A Cognitive Constraint Model of Dual-Task Trade-offs in a Highly Dynamic Driving Task

Duncan P. Brumby[1], Andrew Howes[2], & Dario D. Salvucci[1]

[1]Department of Computer Science
Drexel University
Philadelphia, PA 19104 USA
{Brumby,Salvucci}@cs.drexel.edu

[2]Manchester Business School
University of Manchester
Manchester, M15 6PB UK
HowesA@manchester.ac.uk

ABSTRACT

The paper describes an approach to modeling the strategic variations in performing secondary tasks while driving. In contrast to previous efforts that are based on simulation of a cognitive architecture interacting with a task environment, we take an approach that develops a *cognitive constraint model* of the interaction between the driver and the task environment in order to make inferences about dual-task performance. Analyses of driving performance data reveal that a set of simple equations can be used to accurately model changes in the lateral position of the vehicle within the lane. The model quantifies how the vehicle's deviation from lane center increases during periods of inattention, and how the vehicle returns to lane center during periods of active steering. We demonstrate the benefits of the approach by modeling the dialing of a cellular phone while driving, where drivers balance the speed in performing the dial task with accuracy (or safety) in keeping the vehicle centered in the roadway. In particular, we show how understanding, rather than simulating, the constraints imposed by the task environment can help to explain the costs and benefits of a range of strategies for interleaving dialing and steering. We show how particular strategies are sensitive to a combination of internal constraints (including switch costs) and the trade-off between the amount of time allocated to secondary task and the risk of extreme lane deviation.

Author Keywords

User modeling, multitasking, driver distraction.

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation] User Interfaces: Theory and methods; K.4.1 [Computers and Society] Public Policy Issues: Human safety; H.1.2 [Models and Principles] User/Machine Systems: Human factors.

INTRODUCTION

Today's information- and technology-rich environments are replete with "infotainment" systems for work and pleasure. In many of these environments, interaction with such systems occurs while the person is performing another task — for instance, walking while listening to an iPod, or writing a paper while watching an Internet video broadcast. One multitasking environment that has garnered a great deal of attention is that occupied by the driver of a car. Many empirical studies have found harmful effects of secondary-task interaction (e.g., dialing a cell phone) while driving [2,13,16,17]. Nevertheless, people continue at an increasingly alarming rate to perform distracting tasks while driving — for example, a recent study of over 5,000 American drivers [20] found that 40% of all drivers talk on cell phones, 20% of drivers aged 18-24 select songs on an iPod, and 24% of this younger group send text messages, all while driving.

To better understand multitasking behavior in driving and similarly complex domains, researchers have started to use cognitive modeling as a tool for addressing the many theoretical and practical challenges that arise. One common approach has used production-system cognitive architectures (e.g., ACT-R [1], EPIC [11]) that enable interaction with a simulated complex environment. While models in the ACT-R architecture have accounted for many performance measures of human driver behavior under dual-task conditions [13,14,15,16], these efforts have generally not attempted to capture the rich strategic variability exhibited by humans while multitasking. Another recent approach, cognitive constraint modeling (CCM), posits that predictions of behavior can be derived by finding the set of optimal strategies given an objective function, a set of plausible strategies, and a set of constraints on human performance [7,8]. CCM facilitates exploration of strategic variations in behavior and allows for objective functions to represent desired trade-offs in performance (e.g., between working memory load and task completion time [4]). One feature of the CCM approach is the demand for a quantitative analysis of the constraints imposed by interaction with the task environment. In contrast to simulation-based approaches, constraint models are not intended to interact with a simulated environment; rather, they demand a mathematical understanding of the consequences of action for objective-relevant features of the world. A particular concern in the context of driving is the consequence of inattention for lane deviation.

In this paper, we explore how models of driving and driver distraction can be developed within the context of cognitive constraint modeling (CCM). The most significant challenge in this effort lies in abstracting the driving simulation into a set of computationally efficient formalisms, which can then be utilized for constraint satisfaction and evaluation through an objective function. We develop such an abstraction based on human steering data and integrate the resulting formalism into a CCM model to predict speed/accuracy trade-offs in a driver distraction task.

Although we adopt the example of dialing a cell phone throughout this paper, the analysis is readily generalized to other in-car secondary tasks that contain a high degree of visual/motor interaction. For example, the analysis could be applied to predict driver performance while interacting with a portable digital music player or satellite navigation system. Moreover, there are strong implications for designers or usability evaluators regarding the safety of mobile devices that might be used while driving.

Modeling Human Multitasking

People frequently balance performance between two or more continuous tasks. Cognitive modeling research attempts to provide tools that predict human performance in complex real-world tasks. Models have been used to generate predictions of human performance on single tasks for decades (see http://act-r.psy.cmu.edu/, for example domains). Complexities arise, however, in modeling the combination of tasks [10,14]. These complexities are particularly severe when the tasks share perceptual and/or motor resources.

For instance, Meyer and Kieras [11] point out that if a dual-task involves a limited peripheral resource, such as the eyes, the need to move the eyes from one part of the visual field to another will block any possibility of task parallelism. Dual-task performance requires the allocation of limited resources to be balanced between each of the tasks. Predicting human performance under such conditions presents a significant challenge for cognitive modeling because of the combinatorial explosion in the possible ordering of task operators. In other words, people can adopt multiple behavioral strategies that interleave tasks in different ways. The cognitive modeler is left striving to develop models that account for this strategic variability.

As an example consider the task of dialing a 7-digit phone number while driving. (Throughout, we shall also assume that one “power-on” key-press precedes the 7-digit number and that one “send” key-press follows it — giving 9 key-presses in all.) When entering a cell phone number, the driver could enter all 9 key-presses sequentially without once returning attention to the task of driving. Alternatively, the driver could enter key-presses singly, returning attention to driving after entering each individual key-press. Between these two extreme strategies exist a plethora of alternative ways to complete the task (i.e., whether to switch back to driving between each pair of key-presses). There are at least $2^8 = 256$ possible strategic variations in this space.

