

WHY DRIVING PERFORMANCE MEASURES ARE SOMETIMES NOT ACCURATE (AND METHODS TO CHECK ACCURACY)

Paul Green
UMTRI-Human Factors
University of Michigan
Ann Arbor, Michigan, USA
E-mail: pagreen@umich.edu

Summary: This paper identifies common sources of inconsistency and error in measurements of driving performance and describes methods to determine the size of these errors. Major sources of inconsistency and error discussed include (1) the lack of zeroing procedures (which affects measurements of steering wheel angle), (2) unknown input and output mapping (which affects measurements of throttle position), (3) the failure to control critical factors such as tire pressure, traffic, and wind (which affects measurements of speed), (4) uncertainty about where the lane boundary actually is (which affect measurements of lane position and counts of lane departures), and (5) the failure to define or use consistent definitions for measures such as headway/gap, time to line crossing, and time to collision. The lack of or inconsistency of definitions can lead to multiple interpretations of what could have been measured, and differences between interpretations are of practical significance. The types and magnitude of these inconsistencies and errors vary with the measurement platform, complicating the comparison of driving studies and interfering with building a body of knowledge of driving. By making the driving research community aware of these problems, they can be identified, assessed, and minimized in the future, and published research can be read with a more critical eye.

INTRODUCTION

As the popularity of this conference attests, there has been a renaissance of driving research over the last few years, especially in studies that use driving simulators. The confluence of new people using new technologies, coupled with inattention to basic principles of measurement, has led to research that is not as good as it should be.

First, some researchers are unfamiliar with the all the details of how simulators work and what the data collected actually measure. Assumptions are made that driving performance is being measured one way when the value obtained actually represents something else.

Second, because of the focus on statistically significant differences, verifying the ground truth of performance measures may not occur and the values reported may not be reasonable for real driving. Though analyzing extreme conditions in a simulator can help identify differences in an efficient manner, the results of such research can be unconvincing when extrapolated to apply to normal driving conditions. Further, when unrealistic data are presented too often, it may lead to loss of confidence in human factors driving studies in general.

Third, several measures commonly used by driving researchers are not defined consistently, and some papers do not state which of several definitions is being used.

Accordingly, this paper discusses some common driving measures and statistics and how they are measured, primarily in simulators but on the road as well. This paper concentrates on DriveSafety Vection™ because it is the simulator most commonly used by attendees of this conference, including the author.

MEASURES OF INTEREST

Steering Wheel Angle - Has the Sensor Been Properly Zeroed?

Steering wheel angle refers to the angle of the yoke relative to horizontal, with positive usually being clockwise. In the current version of the DriveSafety simulator, angle is measured using a high-precision optical encoder. Angular accuracy is 5/1000th of a degree, sampling is at 400 Hz, and zero is established by a potentiometer. That sensor should be mounted above the U-joint in the column to reduce measurement error. In older versions of the simulator, zero was determined by turning the wheel from lock to lock, and then the system determined the midpoint. Should the experimenter not turn the wheel fully to the lock during calibration, the determined angular value will be wrong, possibly by several degrees.

With the growth in steer-by-wire systems, steering wheel angle in real vehicles may be measured by vehicle sensors and broadcast on the CAN bus. Generally, manufacturers treat CAN bus codes as highly secret, so obtaining information on accuracy or precision is extremely difficult, and publication is prohibited. One could estimate precision by slowly turning the steering wheel and looking for stair-stepping in the time history of the output. In a driving simulator, one could use the same process or ask the manufacturer for the specifications.

In theory, the mean value for a properly calibrated steering system should be zero. However, when driving a real road, a small steering input may be needed to counteract the crown of the road. In a simulator, the mean should be zero unless the subject is driving a loop. To check the calibration of a real vehicle, find a long stretch of road that is “extremely straight” and has minimal crown and grooving. (Expressways tend to be less crowned than local roads.) “Extremely straight” can be established by standing in the road, and looking down it, either standing upright or squatting down. If the road seems to curve to the eye, find another road.

For a driving simulator, building a long straight section is easy. For example, for a DriveSafety simulator, it is a matter of laying down one straight tile after another. Fortunately, tiles are not crowned and the default for crosswinds is none.

