

Running head: SEQUENTIAL DEPENDENCIES IN DRIVING

Sequential Dependencies in Driving

Anup Doshi

University of California, San Diego

Cuong Tran

University of California, San Diego

Matt Wilder

University of Colorado, Boulder

Michael Mozer

University of Colorado, Boulder

Mohan M Trivedi

University of California, San Diego

Abstract

The effect of recent experience on current behavior has been studied extensively in simple laboratory tasks. We explore the nature of sequential effects in the more naturalistic setting of automobile driving. Driving is a safety-critical task in which delayed response times may have severe consequences. Using a realistic driving simulator, we find significant sequential effects in pedal-press response times that depend on the history of recent stimuli and responses. Response times are slowed up to 100 ms in particular cases, a delay which has dangerous practical consequences. Further, we observe a significant number of history-related pedal misapplications, which have recently been noted as a cause for concern in the automotive safety community. By anticipating these consequences of sequential context, driver assistance systems could mitigate the effects of performance degradations and thus critically improve driver safety.

Sequential Dependencies in Driving

In many naturalistic tasks, it is critically important for an individual to respond quickly to a sequence of cues in a rapidly changing environment. For example, drivers on highways worldwide are increasingly getting stuck in stop-and-go traffic (Chang, 2010), and are forced to engage in a sequence of braking and accelerating actions as dictated by cues from other vehicles. These situations sometimes take a turn for the worse: over 3,500 vehicles were involved in fatal rear-end collisions on highways in the U.S. in 2008 (NHTSA, 2010). In this research, we suggest that the recent sequence of stimuli and responses in driving—specifically, cues to brake and accelerate—have a profound effect on driver response times.

Such *sequential dependencies* have been studied in the psychological literature for over half a century (Anderson, 1960; Estes, 1950; Remington, 1969). A sequential dependency is an influence of one incidental experience on subsequent experience. Sequential dependencies arise when individuals perform a task repeatedly or perform a series of tasks, and one trial influences behavior on subsequent trials. Diverse behavioral measures are affected, including response time, accuracy, type of errors produced, and interpretation of ambiguous stimuli. Recency effects occur across all components of the cognitive architecture, including perception (Maloney, Dal Martello, Sahm, & Spillmann, 2005), selective attention (Kristjansson, 2006), language (Bock & Griffin, 2000), decision making and judgment (Jesteadt, Luce, & Green, 1977), and motor control (Dixon & Glover, 2004, 2009; Wilder, Ahmed, Mozer, & Jones, 2010).

The canonical laboratory task used to study sequential dependencies has been *two-alternative forced choice (2AFC)*, in which individuals are shown a series of trials in which one of two stimuli is presented, and the task is to make a unique response associated with the stimulus. For example, the stimuli might be the letters X and Y and the response a left button press for the X and a right button press for the Y. Theoretical accounts have been developed

for this task that obtain exquisite fits to data, accounting for over 98% of the variance in mean RT conditioned on the recent trial context (Cho et al., 2002; Wilder, Jones, & Mozer, 2010). These theories assume two distinct, additive mechanisms: priming based on what we will refer to as *first-* and *second-degree* statistics. Priming based on first-degree statistics depends on the recent history of stimulus/response identities (Stimuli and responses are confounded in this task so one cannot distinguish on behavioral grounds alone whether the stimulus or the response is producing the priming.) Priming based on second-degree statistics depends on the history of repetitions and alternations of the stimulus/response. Although the first and second degree sequences are in one-to-one correspondence, they can be dissociated in terms of the priming that they predict. For example, the sequence X then Y then X then Y, denoted XYXY, produces first-degree priming for Y because two recent trials have been Y, but strong second-degree priming for X because the stimulus has alternated on the three most recent trials. In fitting data from 2AFC tasks, priming based on first- and second-degree statistics both contribute roughly equally to explaining RT variability (Wilder, Jones, & Mozer, 2010).¹

Sequential dependencies suggest that what might appear to be intrinsic uncertainty in human behavior can instead be understood in terms of a systematic response to a changing world. Recent theoretical perspectives characterize sequential dependencies in terms of an

¹ Traditional analyses of sequential effects has been based a dichotomy between effects of trial $n-1$ (“first order effects”) and effects of trial $n-2$ and beyond (“higher order effects”). Even though these effects are referred to by “order”, the better terminology would have been “lag”: the effect of trial $n-k$ is a *lag k* effect. From the perspective of recent theoretical work (e.g., Jones, Curran, Mozer, & Wilder, 2009; Yu and Cohen, 2009; Wilder, Jones, & Mozer, 2010), analyses using the traditional dichotomy have less value than an analysis based on first- and second-order *statistics* of the sequence. A single decay-based priming mechanism can explain the *lag k* effect for any k . However, as we will discuss in this article, priming based on first- and second-order statistics yields distinct predictions. Because the field has used “order” in the past to refer to lags, we instead use the term “degree”—as in first- and second-degree statistics—to avoid confusing readers who are steeped in the literature.

individual's adaptation to the statistical structure of a nonstationary environment (Jones & Sieck, 2003; Mozer, Kinoshita, & Shettel, 2007; Wilder, Jones, & Mozer, 2010; Yu & Cohen, 2009). Nonstationarity implies that recent events provide a stronger predictor of the future than events further back in time, and therefore that sequential dependencies will continue to be—and are—observed even for overlearned and expert behavior.

