

An RLS-Based Instantaneous Velocity Estimator for Extended Radar Tracking

Nikhil Bharadwaj Gosala¹ and Xiaoli Meng²

¹ETH Zürich, Switzerland ²Hyundai-Aptiv Autonomous Driving JV, Singapore

Abstract—Radar sensors have become an important part of the perception sensor suite due to their long range and their ability to work in adverse weather conditions. However, several shortcomings such as large amounts of noise and extreme sparsity of the point cloud result in them not being used to their full potential. In this paper, we present a novel Recursive Least Squares (RLS) based approach to estimate the instantaneous velocity of dynamic objects in real-time that is capable of handling large amounts of noise in the input data stream. We also present an end-to-end pipeline to track extended objects in real-time that uses the computed velocity estimates for data association and track initialisation. The approaches are evaluated using several real-world inspired driving scenarios that test the limits of these algorithms. It is also experimentally proven that our approaches run in real-time with frame execution time not exceeding 30ms even in dense traffic scenarios, thus allowing for their direct implementation on autonomous vehicles.

I. INTRODUCTION

Autonomous vehicles employ a multitude of sensors to effectively perceive the environment around them. The radar has become an essential part of this myriad of sensors due to its ability to precisely estimate the position and velocity of various targets up to a very long range, and to work reliably in adverse weather conditions like dense fog and heavy rain which trouble other sensor modalities. At the same time, radar has its own set of shortcomings such as the presence of a large amount of noise from sources such as wheel wells, interference from other radar systems, and the extreme sparsity of the point cloud due to its relatively large angular resolution [1, 2, 3, 4].

Traditionally, point trackers have been used to estimate the position and velocity of objects, and to track them over multiple frames in the radar point cloud. However, owing to the development of high-resolution radars that enable sensors to return multiple detections per target, the use of point trackers results in the generation of incorrect and phantom tracks especially when dealing with long objects where the point object assumption is often violated [5]. Additionally, classical filter-based approaches that assume only one return per object are inapplicable to this new problem domain [6]. To address this issue with point trackers, the traditional point object model has been replaced by the extended object model wherein all points belonging to an object are processed together to determine its position, velocity and extents [5]. Extended Object Tracking (EOT) thus refers to the recursive estimation of the position, velocity, and extents of objects that provide multiple noisy returns per frame, and whose shape is unknown and can vary over time [5].

The extended object problem has usually been solved either by using pre-processing techniques that represent

all points from an object using a single meta-point, or by designing models that explicitly account for multiple returns from target objects [6]. The former set of solutions mainly extract the shape by clustering [7, 8, 9] and by shape fitting [10, 11], whereas the latter approaches work by modelling the shape using spatial distributions [12, 13], by adopting physics-based modelling [14, 15, 16], and by modelling target objects as a set of reflection centres [17, 18]. Although the latter set of approaches do not discard useful information by representing multiple points by a single meta-point and are more accurate than the former, they usually are very restrictive in their application, require a lot of human modelling and engineering, and do not run in real-time. Additionally, the former set of solutions are often not perfect and excessively rely on a fool-proof clustering algorithm, use a hard-association algorithm which limits the association of one cluster to only one track, and use insufficient metrics to associate clusters to tracks.

In this paper we present an end-to-end pipeline to perform EOT that runs in real-time, does not extensively rely on fool-proof clustering, and uses both position and instantaneous velocity to perform the one-to-many track-cluster association. Furthermore, as one of the key contributions of this paper, we also propose a novel Recursive Least Squares (RLS)-based algorithm to estimate the instantaneous velocity of a cluster in real-time that allows for accurate track-cluster association, removes the need to perform computationally expensive non-linear state updates, and allows for the estimation of the true velocity even in frames with a large amount of clutter.

The paper is organised as follows: Section II describes our end-to-end EOT pipeline in detail, Section III presents our novel instantaneous velocity estimation algorithm, Section IV evaluates the performance of our velocity estimation algorithm and our EOT pipeline using real-world scenarios, and the paper is concluded in Section V. A video showcasing the real-world performance of our RLS-based instantaneous velocity estimation algorithm can be found online¹.

II. EXTENDED OBJECT TRACKING PIPELINE

In this section we present our real-time end-to-end pipeline to track dynamic extended objects using data from the radar alone, the measurements of which are composed of the position, the Doppler velocity, and the bearing of the reflection point w.r.t. the sensor.