Previous efforts to model human performance during dual-task conditions have tended to focus on a very small fraction of viable strategies — namely, those that could be inferred from empirical data and/or task analysis [10,14]. For example, Salvucci [14] presented a model for dialing while driving that utilized a multitasking general executive. The general executive mechanism provides domain-independent methods for integrating task models. However, as with previous modeling efforts, even this approach requires explicit programming of the points at which control switches between tasks — typically at points that are considered by the modeler to be intuitively convenient or expedient. In this case, the model implemented a strategy that was based on an analysis of inter-key intervals in the human data, which suggested that the chunk structure of the telephone number (e.g., for a 7-digit telephone number this might follow a 3-4 representational structure) provided natural break points to return attention to driving. The model accounted for the increase in dialing time required while driving compared to baseline, and also the degraded steering that resulted from the introduction of the secondary dialing task.

One question that has not been adequately addressed within the literature is *why* people might adopt a particular multitasking strategy over another — for instance, based on their desire to complete the secondary task quickly, or their desire to drive as safely as possible. Both Roberts and Pashler [12] and also Kieras and Meyer [9] have expressed the potential problems of failing to explore the contribution of strategies and architectural constraints on the range of possible models of human performance.

Kieras and Meyer [9] responded by proposing the use of a bracketing heuristic. A bracket was defined by the speed of the fastest-possible strategy for the task, and the slowest-reasonable strategy; thus, observed performance should fall somewhere between the performance of these two strategies. Kieras and Meyer also articulated the importance of exploring the space of strategies to explaining the phenomena being modeled. In this way, bracketing provides a way to construct truly predictive models in complex task domains where the optional strategy optimizations cannot be forecast.

Recent work using cognitive constraint modeling has also addressed the question of modeling strategic variability in complex task domains [7,8]. Eng et al. [4] focused on the question of evaluating two different designs for a flight deck control panel. They demonstrated that performance measures and objective functions could play a key role in modeling behavior: Predictions could be shaped not only by basic task goals and the constraints on the cognitive architecture, but also by the specific performance objectives (e.g., minimizing task completion time or working memory load).

The CCM approach affords the kind of rich strategic analysis we would like for understanding the complex domain of driving and driver distraction. There is, however, a significant limitation in the application of CCM to driving. Previous models of driving [13,14,15,16,18] have tended to interact with a simulation environment that, in essence, provides the mapping from behavior (e.g., steering and acceleration) to changes in the environment (e.g., the vehicle's lane position). Cognitive constraint models in contrast are not intended to interact with a simulated environment; rather, they demand a mathematical understanding of the consequences of action for objective-relevant features of the world.

In order to provide a cognitive constraint model of driving, we developed a computational formalism of the interaction between the driver and the task environment. We constructed this formalism by analyzing trends in empirical driver data and expressing these trends as equations that provide estimates of performance measures. In particular, we attempted to capture how driver performance changes over time based on whether the driver is attending to the task of steering the car. We focused our analysis on established measures of driver performance; namely, the lateral distance (or lateral deviation) of the car from the center of the lane. Focusing on this feature of driving performance enabled a CCM analysis of our illustrative phone-dialing task. In the next section we use this analysis to systematically understand a large space of reasonable multitasking strategies.

A COGNITIVE CONSTRAINT MODEL OF DRIVING IN DUAL-TASK CONDITIONS

One of the components of safely driving a car involves maintaining a central lane position. While driving down even a straight road a driver is required to make many minor corrections to the heading of the car. The driver alters the heading of the car by adjusting the angle of the steering wheel. People rely on visual feedback for this task; removing visual feedback can have disastrous consequences for a person's ability to correctly estimate the heading of the car, especially during lane change maneuvers [5,19]. Essentially, driving can be thought of as a closed-loop motor task that requires the integration of continuous visual feedback. Maintaining a central lane position therefore demands the driver's attention.

In this paper, we consider the consequences of disrupting attention from the task of driving while using a secondary in-car device, such as a cell phone. A high-level model of driving was developed that was based on analyses of human performance data from studies that investigated distracted driving conditions. Observed trends in human performance data were quantified as functions of time and the vehicle's lateral deviation. The aim was to develop a model that predicted changes in lateral deviation under dual-task conditions.

The driver performance data were taken from two experiments that investigated the effect of cell phone use on driving [13,16]. Both experiments were conducted in a fixed-base driving simulator that included the front half of a Nissan 240sx with standard steering and pedal controls. These controls were connected via a hardware interface to a desktop computer that ran the simulation and data collection software. The driving environment in the first study [13] used a single-lane roadway with no other vehicles present. The environment in the second study [16] used a three-lane construction-zone highway environment, where the driver navigated in the center lane behind a rapidly accelerating and decelerating lead vehicle. The secondary task in both experiments involved dialing a 7-digit phone number on a hands-free phone — a simulated cellular phone in the first study and a real Sprint™ cellular phone in the second study. (Both studies also included other dialing conditions that are not included in our analysis below.) In all, we analyzed data from 15 drivers across these two experiments (seven from the first and eight from the second).

Analysis focused on formally abstracting how people's steering behavior affects the lateral deviation of the vehicle. 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). Movement through the center of the lane was discounted by transforming all data by its absolute value (i.e., making all negatives positive). Steering *episodes* were defined within a sequence of steering data as a period of time in which the angle of the steering wheel did not significantly alter.

We assumed that during periods of disruption to normal driving the vehicle's lateral deviation would generally increase over time and would decrease once attention was returned to the primary task of driving. We refer to these as divergent and convergent steering episodes. An episode was considered a *divergent steer* if the lateral deviation of the car at the start of the episode was less than its lateral deviation at the end of the episode. An episode was considered a *convergent steer* if the lateral deviation of the car was less at the end of the episode than it was at the start of the episode. For a given sequence, we collapsed identical contiguous steering episodes. Steering episodes were represented as a tuple {*type, duration, start lane position, end lane position*}. Data from all steering episodes across different participants were pooled. Episodes were aggregated if they were of the same type and had a difference in duration of less than 100 ms. We report an analysis of the mean and standard deviation of these aggregate episode data.