To check the steering wheel angle, have someone drive that road a few times when there is no traffic and minimal crosswinds, instructing them to drive as straight as possible and not to pass. Removing the turnarounds, plot the distribution of wheel angles. If the system is properly calibrated, the mean and the median should both be zero, though the variability will depend on the accuracy of the sensor and subject perturbations. If anything, minor disturbances are missing from driving simulators, so one could drive for 10 or 15 seconds without making any corrections

once the vehicle is aligned, which a driver would not do on the road. For comparison, see Green, Wada, Oberholtzer, Green, Schweitzer, and Eoh (2007) for distributions of steering wheel angle (and for other measures as well) for a fleet of 2002 Buick LeSabres driven on a wide range of roads, from data collected in the advanced collision avoidance system (ACAS) naturalistic driving study.

Throttle Position – What is the Throttle Map?

In the DriveSafety simulator, simulator position (often reported as a percentage) is determined by a linear potentiometer that provides an analog voltage output. In real vehicles, there is likely to be a throttle sensor on the CAN bus if the vehicle has dynamic stability control. Typical accuracy and precision are unknown. If measuring directly using a potentiometer, experimenters need to carefully consider the measurement location, as play in the throttle linkage occurs.

In both cases, a throttle map determines the relationship between throttle input and output, which may not be linear over the entire range. For example, commonly a small initial movement of the throttle leads to a large throttle output, so the vehicle feels responsive. For real vehicles, the throttle map is usually embedded in the engine controller code, which is not published and must be obtained from the manufacturer, if they release it at all.

To create a throttle map, one needs to drive the vehicle on a flat road (real or simulated) and slowly accelerate, measuring throttle position (internally to the vehicle/simulator and independently), engine RPM, and speed. The author has never heard of anyone doing this. Fortunately, for simulators, if one has access to the cab configuration files, this measurement should not be needed. The key point is that throttle input and throttle output may not be the same.

Speed – Are the Tires Properly Inflated? How Do Traffic, Wind, Etc. Perturb Speed?

In a driving simulator, speed is usually determined by dividing the distance traveled by time, specified very accurately by the system clock. In the DriveSafety simulator, the reported accuracy of the clock is 0.00001 s. The highest sampling rate is 60 Hz or 0.01667 s per sample. Considering such, speed should be accurate to the nearest 0.016 km/hr (0.01 mi/hr).

In a real vehicle, the situation is quite different. The ultimate measure of speed is usually determined by counting pulses from the ABS controller, which for some vehicles is four pulses per tire revolution. During each quarter revolution of the tire, the vehicle travels 25% of the tire circumference, assuming no wheel slip. With an accurate clock, speed can accurately be determined. However, the circumference of the tire depends on its inflation pressure. The author does not know of any data that relate tire pressure and vehicle loading to actual tire diameter, but presumably such data exist.

If desired, the ABS-based speed estimate can be compared with estimates based on the engine controller and the transmission controller-measured RPM, the gear and differential ratio, and transmission slip.

One could also use a radar gun to check speeds recorded by on-board instrumentation. This, of course, assumes the radar gun is properly calibrated.

One could also drive on a real, flat road at a fixed speed set by a cruise control when there is no traffic, and time the distance between mile post markers. Older markers should certainly be accurate to within 3 m (10 ft, L. Kostyniuk, April 3, 2007, personal communication), though recently, GPS is used to place markers, so they should be accurate to within a meter. Of course, accuracy can be improved averaging over several miles. Readers should keep in mind that setting a cruise control does not set a single speed, but rather speed over a very limited range, usually less than 1.6 km/hr (1.0mi/hr), that varies in a regular pattern. So, if accurate speed is important, check the tire pressure of the test vehicle regularly.

In an experiment, the speed at which a subject drives depends on traffic, road geometry, sight distance, and road surface condition, as well as tasks other than driving (e.g., using a phone) in which the driver might be engaged. In a driving simulator, the speed cues are different, so subjects tend to drive much faster than normal if they do not attend to the speedometer, though this can vary with subject age, with older drivers creeping along. Often, what is missing in the simulator is following traffic that honks or passes when the driver is going too slowly. In addition, in the real world passing drivers may yell at the subject, give them the finger, or express displeasure in other ways. Most vehicles serving as traffic in simulators have no visible driver, let alone avatars that act.

Thus, because speed in a simulator is a simple geometric calculation, measured speed should be accurate and is not worth checking. However, subject behavior may not lead to realistic driving speeds. On real roads there are invariably slight grades and winds, all of which alter speed. In simulators, roads are perfectly flat and, except for crosswinds to cause lane departures, there are no winds. Hence, speed variance could be less than in actual driving in some circumstances for accomplished subjects.