Although reliable patterns of sequential dependencies have been demonstrated in controlled experimental tasks, it is not known how robust these effects are outside the laboratory. Typical laboratory 2AFC experiments involve trials that are neatly parceled into discrete events. In contrast, driving requires continuous, ongoing behavior, and graded responses that include pedal depression and steering adjustments. In 2AFC, the response-to-stimulus interval is very brief, typically on the order of 500-1000 ms. Recency effects may be short lived and not persist in more naturalistic environments where the events occur irregularly and less frequently. Finally, in 2AFC, participants are focused on a single task, quite different than a situation such as driving, where multiple goals and extraneous stimuli compete to control an individual's attention, and where multiple responses (steering, pedal depression) are required simultaneously.

Here, we examine sequential dependencies in a more complex, naturalistic, time critical task. In our experiment, participants navigate a vehicle in a driving simulator, and are occasionally required to brake or accelerate, indicated by a traffic-light style red or green cue on the windshield. We demonstrate the existence of sequential dependencies in both pedal-press latencies and errors.

Method

Participants

Eighteen volunteers participated in the experiment, 6 female and 12 male. The participants were in their 20s and 30s and varied in driving experience levels from beginner to decades of driving experience.

Apparatus

The experiment was controlled by a PC running an interactive open source simulator, TORCS (*TORCS, The open racing car simulator*, 2010), depicted in Figure 1. The ego-vehicle and road parameters such as friction and vehicle weight were fixed to approach real-world conditions. However they were also constrained in order to facilitate an easy learning process, as some subjects had never used driving simulators before. In particular, the maximum speed of the vehicle was limited, and the dynamics of the vehicle were set such that the car tended to travel at a slow, constant velocity when the pedal was not pressed. This was done in a similar fashion to “idle speeds” on real vehicles, by adjusting the engine RPMs (revolutions per minute). The max RPMs were adjusted so that the maximum speed was approximately 65 mph, and the “idle” (no-pedal) RPMs were set to correspond to about 5 mph.

A 52” screen was placed 3 feet in front of the subject, with a Logitech steering wheel mounted at a comfortable position near the driver, and accelerator and brake pedals placed accordingly. Mounted onto the toy steering wheel was a full-size steering wheel (15”) typical of those found in compact cars. The brake and accelerator pedals, positioned and sized similarly to realistic automobile configurations, registered continuous deflection.

All the driving parameters such as pedal position, lateral position of the vehicle in the road, and steering wheel data were automatically timestamped and logged to disk every 10 milliseconds. Additionally, a camera synchronously captured views of the driver's foot movements at 30 frames per second using an NTSC color camera.

Design

Each participant was run in two conditions: *simple* and *complex*. In the simple condition, the simulator was configured with a wide, straight track that required little steering effort to keep the car on the track. No other vehicles were present, and lane markings, background scenery, and sky features were eliminated. In the complex condition, the track was designed to be a constant S-curve, i.e., a series of switchbacks; the environment was populated with buildings, trees, and other scenery such as mountains in the distance; and the track consisted of three lanes, though no other vehicles were present.

Each condition was run as two continuous driving episodes, with a short break between episodes. The two episodes of the simple condition always preceded the two episodes of the complex condition, to allow the participants to get used to the simulator before increasing the challenge of the task. Trials were embedded within a driving episode and consisted of a cue (red or green disk) and the corresponding driver response (braking or accelerating, respectively). In the simple condition, the interval from response to the next stimulus cue (RSI) was two seconds. In the complex condition, the RSI varied from one to four seconds, drawn from a uniform distribution.

Each driving episode consisted of seven practice trials followed by 128 experimental trials, for a total of 135 trials per episode. Experimental trials were counterbalanced such that all stimulus sequences of length five (e.g., red, green, red, red, green) occurred exactly once per $2^5=32$ trials, resulting in four presentations of each five-stimulus sequence per episode (eight per condition). In the analyses we present, we consider only the influence of the preceding three trials, $n-3$, $n-2$, and $n-1$, on the current trial, n . However, counterbalancing one additional trial back only helps address the possibility that influences of past trials extend further back. The entire experiment took approximately 25 minutes per participant.

Procedure

Participants were told that the experiment examines driver behavior in stop-and-go traffic. They were instructed to respond to a series red and green cues which appear as circular traffic-light style disks in the central field of view; red and green cues required participants to “pump” the brake and accelerator, respectively, with their right foot. Participants were told not to hold their foot on the brake or accelerator position and that the default foot position should be a neutral one in which neither pedal was depressed. Participants were allowed to hover their driving foot over either pedal. Most hovered over the accelerator, as is typical in naturalistic driving, but some hovered over the last pedal that was depressed. Although this hovering requirement is not entirely naturalistic, it is reminiscent of stop-and-go driving and participants had no difficulty following the instruction.

The cues remained visible until a response pedal was pressed. If neither response pedal was pressed within one second of stimulus onset, the trial was terminated by removing the cue and the next stimulus appeared following the RSI.

Data Analysis

A pedal press was defined to occur at the time when either of the pedals was depressed 20% or more of its maximum. Several participants tended to braise one pedal briefly along a trajectory to pressing the other pedal, and we determined criteria—consistent with video observations—that considered the second pedal press as the response.

The video data coming from the foot camera was also analyzed to detect the time of the initial foot movement, and to allow us to decompose the pedal press latency into a foot movement initiation time and a time to move the foot to the pedal. The video processing algorithm used a moving pixel-wise median background to distinguish novel changes in the current frame. Due to the camera perspective, these changes were caused mostly by foot

motion; thus any major foot motion could be detected by placing a threshold on the amount of motion in the frame. Once a foot movement event was detected in this manner, the foot movement signal was traced back to the initial inflection point of the motion, to determine the time of the beginning of the foot motion.

