The Time Course of a Lane Change:
Driver Control and Eye-Movement Behavior

Dario D. Salvucci
Drexel University

Andrew Liu
MIT Man Vehicle Laboratory

Please address correspondence to:

Dario Salvucci
Dept. of Mathematics and Computer Science
Drexel University
3141 Chestnut St.
Philadelphia, PA 19104
+1 215 895 2674
salvucci@mcs.drexel.edu

Running head:

Time Course of a Lane Change

Abstract

In this paper we explore the time course of a lane change in terms of the driver's control and eye-movement behavior. We conducted an experiment in which drivers navigated a simulated multi-lane highway environment in a fixed-base, medium-fidelity driving simulator. We then segmented the driver data into standardized units of time to facilitate an analysis of behavior before, during, and after a lane change. Results of this analysis showed that (1) drivers produced the expected sine-wave steering pattern except for a longer and flatter second peak as they straightened the vehicle; (2) drivers decelerated slightly before a pass lane change, accelerated soon after the lane change, and maintained the higher speed up until the onset of the return lane change; (3) drivers had their turn signals on only 50% of the time at lane-change onset, reaching a 90% rate only 1.5-2 seconds after onset; (4) drivers shifted their eye movements from the start lane to the destination lane immediately after the onset of the lane change. These results will serve as the basis for future development of a new integrated model of driver behavior.

Keywords: driver behavior, lane changes, eye movements, integrated driver model.

INTRODUCTION

Driving is a highly complex task that requires continual integration of perception, cognition, and motor response. Of the various subtasks that comprise the entire driving task, *lane changing* is one subtask that incorporates many of the critical aspects of driving, such as lower-level control (e.g., steering, acceleration), monitoring (i.e., maintaining situation awareness), and decision making (e.g., when to change lanes). Despite the vast attention given to the driving task in general, much less attention has been directed to lane changing, despite its ubiquity in common driving environments — for instance, highway driving, which accounts for roughly 70% of vehicle miles on American roadways (Federal Highway Administration, 1998). Existing work on lane-changing behavior in large part emphasizes the decision-making aspects of the task, particularly gap acceptance and the decision of when to change lanes (e.g., Ahmed et al., 1997; Gipps, 1986). While other studies have addressed different aspects of lane changing, from behavioral aspects such as typical durations (e.g., Finnegan & Green, 1990) to practical development of lane-change collision warning systems (Talmadge, Chu, & Riney, 2000), we have yet to form a complete picture of when and how drivers change lanes including all the various perceptual and motor processes involved.

This paper attempts to flesh out this picture of lane-changing behavior by analyzing the time course of a lane change across several modalities. The paper expands on our recent study of eye movements during lane changing (Salvucci, Liu, & Boer, in press) and examines the integration of steering, throttle, turn signals, and eye movements for left-to-right and right-to-left lane changes. By aggregating and displaying these multimodal together as they occur before, during, and after a lane change, the analysis elucidates several interesting and surprising properties of drivers' lane-changing behavior. These results form a solid foundation on which to develop and validate a rigorous integrated model of driver behavior.

METHOD

This study of lane-changing behavior is based on data collected for a previous study (Salvucci, Liu, & Boer, in press) that focused exclusively on eye-movement data. The experimental task involved driving on a multi-lane highway with traffic in a fixed-base driving simulator. For clarity, we briefly review specific details of the data collection below. We also describe the data processing into time course graphs as needed for the multimodal time-course analysis in the next section.

Subjects

The experiment included a total of 11 participants for final analysis — two women and nine men between the ages of 18 and 31, all with at least two years of driving experience. An additional two participants did not complete the experiment because of serious difficulties in controlling the vehicle in the simulated environment. Another two additional participants were not included in data analysis because their eye movements could not be recorded successfully.

