Comparison of Rectangular and Elliptical Alert Limits for Lane-Keeping Applications

Olivier N. Kigotho and Jason H. Rife, *Tufts University*

BIOGRAPHY

Olivier Ng'weno Kigotho is a staff researcher in the Department of Mechanical Engineering at Tufts University in Medford, Massachusetts. He supports the Automated Systems and Robotics (ASAR) Laboratory, which applies theory and experiment to characterize integrity of autonomous vehicle systems. He received his B.S. in Mechanical Engineering at Tufts University.

Jason Rife is Professor and Chair of the Department of Mechanical Engineering at Tufts University in Medford, Massachusetts, where he directs the ASAR Laboratory. He received his B.S. in Mechanical and Aerospace Engineering from Cornell University and his M.S. and Ph.D. in Mechanical Engineering from Stanford University.

ABSTRACT

This paper introduces an elliptical alert limit (or AL) to characterize horizontal positioning errors for lane-keeping applications. Equations are derived that characterize the maximum dimensions of the AL when traveling on curved roads. On curved roads, the allowable errors in the lateral and longitudinal directions are coupled, so lane-keeping requires both a limit on lateral errors and on longitudinal errors (either or both of which can carry an automated vehicle outside its lane, possibly resulting in a hazardous collision). We apply our equations for characterizing elliptical AL to relevant scenarios involving vehicles of different sizes on representative roadways; in each case, we compare the performance of the elliptical AL to that of a rectangular AL, as has been considered previously by other investigators. Our analyses indicate that the elliptical AL always outperforms the rectangular AL, allowing for larger longitudinal errors, by as much as a factor of two for passenger vehicles on narrow roadways when lateral errors are tightly constrained.

INTRODUCTION

This paper focuses on safety for automobile navigation systems, with a particular emphasis on automated lane-keeping. For lane-keeping on straight roads, the largest tolerable lateral-positioning errors are dictated by lane and vehicle width. For curved roads, errors must be modeled in two dimensions, with the shape of the road coupling the tolerable error levels in the lateral and longitudinal direction. Prior work by Reid *et al.* [1] showed that road curvature reduces the allowable lateral error, with less tolerance for lateral uncertainty when longitudinal uncertainty increases. This paper refines the analysis of Reid *et al*, deriving tighter (less overly conservative) error bounds that may enhance availability in some situations.

To provide context for availability, it is important to note that an ideal navigation system would always be usable. For safety-critical applications, however, it may be necessary to flag the navigation data as unsafe in situations where the probability of a large error is unacceptably high. A screening process can be introduced to signal an alarm if the probability of a large error is too high. This screening process ensures integrity, but at the expense of availability, which is the fraction of the time that the system meets its safety requirements and is therefore usable.

To implement a screen, it is necessary to characterize the errors in the position estimates that the navigation system generates. Positioning errors are nonstationary, meaning they vary in time or space. In GNSS positioning, for instance, errors vary temporally as satellites move across the sky and spatially as the GNSS receiver moves closer to tall buildings, which may block satellites or introduce increased levels of multipath [2], [3]. Even for a stationary error distribution, massive amounts of data must be collected to accurately characterize the probability of rare errors. To limit the amount of data collection required, it is sometimes useful to introduce an approximate error distribution, or *overbound*, that conservatively models the far distribution tails [4]–[7]. A conservative error bound is particularly

valuable for nonstationary errors, since it is sometimes possible to define a stationary overbound that conservatively models a range of distributions for a nonstationary error signal [8]. However, it is always desirable that the overbound be *tight*, meaning that excess conservatism is avoided, because unnecessary over-conservatism causes the screen to activate more often, reducing availability [1]–[3]

A complementary opportunity to enhance availability is to define the AL to be as large as possible, implying the requirements can tolerate larger errors. Our goal is to define a large but safe two-dimensional AL that ensures against lane excursions. Our primary contribution is to show that an elliptical contour can provide a meaningfully larger AL than a rectangular contour, particularly when road curvature is high. We hypothesize that the benefits are especially pronounced when positioning sensors deliver smaller lateral errors than longitudinal ones, as is expected when using a perception sensor like lidar or a camera on a roadway [9]. A series of contours indicating longer longitudinal (along-track) than lateral (cross-track) errors is illustrated in Fig. 1 for a vehicle rounding a curve.

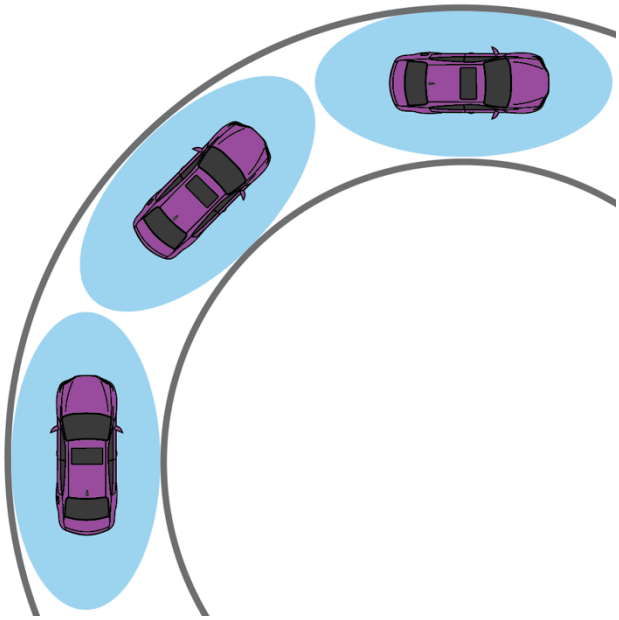

Fig. 1. Roadways tend to be rich in features that have a well defined lateral position but that are not well defined in the longitudinal direction. Examples include lane lines, curbs, barriers, and building facades. As a result, perception sensors like vision and lidar systems produce smaller errors in the lateral direction than the longitudinal direction, as illustrated by the blue error ellipses in the above diagram. As a vehicle progresses through a turn, the positioning-error distribution will tend to rotate to align with the road direction, since the perceptions sensors continually "knock down" lateral errors.

