Salvucci, D. D., Liu, A., & Boer, E. R. (2001). Control and monitoring during lane changes. To appear in *Vision in Vehicles 9*.

CONTROL AND MONITORING DURING LANE CHANGES [*]

Dario D. Salvucci, Drexel University (salvucci@mcs.drexel.edu)
Andrew Liu, MIT Man Vehicle Laboratory (andy@space.mit.edu)
Erwin R. Boer, Erwin R. Boer Consulting (BorromeoRing@aol.com)

This work examines the visual scanning behavior of drivers in preparing and executing a lane change. In an experiment, drivers navigated a naturalistic highway environment with moderate traffic in a fixed-base driving simulator. The data collected in the experiment included steering control data, eye movements, and verbal protocols. Drivers began to exhibit significantly different scanning behavior approximately three seconds before initiation of the lane change, showing increasingly more frequent gazes at the rear-view mirror at the expense of gazes to their current lane. As soon as they decided to make a lane change (as indicated in their verbal protocols), drivers shifted their gaze from salient guiding features of the current lane (e.g., tangent point or lead vehicle) to salient guiding features of the destination lane. In addition, drivers exhibited increased gazes at surrounding vehicles (front and back) before and during lane changes for the purposes of situation awareness and decision making. The results support a dual-purpose view of driver gazes for control and monitoring.

1. INTRODUCTION

In the study of driving and related tasks such as locomotion, researchers have extensively utilized eye movements, or gaze, as a window into how humans execute these tasks. In particular, researchers have studied driver gaze in the context of two critical aspects of the driving task: *control*, or the process necessary to guide steering and speed control; and *monitoring*, or the process necessary to maintain awareness of the external environment as needed for decision making. For control, numerous studies have examined how drivers steer within a single lane and, in particular, what visual cues guide steering (e.g., Land & Lee, 1994). For monitoring, studies have explored drivers' visual scanning strategies in finding and noting important landmarks or potential hazards (e.g., Gipps, 1986).

This paper extends this work by examining driver gaze in a critical component of the driving task: lane changing. Lane changing is an extremely common driver subtask, particularly during highway driving (which accounts for approximately 70% of vehicle miles in the United States (Federal Highway Administration, 1998)), and makes

[*] Please address all correspondence to Dario Salvucci, Dept. of Mathematics and Computer Science, Drexel University, 3141 Chestnut St., Philadelphia, PA 19104, USA

demands on both control and monitoring for successful execution. Surprisingly, relatively little work has been done on lane-changing behavior given its omnipresence in everyday driving. This work has focused primarily on the monitoring and decision-making aspects of lane changing, namely the process of determining when to execute the maneuver (Ahmed et al., 1996; Gipps, 1986). One study has examined the time needed to make a lane change (Finnegan & Green, 1990), and these studies have facilitated the development of real-world systems for lane-change collision avoidance (Talmadge et al., 2000). However, there has been as yet no detailed examination of lane-changing behavior that integrates analyses of driver gaze for the purposes of both vehicle control and monitoring in service of decision making.

This paper describes an empirical study in which we asked drivers to navigate a four-lane highway in a simulated environment with moderate traffic. Using a fixed-base driving simulator, we collected several types of data during the task: standard control data such as steering and acceleration, eye-movement data as a surrogate for drivers' focus of (visual) attention, and verbal protocol data as an indicator of when drivers formed the intentions to perform various actions while driving. Using these data we attempt to build a detailed picture of driver control and monitoring in a lane change — namely, how they prepare for a lane change, what they do during a lane change, and what guides their steering through the lane change.

2. METHOD

In the experiment, participants drove in a simulated highway environment. We were interested primarily in drivers' gaze (i.e., eye movements) and the time-course correspondence of gaze with the execution of lane-change maneuvers. To this end, we constructed a minimalist driving environment in the sense that it included only essential elements — namely, the roadway and other vehicles — with no extraneous scenery. This simplicity reduced the complexity of the downstream eye-movement data analysis while retaining the basic characteristics of the highway driving situation.