Results

Each trial is characterized by three independent variables: (1) the driving environment—simple versus complex, (2) the target response—accelerate versus brake, and (3) the context arising from the three most recent trials. The context is specified in terms of the sequence leading up to the current trial. If we adopt the shorthand notation of X for the {red, brake} stimulus-response pair, and Y for the {green, accelerate} pair, we can represent a sequence like green-green-red-red as YYXX, where the order of trials, left to right, is $n-3$, $n-2$, $n-1$, and n . Because the labels X and Y represent both stimulus and response and the mapping, we are neutral as to whether the effects of context that we describe are due to priming of stimuli, responses, or both.

The context can also be described in terms of the pattern of alternation (A) and repetitions (R) by considering pairs of trials. For example, the trial sequence green-green-red-red, which corresponds to the X-Y sequence YYXX, also corresponds to the repetition/alternation sequence RAR, where the first R denotes that stimuli on trials $n-3$ and $n-2$ are the same, the A denotes that stimuli on trials $n-2$ and $n-1$ are different, and the final R denotes that stimuli on trials $n-1$ and n are the same. We refer to these two representations of context as *first* and *second degree* because the X/Y representation describes individual trials, whereas the A/R representation describes pairs of trials. In laboratory studies of sequential dependencies in 2AFC, the second degree representation of context is

standard. We will present results using both representations, and it should shortly become clear why we include the first-degree representation.

Participants made relatively few response errors. We discuss errors later, but first present an analysis of response times (RTs). For the RT analysis, trial n is discarded if an error has occurred on any of trials from $n-3$ to n , or if the RT is not within three standard deviations of the individual participant's mean RT.

A repeated measures ANOVA with subjects as a random factor yields a main effect of driving environment: RTs are slower in the complex environment than in the simple (654 ms versus 612 ms, $F(1, 17)=27.8$; $p<.001$), expected given the additional distractors and variability present in the complex environment. Participants are slower to brake than to accelerate (659 ms vs. 607 ms, $F(1, 17)=6.97$; $p<.02$), not a big surprise because in normal driving the foot tends to hover over the accelerator pedal. Finally, there is a main effect of context ($F(7, 119)=26.99$; $p<.001$). However, there are no higher order interactions involving any of the independent variables (all $p>.275$). Consequently, we are justified in collapsing across driving environment and response pedal in the response time analyses to follow. Further, because we observe only a main effect of RT on the particular stimulus-response pair identity (i.e., braking versus accelerating), we collapse across complementary sequences, e.g., YYXX and XXYY, and as shorthand, we refer to both sequences using the label that ends in X, here YYXX.

Sequential Dependencies on Response Time

The top row of Figure 2 shows the mean RTs across participants for each three-back context. The left and right graphs encode the context in the first and second degree representations, respectively. Because the two representations are in one-to-one correspondence, the eight data points are the same in the two graphs. However, the data

points are ordered differently. In the first-degree graph, the points are ordered from left to right with recent trials that are most to least similar to the current trial in terms of the stimulus (or response) identity. If RTs are sensitive to priming from recent trials and the priming consists of a decaying trace of stimulus/response identities, the first-degree graph should show a generally increasing trend in the RTs from left to right, because the strongest priming should be obtained for contexts in which recent trials have the same identities as the current trial. In the second-degree graph, the points are ordered from contexts with many repetitions to contexts with many alternations. If RTs are sensitive to priming from recent trials and the priming consists of a decaying trace of the repetition/alternation pattern of trials, the second-degree graph should show a teepee shape, with fastest responses for consistent runs of repetitions or alternations (because recent repetitions prime a repetition, recent alternations prime an alternation), and slowest responses when current trial is a repetition (or alternation) and recent trials have been alternations (or repetitions). By inspection, the first-degree graph appears to show a more meaningful pattern, so we begin with an analysis of the RTs based on the first-degree contexts (left graph of Figure 2), and we return shortly to the second-degree graph.

First-degree context has a robust effect on RT. We can compare each pair of contexts, with a Bonferroni correction for multiple comparisons, yielding several groups of contexts which are significantly different from other groups ($F(7,17)=27.0, p<0.0036$). Grouping the contexts, we have {1:[XXXX], 2:[YXXX], 3:[XYXX], 4:[YYXX], 5:[XXYX, YXYX, XYYX], 6:[YYYX]}. Within these, group 1 is different than 3, 4, 5, and 6; group 2 different than 3, 5, and 6; and group 4 different than 6. These trends are apparent from Figure 2.

To remove the variability due to differences in mean RT and sequential effect magnitudes across participants, we transformed the eight context-specific RTs for each

participant to z-scores, shown in the bottom left graph of Figure 2. The z-score data show an even more reliable dependency on the first-degree context than do the raw data

($F(7,17)=32.3, p<0.0001$). Grouping the contexts, we have {1:[XXXX], 2:[YXXX], 3:[XYXX], 4:[YYXX], 5:[XXYX], 6:[YXYX,XYXX,YYXX]}. Within these, groups 1 and 2 are different than 3, 4, 5, and 6; group 3 different than 6; and group 4 different than 5 and 6.

Decomposing the Response Time

Using automated video-based analysis of the camera aimed at the participant's foot, the overall response time (ORT) can be decomposed into two stages: the time from stimulus onset to initiation of foot movement (FMI), and the time to move from initiation to pedal press (TTM). Trials for which the analysis technique failed (7.87% of the overall trials) were discarded. These cases were caused either by foot motion which began before the stimulus appeared (causing ambiguity in determining response initiation), or by the lack of significant motion prior to the pedal press.