To compare the differences between rectangular and elliptical error bounds, the next section focuses on characterizing the horizontal positioning requirements for safe lane-keeping. The following sections characterize the geometric parameters related to road and vehicle shape and define equations to relate those parameters to maximum tolerable positioning-errors in the lateral and longitudinal directions. An ensuing analysis section explores how these equations must be adjusted to account for heading uncertainty. A final section presents simulation results and discusses the potential impact of using elliptical rather than rectangular AL contours.

HORIZONTAL REQUIREMENTS

The requirements on navigation-system errors are derived from safety requirements for vehicle operations. In this paper, we focus on the lane-keeping operation, and we assume that a potential hazard occurs if the vehicle ever leaves its current road lane. Navigation system errors that cause a lane excursion should be exceedingly rare. In constructing a safety case for an automated vehicle, this requirement can be quantified in terms of a loss-of-integrity probability for lateral-lane excursions P_{lle}. The automated vehicle's navigation data is safe if the probability that a navigation-system error causes a lane excursion is less than P_{lle}. Given an appropriate error-distribution model, an integrity screen can be implemented to check continually that the following condition is met, and to trigger an alarm if otherwise:

$$P\{Navigation\ error\ causes\ lane\ excursion\} < P_{lle} \quad (1)$$

This condition is defined in terms of both an error magnitude (one that causes an excursion) and the probability P_{lle}.

Following the terminology common in aviation applications [5], [10], [11], we will refer to the largest tolerable error as an *alert limit* (or AL). The AL is defined from the operational safety requirements. In the case of a lane excursion, the largest tolerable error is the distance between the road-lane boundary and the vehicle. As shown in Fig. 2a, the lateral distance between the vehicle and the lane boundary is typically the dominant (most constraining) direction when analyzing lane excursions. For straight-line driving, the AL might be quantified as the perpendicular distance from the lane boundary to the nearest point on the vehicle. For straight roads, the lateral alert limit AL_{lat} must be set smaller than half the road width w, accounting for the vehicle width w_v, and the feedback/tracking error FTE:

$$\mathrm{AL}_{lat} < \left(\frac{1}{2}(w - w_v) - \mathrm{FTE}\right) \quad (2)$$

Here the FTE describes an upper bound on the allowed tracking error, noting that vehicle controllers (human drivers or automated feedback systems) do not perfectly match the sensed vehicle position to the reference trajectory, due to limitations of system dynamics and disturbances imposed by the environment.

We could also rewrite the above equation in terms of an expanded vehicle-boundary, where the car width is extended in both lateral directions by the worst possible control and navigation errors. With the expanded vehicle-boundary defined to be $w_e = w_v + 2(\mathrm{AL}_{lat} + \mathrm{FTE})$, the safety criterion (2) could be written in the following alternative form:

$$w_e < w \ \rightarrow \ No\ lateral\ lane\ excursion \quad (3)$$

For lane-tracking on a curved road, a more general two-dimensional analysis is necessary. In work by Reid *et al.* [1], the AL was derived by starting with a rectangular bounding box around the vehicle and expanding the bounding box by margin in both the lateral and longitudinal directions. The result of expanding the original bounding box to account for error margins we will label in this paper as an *expanded bounding box* (EBB). This concept is illustrated in Fig. 2b, where the original bounding box – a rectangle that fully contains the vehicle's projection onto the horizontal plane – is shown in gray. The EBB – which describes the vehicle contour expanded to account for positioning errors – is shown as a red dashed line. Following the approach of Reid *et al*, we analyze the most optimistic case, where the FTE is zero, such that all margin can be assigned to the navigation-system error. In a more refined future analysis, the contour expansion might involve two steps, a first step to expand the contour to account for tracking errors, and a second step expanding the contour to account for navigation errors.

In this EBB framework, the AL is a horizontal area rectangle defined by the lateral and longitudinal error limits. In mathematical terms, another way to represent the EBB is as a dilation of the original bounding box area A_{BB} with the horizontal alert limit area A_{AL}. The dilation operation is used in processing binary images, to expand the contours of a two-dimensional region by a kernel [12]. The dilation outputs all of the points under the kernel as the kernel center is dragged over all points in the original region. Invoking these concepts, the expanded bounding-box A_{EBB} is created by dilating the original bounding box A_{BB} with the kernel, which is the alert-limit A_{AL}:

$$A_{EBB} = \mathrm{dilate}(A_{BB}, A_{AL}) \quad (4)$$

To ensure safety, the expanded bounding box must fit inside the road lane area A_R. We can write this condition in a form analogous to (3) as:

$$A_{EBB} \subset A_R \ \rightarrow \ No\ lateral\ lane\ excursion \quad (5)$$

Invoking dilation allows us to easily generalize the shape of the alert-limit A_{AL}. In the case of Fig. 2b, the EBB area A_{EBB} (dashed red rectangle) was generated by the original bounding box (gray rectangle) with a rectangular alert limit A_{AL}, where the rectangle's width is the tolerable lateral error and its length, the tolerable longitudinal error. To maximize availability, it is desirable to maximize the size of the tolerable errors, so the figure shows an EBB tha touches the inner and outer road-lane boundaries. The aspect ratio of the ellipse is not fixed; rather, the tolerable longitudinal error can be increased by decreasing the tolerable lateral error (and vice versa). Importantly, the vehicle

bounding box is not assumed to be centered in the lane; the reference point can be shifted laterally to maximize the size of the error kernel that satisfies constraint (5).