2.1. Subjects

Fifteen participants were enlisted for the experiment. Of these participants, two were unable to complete the experiment because they could not successfully steer the car in the driving simulator, and another two were excluded because of severe difficulty in tracking their eye movements due to contact lenses. The remaining eleven participants (two women and nine men) were used in the following analyses. Thee participants ranged in age from 18-31 years and had at least two years of driving experience.

2.2. Environment

The simulated highway environment had a total of four lanes with two lanes in each direction (driving on the right side as the American norm). It included no on-ramps or off-ramps and used the standard American lane markings. The highway contained no off-road scenery except for a side wall 2.75 m off the road. Each lane had a width of 4 m while all simulated vehicles (including the driver's vehicle) had a width of 2 m.

With respect to other vehicles, the environment included either no traffic or moderate traffic depending on the condition (discussed shortly). In moderate traffic condition, the environment maintained a constant of ten automated vehicles within a window of 400 m in front or in back of the driver's car — seven cars in the driver's direction and three cars in the opposite direction; when an automated vehicle exceeded this window, it was placed on the opposite side of the window in the same direction and lane. The automated vehicles were assigned to drive at a randomly-chosen desired speed in the range 22.35 – 31.29 m/s (50 – 70 mph) and passed slower vehicles when necessary, avoiding other vehicles in the process. This simple automation provided a natural feel to the environment and fairly realistic traffic patterns.

This environment was integrated into the Nissan Cambridge Basic Research driving simulator (Beusmans & Rensink, 1995). This simulator uses a fixed-base platform with the front half of a Nissan 240SX convertible. The highway environment is projected in front of this platform onto a flat wall resulting in a 70° field of view. A simulated rear-view mirror is inset into the environment display and placed in a position approximately similar to the actual rear-view mirror; we found that subjects could use this simulated mirror quite naturally and effectively. The control data generated by the simulator for the purposes of this study included steering, accelerator, brake, and turn signal data. These data were fed into the simulator computer, which updated the environment according to a vehicle model and provided additional feedback to the driver in terms of steering wheel torque and engine sounds. The simulator uses a three-wheel dynamic model for vehicle dynamics.

Eye movements were collected in the driving simulator using an integrated IScan (Burlington, MA) 60 Hz head-mounted eye tracker attached to a baseball cap. After a short calibration procedure, the tracker computed the pupil center and the reflection point for a small infrared light and estimated the on-screen "point-of-regard" location based on the specifications of the tracking environment. The accuracy of this eye tracker is typically better than 1° of visual angle.

2.3. Materials

The experiment included two primary driving sessions: one with a long straight highway, and another with a highway with shallow curves that could be negotiated comfortably at high speeds (as on normal highways). Different highways were used to allow for possibly different data analysis of driver eye movements, since curves add an extra degree of complexity apart from lane changes. Each session required that the subject drive approximately 25 km (15.5 miles) on the given highway, which took approximately 14 minutes on average. The ordering of these two primary sessions was counterbalanced between subjects. Two additional sessions preceded the primary sessions: a practice session for subjects to become acquainted with the simulator and highway environment, and a session in which subjects were asked to change lanes every five seconds on a long straight road without traffic. Because these preliminary sessions do not involve naturalistic highway driving, they are not of direct interest to this paper and are not discussed further.

A detailed examination of highway driving and lane changing required collection of four types of data, all collated together into a single data file at a sample rate of

approximately 13 Hz. First, we collected standard driver control data as taken from the driving simulator; these data included steering wheel angle, accelerator and brake position, and turn signals. Second, we tracked subjects' eye movements as they fixated various road features and vehicles during driving. The eye tracking produced a sequence of "point-of-regard" locations indicating where the driver was looking on the simulator view (including the simulated rear-view mirror). Third, we recorded all information about the current environment with respect to the position, heading, and speed of each vehicle including the driver's vehicle.

The fourth type of data comprised verbal protocols from subjects indicating when they formed the intention to perform an action and when they completed the action. We recorded five different verbal reports: the intention to perform a passing lane change into the left lane (reported as "pass"), the intention to perform a non-passing lane change ("lane change"), the intention to perform a return lane change back into the right lane ("return" or "lane change"), the intention to follow the vehicle in front ("car follow"), and the completion of the last action ("done"). Each verbal report was keyed in immediately by the experimenter and recorded in the session data file. The distinction between the first two reports, which both involve changes from the right to left lane, arises in the fact that we expected subjects to occasionally change lanes in the absence of the need to pass. We also decided to record car following in the hopes of later analyzing these data, but for our purposes here we do not discuss them further. Subjects occasionally forgot to report the completion of a task; in these cases the experimenter keyed in a "done" when the completion was apparent. It should be noted that while a report of "lane change" could apply to either a left or right lane change, this distinction was obvious to the experimenter given the driver's current lane position.