Driving Simulator and Environment

The Nissan Cambridge Basic Research (CBR) driving simulator (Beusmans & Rensink, 1995) was used as the platform for simulation and data collection. The simulator includes the front buck of a Nissan 240sx convertible that has been instrumented to collect standard driver data (steering position, throttle and brake position, turn signals, etc.). The simulated environment is projected in front of the buck with a field of view of roughly 70°, which we found adequate for a realistic driving experience. The simulator also incorporates an IScan (Burlington, MA) head-mounted eye tracker for collecting eye-movement data with respect to the projected view; this eye tracker maps the pupil center and an infrared reflect point to an estimated point-of-regard with an accuracy of approximately 1° of visual angle. Although the mirrors in the buck were not used, a simulated rear-view mirror was incorporated into the projected environment at

approximately the visual angle of the actual rear-view mirror, and could be utilized effectively by drivers to maintain awareness of surrounding vehicles.

While a number of simulated environments have been developed for the Nissan CBR simulator, we employed a multi-lane highway environment with moderate simulated traffic. The highway included two lanes in each direction, standard lane markings, and a barrier off the side of the road; there were no on- or off-ramps and no extraneous scenery (to focus drivers on the driving and lane-changing task at hand). The environment imposed a three-wheel dynamic model as the translation between driver input and vehicle response. Simulated vehicles moved at a random "desired" speed between 22.35 – 31.29 m/s (50 – 70 mph), passed slower vehicles in the right lane when necessary, and returned back to the right lane after passing as soon as possible.

The simulator in conjunction with the driving environment produced a continual tabular stream of data used for analysis. In particular, four sources of data were collected at a sample rate of 13 Hz and combined to the data file: (1) standard control data, including typical signals of steering-wheel position, throttle and brake position, turn signals, and any other accessible controls within the vehicle; (2) eye-movement data representing where the driver is directing their visual attention — or more accurately, representing gaze as a surrogate for the actual focus of visual attention; (3) vehicle data indicating position, heading, and speed of both the driver's vehicle as well as all other vehicles, to allow for recreation and full analysis of the entire environment; (4) verbal protocol data representing the intention to perform a lane change. This final data type requires further explanation. For our analysis of lane changes, we required some way of partitioning the data stream into segments each representing a single lane change (as well as before- and after-change data). To this end, we asked subjects to report the intention to make a lane change by saying "lane change" or "pass", and the experimenter recorded these reports in the data by means of a keystroke. Subjects also reported the end of a lane change by saying "done"; for the occasional cases in which the subject forgot to report, the experiment recorded the end when it was readily apparent. In summary, these four types of data form the necessary backbone for our time-course analysis of lane-changing behavior.

Procedure

After a general introduction to the simulator and task, the experiment began with two practice sessions in which subjects became acquainted with driving in the simulator and with reporting their lane changes as verbal protocols. Subjects then drove in two consecutive sessions (in a counterbalanced order): one on a straight highway without curves, and one on a highway with shallow curves. For each of these sessions, subjects drove for approximately 14 minutes for a total distance of approximately 25 km (15.5 miles). Subjects had a short break between sessions, during which the eye tracker was re-calibrated if necessary.

RESULTS

Our analysis of the lane-changing data centers around a time course analysis in which we visualize aggregated data sequences from different modalities before, during, and after a lane change. Figures 1 and 2 show the time-course graphs for right-to-left and left-to-right lane changes, respectively. The remainder of this section explains how these graphs were created and what they indicate about the nature of drivers' lane-changing behavior.

<< Insert Figures 1 and 2 approximately here, adjacent to each other if possible >>

The time-course graphs were created in several stages. First, we partitioned all data streams for all subjects and extracted the segments that represented a lane change; these segments were identified using the verbal protocol data, which indicate the start and end of each lane change. Second, we computed a scaled time unit for each lane-change segment as 1/10 of the total lane change time, and re-sampled the data during the lane change by averaging within these time units (thus producing ten data points during each lane change). Third, we extended each segment to include eight scaled time units before and after the lane change. Fourth, we added a single data point before and after the extended segment to include all lane-keeping data that occurred before

and after the segment. Finally, we aggregated all data points for an individual driver, and produced aggregate graphs over these data including both the mean and standard error across the individual driver data. As can be seen in the figures, the graphs thus include five regions of analysis: initial lane keeping, labeled LK; before the lane change, BLC; during the lane change, LC; after the lane change, ALC; and final lane keeping, LK. The data in the graphs include a total of 401 lane changes (36.5 ± 11.2 per driver), and the duration of a single lane change averaged 5.14 ± .86 seconds; thus each scaled unit in the graphs corresponds to approximately one-half second of real time. In the following analysis, we will refer to quantities of time in terms of both seconds and scaled units using this correspondence.

