

Baseline Analysis of Driver Performance at Intersections for the Left-Turn Assist Application

Scott Stevens, Jeffrey Bellone, Philip Azeredo, and Marisol Medri
USDOT Volpe Center

ABSTRACT

This study is aimed at supporting left-turn assist (LTA) applications, which provide warnings to drivers making a left turn across the path of oncoming traffic (LTAP/OD scenarios). The primary goal was to provide much-needed information on typical or “baseline” driving in LTAP/OD scenarios that can be used to refine alert criteria to reduce false and nuisance alerts. A secondary goal was to provide performance data useful for informing practical test procedures, e.g., setting turning speed when evaluating LTA applications on a test track. To accomplish this, LTAP/OD events were identified in the databases of two large-scale naturalistic driving studies. For these events, we estimated the size of the gaps in oncoming traffic into which drivers chose to turn, what factors (environmental, demographic, etc.) affected the choice to turn into a gap of a given size, and the speed profiles throughout each turn. The factors with the largest effect on gap size were age and gender, followed by road wetness, whether or not the turning vehicle stopped before turning, and the number of lanes the turning vehicle had to cross. As a counterpoint to this analysis of safe, typical turning behavior, we also identified instances in a national crash database where LTAP/OD maneuvers led to serious collisions involving vehicles equipped with event data recorders, i.e., scenarios where an alert could have been useful. Speed profiles for turning vehicles did not differ strongly for collisions as compared to baseline driving, but estimated gap sizes were shorter.

CITATION: Stevens, S., Bellone, J., Azeredo, P., and Medri, M., "Baseline Analysis of Driver Performance at Intersections for the Left-Turn Assist Application," *SAE Int. J. Trans. Safety* 5(1):2017, doi:10.4271/2017-01-0033.

INTRODUCTION

Intersections are frequent sites for vehicle collisions. To reduce their frequency and the harm that they cause, applications are being developed that would warn drivers of impending collisions at intersections. However, even if the predictions of these applications are highly accurate, because of the rarity of collisions in everyday driving, drivers will continue to experience false alerts. Additionally, there will be alerts—“nuisance alerts”—that, although true according to the logic of the warning algorithm, are not perceived as helpful by the driver. This combination of false and nuisance alerts can undercut the effectiveness of the warning application by annoying drivers, eroding their trust in the application, and leading them to ignore or even disable the application. It is therefore important to minimize false and nuisance alerts.

This study is aimed at improving our understanding of intersection behavior by identifying metrics that describe driver intersection navigation at both signalized and un-signalized intersections. This supports the overarching goal of knowing when alerts should and should not be issued.

A secondary goal for the study is providing estimates of typical or “baseline” driver behavior that can be used to define conditions in practical test procedures, such as setting the turn speed when testing intersection collision warning systems on a test track.

The focus of this study will be the scenario addressed by left-turn assist (LTA) warning applications, namely left-turn across-path, opposite-direction (LTAP/OD) collisions. In an LTAP/OD collision, a driver is struck by an oncoming vehicle while trying to make a left turn. Different scenarios will be addressed in future work.




Figure 1. Diagram of a Left-Turn Across Path-Opposite Direction (LTAP/OD) scenario

The main contribution of this study is an exploration of the data from existing field-tests of naturalistic driving to learn about baseline driving. This is complemented by another study of LTAP/OD crashes from a national crash database, presented at the end. First, we review existing literature in this area.

LITERATURE REVIEW

The following literature review is aimed at surveying the metrics used by previous researchers for LTAP/OD events and identifying gaps in our knowledge that need to be addressed by future work.

Turning Algorithm

Making a left turn involves crossing a “zone of conflict” (ZOC), which is the point where the driver’s path intersects that of oncoming traffic. Shladover et al. [1] provide a model of a typical left turn in which the driver makes an initial decision to turn 2-5 s from the ZOC and finalizes that decision 1-2 s away. The decision itself involves multiple steps, which have been modelled algorithmically [2, 3]. These steps, which include things like “estimate speed and distance of approaching vehicle” and “can driver stop before reaching intersection” correspond to points where driver error could occur [4]. Bougler et al. [5] distilled these sources into four general categories:

1. Failure to judge safe time gaps correctly
2. Failure to judge speeds of closing vehicles
3. Failure to see oncoming vehicle (i.e., “looked but did not see”)
4. Obstruction of driver’s view by opposing vehicle

Although eye-tracking technology may allow some signs of cognitive distraction, such as narrowing of eye scanning and fixation on the road straight ahead [6], the task is highly complex. Likewise, although map data may indicate the presence of visual obstructions that are always present, such as hills, buildings, or sharp turns, more temporary obstructions, such as vegetation, other vehicles, or bright light, would require more sophisticated technology. Then there is the question of whether a driver will even heed a warning for a vehicle behind a visual obstruction until they can see it, at which point it may be too late.