2.4. Procedure

After a brief introduction to the experimental procedure and simulator setup, the experimenter fitted the eye tracker on the subject and calibrated the equipment using the standard IScan calibration points. Subjects then drove in the two preliminary sessions: in the first, they practiced driving in the highway environment until they felt comfortable controlling the car and generating verbal protocols; in the second, they drove on a straight open road and changed lanes every 5-10 seconds while generating verbal protocols. Subjects then ran through the two primary driving sessions with a short break between sessions. Before each session, the eye-tracking accuracy was checked visually by the experimenter and re-calibrated if necessary. The entire experiment lasted approximately one hour.

3. RESULTS

We now examine drivers' behavior focusing on their gaze patterns before, during, and after lane changes. To this end, we identified the segments of the data protocols that represent lane changes using drivers' verbal protocols as marked in the data; that is, we determined lane-change segments by noting where drivers verbally indicated the start and end of their lane changes. Overall, drivers varied greatly in terms of how many lane changes they made; on average the drivers made 36.5 ± 11.2 lane changes

including both segments of highway. They also varied in the time needed to complete a lane change, averaging 5.14 ± .86 seconds.

3.1. Gaze Pre-Processing

To analyze drivers' gaze behavior, we processed the raw eye-movement data into higher-level chunks to facilitate examination. The process began by associating raw data points to salient visual objects in the external environment. The first set of salient objects pertained to the roadway. Empirical studies (e.g., Land & Lee, 1994; Land & Horwood, 1995) and other work (e.g., Donges, 1978; Rushton & Salvucci, 2001) suggest that drivers utilize several salient road points to control the vehicle: the *vanishing point* for straight road segments, the *tangent point* for curved road segments, and the *near point* (or area) immediately in front of the vehicle. The tangent point and vanishing point can be used to stabilize the vehicle and predict future steering, while the near point can be used to maintain a central position in the lane (Salvucci et al., 2001). We limited the vanishing point to a distance corresponding to four seconds of time headway to better match drivers' typical look-ahead distance, and placed the near point at a distance corresponding to one-half second of time headway. Note that the vanishing point and tangent point are mutually exclusive in the sense that exactly one of them is used depending on the upcoming roadway (i.e., straight or curved road segment, respectively). The second set of visual objects pertained to the other vehicles in the environment. While vehicles are complex objects with potentially many salient features, each vehicle could be treated as a single object for the present purposes.

For each point in time during the simulation, we computed the global position of each visual object and mapped the global position to a corresponding rectangle on the two-dimensional screen image, and occluded rectangles were removed to leave only objects in view. We then assigned each raw eye-movement data point to the nearest visual object in terms of its distance to the two-dimensional image rectangle.

3.2. Aggregate Time Course Analysis

We first explore the time course of drivers' aggregate gaze behavior before, during, and after lane changes. Figure 1 shows the time-course plot for drivers' *gaze dwell ratios* — the mean ratio of time spent looking at objects in the specified areas during the given time window. The time-course plot was generated in several steps. First, for each lane change, we re-sampled the data between the onset and completion of the lane change into ten scaled units (each thus representing 1/10th of the total lane-change time). Second, we extended the window for these data to include eight scaled units before and after the lane change. Third, we averaged together all data points for each individual driver, and finally aggregated these data into aggregate plots showing the mean and standard error across drivers. Fourth, we added data points that represent lane-keeping behavior before and after the lane change. Thus, the figure includes five primary regions: LK, initial lane keeping; BLC, before the lane change; LC, during the lane change; ALC, after the lane change; and LK, later lane keeping. Furthermore, we divided the LC region into two subregions, LC1 and LC2, to represent the stages at which the vehicle is physically in the start lane and end lane, respectively. Because