Analysis of Divergent Steering

It was assumed that periods of driver distraction would lead to an increase in the vehicle's lateral deviation. The size of this increase D in lateral deviation is calculated given a simple mathematical relationship between the duration T

and the lateral velocity *V* of a divergent steering episode ($V = D / T$). We analyzed the episode data in order to gain an estimate of the vehicle's typical lateral velocity during periods of driver inattention.

Divergent steering episode data are presented in Figure 1. (Note that each point in the scatter plot is an average over similar episodes within the pooled participants data.) It can be seen that there is a positive relationship between the duration of an episode and the increase in lateral deviation. Figure 1 also shows the best fitting straight line through this data, which provided a reasonable level of correspondence with the human data ($r^2 = 0.43$), $t\ (62) = 6.90$, $p < .001$. The slope of this best fitting line yields the average lateral velocity of a divergent steering episode (i.e., 0.2833 m/s). Thus, given the duration of a divergent steering episode the associated increase in lateral deviation is predicted by the following,

$$\textit{Lateral Deviation} = 0.2833 \times \textit{Duration} \quad (1)$$

There is another less obvious quality to the data, which is not captured by the aggregate curve: As the duration of a divergent steering episode increases there is an increase in the variance around the mean lateral velocity. Indeed, such an increase in variance should be expected if we consider the mechanics underlying the movement of a car. Essentially, over a period of time the car's position can either remain relatively stable or it can veer from the center of the lane. That is, the driver has either accurately estimated or grossly misestimated the correct steering angle to leave the car in while they engage in a secondary task. Under this assumption, if we consider a short duration between corrective steering updates, then the range of the increase in lateral deviation is quite small. This is because sufficient time has not passed for the car to travel very far from its starting point. But as the duration increases the range of possible end points also increases dramatically because more time has passed for the car to move along its heading. A function was derived using linear regression to predict the relationship between the mean increase in lateral deviation and the standard deviation of the increase,

$$\textit{Standard Deviation} = 0.6824 \times \textit{Lateral Deviation} \quad (2)$$

The function given by Equation 2 is represented in Figure 1. There is a good fit between the function and the human data ($r^2 = 0.66$), $t\ (62) = 11.10$, $p < .001$.

The analysis shows that increases in lateral deviation can be modeled given the duration of time between corrective steering updates. A stochastic performance model can be developed: Equation 1 and 2 provide the mean lateral deviation and standard deviation of the mean that allow sampling from a normal distribution. The resulting model captures the basic assumption that as the time between corrective steering updates increases the more likely the car is to drift from the center of the lane.

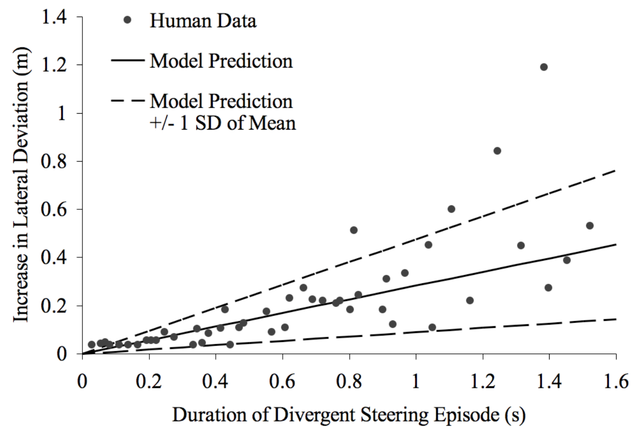

Figure 1. Relationship between the duration of divergent steering episode and increase in lateral deviation.

Analysis of Convergent Steering

We next turn our attention to an analysis of convergent steering episodes. We assume that these episodes are motivated by the driver attempting to move the car back to a central lane position following a period of distraction. We are therefore interested in how the lateral deviation decreases during periods of convergent steering. Indeed, an obvious component of convergent steering is that active adjustments to the heading of the car affects the rate at which its lateral position changes over time — that is, the vehicle's lateral velocity.

Figure 2 presents the relationship between the initial lateral position of the car at the beginning of a convergent steering episode and lateral velocity (i.e., the decrease in lateral deviation / duration of time). The human data suggest that as the car strays closer to the lane boundary during periods of driver inattention, drivers react by making sharper corrective steering movements to return the car to a central lane position. A function was derived using linear regression to predict lateral velocity given the initial lateral deviation of the car at the beginning of the episode,

$$\textit{Lateral Velocity} = 0.1756 \times \textit{Lateral Deviation} + 0.1034 \quad (3)$$

We assumed a maximum lateral velocity of 0.46 m/s. This upper bound was determined by the upper confidence interval on the distribution of observed velocities. The function given by Equation 3 is represented in Figure 2, and it provided a reasonable degree of correspondence with the human data ($r^2 = 0.42$), $t\ (32) = 4.86$, $p < .001$. Further analysis of the human data revealed a constant standard deviation of the mean of 0.09 m/s. This is also represented in Figure 2. In order to develop a stochastic model of the lateral velocity of a convergent steering episode, we sampled values from a normal distribution given the mean lateral velocity (Eq. 3) and the standard deviation observed in the human data.