The presence of such risk factors could be used to increase the likelihood of issuing a warning, but it may not be necessary to diagnose these causes specifically: the system could simply warn when the gap into which the vehicle is turning is too short. However, doing so requires accurately identifying when the driver initiates a turn. Brake release and forward acceleration, for example, would be insufficient to distinguish turning from creeping forwards for a better view of traffic. The problem is profound, as in some cases even the driver may not yet know what they are going to do: as noted by Bougler et al. [5], the decision to turn is “fluid and constantly being re-evaluated with an expectation of a change in conditions (e.g., the signal may change from green to amber). Drivers are, in essence, prepared to change their decision and stop at a moment’s notice.”

Aggressive driving raises another issue: someone who is pushing into a small gap, counting on the oncoming vehicle to slow, might regard an alert to be a nuisance, but should it be suppressed for that reason even if the maneuver is risky?

Intersection Approach

The first information available to a warning application about a driver’s intent at an intersection occurs well before the turn is initiated, when the vehicle is still approaching the intersection. Referring to an automatic braking system, Chovan et al. [4] list three factors present during the approach that could be used to identify an impending left turn at an intersection:

1. Turn indicator activation
2. Deceleration with a green light or right of way
3. Merging into a left-turn lane

Typical speed profiles during left turns have been characterized based on observing naturalistic driving [5, 7]. These sometimes reveal consistent patterns, such as these observed by Bougler et al. [5]:

- Turns without stopping and with no lead vehicle included minimum speeds between 11.2 and 15.6 mph.
- Turns where the vehicle slowed to let an oncoming vehicle pass included minimum speeds between 4.5 and 8.9 mph.
- Vehicles generally entered the left-turn lane at 20.1-29.0 mph.
- Drivers constantly reevaluated their decision, prepared to stop if another vehicle appeared or the light changed, with reaction times around 0.4 s.
- In general, driver decisions to turn without stopping or to stop can be distinguished beginning 17 m (or 2 s), from the stop bar (the decision point was therefore reached 0.5 s before that).

Nonetheless, the authors note that “it was very difficult to accurately predict from the [driver’s] approach speeds whether the driver intended [to] turn or stop to let the [oncoming vehicle] clear.”

Although not specific to LTAP/OD, a study found that red-light violations could be predicted by models incorporating range to intersection, speed, and longitudinal acceleration, as well as possibly estimated time of arrival and the required deceleration parameter [8]. This study was based on data from the U.S. DOT Cooperative Intersection Collision Avoidance System for Violations (CICAS-V).

In regards to driver preferences for the timing of warnings during the approach, a study conducted at a controlled intersection on a test track showed little difference between receiving the warning two, three, or four seconds ahead of time (there was a slight preference for three, which was half a second before drivers typically applied the brakes in this instance). Additionally, younger drivers were more amenable to a later warning and older drivers to an earlier warning.

One study determined that a last-second warning allowing enough time to stop should be issued around 14 m from the collision point [9]. Unfortunately, the speed profile distributions of drivers intending to stop and those intending to proceed were indistinguishable at that distance, meaning alerts issued that early would produce a large number of nuisance alerts.

Another key aspect of intersection approach measured by Bougler et al. is the distribution of where drivers tend to stop relative to the stop bar. This was compared between different age and gender groups and appeared to differ by intersections.

Turn Indicators

Turn signals not only provide useful information to the oncoming vehicle [4], but can indicate driver intent to the warning application—in fact they have already been used as a prerequisite condition for several warning applications. The primary difficulty with using turn indicator status is unreliability, since not all drivers use them. Furthermore, when drivers do signal, there is likely a wide variation in how long before a turn and at what distance from the intersection they switch it on. There is consequently a need for researching how drivers use turn indicators, including the timing of when drivers activate their turn indicator.