Lane Position

The first panels in Figures 1 and 2 show the lane position as a function of time. The general picture is not surprising: the vehicle begins moving from the center of the start lane at the onset of the lane change, and it begins to level out to the center of the destination lane by the end of the lane change. One interesting aspect of these data occurs in the 5 seconds (10 scaled units) before the start of the lane change — namely, there is a subtle shift away from the destination lane approximately 2-3 seconds (4-6 units) before lane-change onset. This shift could be attributed to drivers compensating for potential danger while visually scanning their environment. In the five seconds prior to the lane change, drivers begin to look around their vehicle to ensure a safe lane change (as we will see in the eye-movement data). It seems that in these five seconds, drivers move slightly away from the destination to avoid possible vehicles in that lane. Of course, assuming the lane is clear, the driver can then proceed to steer the car into the destination.

Steering

The steering profiles for both lane-change directions have similar characteristics. First, the initial steering from center occurs approximately a half second (one scaled unit) before the reported initiation of the lane change. This latency is likely due to a slightly delayed reporting

the verbal protocol. From there, the steering profile takes on a rough sine-wave shape. Drivers first steer toward the destination and then back to center to level the vehicle on a steady path to the destination lane. Then, in a continual smooth movement, drivers steer in the other direction and back to center to straighten the vehicle in the destination lane. Note that the second peak has a slightly larger duration and smaller maximum magnitude, indicating a smoother leveling-off process.

We should note one aspect of the steering data that we noticed in viewing individual lane-change protocols in replay. In the aggregate graphs in the figures, the steering profiles show only a single sine-wave like function. However, many individual protocols seemed to show several sine-wave like functions that occur one after another — more specifically, a larger sine wave followed by one or more smaller ones. In essence, drivers seemed to perform a sequence of small "lane changes" that move the vehicle incrementally toward the center of the destination lane. First, drivers would perform a single, larger steering maneuver to move the vehicle roughly into the next lane. Then, they would execute smaller sine-wave maneuvers to adjust the position of the vehicle within the lane. The graphs in the figure aggregate these many sine waves together, with the final appearance of a single maneuver. Given the variability in our data, we were unable to perform a more systematic analysis to determine the presence of these smaller maneuvers, but we hope to explore this issue further in future studies.

Throttle Position

In examining throttle position, it is essential to look at the two types of lane changes separately, given that they represent very different situations. Figure 1 shows the throttle-position profile for right-to-left lane changes in which the driver is generally passing a slower lead vehicle. In the 5 seconds (10 scaled units) before the lane change, the driver decelerates gradually, presumably to avoid colliding with the slower vehicle. However, soon after lane-change onset, the driver accelerates to passing speed and maintains that speed through the rest of the lane

change and beyond. Of note, the average speed remains high even long after the lane change (the right-most data point), indicating an overall speed change during the pass.

Figure 2 shows the profile for left-to-right lane changes in which the driver returns to the right lane after a pass. Generally, the driver does not need to decelerate in this case; if there is a slower vehicle in the right lane ahead of the driver's car, most drivers simply continue the pass maneuver. However, drivers do indeed decelerate immediately upon starting the lane change, reducing quickly to slow-lane speed and remaining at this speed from then on.

Turn Signals

The turn-signal ("indicator" or "blinker") profiles in the figures show the ratio of protocols in which the turn signal was on at the given time. The profiles for lane changes in both directions look relatively similar. First, drivers began turning on the signals approximately 1.5 seconds (3 scaled units) before the start of the lane change. However, at the time of lane-change onset, the signal was on only half of the time (with a ratio of approximately .50); only after an additional 1.5–2 seconds (3-4 units) did the ratio reach its peak of approximately .90. Drivers also generally turned off the signals well before the end of the lane change. The smoothness of this profile is quite striking, given that one might expect more of a step-function profile, with a rapid increase at onset and rapid decrease at offset.

It should be noted with these turn-signal data that drivers may have used signals in two distinct ways, just as in normal driving: (1) to indicate the *execution* of a lane change, or (2) to indicate the *intention* to execute a lane change. For the latter case, drivers may turn on the signal before looking around their environment, to warn other drivers of the imminent maneuver or possibly to "ask" for help in creating space for the maneuver (as might be common on highly congested roadways). However, given that in this experiment drivers turned on the signal around onset time or even later, it seems that drivers generally followed the first interpretation.