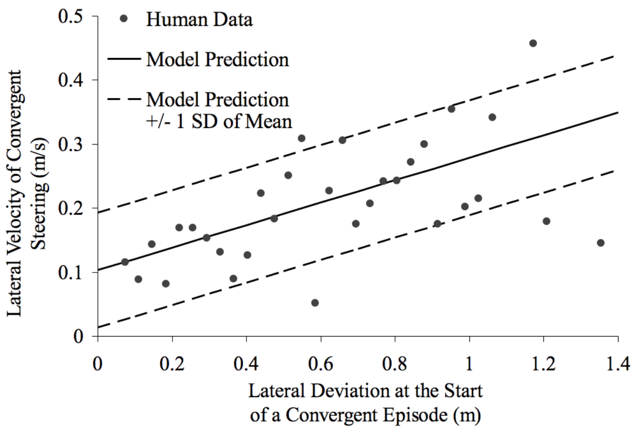

Figure 2. Relationship between lateral deviation at the start of a convergent steering episode and lateral velocity.

The decrease in lateral deviation for a convergent steering episode is given by multiplying lateral velocity by the duration of time that is given up to steering. As should be apparent at this stage, lateral deviation must be greater than zero because a value of zero causes problems with multiplications in the model; therefore, we assume an arbitrary minimum lateral deviation of 0.001 m. Moreover, this assumption reflects the idea that under dual-task conditions, the decrease in lateral deviation associated with ever increasing time given up to steering asymptotes. We return to this issue shortly.

The method described above captures some important qualities of the convergent steering data: 1) The more time that is given up to a convergent steering episode, the closer the car returns to a central lane position; 2) as the car strays closer to the lane boundary, the greater the lateral velocity of the convergent steering episode; 3) the greater the lateral movement of the car, the more rapidly the car returns to a central lane position; 4) velocity only increases up to a maximum rate, however, at which point the car returns to the center of the lane at a constant rate.

A Cognitive Constraint Model of Dialing while Driving

The analyses of divergent and convergent steering behavior outlined in the previous section were integrated to develop a cognitive constraint model of driving under dual-task conditions. There were several operating criteria for specifying this model:

1. The model presents multitasking between an ongoing dynamic task (steering) and a secondary task (dialing).
2. Allocation of resources between the two tasks should be predicted given psychological constraints and an objective function.
3. An objective function that quantifies performance in the overall task should reflect a trade-off between steering task performance (i.e., minimizing divergence from the center of the lane) and dial task performance (i.e., minimizing task completion time).

The cognitive constraint model made minimum commitments to representing task processes for the driving task. In the model, driving performance was determined by the duration of distinct periods of distraction, during which the secondary dial task was processed, and subsequent periods of corrective steering. Given the previous analysis of the human steering data we developed a computational formalism that abstracted over the fine-grained details of the driving task.

For the dialing task a 7-digit number was entered. 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. The duration of each key-press was set at 310 ms, which assumed 50 ms for recalling the digit, a 50 ms step of cognition, where the motor response is initiated, and 150 ms motor preparation and 60 ms motor execution for the key press. In addition, it was assumed that the dial task was preceded by a hand movement to the phone from the steering wheel, and that it was followed by a hand movement back to the steering when the dial task was completed. We assumed that each of these hand movements took 800 ms. All of the assumptions regarding timing estimates for the dial task were based on a previously reported model in the literature [14].

A further assumption of the model was that switching between tasks would carry a cost overhead (what we shall refer to as a switch cost). More specifically, it was assumed that a time cost would be incurred by moving visual attention between the outside 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 used 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 [1].

The model assumed that only a single task operator could be processed at any given time (e.g., a convergent steering episode could not occur at the same time as a key-press). Switches of attention between tasks could only occur at unit task boundaries (e.g., on the completion of key-press or at the end of a convergent steering episode). Although it might appear that this high-level representation of task operators is an over simplification, it was functionally equivalent to at least one fine-grained model of driver distraction in the literature: Salvucci's [14] model developed within the ACT-R cognitive architecture that used a queuing mechanism to switch between tasks at operator boundaries.

Figure 3 provides a schematic overview of the cognitive constraint model. The upper panel of Figure 3 represents the dial and steer operators. Switch costs operators are represented as a switch-operator that links transitions between dial and steer operators. Operators are represented as a PERT-like representation depicting the flow of information in separate task streams over time. Time is represented on the horizontal axis (moving left to right) and resource streams on the vertical axis.

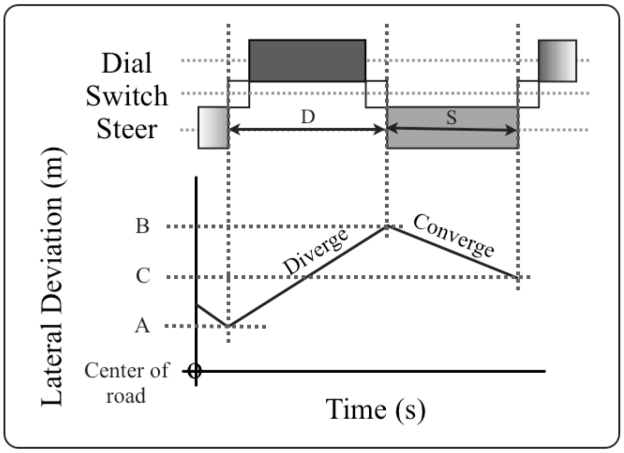

Figure 3. Behavior graph of interleaved steering and dialing over time, and corresponding data plot of shifts in lateral deviation over time.

The lower panel of Figure 3 represents how the operator ordering specified by a particular multitasking strategy was used to model changes over time in the lateral deviation of the car. As an example, we can define an initial lateral deviation of *A* on the *y*-axis in the lower panel of Figure 3. While dialing, lateral deviation increases. The extent of this increase is dependent on the amount of time that attention was directed away from the task of steering. For instance in the figure, arrow *D* defines a period of time between the end of one steering operator and the start of the next steering operator, which followed the completion of a key-press. The lateral position of the car at the end of this period is represented by point *B*, which is given by the sum of the increase in lateral deviation and the initial position *A*. As we outlined earlier, the size of the increase is defined by a randomly sampled value from a normal distribution with *mean* = 0.2833 x *D* (Eq. 1) and *SD* = *mean* x 0.6824 (Eq. 2). It is important to note that the duration of any switch costs were added to the duration of divergent steering episodes and not to the duration of convergent steering episodes.