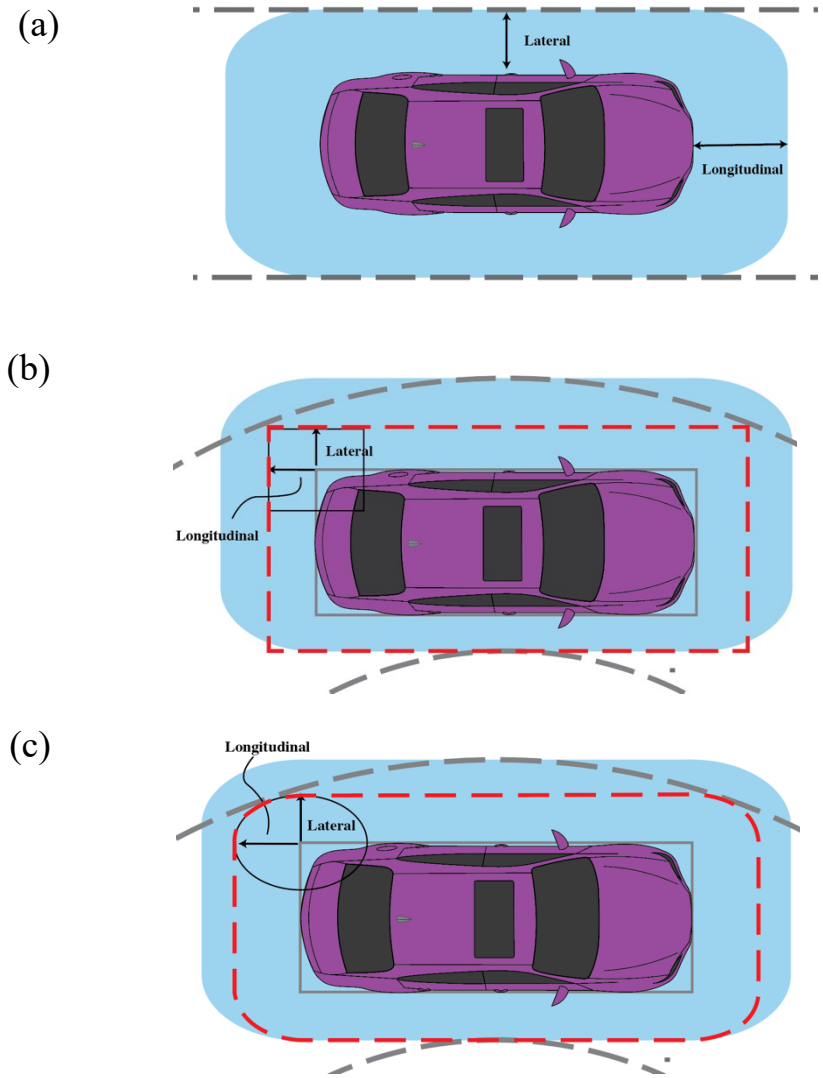

Fig. 2. Positioning Requirements for an Autonomous Vehicle. On a straight road, as shown in (a), the lateral and longitudinal AL are decoupled. However, on a curved road, there is a relationship between the lateral and longitudinal directions as shown in (b) and (c). In (b) the rectangular AL used in [1] is depicted and in (c) our elliptical AL, which provides a tighter bound, is shown.

Instead of defining A_{AL} to be a rectangle, it could instead be an ellipse. For an elliptical A_{AL}, the EBB is a rounded rectangle, as shown in Fig. 2c by the red-dashed line. The ellipse in this case is again defined by two parameters, a maximum tolerable lateral error and a maximum tolerable longitudinal error. One of the major goals of this paper is to determine if the elliptical AL offers an advantage over the sharp-cornered rectangular AL.

QUANTIFYING ROAD AND VEHICLE GEOMETRY

This section parameterizes the areas used in (4) and (5), specifically, the geometry of the vehicle bounding box A_{BB} and the geometry of the roadway A_R. In quantifying vehicle and road geometry, we appeal to the approach of [1]. This earlier work offered standard dimensions for different classes of vehicle and for different classes of road, which are repeated in the two tables below.

Standard vehicle dimensions are summarized in Table 1. These include a range of vehicles from midsize sedans to 6-wheel pickups. Not all vehicles are designed to be driven on all roads. For instance, semi-trailers are not designed to negotiate narrow roads. Therefore in this analysis, as in [1], we consider only passenger vehicles.

Table 1: Standard Passenger Vehicle Dimensions [13], [14]

Vehicle Type	Length [m]	Width [m]	Height [m]
Mid-size	4.87	1.85	1.48
Full-size	5.15	1.94	1.54
Standard pickup	5.32	2.03	2.06
US standard passenger	5.8	2.1	1.3
6-Wheel pickup	6.76	2.43	2.06

Typical road dimensions are parameterized in Table 2. The design of road geometry is usually dependent on the intended vehicle speed [14]. The speed and curvature in a corner can be evaluated from the road's side friction factor and superelevation [14], both of which provide centripetal force for the car to carve its turn.

Table 2: Passenger Vehicle Road Design Geometry Based on [14]–[16]

Road Type	Speed [km/h]	Width [m]	Radius of Curvature [m]
Freeway	80 – 130	3.6	195
Interchanges	30 – 110	3.6 – 5.4	150 – 15
Arterial	50 – 100	3.3 – 3.6	70
Collector	50	3.0 – 3.6	70
Local	20 – 50	2.7 – 3.6	10
Hairpin/ Cul-de-Sac	< 20	6.0	7
Single Lane Roundabout	< 20	4.3	11

METHODOLOGY: LARGEST TOLERABLE ELLIPSE

To complement the equations for the largest tolerable rectangular AL, as provided in [1], this section provides equations for the largest tolerable elliptical AL. Our approach dilates the vehicle bounding box, as quantified by Table 1, until at multiple points the EBB becomes tangent to the road-lane boundary, as quantified by Table 2. The ellipse used to dilate the bounding box is characterized by two axes, the principal axes in the lateral and longitudinal directions, Δ_{lat} and Δ_{long}, respectively. These principal axes are assumed to be aligned with the road lateral and longitudinal directions.