Eye Movements

In an earlier paper (Salvucci, Liu, & Boer, in press), we examined drivers' eye movements during lane changes as a single profile aggregated for both directions of lane changes. Figures 1 and 2 show the eye-movement data for each direction as *gaze dwell ratios* — that is, the ratio of time spent looking at objects in the start lane, the end lane, or the rear-view mirror. In Figure 1, drivers spend most of their gaze time before the lane change looking at the start lane (i.e., their current lane). As onset approaches, more gaze time is directed to the end lane and the mirror, since drivers are checking to the side and rear of the vehicle to ensure safe passage. When drivers initiate the lane change, the dwell ratios for both start and end lane converge to approximately the same value while the ratio for the mirror trails off. These values remain steady through the lane change and beyond; essentially, drivers need to look at the destination lane to control the vehicle, but keep looking at passed car in the start lane to monitor its position.

In Figure 2, drivers begin as they ended above, with equal dwell ratios to the start and end lanes. Again, as onset approaches, drivers direct more of their gaze to the mirror and less to the start lane. Somewhat by onset time, and certainly by the midpoint of the lane changes, drivers direct the vast majority of their gaze to the destination lane. Thus, drivers have come full circle and the end of the profile matches the start of that in Figure 1. The most striking aspect of both profiles is how quickly the gaze distributions change at the onset of the lane change: within only 1-2 seconds (2-4 scaled units) after onset, drivers focus their gaze on the destination lane, even when the vehicle remains (at least partially) in the start lane.

Summary

To summarize the results above, we noticed several aspects of the time course of driver lane changes that stood out in our analysis:

1. Drivers exhibited the expected sine-wave steering pattern except for a longer and flatter second peak as they straightened the vehicle. In addition, drivers exhibited a slight

tendency to steer away from the destination lane 2-3 seconds before the onset of the lane change.

2. Drivers decelerated slightly before making a pass lane change, accelerated soon after the lane change to a higher passing speed, and maintained this speed up until the onset of the lane change back to the slow lane.

3. Drivers had their turn signals on approximately 50% of the time at lane-change onset, reaching approximately 90% of the time 1.5-2 seconds into the lane change.

4. Drivers began shifting their gaze away from the start lane approximately 5 seconds before lane-change onset, and shifted their gaze toward the destination lane almost immediately after lane-change onset.

Given the smaller scope of our study and the use of a driving simulator rather than a real vehicle, these results should be taken as preliminary observations into the nature of driver behavior during lane changes. However, these data do provide a baseline for further, more rigorous studies on lane changing as well as analogous time-course analyses for similar driver maneuvers.

DISCUSSION

In addition to augmenting our general knowledge about lane-changing behavior, we are in the process of using these data to develop and validate an integrated model of driver behavior (Salvucci, Boer, & Liu, 2001; for similar models, see also Aasman, 1995; Levison & Cramer, 1995). Our initial efforts have attempted to formalize a simple control mechanism that enables basic steering, lane changing, and curve negotiation in as straightforward a way as possible. To this end, we have constructed an integrated driver model (Salvucci, Boer, & Liu, 2001) that utilizes a simple two-level steering model (see Donges, 1978; Land & Horwood, 1995; McRuer et al., 1977) based on two control points: (1) a *far point* that guides predictive steering, possibly the road tangent point (Land & Lee, 1994), road vanishing point, or lead vehicle; and (2) a *near point*

that guides steering to lane center. With this steering model, lane changing is implemented in an extremely simple way: to initiate a lane change, the model simply switches from using the far and near points of the start lane to using the far and near points of the destination lane. The only minor adjustment to this method was the addition of a control parameter to limit the speed of the lane change, modeling individuals' variability in their desire to execute the maneuver more quickly or slowly. We have validated this model's behavior with parts of the collected data, namely the steering and gaze data. In future work we hope to capture additional aspects of behavior presented here, including throttle and turn-signal data.