The only study we found was by Ponziani [10], who himself noted that he could find no other formal research into turn indicator use. In his study, a single observer tallied turn signal use and neglect while driving in an unmarked vehicle around Dayton, OH, aiming for a mix of weather, time of day, day of week, and types of turn, and excluding emergency vehicles, bicycles and mopeds. Indicators were used in roughly 75 percent of 10,000 turns.

Gap Length

Having reached an intersection, a driver wishing to make an LTAP/OD maneuver will typically perform one of three actions [5]:

1. Turn without slowing to below turning speed (when oncoming vehicle is distant or absent)
2. Slow to below turning speed in order to allow an oncoming vehicle to pass and then turn (when oncoming vehicle is close)
3. Stop and turn after the oncoming vehicle has passed

The first action is not an LTAP/OD, strictly speaking, since the oncoming vehicle is not present or close enough to influence the driver's behavior. In the second two, the key factor is the size of the gap in oncoming traffic (generally the time from the first opportunity to turn until the oncoming vehicle reaches the ZOC). Since the driver making a turn in an LTAP/OD scenario generally does not have the right of way, their task boils down to assessing the length of those gaps [11]. For this reason, gap length has received much attention in research and will be the primary focus here as well.

Whether a driver turns into a given gap is not dependent solely upon the length of that gap. For example, drivers will accept shorter gaps the longer they have had to wait (either in terms of the total waiting time [12] or the number of untaken gaps that have passed [13]), or the longer they expect to have to wait after the coming gap [14].

Two studies of gap choice for drivers merging into cross traffic may be relevant: in the first, estimation of gap sizes were affected by “a bias toward underestimation of high speeds and overestimation of lower ones” [15]. In the second study, conducted in a driving simulator, the likelihood of turning into a gap increased the faster the

speed of the oncoming vehicle [16]. These results likely apply to oncoming traffic as well—perhaps even more so given the weaker visual cues used to judge the speed of an approaching vehicle. Knowing that a driver may be more likely to turn into a dangerously short gap under certain conditions could be used to tailor alert thresholds to those conditions.

Rather than basing warnings strictly on the kinematics of what is technically possible in avoiding collisions, alert thresholds could be set to omit psychologically uncomfortable “close calls” as well [4]. Or thresholds could be based simply on the size of gaps that drivers typically choose—Ragland et al. [11] looked at five intersections in the San Francisco Bay area and found available gaps (here defined as the time between oncoming vehicles) was a right-skewed distribution with 30 percent of gaps under two seconds, 18 percent were three seconds long, and 12 percent were 4 seconds (gaps over 12 seconds were omitted). Drivers never turned when the gap was under 3 s and always turned when it was over 12 s. The percentage of times drivers turned increased step-wise for gap lengths between 3 and 12 s: 15 percent turned when the gap was under 4.1 s, 50 percent when the gap was under 6.0 s, and 85 percent when it was under 8.6 s. Given the variation seen between intersections and between drivers, and since they expect to see variation between different genders and between different weather or lighting conditions, the authors suggest tailoring warning applications to these parameters: “parameters of the logistic function could be used to characterize differences across intersections, across conditions within the same intersection, and across different categories of drivers.”

In another field test, mean gaps were $4.3 \text{ s} \pm 3 \text{ s}$ when drivers turned and $0.5 \text{ s} \pm 2.5 \text{ s}$ when they did not [5]. Results from the test track study were more narrowly distributed (as expected given the controlled situation), with only three seconds between zero and 100 percent of drivers turning.

Bougler et al. noted that predicting collisions based on gap length faces three challenges: determining the best way to measure gaps, better justification for threshold settings, accounting for variation in thresholds based on differences between intersections [5]. In a subsequent study, the authors set a threshold at 1.5 s, since they observed less than five percent of drivers to turn with a gap that small at a series of intersections observed via roadside radar [1].

Crossing Behavior

Studies often predict vehicle movement through the intersection as a clean arc, but it is important to model the actual turning radii, speeds, and accelerations, as well as any other variables that drivers use in the real world. Good profiles for turn initiation may allow applications to distinguish, for example, between initiating a turn and creeping into an intersection for a better view.