A diagram of how the road geometry determines the maximum tolerable lateral and longitudinal dimensions of A_{AL} is shown in Fig. 3. The alert-limit ellipse, with principal axes Δ_{lat} and Δ_{long}, can be grown until the EBB (blue shaded region) contacts the road-lane boundary (solid black) in three places: near both outer corners and near the midpoint of the vehicle's inner face. At present, the vehicle's orientation is assumed to be aligned with the road; this assumption is relaxed in the next section.

The largest possible ellipse axes Δ_{lat} and Δ_{long} can be obtained by solving for the point where the EBB contacts the road boundary such that the ellipse and road boundary are locally tangent at the point of contact. To define this contact condition, we first define the elliptical bound as the locus of all points $\mathbf{x}_{AL}$ for values of the parameter $\gamma \in [0,2\pi)$.

$$\mathbf{x}_{AL} = \begin{bmatrix} \Delta_{long} \sin\gamma \\ \Delta_{lat} \cos\gamma \end{bmatrix} \tag{6}$$

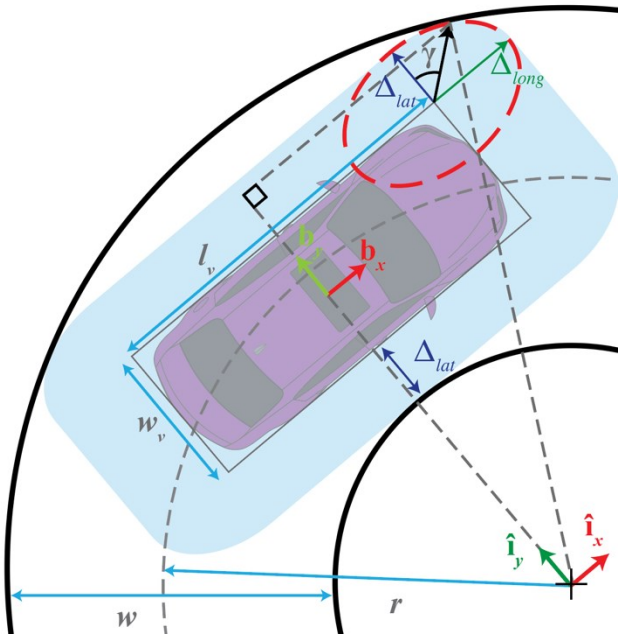

Fig. 3. Geometry of the elliptical AL used to model the uncertainty in the localization of the vehicle. The ellipse shown (red-dashed contour) extends from the outer-front corner of the vehicle's bounding box. The ellipse is just large enough to be tangent to the outer lane-boundary.

The corners of the EBB are formed from quarter-sections of the ellipse (6), translated so its center lies at each corner of the original bounding box. The flat edges of the EBB are defined by expanding the original bounding-box by Δ_{lat} on each side and by Δ_{long} front and back. The distance to the outer front corner of the vehicle from the vehicle's center can be defined by the vector $\mathbf{x}_{OF}$ where

$$\mathbf{x}_{OF} = \begin{bmatrix} l_v/2 \\ w_v/2 \end{bmatrix} \tag{7}$$

For the EBB to touch the inner lane boundary (as shown in Fig. 3), the center of the inner bounding box must be shifted by a distance of Δ_{lat} from the inner lane boundary, which lies at a radius of $r - \frac{w}{2}$ from the road's center of curvature, where r is the radius of curvature of the lane and w, the width of the lane. Tracing from the inner lane-boundary to the car center, we must add an additional radial distance of $\Delta_{lat} + w_v/2$. We can therefore define the vector to the vehicle center from the road center-of-curvature to be $\mathbf{x}_{VC}$, where

$$\mathbf{x}_{VC} = \begin{bmatrix} 0 \\ r - \dfrac{w}{2} + \Delta_{lat} + \dfrac{w_v}{2} \end{bmatrix} \tag{8}$$

The rounded outer-front corner of the expanded bounding box is described by

$$\mathbf{x}_{EBB} = \mathbf{x}_{AL} + \mathbf{x}_{OF} + \mathbf{x}_{VC} \tag{9}$$

for the quarter ellipse with $\gamma \in \left[0, \frac{\pi}{2}\right]$. The two coordinate systems shown in Fig. 3 (vehicle-fixed, road-fixed) are aligned in this case, so there is not need to distinguish between the two in summing vectors in (9). The size of the AL is maximized when the quarter ellipse is tangent to the outer lane-boundary, so we set the distance $\mathbf{x}^*_{EBB}$ (where the star denotes the contact point) equal to the radius of the outer lane-boundary.

$$\|\mathbf{x}^*_{EBB}\| = r + \frac{w}{2} \tag{10}$$

To be tangent at the contact point, the slope of the quarter ellipse (6) at $\mathbf{x}_{EBB}^{*}$ must match the slope of the outer lane boundary. Thus we obtain a second constraint equation for matching slopes:

$$-\frac{\mathbf{x}_{EBB}^{*}\cdot\hat{\mathbf{i}}_x}{\mathbf{x}_{EBB}^{*}\cdot\hat{\mathbf{i}}_y}=-\frac{\frac{d}{\mathrm{d}\gamma}(\mathbf{x}_{AL}\cdot\hat{\mathbf{i}}_y)}{\frac{d}{\mathrm{d}\gamma}(\mathbf{x}_{AL}\cdot\hat{\mathbf{i}}_x)} \tag{11}$$

