In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 121 - 126. Oxford, UK: Taylor & Francis/Psychology Press.

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 of 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 of 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 of 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 at 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 function[1],

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

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 vehicles 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.