In the Bougler et al [5] study, turning times, defined as the time from when the front bumper crosses the outside line of a crosswalk until the rear bumper exits the outside line of the next crosswalk, varied between 5.2 and 7.8 s. Some of this variation was due to whether the turning vehicle was accelerating from stopped, to gaps that vary from differences between intersections, which vary in size, the

number of lanes, the number of crossing lanes, whether or not there is a dedicated turn lane, the geometry of the intersection (straight or curved), and whether there are any visual obstructions such as hills or turns. Even potholes can change how a driver will move through an intersection in terms of speed or course.

Environmental Conditions

Environmental conditions can affect intersection behavior and introduce further complexity into an already complex area. For example, visibility can be affected by light levels, fog, rain, or snow. Road surface conditions, such as rain and ice, can affect acceleration and braking times, potentially changing when a warning should be issued or, since many people may underestimate how the conditions affect their vehicle's performance, the level of the response required by the driver to avoid colliding.

Left-turn maneuvers in general require higher-than-normal cognitive workloads and are often made under stress due to time pressure and perhaps the sense that others are waiting on you [4]. In addition to the above effects on vehicle performance, environmental conditions may further increase the driver's stress.

There is little research into how weather affects intersection behavior in particular, although one showed that drivers avoided smaller gaps with increasing amounts of rain [17].

Driver Demographics

Older drivers are more likely to be involved in intersection collisions due to decreased reaction times, visual acuity, agility, ability to handle complex tasks such as interactions with other road users, and make more mistakes in gap selection [18]. Their reduced ability to inhibit the primary task of monitoring the vehicle's travel path may result in checking for other vehicles at intersections less frequently [19]. Older drivers are more likely to get into collisions while making left turns [20], leading to the official advice given by some state DOTs for older drivers to make a series of right turns instead. All in all, a driver is eight percent more likely to crash while making a left turn for every year of age past 65 [21].

That said, older drivers tend to compensate for many of these factors by driving within their abilities, avoiding rush hour, driving at night, and some follow the advice of avoiding left turns where possible and waiting for larger gap sizes. Active simulator training can even compensate for the reduction that older drivers show in looking for threats when making turns [22].

However, given these age-related effects, differences do not show up in all variables, e.g., Bougler et al. found no effect of age on turning time or stopping location relative to the stop bar [5].

Bougler et al. also saw no differences in gender, which is the other demographic factor usually considered. On the other hand, two simulator studies did find gender differences, both with women selecting longer gaps [23, 24].

Ragland et al. propose the most obvious solution which is to use logistic regression models to optimize acceptable gap lengths for subpopulations [11]. Of course, the same could be done for individual intersections or environmental conditions. Another option is to allow drivers to customize their warning applications by setting the alert sensitivity or even selecting for larger gap sizes.

In a theoretical paper, Dabbour & Easa proposed a solution in more detail [25]. Noting that mistakes in judging gap sizes are more common with certain age groups and people with health conditions, they proposed a warning application that takes into account an individual driver's perception-reaction times and acceleration behavior. Drivers would input their age and gender and the system would use that data to calculate turning time to compare to the available gap size.

Driver Distraction and Inattention

Distraction and inattention [26] are primary targets for warning applications, and, as of 2008, distraction was a contributing factor in 21 percent of LTAP/OD collisions at signalized intersections and 29 percent of non-signalized intersections (for the oncoming vehicle, distraction was a factor in only four percent of collisions at both types of traffic control devices).

One difficulty in measuring inattention in LTAP/OD scenarios is that it is more difficult to use eye glances and head turns since the angle of oncoming traffic is often directly ahead (no conspicuous turns left or right). The only study of eye or glance behavior that specifically addresses LTAP/OD scenarios was a comparison between motorcyclists and car drivers [27]. That study found that motorcyclists were less likely to make a last glance in the direction of maximum threat (oncoming traffic)—30 percent of the time (± 29 percent) compared to 48 percent of the time (± 33 percent) for cars.

Cognitive distraction is more difficult to detect since the driver's eyes may remain on the road. As mentioned above, eye-tracking technology is one way to address this [28].