Together (10) and (11) describe the matched radius and slope at $\mathbf{x}_{EBB}^{*}$; these two equations can be used to constrain the relationship among three unknown variables. The set of unknowns is $\{\Delta_{lat},\Delta_{long},\gamma^{*}\}$, where γ^{*} identifies the value of the ellipse parameter γ at the tangent point. For this problem, it is convenient to consider all of the possible ellipses that we might use, sweeping γ^{*} over its range from 0 to $\pi/2$. One extreme, with $\gamma^{*}=0$, results in the widest possible ellipse, a degenerate case (a line segment) with width Δ_{lat} and length 0. The other extreme, with $\gamma^{*}=\frac{\pi}{2}$, results in the widest possible ellipse, a degenerate case with width 0 and length Δ_{long}. Scanning over the range between the extremes gives us the Pareto optimal trade-off for Δ_{long} as a function of Δ_{lat}, for the case of an elliptical AL. Earlier work by Reid *et al* [1] computed a similar trade-off for a rectangular AL.

By parameterizing over all possible values of γ^{*}, we can now directly solve (10) and (11) to obtain $\{\Delta_{lat},\Delta_{long}\}$ in each case. For numerical efficiency, we expand (11) and use the quadratic formula to solve for Δ_{long}:

$$\Delta_{long}=\frac{-l_v/2+\sqrt{l_v^2/4+4\Delta_{lat}(r+w/2-w+\Delta_{lat}(1+\cos\gamma+w_v)\tan\gamma^{*}\sin\gamma^{*}}}{2\sin\gamma^{*}} \tag{12}$$

Plugging the result into (10) gives the following quadratic equation:

$$\left(\frac{l_v}{4}-\frac{1}{2}\sqrt{\frac{l_v^2}{4}+4\Delta_{lat}(r+\frac{w}{2}-w+\Delta_{lat}(1+\cos\gamma^{*}+w_v)\tan\gamma^{*}\sin\gamma^{*}}\right)^2 + \left(r-\frac{w}{2}+\Delta_{lat}(1+\cos\gamma^{*})+w_v\right)^2-\left(r+\frac{w}{2}\right)^2=0 \tag{13}$$

Discretely sampling over the range of γ^{*} values, (13) can be solved for Δ_{lat} using a numerical solver (e.g. `fzero` in MATLAB). Using this solution for Δ_{lat}, the positive root of (12) can be computed to obtain Δ_{long}.

ATTITUDE SENSITIVITIY

In the prior section, we assumed the basis vectors for the vehicle (labeled $\hat{\mathbf{b}}_x$ and $\hat{\mathbf{b}}_y$) were aligned with the basis vectors for the road (labeled $\hat{\mathbf{i}}_x$ and $\hat{\mathbf{i}}_y$), but this assumption may not be true in general. This section generalizes the elliptical AL calculations from the prior assumption, allowing for the vehicle to be rotated relative to the road.

In this section, we consider only attitude changes due to changes in yaw heading (the ψ coordinate). We could also consider more general attitude changes (including pitch and roll), as were considered in the paper by Reid [1]; however, we reduce the dimensions of the problem by asserting that the lane-excursion requirement applies only in the horizontal direction (and not in the vertical direction). Since we consider only a horizontal AL, vertical motion of the car is not relevant. Moreover, pitch and roll rotations have minimal impact on the horizontal bounding box, since the suspension keeps these angles relatively small during normal driving.

Another key assumption of our analysis is that the error distribution is defined in the frame of the roadway (sometimes called the *navigation frame*). For instance, if a lidar system detects lane markers, the resulting position errors will be smaller in the direction perpendicular to the lane, a coordinate aligned with the roadway direction (e.g. the $\hat{\mathbf{i}}_y$ direction in Fig. 4) and not necessarily with the vehicle (e.g. the $\hat{\mathbf{b}}_y$ direction in Fig. 4).

With these assumptions, we can repeat the analysis of the prior section, considering the rotation of the vehicle relative to the road direction. Although it is not hard to develop an analysis for larger angles, we present the small angle result for compactness. The first step of analysis is to consider that the vehicle center moves slightly outward as the vehicle rotates subject to the condition that the EBB remains tangent along the inner boundary. The change to $\mathbf{x}_{VC}$ can be shown to be negligible for small angles (meaning a small change in yaw ψ between the bases in Fig. 4 causes no appreciable change to $\mathbf{x}_{VC}$). The second step in the analysis is to consider the moment-arm rotation of the vehicle corner about the vehicle center. In this case we consider a positive yaw angle about the upward axis, rotating the vehicle out of alignment with the road. For this rotation, the size of the AL is limited by the tangent condition of the EBB near the outer-front corner of the vehicle, as shown in Fig. 4. The effect of rotation can be computed by introducing a 2D rotation matrix $\mathbf{R}_{\psi}$ into (9):

$$\mathbf{x}_{EBB} = \mathbf{x}_{AL} + \mathbf{R}_{\psi}\mathbf{x}_{OF} + \mathbf{x}_{VC} \tag{14}$$

For a small-angle approximation, the diagonal terms of the rotation matrix have a magnitude of 1 and the off diagonal entries have a magnitude of ψ. By replacing (9) with (14), it is possible to re-evaluate (10) and (11) to account for the impact of a small-angle heading error on the horizontal AL.