Lateral Position – So, Really, Where Is the Lane Edge? And Do Not Forget, People Cut Curves

In most simulators, lane position is usually the distance from the centerline of the vehicle to the centerline of the lane marking, but it could be to an edge of the lane marking. Keep in mind that the painted line has a finite width (in the U.S., typically 10 cm (4 in)). For double lines, the situation is more complex.

An even larger problem is dealing with curves. In highway engineering, a curve consists of a straight section (the tangent), a true curve (with a radius), and a transition section. The transition section is needed to avoid the sudden change from zero steering angle to the angle required by the curve. In a simulator, curves can be generated using splines or as a series of chords. If the chords are short enough, drivers do not notice the angular change, especially if the scene has texture. If chords are used, errors in lane position can be substantial since drivers follow the curve, not the chords. A rather tight 200-m radius curve simulated with 10-m chords would have about 125 chords per circle. The difference between where the chord contacts the radius and where the distance is maximum (the midpoint of the chord) is 0.125 m. Considering that the

standard deviation of lane position of 0.2 m is typical (Green, Cullinane, Zylstra, & Smith, 2003), this is a significant potential error. For 400-m curves, the value is about half of that, 0.0625 m (6 cm), which is still a significant amount. In the DriveSafety simulator, curves are approximated from a series of points, but those points are very close together. Thus, for a simulator, it is important to know how curves are generated and how lateral position is determined.

For real roads, the assumption is that the pavement and the paint (the lane markings) are perfectly aligned. This is not true, as may readily be observed when driving on a concrete road. An observer, driven by someone else, should rest their head on the window edge to provide a fixed reference, and watch the road flow by. For concrete roads, because of the design of paving machines, the width varies very little. For many roads, the painted marking will actually wander back and forth from the edge of the concrete pavement by a few inches. This variation is important because most lane position systems use the lane markings to determine lane position, not the edges of the pavement. If drivers used the pavement to guide the vehicle, the vehicle instrumentation used the paint to measure location, and the alignment of the painted marking was poor, then the instrumentation could report considerable variation in lane position when in fact lane position hardly varied at all.

Further, a lane tracker can be fooled. Sometimes longitudinal cracks are mistaken for lane markings, as are rivulets of water and tar strips when driving into the sun. Fortunately, lane trackers are getting better at rejecting these objects as lane markings, though lane trackers based on lane departure warning systems tend to have lags and do not work well when drivers make quick maneuvers laterally or when pitching occurs.

Finally, real drivers sometimes do not drive down the center of a lane, and in fact, they cut curves. So, if lane position across straight and curved sections is averaged, curve cuts can increase the estimated standard deviation of lane position even though the driver is not deviating from the intended path. Such calculations are greatly affected by the inclusion or exclusion of lane changes (D. LeBlanc, April 4, 2007, personal communication).

A very common derivative measure from lane position is lane departure. Certainly, if drivers make an unintended departure from their lane, they could strike another vehicle or a fixed object, and for that reason lane departures serve as a primary safety surrogate.

However, there is no consistency in determining when a departure occurs. At issue is what part of the vehicle is considered (outside edge of the wing mirror, curve of the body, outside wall of the tire, or the contact patch or the centerline of the tire) and which part of the lane marking is considered (inside edge, centerline, outside edge). This is problematic.

Headway – Is the Highway Engineering or the Human Factors Definition Being Used?

Headway is generally defined, when it is defined, in two ways. Highway engineers think of headway in terms of passage of two vehicles past a point, usually the difference in arrival time between the front of a vehicle and the front of a following vehicle. The inverse of this difference indicates the traffic-carrying capacity of that lane (measured in vehicles/hour).

In human factors, headway (or more properly, gap) refers to the distance or difference in arrival time between the back of a lead vehicle and the front of a following vehicle. This difference represents the time or distance in which a driver can respond to some action of a lead vehicle. In the DriveSafety simulator, this is the HeadwayDist variable, and is computed as the distance along the path. (If computed as Pythagorean distance in the x,y positions of the two vehicles, it will be less.)

So for headway/gap, the major problem is that the measure is often not defined, and the difference in definitions can be of practical importance. For example, suppose a driver is following a lead vehicle fairly closely with a 1.0 s gap. At 96 km/hr (60 mi/hr), the gap between the two vehicles would be 26.8 m (88 ft). If the vehicles in question were 4.7-m (15-ft) long, the headway gap would be about 17% greater or 1.17 s, a difference of practical significance. Obviously, if some of the vehicles amongst the traffic are tractor-trailers (roughly 18.3-m (60-ft) long), the difference can be much larger.