Lane changing is not simply a control problem, however; it involves both monitoring to maintain situation awareness and higher-level decision making to determine when to execute. Thus, we would like our integrated driver model to capture these aspects of behavior as well. The recently-developed model has been developed in a *cognitive architecture*, or computational framework that incorporates the abilities and limitations of human behavior. By utilizing a cognitive architecture (specifically the ACT-R architecture of Anderson & Lebiere, 1998), the model operates over declarative memory structures that embody current situation awareness but fade over time. Thus, by "gazing" at its environment with simulated "eyes", the model can gather information about its surroundings and, for instance, can recall or forget a vehicle in the driver's "blind spot". Also, the model incorporates higher-level decision making based on both the contents of its declarative memory structures as well as its current perceptual information. The integration of higher-level cognition and lower-level control offers the integrated model the potential to represent driver behavior in both coarse and fine-grained levels of details, and also predict the effects of secondary tasks on driver performance (e.g., Salvucci, 2001). Time-course analyses for lane changing and other maneuvers and tasks are crucial to providing insight into how to build such models and providing data with which to validate these models in a rigorous manner.

REFERENCES

Aasman, J. (1995). *Modelling driver behaviour in Soar*. Leidschendam, The Netherlands: KPN Research.

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.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

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: The 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.

Levison, W. H., & Cramer, N. L. (1995). Description of the integrated driver model (Tech. Rep. No. FHWA-RD-94-092). McLean, VA: Federal Highway Administration.

McRuer, D. T., Allen, R. W., Weir, D. H., & Klein, R. H. (1977). New results in driver steering control models. *Human Factors*, *19*, 381-397.

Salvucci, D. D. (2001). Predicting the effects of in-car interface use on driver performance: An integrated model approach. *International Journal of Human-Computer Studies, 55*, 85-107.

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

Salvucci, D. D., Liu, A., & Boer, E. R. (in press). Control and monitoring during highway driving. To appear in *Vision in Vehicles 9*. Amsterdam: Elsevier.

Talmadge, S., Chu, R. Y., & Riney, R. S. (2000). Description and preliminary data from TRW's lane change collision avoidance testbed. In Proceedings of the Intelligent Transportation Society of America's Tenth Annual Meeting and Exposition.

ACKNOWLEDGMENTS

FIGURE CAPTIONS

Figure 1. Time-course graphs for right-to-left lane changes. Note that one scaled time unit represents approximately 0.5 seconds.