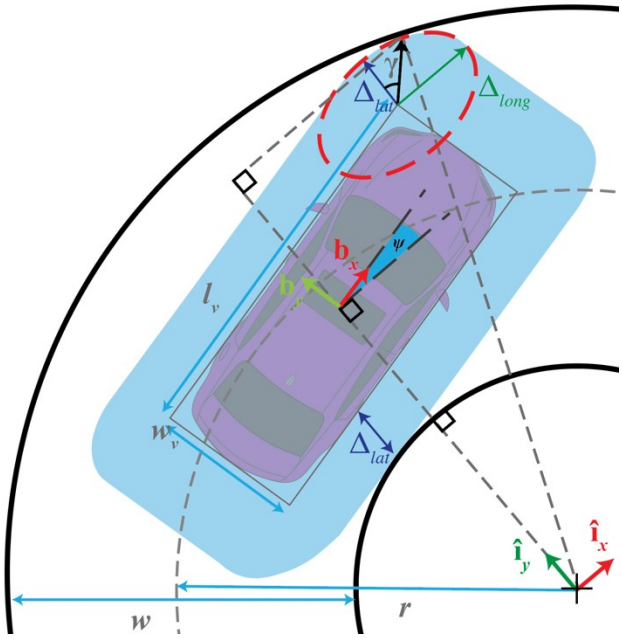

Fig. 4. The effect of an attitude error of ψ on the EBB, which is drawn as tangent to the outer lane-boundary near the outer-front corner of the vehicle..

RESULTS: TRADEOFF OF LATERAL AND LONGITUDINAL DIMENSIONS

From the point of view of performance, generally speaking, larger AL dimensions are better. The idea is that the AL is a specification that determines the worst-allowable error. At any given time, the confidence bound for the worst-possible error should fit inside the AL. This confidence bound, sometimes called a protection level or PL [5], must be assessed at each moment, considering the sensor signals being used and the environmental conditions. If the PL ever grows larger than the AL, an alert must be triggered indicating that the system is not operating safely.

This section compares the dimensions of rectangular and elliptical ALs, to assess their relative dimensions. As part of this analysis, it is important to remember that the lane-excursion conditions described above are characterized by a Pareto optimal surface. In other words, there is no specific "best" value of the pair (Δ_{long}, Δ_{lat}), but rather a trade-off curve describing the best Δ_{long} as a function of Δ_{lat}.

For the purposes of our comparison, the tradeoff for the rectangular AL was computed using the methods of [1], and for the elliptical AL, using equations (10) and (11).

Vehicle Aligned with Roadway

We first consider the case of the vehicle aligned with the roadway, such that the relative heading between the vehicle and the road is $\psi = 0$. For this analysis, the vehicle was assumed to be a US standard passenger vehicle. Six possible roadways were considered (see Table 2), including a freeway, an interchange, two local streets, a roundabout, and a hairpin/cul-de-sac. These assumptions and roadway conditions were chosen to match test cases selected in [1].

Pareto optimal tradeoff curves (Δ_{long} as a function of Δ_{lat}) are plotted for six possible roadways, as shown in Fig. 55. As can be seen in the plots in Fig. 5, there is a modest but clear increase in the longitudinal dimension (vertical axis) for any given lateral dimension (horizontal axis) when using an elliptical AL rather than a rectangular AL.

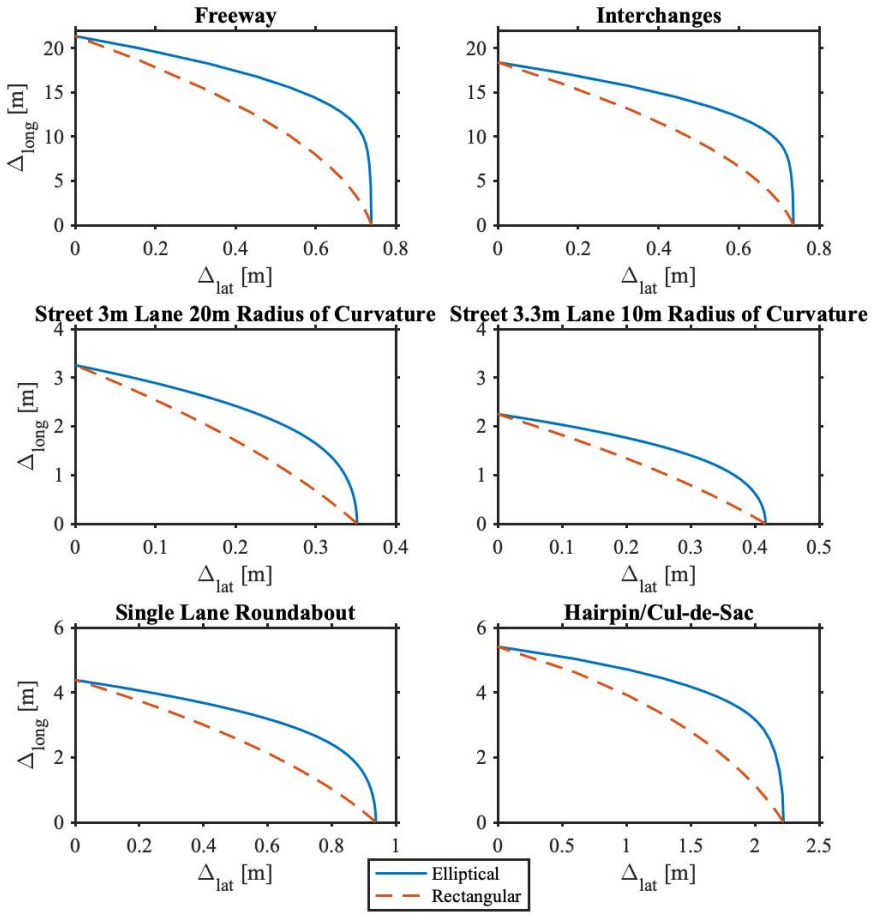

Fig. 5. Tradeoffs between tolerable lateral and longitudinal errors for elliptical a rectangular ALs. The elliptical AL provides larger tolerable errors than the rectangular AL.