Figure 3 shows sequential effects indexed by first-degree sequence for the two components of ORT. Systematic variation is observed in TTM but not in FMI. Indeed, the sequential dependencies in TTM appear to be as robust as those in ORT. Pairwise comparisons across the different contexts, with a Bonferroni correction for multiple comparisons, yielded several groups of contexts which were significantly different from other groups ($F(7,17)=23.7, p<0.0036$). Grouping the contexts, we have {1:[XXXX], 2:[YXXX], 3:[XYXX,YYXX], 4:[XXYX], 5:[YXYX,XYXX,YYXX]}. Within these, group 1 is different than 3, 4, and 5; group 2 different than 4 and 5; and group 3 different than 5. These trends are apparent in Figure 3. For example, contrasting the extreme contexts of XXXX and YYYX, the TTM difference is 97 ms, comparable to the difference in ORT.

These strong patterns are all the more impressive considering that the quantization of the video stream is relatively coarse; our timing accuracy is limited by the video frame rate, 33 ms. One potential explanation for why sequential dependencies are observed in TTM is that our experiment did not control for foot position at the start of a trial, and if participants rested their foot over the pedal that was last depressed, then one would expect TTM to be faster for an XX sequence than a YX sequence. Although foot position might explain why sequential effects are observed in TTM at lag 1, this account would not explain effects beyond lag 1, which are found in the data.

Whereas FMI is not susceptible to sequential effects, it is affected by the interval from the preceding response to the current stimulus (RSI). In the *complex* condition, we divided trials into shorter and longer RSIs, with the dividing line being the median RSI of 2.52 seconds. In other words, we labeled those trials whose RSI was between 1 – 2.52 seconds as “short”, and those between 2.52 – 4 seconds as “long”. FMI was reliably faster following short RSIs (FMI of 343.8 and 365.7 ms for short and long RSI, respectively; $F(1, 17)=14.55$; $p<0.01$). However, RSI had no systematic effect on TTM (298.6 and 301.1 ms for short and long RSI; $F(1, 17)=0.28$; $p>0.25$). Thus a stimulus readiness effect is observed in the early stages of processing that determine FMI, whereas sequential dependencies are observed in later stages reflected in the TTM.

Sequential Dependencies Based on First- And Second-Degree Statistics

To investigate the relative contribution of first- and second-degree statistics to priming, the data were fit to models that assume that priming from recent trials decays exponentially, and is based on first-degree statistics, second-degree statistics, or both.

We model the probability of the stimulus on trial $n+1$, X_{n+1} , based on first-degree statistics as

$$P_1(X_{n+1} = x) = c_1 \sum_{k=0}^3 \lambda_1^k \cdot \mathbf{I}[X_{n-k} = x];$$

with λ_1 representing the first-degree decay parameter and $c_1 = (1 - \lambda_1)/(1 - \lambda_1^4)$ serving to scale the weights λ_1^k such that their sum is 1. $\mathbf{I}[\cdot]$ represents the indicator function and takes on the value 1 if the statement in the brackets is true. The probability of the stimulus on trial $n+1$ based on second-degree statistics is given by

$$P_2(X_{n+1} = x) = \mathbf{I}[X_n = x]P_R + (1 - \mathbf{I}[X_n = x])(1 - P_R);$$

$$\text{with } P_R = P(X_{n+1} = X_n) = c_2 \sum_{k=0}^2 \lambda_2^k \cdot \mathbf{I}[X_{n-k} = X_{n-k-1}].$$

Here λ_2 is the second-degree decay parameter and $c_2 = (1 - \lambda_2)/(1 - \lambda_2^3)$ is again a normalizing term. The probability for the combined model is obtained by taking a weighted combination of the first- and second-degree probabilities,

$$P_c(X_{n+1} = x) = wP_1(X_{n+1} = x) + (1 - w)P_2(X_{n+1} = x).$$

The pure models have one free parameter – the memory decay rate (λ); the combined model has three – two decay parameters (λ_1 and λ_2) and one additional parameter w for a mixture weight (See Wilder, Jones, & Mozer, 2010, for details of these models.) Parameters for each model are obtained by minimizing the squared prediction error.

Table 1 reports the percent of variance explained by the three different models for the three different RT measures – ORT, FMI, and TTM. As expected, no model does a good job of fitting the FMI, which has no systematic relationship with the recent history. The first-degree model does a good job of fitting the patterns of ORT and TTM, consistent with what one observes in Figures 2 and 3.

Although the second-degree model obtains a poor fit, priming based on second-degree statistics may be present but dominated by first-degree statistics, as suggested by the fact that the combined model shows an improvement in fit over the first-degree model for all RT measures. However, the improvement could be due only to the combined model's additional degrees of freedom. Comparing the first-degree and combined models, a likelihood-ratio test fails to reject the null hypothesis that the two models obtain a comparable fit (for all RT measures, $p > 0.25$). Thus, we conclude that a model of first-degree priming (i.e., priming of the stimulus or response identities) with exponential memory decay does as well as any account in explaining both the ORT and TTM data. Figure 4 shows the model fit to ORT data. The model has one free parameter, and requires an affine transform—two additional parameters—to predict RTs, although this affine transform is irrelevant to the percent of variance explained by the model.

Another way of assessing the magnitude of sequential effects is to construct a *lag profile* that shows the influence of the n th trial back on the current trial. To construct a lag profile based on first-degree statistics, the RT data is partitioned by whether the current trial is the same or different as the n th-back trial. The difference in mean RT between these two conditions reflects the extent of priming. The upper graph of Figure 5 shows the first-degree lag profile for the three response time measures. For example, the upper graph in Figure 5 shows that at lag 1, overall RT is on average about 60ms faster when the current trial is the same as the previous trial, compared to when it is different than the previous trial. Because the values in a lag profile collapse over more trials and contexts than measures based on context-specific history, the lag profile offers the opportunity to detect small but statistically reliable effects.