Summary

- There is strong evidence that variables such as intersection approach, the size of gaps when drivers turn, and crossing behavior, vary substantially between individual intersections.
- There is strong evidence that individual trajectories vary substantially due to conditions such as a lead vehicle being present and what traffic signal (e.g., red or green light) the driver encounters.
- Age effects have been seen in gap selection behavior, but do not show up in all variables. They were not seen in turning time or stopping location, for example.
- Gender effects appear in some studies but not others.
- Models of driver behavior need to be optimized to include an array of variables in order to accurately predict driver behavior.

BASELINE LTAP/OD ANALYSIS

This section describes an analysis of LTAP/OD driving aimed at capturing normal or “baseline” behavior. The analysis is aimed at two questions:

1. What characterizes a typical driver turn?
2. What variables affect/indicate a driver's decision to turn?

The data comes from two recent naturalistic driving studies: the Safety Pilot Model Deployment [29] and Driver Adaptation [30]. Both of these studies involved giving volunteers cars with integrated collision-warning systems for everyday use. In both cases, the warning systems utilized vehicle-to-vehicle (V2V) technology. In Safety Pilot, 127 participants had their vehicles for six months in and around Ann Arbor, Michigan. In Driver Adaptation 37 participants had their vehicles in the Washington, DC, area for either three, nine, or twelve months. In both studies the vehicles contained cameras showing views looking both in and out of the vehicle. Additionally, multiple data channels recorded vehicle dynamics (acceleration, speed) and inputs (throttle and brakes). Safety Pilot drivers were spread evenly across different age groups and Driver Adaptation drivers were all under 30. Both studies were balanced for gender.

Methodology

The general approach involved querying the databases for LTAP/OD events by first identifying left turns using steering wheel rotation and vehicle yaw. The resulting events were reviewed by watching the video from the vehicle's onboard cameras using a custom-built data analysis tool that displayed all cameras and selected sensor data on a single screen. Manual review was necessary to exclude turns made with no oncoming vehicles, which included both low-density areas and intersections with protected left-turn lights.

Video review also allowed us to record aspects of the turn not available from the vehicle's sensor data, such as the time when the driver first had an opportunity to turn, when the driver cleared the intersection, and environmental conditions.

Using both the onboard sensor data and the data collected from the video review, we collected the following metrics for analysis:

- Gap lengths into which drivers turned
- Age group
- Gender
- Presence of a visual obstruction
- Distraction status of the driver
- Presence of a dedicated left-turn lane
- Speed, acceleration, and braking data throughout a turn
- Information on the intersection layout
- Timing relative to the oncoming vehicle
- Weather, road conditions, day or night status
- Turn signal use and timing

Gaps into which drivers turned (“go” events) were here defined as the amount of time available for a driver to complete a turn, from the start of the turn until the oncoming vehicle entered the ZOC. Gaps where the driver did not turn (“no go”) were defined as the time from when the driver reached the intersection and was first able to turn (after any vehicles ahead in a queue had gone) until the oncoming vehicle reached the ZOC.

All of these metrics were explored but the primary metric was gap length, since the question, “do I have time to turn?” is at the heart of the driver's dilemma in an LTAP/OD scenario. The question then became what variables affected the choice of gap length and how.

Statistical Analyses

Since this analysis is a *post hoc* exploration in which we identified potentially meaningful results from a larger set of results, the probabilities associated with *p*-values do not apply. Consequently, effect size with 95 percent confidence intervals (CI) were reported instead. Ensuring that these results reflect real patterns and not just anomalies in this data set will require corroboration with future tests.

There are multiple ways to measure effect sizes. In this analysis we used two methods. For comparisons between two groups, we used Cohen's *d*, which describes the magnitude of change in one variable produced by a change in the other. This change in mean is normalized by dividing by the standard deviation (SD). Therefore *d* = 0 means no effect and *d* = 1 means an effect size of one SD. Although values can range from -3 to 3 (at which point 100 percent of one group will be above or below the other group's mean), *d* is usually interpreted with values around 0.2 as small, around 0.5 as medium, and over 0.8 as large [31]. Where possible, results were also described in real-world terms to aid in clarity.