The lateral deviation of the car decreased during periods of convergent steering. In the figure, arrow *S* defines the duration of a steering operator during which the lateral position of the car decreases from point *B* to point *C*. The lateral velocity *V* of the corrective steering movement was determined by randomly sampling a value from a normal distribution with *mean* = 0.1756 x *B* + 0.1034 (Eq. 3) and *SD* = 0.09. The updated lateral position *C* was determined by the initial lateral position *B*, lateral velocity *V*, and duration *S* of the steering movement ($C = B - V \times S$, where *C* is constrained to be greater than 0.001 m).

In order to fully evaluate the strategy space we generated all permutations for ordering the task operators and, over each of these permutations, we further enumerated over a set of steering durations. That is, for a given interleaving strategy we explored a range of discrete values for the duration of each of the steering operators. Following a validated control model of steering [15], which posits 150 ms processing time for each steering update, we explored multiples of 150 ms for each steering operator. We set a conservative upper limit of 1500 ms for steering updates that were evaluated. This upper limit was chosen because we found that using steering updates greater than 1500 ms predicted asymptotic steering performance. (We provide evidence to support this claim in the Results section.) For each steering operator in a given strategy all permutations of duration 150, 300, 450, 600, 750, 900, 1050, 1200, 1350, 1500 ms were explored. In total this meant that some 262,701 distinct strategies were evaluated.

For each strategy, we modeled changes in lateral deviation over the duration of the multitasking episode. For each run, the lateral deviation of the car was initiated at 0.05 meters from the center of the lane. This estimate was based on that observed in the human data. As we have said, lateral deviation increased during periods of driver distraction and decreased during periods of convergent steering. For each strategy changes in lateral deviation were recorded. Each strategy was run for 50 trials; note that trials varied in terms of lateral deviation because of the random sampling from distributions specified for steering episodes. Trials did not vary in terms of dial time, however.

RESULTS

For each strategy, we were primarily interested in the time to complete the dial task and the average lateral deviation of the car while the dial task was being completed. We analyzed the performance of each of the strategy permutations, which essentially ranged from doing the entire dial task without driving (all 9 key-presses in sequence) to maximally interleaving driving (with steering updates between each pair of key-presses). Each strategy was evaluated using the cognitive constraint model described in the previous section. Figure 4 presents the performance prediction for each of the 262,701 distinct strategies, where the *x*-axis denotes the time to complete the dial task and the *y*-axis denotes the mean lateral deviation of the car from the center of the lane. Each of the model data points in the figure represents an average over 50 trials. It is clear that as the duration of each dial episode decreases, total dial time increases and, as a consequence, lateral deviation decreases. Moreover, Figure 4 clearly shows the speed/accuracy trade-off that exists between dialing quickly and driving safely: The upper-left portion of the plot represents faster but less safe performance, whereas the bottom-right portion represents slower but safer performance.

It is possible to highlight interesting and illustrative example strategies within the strategy space shown in Figure 4. The details of each of these example strategies are presented in Figure 5 using a PERT-like representation to depict the flow of activity in separate task streams over time. The duration of each steering operator is given in the figure. It is worth noticing that none of these illustrative strategies, in particular the *safest* strategies, incorporated a

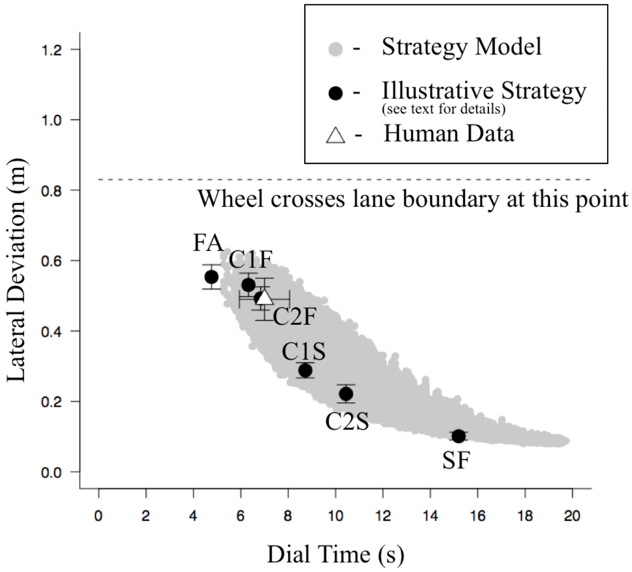

Figure 4. Data plot of dial time and average lateral deviation for each of the 262,701 strategy models and human data. Illustrative strategies represent: *FA* = Fastest; *C1F* = fastest 3-4 chunking; *C2F* = fastest 3-2-2 chunking; *C1S* = safest 3-4 chunking; *C2S* = safest 3-2-2 chunking; *SF* = safest. All error bars represent 95% confidence interval of mean.

steering update greater than 1.35 s. We show shortly that giving up any more time to steering was redundant because it did not bring a significant decrease in lateral deviation.

Figure 5 also includes, for each strategy, an estimate of the average time required to return the car to a threshold lateral deviation of 0.10 m. Although not explored in a more detail here, this *recovery time* is interesting because it points to another potential trade-off between 1) the time required to complete the dial task and 2) the overall time needed to complete the dial task and "reset" the vehicle back to its pre-secondary task state.

Another interesting facet to the lateral deviation data was that although some strategies maintain a low average lateral deviation, they nonetheless predict occasional lane departures. Assuming a standard lane width of 3.66 m and a mid-sized car width of 2 m, the outer wheel of the car would be expected cross the lane boundary at 0.83 m lateral deviation from the center of the lane. For each strategy we therefore calculated the frequency of trials with which it committed a lane departure while completing the dial task. It was interesting to find that all strategies that made fewer than two steering updates during the dial task posed a significant risk of committing a lane violation, in that on fewer than 5% of trials they committed a lane violation. With this in mind, we return to the modeling results for each of the illustrative strategies.