A. Overall Pipeline

An EOT pipeline enables the estimation of the length and width of tracked objects in addition to parameters like

¹<https://www.youtube.com/watch?v=B5SPXXkZpz8>

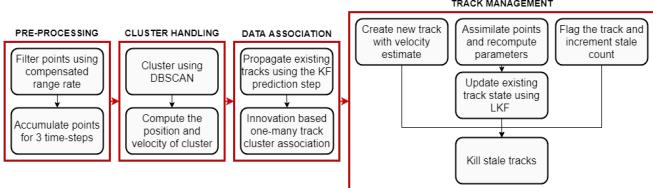


Fig. 1: The overall extended object tracking pipeline. Parts of the pipeline have been grouped together for ease of understanding. Image best viewed on a computer.

position and velocity which allows for the precise representation of the world around an autonomous vehicle, enabling accurate path planning and robust collision avoidance. The tracks generated by this tracker are used in conjunction with those from other sensor modalities to generate very precise final tracks. However, when all the other sensors malfunction, the tracks from radar can be used either to continue operation at a reduced speed, or to allow for the safe stoppage of the vehicle. Our EOT pipeline, depicted in Figure 1, can broadly be classified into four categories namely, *Pre-processing*, *Cluster Handling*, *Data Association*, and *Track Management*, each of which is detailed below.

B. Pre-processing

The pre-processing phase encompasses two phases namely, *Static Point Removal* and *Frame Accumulation*. In the former, only dynamic points, i.e., points having a compensated range rate of at least 0.5 m/s are retained which facilitates the removal of noise and other indiscernible points from the input point cloud. However, this filtration also removes non-noise points with a range rate less than this threshold, but it has experimentally been noted by us that their omission does not affect the performance of the tracker. Furthermore, we increase the density of the radar point cloud by accumulating it over 3 time steps because the extreme sparsity of the point cloud makes detection and tracking of objects from only one time step very challenging. Doing so not only makes detection and tracking easier but also provides a good prior for estimating the instantaneous velocity of clusters, as will be discussed in section III. However, accumulation introduces significant distortion in an object which can skew its position and extent estimates. For instance, when the radar refreshes at 14 Hz and the object's speed is 14 m/s, the distortion over 3 time steps is as large as 3 m. This distortion is corrected using its velocity estimate computed in section II-C.2 as shown in equation (1), and this correction is performed after the clustering phase, detailed in section II-C, is completed.

$${}^{t_1}T_n^{t_2} = -v_n(t_2 - t_1) \quad (1)$$

Here, ${}^{t_1}T_n^{t_2}$ is the transformation that corrects the distortion between time steps t_1 and t_2 in the n^{th} cluster, v_n is its velocity as computed in section II-C.2.

C. Cluster Handling

In the second phase of the pipeline, all points from the accumulated frames are clustered using the DBSCAN algorithm [19], and the position and velocity metrics are

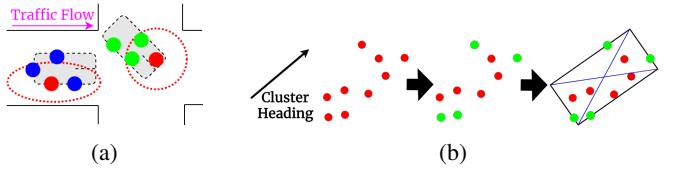


Fig. 2: (a) The elliptical and circular grouping when the point, shown in red, is in an area with a defined and ambiguous traffic flow direction respectively. (b) Steps involved in estimating the centre of a cluster. The extreme points (green) are computed using the cluster, after which a rectangular bounding box is drawn through these points. The centre is then the geometric centre of this bounding box. The arrow shows the heading of the cluster.

computed for each cluster. DBSCAN is a density-based clustering algorithm that groups points that are densely packed together and marks points lying in low density regions as outliers. DBSCAN is preferred to other clustering algorithms because of its adaptability to varying cluster shapes, and its efficient handling of outliers and clutter in the input dataset. To allow for an accurate grouping of vehicles on the road wherein most of them have their lengths aligned to the flow of traffic, we use the available pre-annotated high-definition (HD) map to group points within an elliptical region whose major axis is aligned with the flow of traffic instead of the traditional circular grouping of points. When the flow of traffic is unclear, e.g. in intersections, we revert to using the traditional circular region to group points as shown in Figure 2a. Lastly, as detailed later in section II-D, because we employ a many-one data association algorithm to associate clusters to existing tracks, it is acceptable for the clustering phase to not be accurate and not cluster all returns from an object together. Such an approach also ensures accurate tracking even when the vehicle is not aligned to the traffic flow direction, e.g. when the vehicle is changing lanes or is pulling out onto the road.

Once the grouping is complete, the position and velocity of the clusters are estimated in the world frame and the procedure to compute them is detailed below.