The lag profile confirms that for ORT and TTM, priming based on first-degree statistics are statistically reliable at lags 1, 2, and 4 ($p < 0.01$ by one-tailed t-test), though not at lag 3

(ORT lag 3 $t(17)=1.20$; $p=0.248$; TTM lag 3 $t(17)=1.79$; $p=0.092$). There is no significant first-degree priming of FMI ($p>0.10$ for all lags).

In the lag profile based on second-degree statistics (lower graph of Figure 5), each trial is characterized as a repetition or alternation of the previous, and the lag n priming is the speed up that occurs when the current trial matches the n th trial back. For example, the lower graph in Figure 5 shows that at lag 2, ORT is about 5 ms faster when the current trial n and trial $n-2$ are both repetitions or both alternations, compared to when trial n is an R and $n-2$ an A or vice versa. Second-degree priming at lag 1 is identical to first-degree priming at lag 2. Consequently, evidence for effects based on second-degree statistics comes from reliable priming at lags 2 or 3. Second degree priming at lags 2 and 3 is statistically reliable for ORT (lag 2: $t(17)=3.73$; $p<0.01$; lag 3: $t(17)=2.70$; $p=0.015$), albeit small in magnitude. For TTM, second-degree priming is not statistically reliable (lag 2: $t(17)=1.27$; $p=0.221$; lag 3: $t(17)=1.83$; $p=0.086$). Thus, there is evidence for second-degree priming in the lag profile, but it has a small absolute effect on RTs.

Pedal Misapplications

Sequential effects typically arise in error rates as well as in latencies (Cho et al., 2002), with the slower responses being more error prone, indicating that sequential effects are not merely a speed-accuracy trade off.

Trials in which the participant initially fully pressed the wrong pedal were rare in the experiment, occurring on only 2.33% of trials. In nearly every case, participants detected their error and corrected it afterward (e.g., Figure 6). Due to the small number of errors, instead of comparing error patterns in three-back contexts, we examined errors based only on the previous trial context.

Pedal misapplications occurred on 3.30% of trials when the previous stimulus differed from the current, but only on 1.46% when the previous stimulus was the same as the current (reporting a subjectwise median). The difference is significant by a Wilcoxon signed-rank test ($Z=2.46$; $p<0.05$). On 0.39% of trials, participants made no response. When these no-response trials are considered to be misapplications, the effect of the previous trial is still evident ($Z=2.16$; $p<0.05$).

Discussion

In a simulated driving task, we observe a robust effect of the recent sequence of pedal presses on both response latencies and error rates. The task is perhaps the most realistic, naturalistic scenario in which recency effects have been explored. We obtained evidence for effects well beyond the immediately preceding trial, extending up to four trials back. The effects are primarily what we characterize as first degree, corresponding to the traditional conception of priming of the response (or stimulus, as the two are confounded). In contrast to laboratory 2AFC tasks, second degree effects—effects of the repetition and alternation sequence—are minimal in magnitude.

Our experiment extends the canonical laboratory 2AFC task toward a more naturalistic scenario in three respects. First, the visual environment is cluttered, heterogeneous, constantly changing, and contains elements that distract from the pedal pressing task, such as buildings, trees, mountains, and grass. Second, driving requires multitasking: in addition to braking and accelerating as demanded by the pedal pressing task, the driver needs to constantly monitor and adjust steering to remain in the center lane of a regularly curved road. Third, the response-to-stimulus intervals are variable, ranging from one to four seconds, and are larger than the typical interstimulus intervals of less than a second.

Even the simple condition in our experiment—which offered little in the way of distractor stimuli and required minimal multitasking—moves in the direction of ecological validity over the canonical 2AFC task. However, the finding that sequential effects are robust and comparable in the complex condition supports the notion that environmental and task complexity—via additional variability of the stimulus onset, additional visual distractors, and additional difficulty of the secondary (driving) task—do not modulate the nature of sequential effects.

Although it's not altogether unexpected to find priming in a complex, multitasking environment, a further analysis of behavior offers an intriguing suggestion concerning the stage of processing at which priming occurs. We decomposed the overall response time into the time from stimulus onset to foot movement initiation (FMI), and the time from movement initiation to pedal press (TTM). We find that only TTM is affected by the trial sequence, and only FMI is affected by the response-to-stimulus interval, implying that sequential dependencies occur at a distinct stage of processing from stimulus readiness.

Our motivation for the RT decomposition is a related decomposition that has been performed via event-related potentials to analyze sequential dependencies. In a 2AFC task with the two responses mapped to the two hands, the start of response selection can be inferred as the onset of the lateralized readiness potential (LRP). Jentzsch and Sommer (2002) and Jones, Curran, Mozer, and Wilder (2009) have analyzed sequential effects on RT, decomposing the overall RT into stimulus processing (pre-LRP) and response processing (post-LRP) stages. Wilder, Jones, and Mozer (2010) have shown that the stimulus processing component of RT is well fit by a second-degree model of sequential dependencies, and the response processing component is well fit by a first-degree model. That is, sequential dependencies are based on frequency of repetitions and alternations of the *stimuli*, and on the frequency of the two alternative *response* identities.

The presence of priming based on first-degree statistics in the TTM in the driving task is predicted by the dual priming theory (Wilder, Jones, and Mozer, 2010), because TTM reflects a late stage of processing. However, on its face, the theory predicts second-degree dependencies at an early stage of processing, which are not observed in the FMI. It remains to be determined whether the absence of second-degree priming is a failure of the theory or is attributable to noise in the data or another factor that differentiates the driving task from canonical 2AFC tasks, e.g., the complexity of the visual environment, multitasking, or the asymmetry of pedal presses.

