more or less digits per episode) and instead focuses on *dividing single-task dial time in to more or less equal chunks of time*; thus, abstractly representing points in the strategy space of more or less interleaving.

Figure 4 presents a scatter plot of total time to complete the dial task and RMSE lateral deviation for strategies that systemically vary in the degree of task interleaving. In particular, at each point in the space we divide baseline dial-time (5.21 sec) by N , where N varies between 1 (no-interleave strategy) and 9 (maximum-interleave strategy). Given an estimate of the amount of time between steering updates that a particular strategy affords, Equations 2 and 3 are used to derive predictions of lateral deviation. For instance, if we consider adopting a no-interleave strategy, which completes the dial task without once returning attention to the primary task of driving, then the interval between steering updates would be 5.73 seconds (i.e., $5.21 + 0.185 \times 2 + 0.15$). It is clear from Figure 3 that updating steering control at this interval would likely have catastrophic consequences for the driving task, with the car being likely to cross over the lane boundary.

In contrast, if the dial task were conducted with steering updates occurring after each and every individual digit was entered, what we shall refer to as a maximum-interleave strategy, then the interval between steering updates would be only 1.09 seconds (i.e., $5.21 / 9 + 0.185 \times 2 + 0.15$). It can be seen in Figure 3 that updating steering control at this interval would not likely lead to an egregious lateral deviation. However, this strategy would incur 4.70 seconds of additional time costs because of frequently switching between tasks and updating steering control (i.e., $9 \times (0.185 \times 2 + 0.15)$).

Figure 4 represents the speed/accuracy trade-off that clearly exists between dialing quickly and driving safely: The upper-left portion of the plot represents faster but less safe performance resulting from less interleaving, while the bottom-right portion represents slower but safer performance resulting from more interleaving. There are diminishing returns for interleaving, however. Such that, while interleaving tasks more often generally leads to safer performance there is a point in the space where further interleaving gives only small improvements in safety.

We compare these model-based predictions shown in Figure 4 with previous empirical data. In particular, Salvucci (2001) reports dual-task performance of 7 sec ($SD = 1.77$ sec) for the dialing task and RMSE lateral deviation of 0.49 m

($SD = 0.10$ m) for the driving task. These human data are also presented in Figure 4.

It is interesting that the human data lie close to the “turning point” where lateral deviation starts to increase dramatically within the modeled strategy space. This suggests that any less interleaving between tasks would likely result in a dramatic increase in lateral deviation, but also that more interleaving between tasks would not likely result in a significant reduction in lateral deviation given the additional time costs.

The model-based predictions demonstrate that adopting a strategy that returns attention to driving after entering three digits at a time is particularly efficient. This strategy does not allow driving performance to become too egregious because the interval between steering updates increases to only 2.26 seconds (i.e., $5.21 / 3 + 0.185 \times 2 + 0.15$). But at the same time the strategy keeps the additional time costs incurred as a result of interleaving tasks down to only 1.56 seconds (i.e., $3 \times (0.185 \times 2 + 0.15)$).

Finally, notice that predictions for lateral deviation in Figure 4 were derived using the exponential loss function derived from running the model over a single-event (Eq. 2) and also over a long-term measurement interval (Eq. 3). It is interesting that for the most part the single-event model and the long-term model gave fairly consistent predictions for strategies that interleaved tasks more often. However, the model predictions differed fairly significantly for strategies that completed the dial task in only one or two bursts (i.e., the no-interleave strategy). The reason for this discrepancy is that when a particular strategy was maintained for a longer time (i.e., 60 sec), the average deviation could grow quite large at longer intervals between updates.




Figure 4: Data plot of dial time and average lateral deviation across strategies of varying task interleaving.

General Discussion

The question addressed at the start of this paper was why people return attention to steering control after dialing every three or four digits of a telephone number. To address this question, a control model of steering was developed from an analysis of driver performance data. The model made minimal commitments to human cognitive architecture and minimal assumptions about the constraints imposed by the environment. The model was used to predict the average rate at which the lateral deviation of the vehicle from the lane center increases with increasing time between updates

of steering control. This bounded rational analysis suggests that the *rate* of decline in driving performance with time away from steering control might determine the amount of time that people are prepared to give up to focus on a secondary task while driving and, by consequence, the degree of task interleaving. We demonstrate that returning attention to driving after entering bursts of three digits at a time is a particularly efficient strategy for completing the dial task while driving because it does not allow driving performance to become too egregious, while at the same time it keeps the time costs incurred from switching between tasks minimal. Moreover, we show that any less task interleaving would result in a dramatic increase in lateral deviation, with possibly unacceptable consequences for safety, and that any more interleaving would incur additional time costs while not affording a significant improvement for driver safety.