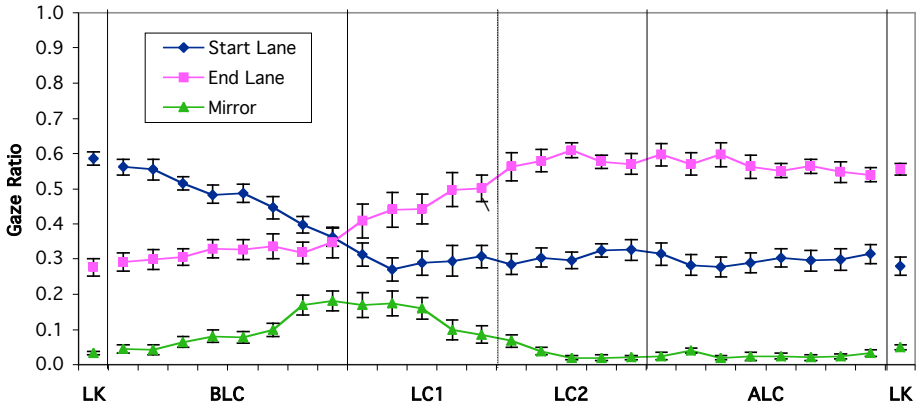

Figure 1. Gaze dwell ratios for the start lane, end lane, and mirror during initial lane keeping (LK), before the lane change (BLC), during the lane change in the start lane (LC1) and end lane (LC2), after the lane change (ALC), and later lane keeping (LK).

drivers averaged 5.14 seconds per lane change, each scaled unit represents roughly a half second of real time.

Figure 1 shows gaze dwell ratios for three areas: the start (or current) lane, the end (or destination) lane, and the (rear-view) mirror. The time-course plot reveals a great deal about driver control and monitoring behavior for lane changes. During initial lane keeping, drivers look at the start (current) lane approximately twice as much as the other lane, with very little time spent looking at the mirrors. Before a lane change, however, drivers begin to shift a great deal of their dwell time from the start lane to the mirror. As soon as they begin the lane change, dwell time for the end lane surpasses that of the start lane. Dwell time for the mirror tails off during the lane change, while dwell time for the end (destination) lane gradually increases. By the end of the lane change, drivers look at the end lane (now the current lane) twice as much as the start lane — that is, the complement of the initial state, as we would expect given the change of current lane.

These results contain two interesting findings. First, the figure shows a clear symmetry around the onset of the lane change, particularly with respect to the dwell ratios for the start and end lanes. This symmetry suggests that the lane-change onset marks a switching of roles between the two lanes — namely, that when drivers initiate a lane change, drivers stop controlling the vehicle using features of the start lane and begin controlling using features of the end lane. Second, the results give an initial indication into the demands of control versus monitoring during highway driving. During lane keeping, drivers must look at both lanes to monitor vehicles but only at their current lane to control their own vehicle (for the most part, at least). Both before and after lane keeping, we see a clear trend that the current lane draws twice as much dwell time as the other lane. Thus, for similar environments with two-lane highways and moderate traffic, we can roughly characterize drivers' gaze dwell time as being divided equally among three tasks: vehicle control, monitoring the current lane, and monitoring the other lane. Of course, this rough sketch does not consider the many

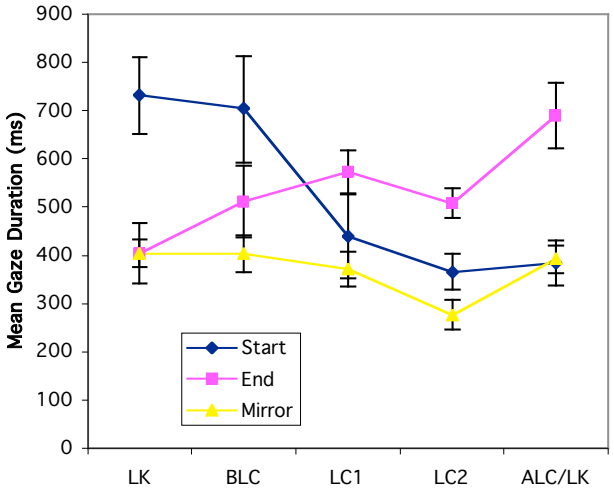

Figure 2. Mean gaze duration for the start lane, end lane, and mirror during initial lane keeping (LK), before the lane change (BLC), during the lane change in the start lane (LC1) and end lane (LC2), and after the lane change (ALC) with later lane keeping (LK).

possible situations that can arise or the fact that sometimes a single gaze may serve the purposes of both monitoring and control. However, it does give a rough estimate of how much gaze dwell time is spend on control versus monitoring.