Illustrative Strategies

It can be seen in Figure 4 that the fastest strategy (*FA*) and the safest strategy (*SF*) show the extremes of the speed/ accuracy trade-off space. As expected, the fastest strategy completes the dial task in a single contiguous block without once returning attention to driving. The predicted dial-time for this strategy is 4.76 s, which is simply the sum of the dial-task operators (i.e., 9 x 0.31 s for key-presses, 2 x 0.185 s for switch costs, and 2 x 0.8 s for hand movements to and from the phone). As would be expected, completing the dial task without returning attention to driving resulted in a relatively large lateral deviation (*M* = 0.554 m, *SD* = 0.122 m) compared to other strategies. Moreover, this strategy frequently resulted in lane departures (38/50).

We might assume that the safest strategy would maximally-interleave dialing with frequent steering updates. Indeed, strategies with steering updates of up to 1.5 s each were explored; thus, the maximum dial-time for a strategy was 19.72 s (i.e., 9 x 0.31 s for key-presses, 8 x 1.5 s for steering updates, 18 x 0.185 s for switch costs, and 2 x 0.8 s for hand movements to and from the phone). As would be expected, this strategy resulted in a relatively small lateral deviation (*M* = 0.088 m, *S.D.* = 0.029 m) compared to other strategies. However, it is clear from Figure 4 that lateral deviation asymptotes with dial-times greater than approximately 14 s.

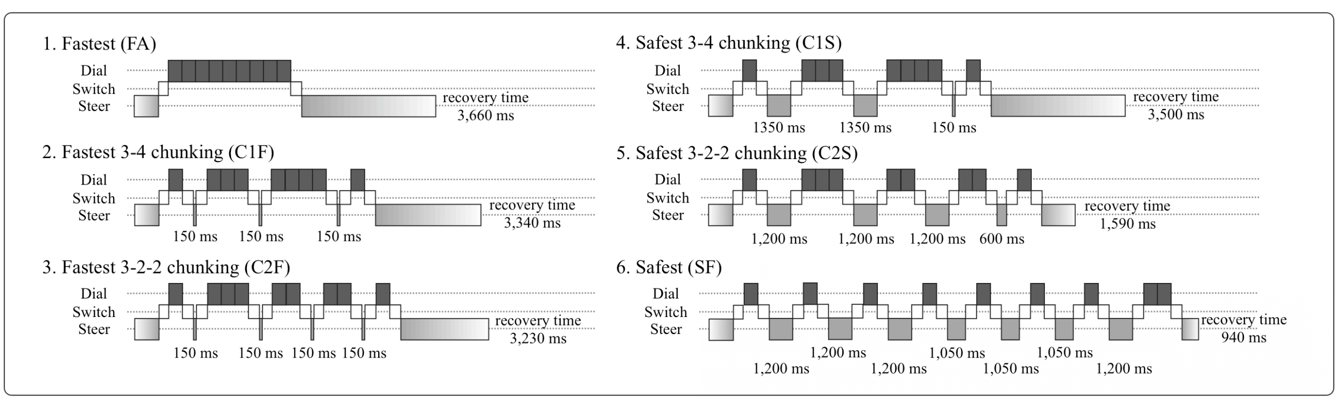

Figure 5. Behavior graphs representing interesting and illustrative strategies within the space of possible multitasking strategies.

We preset an analysis that defines the safest strategy that also completed the dial task in a timely manner. Because the driving model was stochastic, the average lateral deviation of a strategy differed between model runs. It was possible to utilize this variance in order to identify a subset of strategies that did not significantly differ from the strategy that gave up the most amount of time to driving while completing the dial task. To this end, a series of t-tests were conducted to reject strategies. A strategy was rejected if it produced a lateral deviation that was significantly greater ($p < 0.05$) than the lateral deviation of the maximum dial-time strategy (see above). From this set of *safe-driving* strategies, the strategy that completed the dial task in the least amount of time was defined as the safest strategy (*SF*). In other words, although strategies with steering updates of up to 1.5 s were explored, we identified a strategy that proved to be as safe, in terms of lateral deviation (M = 0.090, SD = 0.036 vs. M = 0.088 m, $S.D.$ = 0.029 m, respectively), but completed the dial task in a more timely manner (16.58 s vs. 19.72 s, respectively). The safest strategy (*SF*) never predicted a lane departure (0/50).

Moreover, Figure 5 shows that the safest strategy did not actually need to maximally-interleave dialing with steering updates. This is a somewhat unexpected finding: Dialing up to two digits in quick succession, at least in some conditions, can be as safe as dialing digits singly.

We next consider a set of strategies that correspond to North American phone number chunking structures — that is, a 3-4 or 3-2-2 breakdown of a phone number (plus an initial "power-on" and final "send"). For each of these two different chunking strategies we highlight only the fastest and safest alternatives (these variations were identified in same way as strategies *FA* and *SF*). It is clear from Figure 4 that giving up more time to driving between dial bursts provides much safer performance than when dial time is minimized by giving up less time to driving between dial bursts. Moreover, the fastest strategy variants resulted in many more lane departures than the safest strategy variants: *C1F* = 33/50; *C2F* = 33/50; *C1S* = 1/50; *C2S* = 0/50.

Comparison with Human Data

A comparison is offered between the model-based predictions and previous empirical results. Salvucci [13] reports that participants required 7 s (SD = 1.77 s) for an equivalent dialing task, and that their average lateral deviation was 0.49 m (SD = 0.10 m) during secondary-task performance. These human data are presented in Figure 4. It is clear that the human data are within the performance limits predicted by the cognitive constraint model. In particular, one of the fast chunking strategy variants (*C2F*) predicted performance, in terms of both dial-time and lateral deviation, that was very close to that observed in the human data (see Figure 4). Indeed, this result corroborates previous modeling efforts that have assumed that people adopt a strategy that corresponds to the chunking structure of the telephone number. The major difference here though is that this correspondence has been reached following a comprehensive evaluation of the plausible strategy space.