An open empirical question that is posed by the analysis presented here is that if the rate of decline in driving performance with time away from task were different, then people might interleave task differently. For instance, imagine if driving performance were to decline much more gradually with time away from task (i.e., when driving at a slower speed), then there would be little value in interleaving tasks: People may as well complete the dial task in a single contiguous burst in order to avoid incurring the costs of switching between tasks. Whereas, if driving performance were to decline at a much more rapid rate (i.e., when driving at a faster speed), then people might be prepared to give up less time per visit to the secondary task and consequently interleave more frequently. That is, driving speed should have an effect on both dial time and lateral deviation in dual-task conditions (see Brumby, Salvucci, & Howes, submitted, for an initial investigation into this question). Moreover, there is evidence that drivers tend to slow down on their own accord when engaging in a secondary dialing task (Salvucci, 2001; Salvucci & Macuga, 2002). This slowing behavior might reflect active attempts to reduce the consequences of directing attention away from the road for driving performance.

We might also consider applying the analysis presented here to some other secondary in-car task that demands a series of longer interaction episodes than a simple dialing task (e.g., selecting media content on an Apple iPod). The analysis presented here clearly suggests that lateral deviation should increase as the amount of time spent on the secondary task increases. An interesting question there emerges from considering a longer task, where the vehicle is more likely to drift from the lane center, is whether people give up more time to steering control (i.e., by conducting a series of steering updates in succession). The approach taken here for running the model over a long-term measurement interval was to assume that only a single corrective steering update is performed, regardless of how far from the lane center the vehicle has become. This seems like a rather implausible assumption, however. An alternative assumption is that people only resume the secondary task return when the vehicle has been to returned to a stable lateral position in the roadway (as in Salvucci's, 2005, 2001, driver models). Further work is required to explore techniques for enumerating over various durations of time given up to

steering control for each of the possible multitasking strategies discussed here (see Brumby, Salvucci, Mankowski, & Howes, 2007, for some more recent progress on this issue).

Acknowledgments

This research was supported by National Science Foundation grant #IIS-0426674. We would like to thank three anonymous reviewers for providing comments for improving this paper.

References

- Allport, A., Styles, E.A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà, & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 421-452). Cambridge, MA: MIT Press.
- Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of mind. *Psychological Review*, *111*, 1036-1060.
- Brumby, D.P., Salvucci, D.D., & Howes, A. (submitted). An empirical investigation into dual-task trade-offs while driving and dialing. Paper submitted to the *British HCI Group Annual Conference*. Lancaster University, UK.
- Brumby, D.P., Howes, A., & Salvucci, D.D. (2007). A cognitive constraint model of dual-task trade-offs in a highly dynamic driving task. To appear in *Human Factors in Computing Systems: CHI 2007 Conference Proceedings*. New York, NY: ACM Press.
- Brumby, D.P., Salvucci, D.D., Mankowski, W., & Howes, A. (2007). A cognitive constraint model of the effects of portable music-player use on driver performance. To appear in the *Proceedings of the Human Factors and Ergonomics Society 51st Annual Meeting*. Santa Monica, CA: Human Factors and Ergonomics Society.
- Gray, W.D., Neth, H., & Schoelles, M.J. (in press). The functional task environment. In A. Kramer, A. Kirlik, & D. Wiegman (Eds.), *Applied Attention*. New York, NY: Oxford University Press.
- Green, R. F. (1984). Stopping rules for optimal foragers. *American Naturalist*, *123*, 30-43.
- Hildreth, E.C., Beusmans, J.M.H., Boer, E.R., & Royden, C.S. (2000). From vision to action: experiments and models of steering control during driving. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 1106-1132.
- Howes, A., Vera, A., & Lewis, R.L. (2007). Bounding rational analysis: Constraints on asymptotic performance. In W.D. Gray (Ed.) *Integrated Models of Cognitive Systems* (pp. 403-413). New York, NY: Oxford University Press.
- Levy, J., Pashler, H., & Boer, E. (2006). Central interference in driving: Is there any stopping the psychological refractory period? *Psychological Science*, *17*, 228-235.
- Payne, S.J., Duggan, G.B., & Neth, H. (in press). Discretionary task interleaving: Heuristics for time allocation in cognitive foraging. *Journal of Experimental Psychology: General*.
- Salvucci, D.D. (2001). Predicting the effects of in-car interface use on driver performance: An integrated model approach. *International Journal of Human-Computer Studies*, *55*, 85-107.
- Salvucci, D.D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science*, *29*, 457-492.
- Salvucci, D.D., & Macuga, K.L. (2002). Predicting the effects of cellular-phone dialing on driver performance. *Cognitive Systems Research*, *3*, 95-102.
- Son, L.K., & Sethi, R. (2006). Metacognitive Control and Optimal Learning. *Cognitive Science*, *30*, 759 -774.