What are the practical implications of recency effects in driving? Consistent with previous studies using simple 2AFC tasks, we find that recent events can result in response delays of nearly 100 ms (e.g., for the most extreme contrast—XXXX versus YYYY—the difference in ORT is 91 ms, and the difference in TTM is 97 ms). Because sequential dependencies are found in complex laboratory conditions that approximate real-world driving, it seems altogether plausible that they will be observed in situations such as stop-and-go traffic, or a series of traffic lights in an urban environment.

At 20mph, a delay of up to 100 ms corresponds to an increased stopping distance of nearly one meter, which could be the difference between running into a leading vehicle and stopping safely. At 65mph this 100 ms difference corresponds to an extra three meters in stopping distance. Even smaller differences in response times have been shown to have severe consequences in both the probability and severity of vehicle crashes (Evans, 1991).

Beyond delays in responding, we have shown that recent events can affect the likelihood of a pedal misapplication. In our study, the probability of a pedal misapplication more than doubled if the target pedal differs from the previous pedal versus when the two are the same.

Many recent unintended-acceleration-related accidents in the U.S. may have been a result of pedal misapplications (Healey & Carty, 2010). The context-dependent pattern of driver response delays and errors could potentially be exploited in a holistic Advanced Driver Assistance System, or ADAS (Doshi, Cheng, & Trivedi, 2009; Trivedi, Gandhi, & McCall, 2007). Given a particular observed history of pedal presses and familiarity with a particular driver, the ADAS could predict real-time performance. These predictions could be used to provide additional or earlier alerts to drivers in situations where delays/errors were likely. In more critical circumstances, the vehicle could take action to reduce impending accident severity by priming the brake system. In less critical situations, if the upcoming context is ripe for higher response times (e.g., more than 600 ms), an urgent alert could help to refocus the driver to the potentially appropriate response. An adaptive intelligent ADAS design could thus utilize sequential context to help the driver and mitigate dangerous or uncomfortable circumstances.

References

- Anderson, N. (1960). Effect of first-order conditional probability in a two-choice learning situation. *Journal of Experimental Psychology*, 59(2), 73-93.
- Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General*, 129, 177-192.
- Chang, A. (2010, August 24). *China's massive traffic jam could last for weeks*. Associated Press. Available from http://hosted.ap.org/dynamic/stories/A/AS_CHINA_TRAFFIC_JAM
- Cho, R. Y., Nystrom, L. E., Brown, E. T., Jones, A. D., Braver, T. S., Holmes, P. J., et al. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive Affective Behavior Neuroscience*, 2(4), 283-299.
- Dixon, P., & Glover, S. (2004). Action and memory. In B. Ross (Ed.), *The psychology of learning and motivation* (p. 143-174). San Diego, CA: Elsevier.
- Dixon, P., & Glover, S. (2009, May). Perseveration and contrast effects in grasping. *Neuropsychologia*, 47(6), 1578-1584.
- Doshi, A., Cheng, S. Y., & Trivedi, M. M. (2009). A novel, active heads-up display for driver assistance. *IEEE Transactions on Systems, Man, and Cybernetics Part B*, 39(1), 85-93.
- Estes, W. (1950). Toward a statistical theory of learning. *Psychological Review*, 57, 94-107.
- Evans, L. (1991). *Traffic safety and the driver*. New York: Van Nostrand Reinhold.
- Healey, J. R., & Carty, S. S. (2010, July 14). Driver error found in some Toyota acceleration cases. *USA Today*. Available from http://www.usatoday.com/money/autos/2010-07-14-toyota14_ST_N.htm
- Jentzsch, I., & Sommer, W. (2002). Functional localization and mechanisms of sequential effects in serial reaction time tasks. *Perception and Psychophysics*, 64(7), 1169-1188.

- Jesteadt, W., Luce, R. D., & Green, D. M. (1977). Sequential effects in judgments of loudness. *Journal of Experimental Psychology Human Perception and Performance*, 3(1), 92-104.
- Jones, M., Curran, T., Mozer, M., & Wilder, M. H. (2009, November). Sequential effects reflect learning of temporal structure. *Presented at the 50th Annual Meeting of the Psychonomic Society*.
- Jones, M., & Sieck, W. R. Learning myopia: An adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 29, 626-640.
- Kristjansson, A. (2006). Rapid learning in attention shifts: A review. *Visual Cognition*, 13(3), 324-362.
- Maloney, L. T., Dal Martello, M. F., Sahm, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences*, 102, 3164-3169.
- Mozer, M. C., Kinoshita, S., & Shettel, M. (2007). Sequential dependencies offer insight into cognitive control. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 180-193). Oxford: Oxford University Press.
- NHTSA. (2010). National Highway Traffic Safety Administration - Fatal Analysis Reporting System (FARS). Available from www-fars.nhtsa.dot.gov
- Remington, R. (1969). Analysis of sequential effects in choice reaction times. *Journal of Experimental Psychology*, 82(2), 250-257.
- TORCS, the open racing car simulator*. (2010). Available from <http://torcs.sourceforge.net>

Trivedi, M. M., Gandhi, T., & McCall, J. C. (2007, January). Looking in and looking out of a vehicle: Computer vision-based enhanced vehicle safety. *IEEE Transactions on Intelligent Transportation Systems*, 8, 108-120.

Wilder, M., Ahmed, A., Mozer, M., & Jones, M. (2010). Sequential effects in motor adaptation: The importance of far back trials. *Poster presentation at Society for Neuroscience*. San Diego, CA, November 15, 2010.

Wilder, M., Jones, M., & Mozer, M. C. (2010). Sequential effects reflect parallel learning of multiple environmental regularities. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, & A. Culotta (Eds.), *Advances in Neural Information Processing Systems* (Vol. 22, pp. 2053-2061). La Jolla, CA: NIPS Foundation.

Yu, A. J., & Cohen, J. D. (2009). Sequential effects: Superstition or rational behavior? *Advances in Neural Information Processing Systems*, 21, 1873-1880.