It is also clear that a whole range of strategies actually fit the human data approximately as well: We found that some 9,479 strategies gave performance predictions within the 95% confidence intervals of the human data. These strategies differ widely on dimensions other than their correspondence to the data, and only a small proportion of these strategies employed a chunking structure. Strategies varied in the number of steering updates: A strategy that made only a single, but very long driving update, also fit these data, as did a strategy that made as many as six shorter driving updates. These strategies make very different predictions for the total time required to complete the dial task and return the vehicle to its pre-secondary task state. Moreover, strategies differed in terms of the proportion of trials that a lane departure occurred (e.g., seven of the strategies committed 4/50 lane departures, which is far fewer than the 33/50 made by strategy *C2F*).

GENERAL DISCUSSION

In this paper we have described an approach to modeling the space of strategies available for performing secondary tasks while driving. For this purpose, we described the constraints on the interaction between driver steering episodes and the local task environment, and analyzed the speed and safety of the set of possible strategies given these constraints. We have shown how various strategies for interleaving dialing and driving result in different speed/ accuracy trade-offs. It was found that human performance data in the literature [13] fit within the bounds on the predicted dual-task lateral deviation and dial-time.

We found that strategies with a familiar chunking structure gave performance predictions that were very close to the human data (e.g., adopting a 3-4 or 3-2-2 dialing pattern). This strategy account is consistent with others in the literature. For instance, based on an analysis of inter-key intervals in the human data, Salvucci [14] developed an ACT-R model that assumed attention was ceded away from steering while a chunk of digits was dialed. The resulting model was found to fit the human data. What is novel about the work presented here is that all reasonable strategies are evaluated and their performance characteristics generated. The analysis offers an explanation for *why* people behave the way that they do in terms of the trade-off between dial time and safety. It can be seen how participants might view other strategies as inferior in terms of this trade-off. In this respect the approach presented here has similarities to bracketing [9].

The work that we have reported differs from other approaches to modeling driver behavior [13,14,15,16,18] that have focused effort on programming software that links a computational cognitive architecture to a virtual reality simulation of a driving environment. A feature of this approach has been a concern with reproducing, by simulation, the internal psychological processing mechanisms, the mechanics of vehicle motion, and the interaction between the two. In contrast, we have focused on understanding the constraints on the interaction between the driver and the task environment that are important to some critical performance variables (e.g., lateral deviation and dial-time). Two potential advantages of our approach are (1) parsimony of expression and (2) explanations for why people prefer the strategies that they to choose in terms of objective performance criteria.

An interesting empirical question that is posed by the current analysis is whether manipulation of the speed/ accuracy trade-off would result in behavioral changes equivalent to those predicted. For instance, if participants are encouraged to dial as quickly as possible, then we might expect this to result in a relatively large lateral deviation. In contrast, if participants are encouraged to drive as conservatively as possible while dialing, then we might predict that the duration of dial episodes would decrease and consequently the total dial time would increase. Further empirical work is required to address these questions. Moreover, the analysis revealed that a whole range of strategies fit the human data equally as well, and that these strategies vary along other performance dimensions (e.g., the number of steering updates, the frequency of lane departures, and the total time required to complete the dial task and return the vehicle to a central lane position). Data are not currently available to discriminate between these accounts. There is an open empirical question regarding the type(s) of objective function that people might be sensitive to. Given experimental control over an objective function, we should be able to move beyond bracketing to make more precise predictions of human performance.

Limitations

In the current analysis it was assumed that steering and dial operators were limited by a serial bottleneck. This is clearly an oversimplification of human performance. Future work should extend the cognitive constraint analysis presented here to one that incorporates more detailed assumptions of the resource constraints of the human cognitive architecture. Assumptions of the human cognitive architecture have been well specified elsewhere (e.g., ACT-R [1], EPIC [11]). Although these account differ, it is nonetheless commonly assumed, for instance, that people can move their eyes to a visual location at the same time that their finger punches a key. Presumably incorporating such assumptions about perceptual/motor parallelism would have consequences for predicting dual-task performance here.

One avenue would be to develop a set of task models within a task description language, called Information Requirement Grammar (IRG, [6]). IRG is motivated by the theory that higher-level task performance is constrained by the information requirements and resource demands that operate on lower level task processes. A Prolog-based tool, called CORE [6,7], can be used to expand the task description specified in the IRG to determine an optimal schedule of the start times for each low-level process. Previous research has found that this approach allows for exploration over the set of permissible strategies within the set of task and architectural constraints (Eng et al. [4]; Howes et al. [7], but see Brumby & Salvucci [3] for an initial report on progress towards this goal in the domain of driver distraction).

In the current analysis all strategy permutations were exhaustively evaluated. There is redundancy in this analysis though because many of the strategies do not meaningfully differ, both in terms of the specified task interleaving and the performance predictions that are derived. For example, it might be argued that the comparison between strategies 3-3-4 and 3-4-3 is not particularly meaningful or useful. Moreover, evaluating all of the strategy permutation suffers the problem of exponential complexity: As the number of operators increases, the number of solutions increases exponentially. Search heuristics could be used to avoid this redundancy in the analysis by terminating the evaluation of strategy variants that do not significantly differ from one another.

Implications

The work presented here further develops the cognitive constraint modeling approach [4,6,7,8] in understanding the implications of a space of strategies for behavior. The work in the current paper demonstrates that it may be viable to utilize this approach in order to inform the development of interfaces that optimize human performance under dynamic dual-task (or multitasking) conditions. This method could also be used by usability evaluators to compare the methods that a proposed or current design makes available to the user to the best possible set of methods.