As a practical matter, the problem is determining what is the lead vehicle. In real vehicles, the headway sensor may lock on vehicles in other lanes, especially in curves. In a simulator, it is usually a matter of deciding which vehicle to choose. For example, suppose the driver is changing lanes. In their lane is a vehicle with a gap of 30.5 m (100 feet). In the lane they are entering is a vehicle that is closer, say 324.4 m (80 ft) down the road. At what point does the vehicle in the other lane become the lead vehicle—when any part of the subject vehicle is in the adjacent lane, when more than half of the vehicle is in that lane, when it is completely in that lane, or something else?

Time to Line Crossing (TLC) – Which Method Is Being Used to Estimate It?

Time to line crossing (Godthelp, Milgram, & Blaauw, 1984; van Winsum, Brookhuis, & de Waard, 2000; Mammari, Glaser, Netto, & Blosseville, 2004) is a measure of the driver's lateral control performance. Interestingly, van Winsum et al. (2000, p. 24) note, "TLC represents the time available for a driver until the moment at which any part of the vehicle reaches one of the lane boundaries." Is reaching the boundary when some part of the car (e.g., a wing mirror) is above the lane boundary or does reaching the boundary require a tire to touch it? (As a footnote, in DriveSafety simulators, "crossing" the line occurs when a tire touches the line.) Van Winsum et al. (2000) describe two methods to estimate TLC: (1) dividing the lateral distance by the lateral velocity and (2) also including the lateral acceleration. Based on experimental data for several maneuvers and TLC values computed using both methods, Van Winsum et al. (2000) note that the preferred calculation method depended on the maneuver. Often, the computation procedure used by a particular simulator (and for on-road studies as well) is unspecified.

Time-to-Collision (TTC) – Which Method Is Being Used to Estimate It?

Time to collision is defined as "The time required for two vehicles to collide if they continue at their present speed and on the same path" (Hayward, 1972, as cited in van der Horst and Hogema, 1993). In many ways, the issues for TTC are the same as those for TLC—namely, whether acceleration is considered in the estimate and what assumptions are made about changes in the path of the subject vehicle and the lead vehicle to be struck.

CONCLUSIONS

When conducting human factors studies of driving, be sure to consider the following factors: (1) how steering is zeroed when measuring steering wheel angle, (2) the throttle map used for measuring throttle position, (3) if the tires are properly inflated and how traffic, wind, and so forth perturb measurements of speed, (4) how lane position and lane departures are determined, (5) which definition of headway/gap is being used, the highway engineering or human factors definition, (6) how time to line crossing is being computed, and (7) how time to collision is being computed. To do quality science, well-defined measures are needed, as well as methods for collecting them. In driving research to date, that has not always been the case.

ACKNOWLEDGMENTS

I would like to thank Doug Evans of DriveSafety for technical input to this paper.

REFERENCES

- Godthelp, J., Milgram, J., & Blaauw, G.J. (1984). The development of a time-related measure to describe driving strategy. *Human Factors*, 26, 257-268.
- Green, P., Cullinane, B., Zylstra, B., & Smith, D. (2003). *Typical values for driving performance with emphasis on the standard deviation of lane position: A summary of literature*. (Technical Report UMTRI-2003-42). Ann Arbor: The University of Michigan Transportation Research Institute.
- Green, P.E., Wada, T., Oberholtzer, J., Green, P.A., Schweitzer, J., & Eoh, H. (2007, In press). *How do distracted and normal driving differ: An analysis of the ACAS naturalistic driving data*. (Technical Report UMTRI-2006-35). Ann Arbor: The University of Michigan Transportation Research Institute.
- Hayward, J.C. (1972). *Near miss determination through use of a scale of danger* (Report # TTSC 7115), University Park: The Pennsylvania State University.
- Mammar, S., Glaser, S., Netto, M., & Blosseville, J.-M. (2004). Time-to-line crossing and vehicle dynamics for lane departure avoidance. In *Proceedings of the 7th International IEEE Conference on Intelligent Transportation Systems*. Piscataway, NJ: IEEE Intelligent Transportation Systems Council, 618-623.
- van der Horst, R., & Hogema, J. (1993). Time-to-collision and collision avoidance systems. In *Proceedings of the 6th ICTCT Workshop*. Vienna: International Cooperation on Theory and Concepts in Traffic Safety, 15-22.
- van Winsum W., Brookhuis K.A., & de Waard D. (2000). A comparison of different ways to approximate time-to-line crossing (TLC) during car driving. *Accident Analysis and Prevention*, 32(1), 47-56.