Author Note

This research was supported by a University of California Transportation Center PhD Dissertation Fellowship, UC-Discovery, and NSF grants BCS-0339103, BCS-720375, and SBE-0518699.

Table 1

Percent Variance Explained by Three Models for Three Different RT Measures

<i>RT Measure</i>	<i>Model</i>		
	1st Degree	2nd Degree	1st + 2nd Degree
FMI	9.81	8.23	12.61
TTM	94.32	17.56	95.32
ORT	87.91	16.97	89.68

Figure Captions

Figure 1. Simulator layout and experimental setup. The screen is shows a typical driver view in the complex condition; the left inset shows a typical view in the simple condition. The right inset shows a sample frame of the camera used for initial foot motion detection.

Figure 2. Sequential Effects in Pedal Press Response Times (top) and Z-scores (bottom). The left side show effects with respect to the First-degree stimulus history, and the right side shows the same data with respect to the Second-degree history (repetitions/alternations). Error bars indicate one standard error of the mean.

Figure 3. Sequential effects on response time, broken into foot movement initiation time (top) and time to move to pedal (bottom).

Figure 4. Mean overall reaction time (ORT) as a function of trial context (circles) and fit of first-order model with exponentially decaying memory (diamonds)

Figure 5. Lag plots of first- and second-degree sequential effects on overall RT (ORT), foot movement initiation Time (FMI) and time to move to pedal (TTM). The error bars represent +/-1 standard error of the mean. For example, for the first degree graph (top), at lag 1, ORT is on average about 60ms faster when the current trial is the same as the previous trial, compared to when it is different than the previous trial. At lag 2, ORT is about 30ms faster when the current trial n is the same as trial n-2, compared to when the current trial is different than trial n-2. For the second degree graph (bottom), at lag 1, ORT is about 30ms faster when current trial n and trial n-1 are both repetitions (or both alternations), compared to when trial n is a repetition and trial n-1 is an alternation (or vice versa).

Figure 6. Sample series in the complex condition of the experiment, showing an example of pedal misapplication. After a number of accelerate cues in the recent history, at the 88-second mark, the driver sees a brake cue but accidentally steps on the accelerator instead.

To Appear, Cognitive Science, 2011



Figure 1. Simulator layout and experimental setup. The screen is shows a typical driver view in the complex condition; the left inset shows a typical view in the simple condition. The right inset shows a sample frame of the camera used for initial foot motion detection.

To Appear, Cognitive Science, 2011

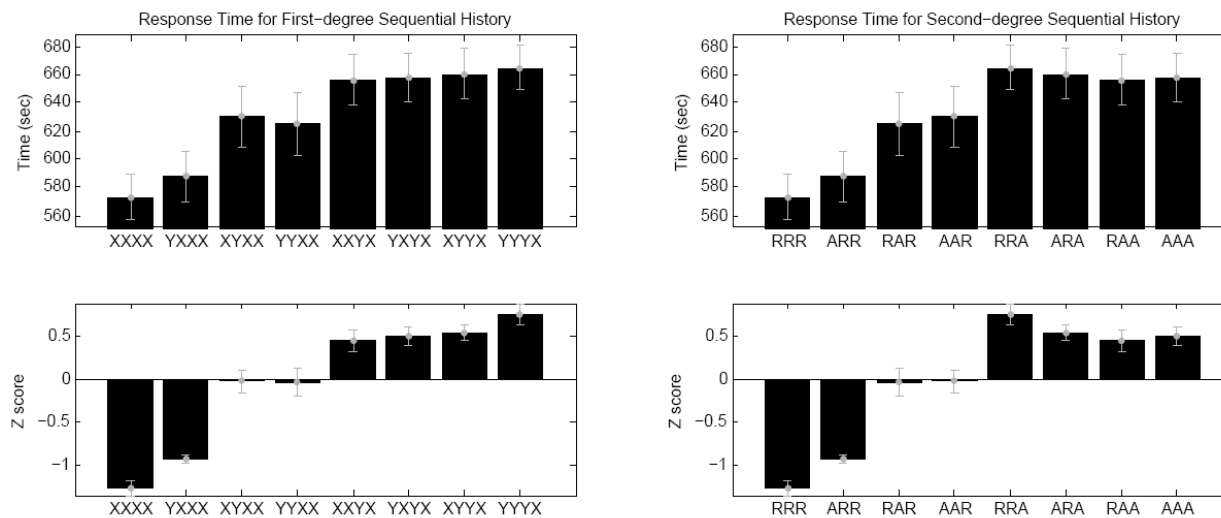


Figure 2. Sequential Effects in Pedal Press Response Times (top) and Z-scores (bottom). The left side show effects with respect to the First-degree stimulus history, and the right side shows the same data with respect to the Second-degree history (repetitions/alternations). Error bars indicate one standard error of the mean.