As a follow-on analysis, we consider one specific roadway and vary vehicle type. In this second analysis, we choose a road of width of 3 m and radius of curvature of 20 m. We consider four vehicle types listed in Table 1 (mid-size, full-size, standard pickup, US standard passenger). We also consider the special case in which the dimensions of the AL are equal, that is $\Delta_{long} = \Delta_{lat}$, which represents a case where GPS/INS sensors are dominant for navigation, and where perception sensors (e.g. LIDAR or camera sensors that detect lane markers) are not used. The results for this second analysis are shown in Table 3. For this analysis, the benefits of the elliptical (or circular) AL are very modest

relative to a rectangular (or square) AL bounds, with dimensions about 5% larger for the circular as compared to the square AL.

Table 3. A Comparison Between Square and Circular Limits on a 3m Wide Lane With 20m Radius of Curvature

Vehicle Type	Square Limit [m]	Circular Limit[m]
Mid-size	0.48	0.51
Full-size	0.42	0.45
Standard pickup	0.38	0.40
US standard passenger	0.33	0.35

Unequal Lateral and Longitudinal Bounds

Whereas the first two-analyses mirror cases studies by Reid [1], we consider an additional case that is unique to this paper. In particular, we consider that the sensor error distribution is significantly longer in the longitudinal direction than the lateral. This is potentially the case for a navigation system that leverages perception sensors like LIDAR or a camera to detect lane boundaries. In this case, it is expected that the localization in the lateral direction is more precise than in the longitudinal direction because most features in the environment are aligned with the road (e.g. buildings, lane markers), allowing accurate localization perpendicular to those features. In such cases, the nominal error ellipse is expected to be longer in the along-track direction, as shown in Fig. 1. To analyze this case, we introduce a third analysis, identical to the second but constrained to allow for a lateral bound Δ_{lat} = 0.3m. The longitudinal bound Δ_{long} is then a specific value. Considering the four vehicle types of the prior analysis and the same road (3m Wide Lane with 20m radius of curvature), Δ_{long} values were computed and compiled in Table 4. Results for the same four vehicle types on a tighter turn (3.3m Wide Land with 10 m radius of curvature) were also computed and compiled in Table 5. In all cases, there is a substantial improvement in Δ_{long} for the elliptical AL as compared to the rectangular AL.

Table 4. A Comparison Between Rectangular and Elliptical Limits on a 3m Wide Lane With 20m Radius of Curvature

Vehicle Type	Lateral [m]	Rectangular Longitudinal [m]	Elliptical Longitudinal [m]
Mid-size	0.30	2.40	3.39
Full-size	0.30	1.85	2.87
Standard pickup	0.30	1.31	1.98
US standard passenger	0.30	0.68	1.65

Table 5. A Comparison Between Rectangular and Elliptical Limits on a 3.3m Wide Lane With 10m Radius of Curvature

Vehicle Type	Lateral [m]	Rectangular Longitudinal [m]	Elliptical Longitudinal [m]
Mid-size	0.30	1.93	2.53
Full-size	0.30	1.56	2.18
Standard pickup	0.30	1.23	1.69
US standard passenger	0.30	0.79	1.40

Vehicle Rotated Relative to Roadway

It is also relevant to consider the impact of a car orientation error on the size of the AL. To assess orientation sensitivity, we consider the same cases described by Tables 4 and 5, with the only change being the introduction of a heading error. Again, we consider conditions discussed in [1], and so we select a heading error of $\psi = 1.5^{\circ}$ for purpose of comparison. Again Δ_{long} was computed assuming $\Delta_{lat} =$ 0.3m. Values of Δ_{long} were computed for each case and compiled in Tables 6 and 7.

Table 6. A Comparison Between Rectangular and Elliptical Limits on a 3m Wide Lane With 20m Radius of Curvature and 1.5˚ Attitude Error

Vehicle Type	Lateral [m]	Rectangular Longitudinal [m]	Elliptical Longitudinal [m]
Mid-size	0.30	1.39	3.15
Full-size	0.30	0.85	2.58
Standard pickup	0.30	0.33	1.94
US standard passenger	0.30	N/A	0.95

Table 7. A Comparison Between Rectangular and Elliptical Limits on a 3.3m Wide Lane With 10m Radius of Curvature and 1.5˚ Attitude Error

Vehicle Type	Lateral [m]	Rectangular Longitudinal [m]	Elliptical Longitudinal [m]
Mid-size	0.30	1.37	2.40
Full-size	0.30	1.00	2.02
Standard pickup	0.30	0.67	1.58
US standard passenger	0.30	0.22	1.17

In all cases, the size of Δ_{long} was much larger (more favorable) for the elliptical AL as compared to the rectangular AL. In the most extreme case, with the US standard passenger vehicle on the 3m Wide road, the EBB formed with an elliptical AL fit comfortably inside the lane boundaries, but the EBB formed with a rectangular AL did not.

DISCUSSION

This paper shows that in all cases an elliptical AL is an improvement over a rectangular AL. Admittedly, the benefits are modest for the case of bounds with equal dimensions ($\Delta_{long} = \Delta_{lat}$), with Table 3 showing only approximately a 5% increase in the size of the elliptical AL relative to the rectangular one. The benefits are much stronger when the aspect ratio is not one, however, as shown by Tables 4 through 7. If lateral sensing (e.g. with LIDAR or camera) is accurate enough to achieve a confidence bound (or PL) no wider than 0.3m, then significant margin is available for longitudinal error. In Table 4, for example, there is always at least 40% additional margin in the longitudinal direction for any of the vehicles considered, and as much as 142% additional margin (for the US standard passenger vehicle).