Figure 2. Time-course graphs for left-to-right lane changes. Note that one scaled time unit represents approximately 0.5 seconds.
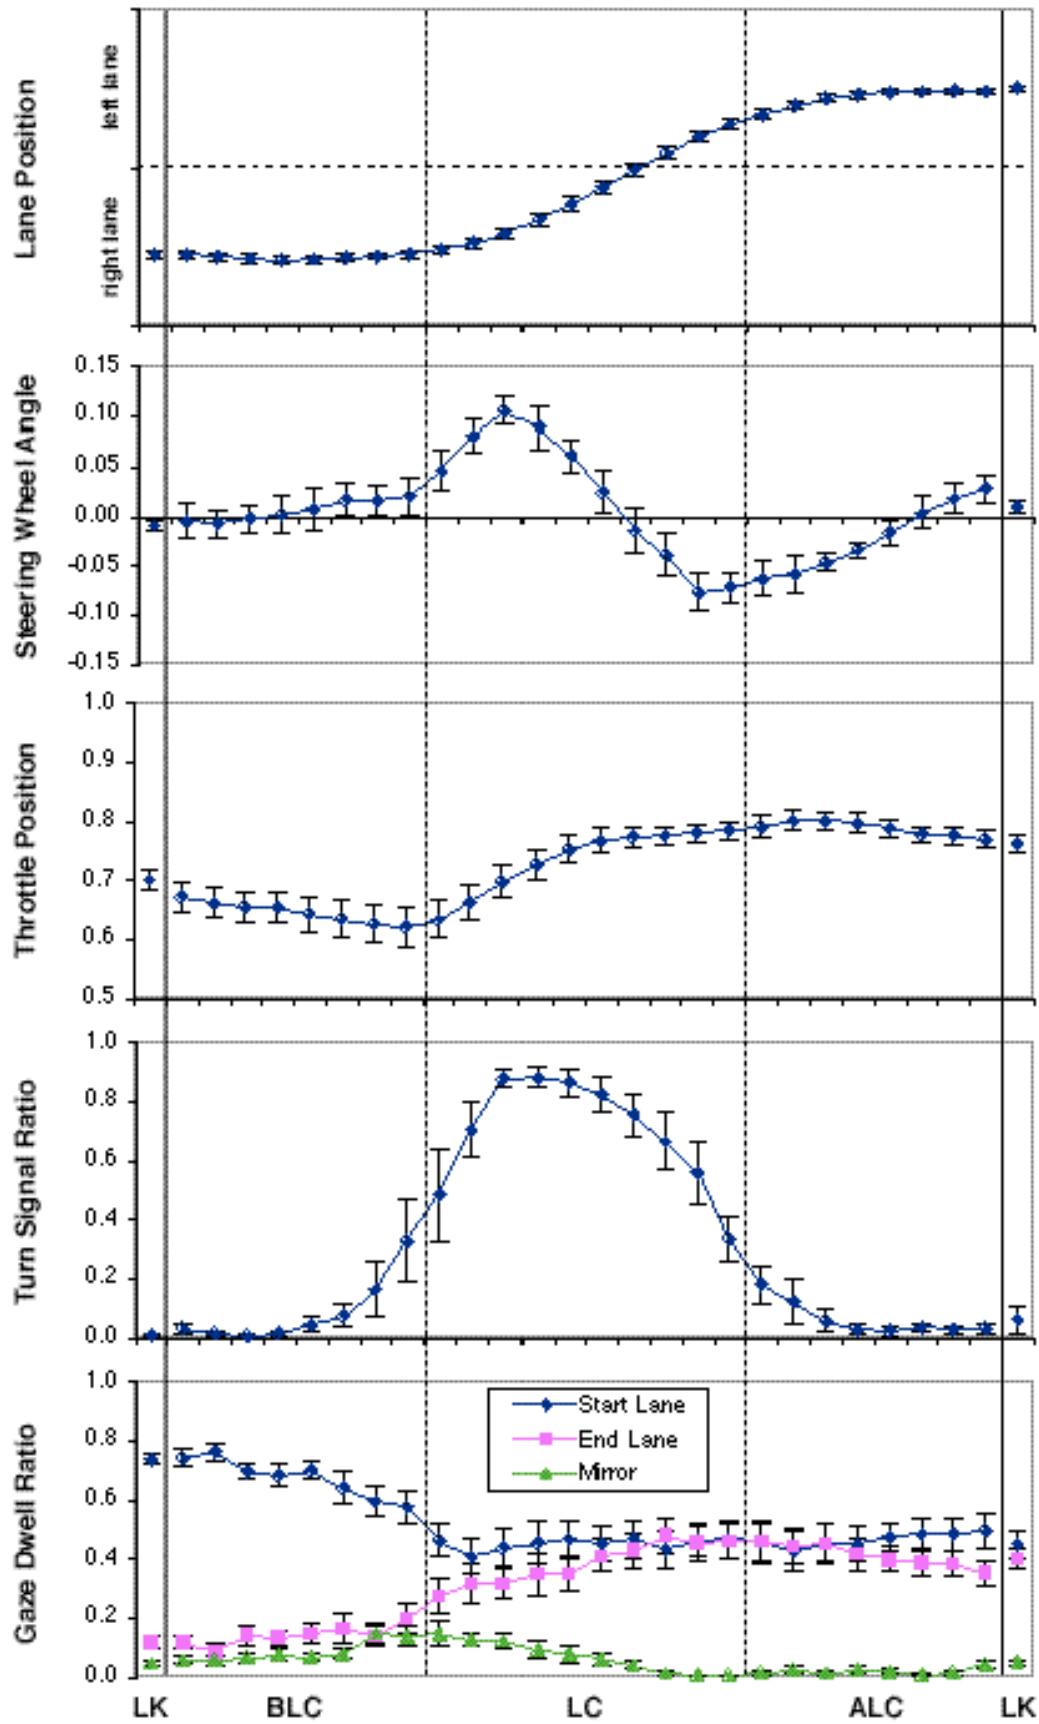
Figure 1.

Figure 2.