To Appear, Cognitive Science, 2011

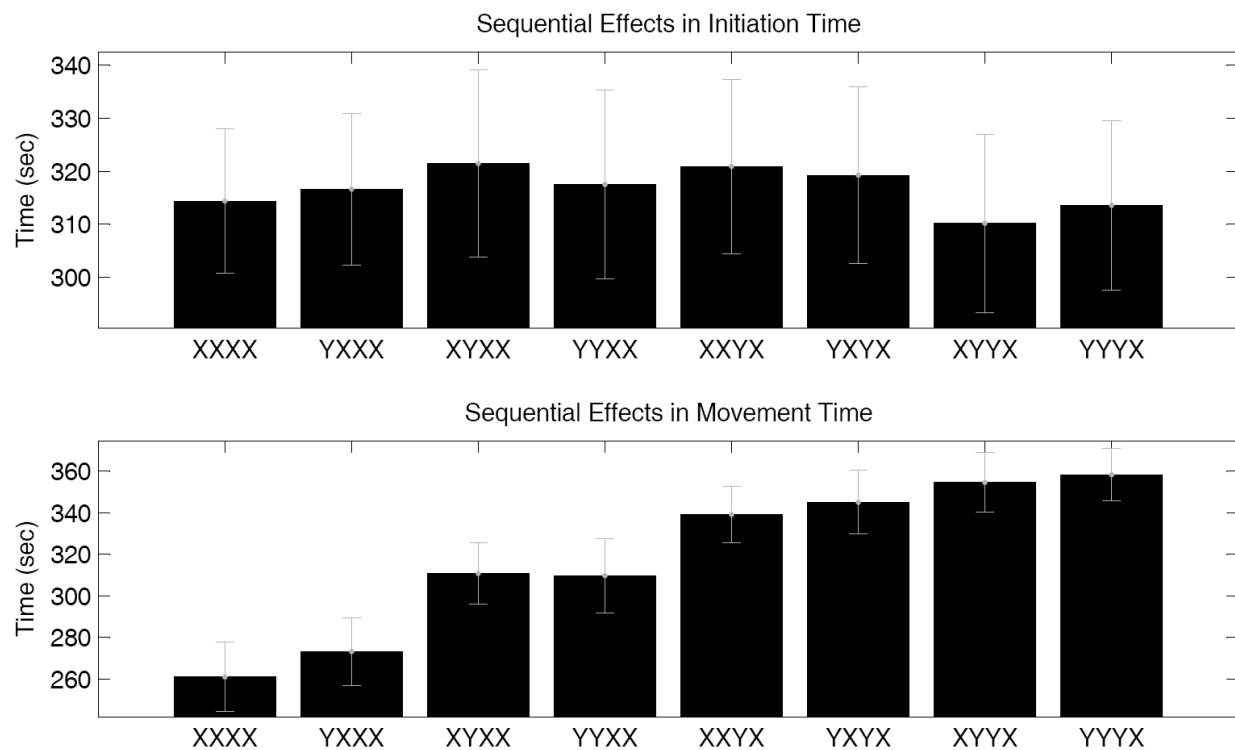


Figure 3. Sequential effects on response time, broken into foot movement initiation time (top) and time to move to pedal (bottom).

To Appear, Cognitive Science, 2011




Figure 4. Mean overall reaction time (ORT) as a function of trial context (circles) and fit of first-order model with exponentially decaying memory (diamonds)

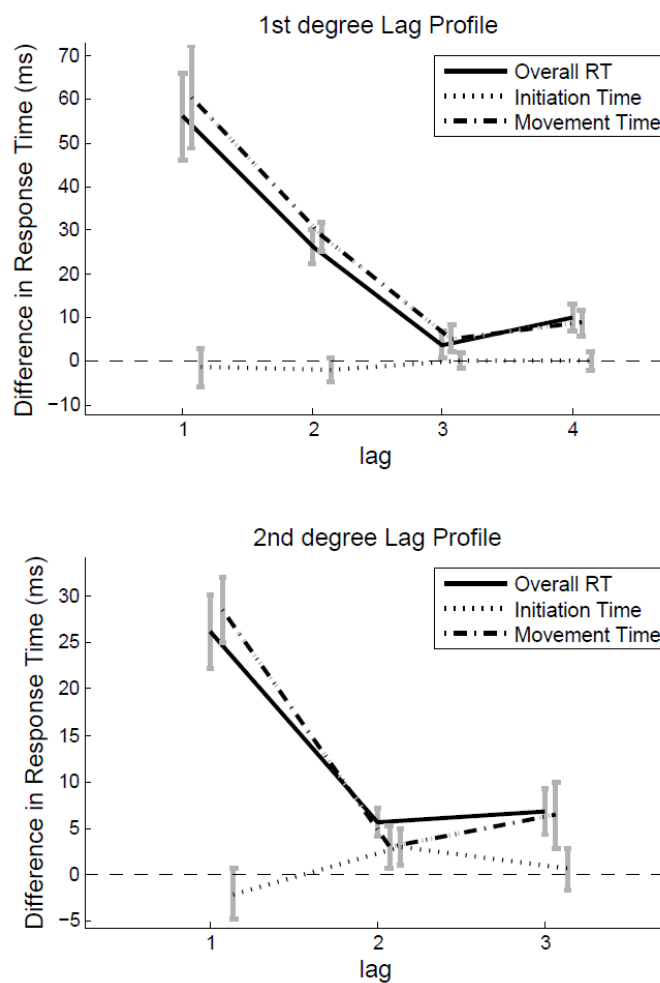


Figure 5. Lag plots of first- and second-degree sequential effects on overall RT (ORT), foot movement initiation Time (FMI) and time to move to pedal (TTM). The error bars represent ± 1 standard error of the mean. For example, for the first degree graph (top), at lag 1, ORT is on average about 60ms faster when the current trial is the same as the previous trial, compared to when it is different than the previous trial. At lag 2, ORT is about 30ms faster when the current trial n is the same as trial $n-2$, compared to when the current trial is different than trial $n-2$. For the second degree graph (bottom), at lag 1, ORT is about 30ms faster when current trial n and trial $n-1$ are *both* repetitions (or both alternations), compared to when trial n is a repetition and trial $n-1$ is an alternation (or vice versa).

To Appear, Cognitive Science, 2011




Figure 6. Sample series in the complex condition of the experiment, showing an example of pedal misapplication. After a number of accelerate cues in the recent history, at the 88-second mark, the driver sees a brake cue but accidentally steps on the accelerator instead.