Though we did not conduct a detailed availability analysis, we note that these enhancements are particularly significant because the AL dimensions are about 1m. It is very difficult to use GNSS to achieve integrity when the AL is smaller than 1m. Achieving a 1m AL has been a very challenging target for other safety systems, such as the Joint Precision Approach and Landing System (JPALS) [17]. Even with additional sensors (e.g. perception sensors), it is harder to enhance longitudinal accuracy than lateral accuracy. Doubling the length of the AL from 1m to nearly 2m, as is approximately the case of the standard pickup and US standard passenger vehicles in Tables 4 and 5, has an enormous potential benefit to enhance availability in that current sensors support a 2m AL much more easily than a 1m AL. The potential advantages of an elliptical AL are even more pronounced when heading errors must be tolerated. For a rectangular AL, the sharp corners of a rectangular EBB are problematic. The rounded EBB corners for the elliptical AL, accounts for dramatically larger Δ_{long} in the cases of orientation error (Tables 6 and 7), where Δ_{long} for the elliptical AL has a minimum improvement of 127% over the rectangular AL.

Conducting a detailed availability analysis for realistic driving scenarios is an important topic of future work. Because the availability of the system is dependent on road geometry, an increase in availability implies an increase in the locations where autonomous navigation can work, such that the autonomous vehicle can travel more places. Such studies would confirm the degree to which elliptical bounds could improve the efficiency of the transportation.

.

CONCLUSIONS

This paper quantifies the tolerable errors for lane-keeping on straight and curved roads. Specifically, we characterize the largest tolerable error (or AL) in terms of an ellipse on the horizontal plane. Our approach contrasts with earlier work by Reid *et al.*, which characterized the AL as a rectangle.

Through analyses of various vehicles operating on various roads, we show that the elliptical AL for lane-keeping always outperforms the rectangular AL when operating on a curved roadway. The curved roadway is an interesting case, in that the shape of the lane boundaries couples together the maximum tolerable lateral and longitudinal errors. Whereas the benefits of the elliptical AL are modest when the navigation-error distribution is uniform, the benefits of the elliptical AL are substantial when lateral errors are smaller than longitudinal errors, as might be expected for roadway navigation systems relying heavily on LIDAR or cameras. The larger elliptical AL tolerates larger positioning errors, which in turn offers potential to augment navigation availability. Though availability was not simulated directly, it is notable that elliptical AL bounds of 0.3m in width were roughly twice the length of rectangular bounds on mildly curved local roads.

ACKNOWLEDGEMENTS

We gratefully acknowledge NSF grant CNS-1836942 and DOT contract DTRT57-17-D-30011, Task Order 1 (through SAIC subcontract P010253814), which supported different aspects of this research. Opinions discussed here are those of the authors and do not necessarily represent those of the NSF, the DOT, SAIC, or other affiliated agencies.

REFERENCES

[1] T. G. R. Reid *et al.*, "Localization Requirements for Autonomous Vehicles," *SAE Int. J. Connect. Autom. Veh.*, vol. 2, no. 3, pp. 1–16, 2019.

[2] P. D. Groves, Z. Jiang, M. Rudi, and P. Strode, "A portfolio approach to NLOS and multipath mitigation in dense urban areas," *26th Int. Tech. Meet. Satell. Div. Inst. Navig. ION GNSS 2013*, vol. 4, no. September, pp. 3231–3247, 2013.

[3] P. Xie and M. G. Petovello, "Measuring GNSS Multipath Distributions in Urban Canyon Environments," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 2, pp. 366–377, 2015.

[4] B. DeCleene, "Defining Pseudorange Integrity - Overbounding." pp. 1916–1924, 22-Sep-2000.

[5] J. Rife and S. Pullen, "The Impact of Measurement Biases on Availability for CAT III LAAS," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, pp. 1386–1395, Jun. 2005.

[6] J. Rife, S. Pullen, P. Enge, and B. Pervan, "Paired overbounding for nonideal LAAS and WAAS error distributions," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 4, pp. 1386–1395, 2006.

[7] J. D. Larson, D. Gebre-Egziabher, and J. H. Rife, "Gaussian-Pareto overbounding of DGNSS pseudoranges from CORS," *Navig. J. Inst. Navig.*, vol. 66, no. 1, pp. 139–150, Mar. 2019.

[8] J. H. Rife, "Robust Chi-square Monitor Performance with Noise Covariance of Unknown Aspect Ratio," *Navig. J. Inst. Navig.*, vol. 64, no. 3, pp. 377–389, Sep. 2017.

[9] J. Rife, "Collaborative Vision-Integrated Pseudorange Error Removal: Team-Estimated Differential GNSS Corrections with no Stationary Reference Receiver," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 1, pp. 15–24, 2012.

[10] S.-S. Jan, W. Chan, T. Walter, and P. Enge, "Matlab Simulation Toolset for SBAS Availability Analysis." pp. 2366–2375, 14-Sep-2001.

[11] C. A. Shively and T. T. Hsiao, "Availability Enhancements for CAT IIIb LAAS." pp. 474–489, 25-Jun-2003.

[12] R. Gonzales and R. Woods, *Digital Image Processing*. Addison-Wesley Publishig Company, 1992.

[13] *U.S. Department of Transportation Federal Highway Administration*. Federal Size Regulations for Commercial Motor Vehicles, Washington DC., 2017.

[14] *A Policy on Geometric Design of Highways and Streets*. American Association of State Highway and Transportation Officials, 2001.

[15] *Roundabouts: An Informational Guide*, vol. FHWA-RD-00-067. US Department of Transportation Federal Highway Administration.

[16] *Design Manual*. Washington State Department of Transportation, Olympia, WA., 2011.

[17] J. Rife *et al.*, "Navigation, Interference Suppression, and Fault Monitoring in the Sea-Based Joint Precision Approach and Landing System," *Proc. IEEE*, vol. 96, no. 12, pp. 1958–1975, Dec. 2008.