Cohen's *d* can only be used to compare two groups, so where there were more than two levels to a comparison, eta-squared (η^2) was used instead. Unlike *d*, which gives a standardized measure of a difference between means, η^2 is like r^2 in that it measures the proportion of variation in one variable explained by the other. A one-way analysis of variance (ANOVA) test is used to calculate η^2 , and it is loosely interpreted on a scale from zero to one, with 0.02 as small, 0.13 as medium, and 0.26 as large. A downside is that when sample sizes are small, η^2 will slightly overestimate effect sizes. In an exploratory study, though, this bias can be considered an increase in sensitivity since if no effect is detected, the evidence that it does not exist is stronger. Furthermore, η^2 has an advantage over alternatives in that it can be used to compare unequal sample sizes. For η^2 values only, the 90 percent CI was used instead of 95 percent CI since values cannot be less than zero.

In spite of the focus on effect sizes, there were a few instances where standard inferential statistical tests were used. These included preplanned comparisons such as comparing different databases before pooling. In these cases, parametric tests were used once normality was established. When the data were non-normal, non-parametric tests were used and medians and interquartile ranges (IQR) were reported instead of means and SDs.

In the final data set, some drivers only had one event and others had many. To avoid overweighting drivers with multiple events (pseudoreplication) and to preserve assumptions of independence, events were averaged per driver so that driver, rather than event, was the unit of analysis. Within-subject comparisons were conducted whenever $n \geq 8$ to take advantage of the greater statistical power and since behavior tends to show high inter-subject variation. To avoid throwing out the data from non-matched drivers, plots and tables differentiate between the paired data in the comparison and those that were omitted from the test.

All statistical tests were conducted using RStudio, version 0.99.486 [32], which ran R version 3.2.2 (2015-08-14).

Results

With both the Safety Pilot and Driver Adaptation databases combined, a total of 220 LTAP/OD events were found. This included 83 events where the driver did not turn (“no go”) and 137 events where the driver did (“go”) (Table 1). These events included 66 different drivers. The distribution of events per driver was uneven, ranging from one to 43 events per driver, but with a median of two events (Figure 2). Of the drivers, 54 had “no go” events and 36 had “go” events.




Figure 2. Average number of events per driver

Table 1. Number of events and drivers

	Safety Pilot		Driver Adaptation		Total	
	Events	Drivers	Events	Drivers	Events	Drivers
Go	47	31	36	5	83	36
No go	91	46	46	8	137	54
Total	138	56	82	10	220	66

Gap Lengths

There was no statistically significant difference between the average gap lengths between the two databases (independent t -test, $t = -0.37$, $df = 34$, $p = 0.72$, two tailed) so the data were pooled for all subsequent analyses.

As expected, the average length of gaps into which drivers turned (5.0 s) was longer than those into which they did not turn (3.3 s) (Table 2, in which “WS” refers to the within-subject comparison data). This is important because any inferences based on differences between gap lengths would be invalid if we could not demonstrate a relationship between gap length and whether or not a driver turned. Comparing within subjects, Cohen’s d for the difference between the groups was 1.3 (95

percent CI: 0.62 to 1.93)-a large effect. In other words, 90 percent of gaps where drivers turned were longer than the average gap where drivers did not turn. Since not all drivers had both go and no-go events to compare, these statistics are for a subset of the total data. Figure 3 shows the distributions for both all drivers (the dark histograms in the background) and for the subset of drivers with both go and no-go events (the transparent, white histograms in the foreground, representing the subset of data used in the within-subject comparison). All events are averaged by driver. It should be repeated, however, that the methodologies for measuring gaps for “go” (from turn start) and “no go” (from first opportunity to turn) events were slightly different, which lessens the precision of this difference. Nonetheless, the observed differences in gap sizes were large enough that this likely had a small effect.




Figure 3. Length of gaps when drivers did and did not turn

Table 2. Length of gaps when drivers did and did not turn

	Mean	SD	Min	Max	n
Go	5.0	0.8	3.6	6.6	36
No go	3.3	1.3	0.8	6.3	54
Go (WS)	5.0	0.8	3.6	6.6	24
No go (WS)	3.3	1.2	0.8	6.2	24

Another way to show these results is with a cumulative density distribution showing the percent of drivers who either have or have not turned by a certain length of gap. In Figure 4, you can see that 80 percent of drivers turned by the time the gap was increased to 5.8 s.




Figure 4. Cumulative density percentiles by gap length for drivers who turned