We believe that it would be helpful, from a cognitive engineering perspective, to develop a high-level driver model that is capable of abstracting across the effects of a wide range of different in-car devices that might vary, for instance, in the amount of visual or motor interaction that is required to use the device. Initial, unreported efforts suggest that there is a high degree of correspondence between the best-fitting values given in Eq. 1-3 and those derived from a similar analysis of data collected from an iPod distraction study [17]. This suggests that the cognitive constraint model of distracted driving developed here is likely to be robust across different types of secondary tasks.

In contrast, one important area where different types of mobile devices differ is the amount of interaction that is required to complete a task. The analysis presented here clearly suggests that as the duration of an interaction episode increases, the associated risk to driving increases considerably. Allowing shorter interaction episodes with the device is therefore vital, if it is at all likely that the driver of a car will use a device. For example, scrolling through a list of contacts on a cellular phone or a list of artists on an iPod might demand longer interaction episodes with the device than manually entering the friends telephone number or playing music through the “shuffle” option. Task time is also related to conflicts between input and output modalities; for instance, some studies have reported that speech-based dialing can consume more time than manual dialing but result in no significant degradation in performance [13]. The modeling approach developed here could potentially serve the foundation for finding particular interaction-styles that could be designed into a device so as to minimize potential hazardous effects of use while driving.

SUMMARY

Despite the obvious risk factors, people continue to use mobile devices while driving [20]. Analyses of driving performance data revealed that a set of simple equations accurately modeled changes in the lateral position of the car within the lane under dual-task conditions. The model quantified how the vehicles lateral deviation increased during periods of driver inattention, and decreased during periods of active steering. The benefit of this approach was demonstrated by modeling the dialing of a cellular phone while driving. It was shown that understanding, rather than simulating, the constraints imposed by the task environment can help to explain the costs and benefits of a range of multitasking strategies. For instance, each multitasking strategy was sensitive to a combination of internal constraints (including switch costs) and the trade-off between the amount of time allocated to secondary task and the risk of extreme lane deviation.

The work presented is of value to designers wishing to predict the extent of disruption caused by using mobile devices while driving. Abstracted away from dialing, the analysis suggests that the interaction strategy adopted by a user has an effect on the level of disruption caused by using a mobile device while driving. Care should be taken in considering the type of interaction strategy that a device demands, if there is a chance the device may be used while driving or during other complex tasks.

ACKNOWLEDGMENTS

This research was supported by National Science Foundation grant #IIS-0426674 and Office of Naval Research grant #N00014-03-1-0087. We would like to thank the six anonymous reviewers for their astute comments and suggestions for improving the initial draft of this manuscript. We would also like to thank Mark Zuber and Dan Markley for their help with data analysis.

REFERENCES

1. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of mind. *Psychological Review, 111*, 1036-1060.
2. Briem, V., & Hedman, L.R. (1995). Behavioral effects of mobile telephone use during simulated driving. *Ergonomics, 38*, 2536–2562.
3. Brumby, D.P., & Salvucci, D.D. (2006). Exploring Human Multitasking Strategies from a Cognitive Constraint Approach. In *Proc. Cognitive Science Society* (pp. 2451). Lawrence Erlbaum Associates.
4. Eng, K., Lewis, R.L., Tollinger, I., Chu, A., Howes, A., & Vera, A. (2006). Generating automated predictions of behavior strategically adapted to specific performance objectives. In *Proc. CHI 2006* (pp. 621-630). ACM Press.
5. Hildreth, E., Beusmans, J., Boer, E., & Royden, C. (2000). From vision to action: experiments and models of steering control during driving. *Journal of Experimental Psychology: Human Perception and Performance, 26*, 1106–1132.
6. Howes, A., Lewis, R.L., Vera, A., Richardson, J. (2005). Information-Requirements Grammar: A theory of the structure of competence for interaction. In *Proc. Cognitive Science Society* (pp. 977-982). Lawrence Erlbaum Associates.
7. Howes, A., Vera, A., Lewis, R.L., & McCurdy, M. (2004). Cognitive constraint modeling: A formal approach to supporting reasoning about behavior. In *Proc. Cognitive Science Society* (pp. 595-600). Lawrence Erlbaum Associates.
8. Howes, A., Vera, A., & Lewis, R.L. (In press). Bounding rational analysis: Constraints on asymptotic performance. In W.D. Gray (Ed.) *Integrated Models of Cognitive Systems*. Oxford University Press.
9. Kieras, D.E., & Meyer, D.E. (2000). The role of cognitive task analysis in the application of predictive models of human performance. In J.M.C. Schraagen, S.F. Chipman & V.L. Shalin (Eds.), *Cognitive Task Analysis* (pp. 237-260). Lawrence Erlbaum Associates.
10. Kieras, D.E., Meyer, D.E., Ballas, J.A., & Lauber, E.J. (2000). Modern computational perspectives on executive mental processes and cognitive control: Where to from here? In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes: Attention and Performance XVIII* (pp. 681-712). MIT Press.
11. Meyer, D.E., & Kieras, D.E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review*, *104*, 3-65.
12. Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107*, 358-367.
13. 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.
14. Salvucci, D.D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science, 29*, 457-492.
15. Salvucci, D.D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors*, *48*, 362-380.
16. Salvucci, D.D., & Macuga, K.L. (2002). Predicting the effects of cellular-phone dialing on driver performance. *Cognitive Systems Research*, *3*, 95-102.
17. Salvucci, D.D., Markley, D., Zuber, M., & Brumby, D.P. (In press). iPod distraction: Effects of portable music-player use on driver performance. To appear in *Proc. CHI 2007*. ACM Press.
18. Tsimhoni, O. & Liu, Y. (2003). Modeling steering using the queueing network – model human processor (QN-MHP). In *Proc. Human Factors and Ergonomics Society* (pp. 1875-1879). Human Factors and Ergonomics Society.
19. Wallis, G., Chatziastros, A., & Bülthoff, H. (2002). An unexpected role for visual feedback in vehicle steering control. *Current Biology, 12*, 295-299.
20. "2006 GMAC Insurance National Drivers Test", available at http://www.gmacinsurance.com/SafeDriving/2006/.