3.3. Gaze Duration Analysis

The time-course plot above indicates the total gaze duration spend on various areas. Another way to analyze these data examines the mean duration per gaze — that is, the amount of time spent looking at an object during a single gaze (before looking at another object). Figure 2 shows the mean gaze durations for the start lane, end lane, and mirror in the five stages of the lane change discussed earlier. Gazes at the mirror maintain a fairly constant duration of approximately 350 ms. Gaze durations for the start lane are roughly 700 ms before the lane change, but afterwards drop sharply to the level of the mirror gaze durations. In contrast, gaze durations for the end lane start at the level of the mirror durations but rise (gradually but significantly) to roughly 700 ms.

As in the dwell-time plot, these results show a reversal at the onset of a lane change; however, while Figure 1 shows a reversal of total dwell time on the different lanes, Figure 2 shows a reversal of total time *per gaze*. These results again suggest a separation of control versus monitoring. Gazes at the lane guiding control (i.e., the start lane before LC, the end lane during and after LC) are relatively long, roughly 700 ms. Gazes at the monitoring areas (i.e., the mirror and the lane not guiding control) are approximately half as long, roughly 350 ms. Thus, drivers tend to focus more intently on objects guiding control and only briefly glance at objects for purposes of monitoring.

4. CONCLUSIONS

The main findings of this study should aid in the development of future theories and models of driver behavior in two ways. First, the study provides a parsimonious way of understanding the execution of a lane change: at the onset of the lane change, drivers simply switch their control from using features of the current lane to features of the destination lane. Such a technique has been implemented in an integrated computational model of driver behavior with good success (Salvucci et al., 2001). Second, the study elucidates the time course of gazes for the dual purposes of control and monitoring. The results can be used to augment and generalize current theories and models of vehicle control to include behaviors for monitoring and situation awareness.

REFERENCES

Ahmed, K. I., Ben-Akiva, M. E., Koutsopoulos, H. N., & Mishalani, R. G. (1996). Models of freeway lane changing and gap acceptance behavior. In J.-B. Lesort (Ed.), *Transportation and Traffic Theory*. New York: Elsevier Science Publishing.

Beusmans, J., & Rensink, R. (Eds.) (1995). Cambridge Basic Research 1995 Annual Report (Tech. Rep. No. CBR-TR-95-7). Cambridge, MA: Nissan Cambridge Basic Research.

Donges, E. (1978). A two-level model of driver steering behavior. *Human Factors*, *20*, 691-707.

Federal Highway Administration (1998). Our nation's highways: Selected facts and figures (Tech. Rep. No. FHWA-PL-00-014). Washington, DC: U. S. Department of Transportation.

Finnegan, P. and Green, P. (1990). The time to changes lanes: A literature review (Tech. Rep. No. UMTRI-90-34). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Gipps, P. G. (1986). A model for the structure of lane-changing decisions. *Transportation Research – Part B*, *5*, 403-414.

Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature*, *369*, 742-744.

Land, M., & Horwood, J. (1995). Which parts of the road guide steering? *Nature*, *377*, 339-340.

Rushton, S. K., & Salvucci, D. D. (2001). An egocentric account of the visual guidance of locomotion. *Trends in Cognitive Sciences*, *5*, 6-7.

Salvucci, D. D., Boer, E. R., & Liu, A. (2001). Toward an integrated model of driver behavior in a cognitive architecture. *Transportation Research Record*, No. 1779.

Talmadge, S., Chu, R. Y., & Riney, R. S. (2000). Description and preliminary data from TRW's lane change collision avoidance testbed. Presented at the Intelligent Transportation Society of America's Tenth Annual Meeting and Exposition; Boston, MA; May 1-4, 2000.