1) *Position*: All points in the cluster are represented by a meta-point that is assumed to be in the geometric centre of the cluster. The geometric centre of the cluster is obtained by computing the centre of the tightest possible rectangular bounding box around all points in the cluster as shown in Figure 2b. The geometric centre is chosen over the centroid because the distribution of points over the target object is generally not uniform resulting in distorted centroid estimates. The position innovation (I_{pos}) is then the squared distance between the predicted position of the track and the geometric centre of the cluster. Mathematically,

$$I_{\text{pos}} = (x_{\text{track}} - x_{\text{cluster}})^2 + (y_{\text{track}} - y_{\text{cluster}})^2 \quad (2)$$

where x_{track} and y_{track} denote the estimated position of the centre of the track, and x_{cluster} and y_{cluster} denote the coordinates of the centre of the cluster.

2) *Velocity*: Unlike position which is directly obtained from the radar sensor, the velocity of the cluster needs to be explicitly computed using the Doppler velocity and the bearing w.r.t. the sensor of each point in the cluster. We use a novel RLS-based algorithm, presented in detail in section III, to compute the instantaneous axis-aligned velocity of the

cluster. The velocity innovation (I_{vel}) is then the squared difference between the velocity of the track and the velocity of the cluster. Mathematically,

$$I_{\text{vel}} = (v_x^{\text{track}} - v_x^{\text{cluster}})^2 + (v_y^{\text{track}} - v_y^{\text{cluster}})^2 \quad (3)$$

where v_x^{track} and v_y^{track} , and v_x^{cluster} and v_y^{cluster} represent the axis-aligned velocity of the track and cluster respectively.

D. Data Association

Once the clustering phase is complete, the next part of the pipeline is to ascertain whether any of the newly computed clusters are point returns from objects that are already being tracked by the tracker. This phase is called *Track-Cluster Association* and is not trivial because we do not know which data points come from which object and vice-versa. We use two parameters - position and velocity - to compute the association between the tracks and the clusters. A cluster is associated to a track if the innovation between the parameters of the track predicted using the constant velocity (CV) motion model and those of the cluster is less than 9 m^2 and $9 \text{ m}^2/\text{s}^2$ respectively, and the sum of innovations of the two parameters is the least amongst all potential tracks. It is worth noting that because we use a one-many data association algorithm, multiple clusters can be associated to the same track which improves the performance in case of long objects and reduces the reliance on the need for an accurate clustering algorithm, as depicted in Figure 3a.

E. Track Management

The track-cluster associations generated from the previous phase are used to either create, update, or delete tracks based on their association history. This section explains how the track state is maintained and how the track-cluster association is handled.

1) *Track State*: The state of a track is defined using a 6-dimensional state vector that keeps track of its position, velocity, and extents. Mathematically, it is represented as $\{x, y, v_x, v_y, l, w\}$, where x and y denote the position of the centre of the track in the XY plane, v_x and v_y denote the axis-aligned velocity of the track, and l and w denote the length and width of the track respectively. The track is maintained using a Linear Kalman Filter (LKF) with the propagation step using a CV motion model to propagate the tracks. An LKF is used because all the state parameters can be updated in a linear fashion directly from the raw or derived measurements and is much more computationally efficient as compared to its non-linear counterparts. Our LKF uses the standard update equations and they have been omitted in the interest of space.

2) *Cluster not associated to any track*: When a cluster is not associated to any track, a new track is created with the state parameters initialised using the corresponding values of the cluster. Since the majority of the moving objects on roads are vehicles, we initialise the l and w to 4 m and 2 m respectively. To prevent noise and other random reflections from interfering with the track list, the new track is initially set to an *invalid* state which becomes valid when the track is associated to at least one cluster for 3 iterations.

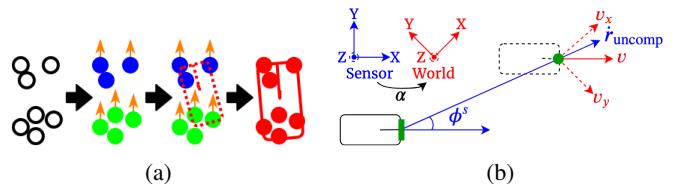


Fig. 3: (a) One-many track-cluster association. All points in the image belong to the same object, but DBSCAN initially clusters the top and bottom sets of points into two different clusters. Their computed velocity vectors are shown in orange. When track-cluster association is performed between the clusters and the propagated track (dashed red rectangle), both clusters are associated to the same track resulting in the accurate grouping of points. The heading and extents of the existing track are also corrected in the measurement update and the updated track is shown using a solid red rectangle. (b) Geometry to compute the axis-aligned velocity of the tracked object. Items in the sensor and world frames are depicted using blue and red respectively, and the rotation between them is denoted by α . ϕ^s is the bearing of the point w.r.t. the radar in the sensor frame, r_{uncomp} is the uncompensated range rate as measured by the radar and does not account for the motion of the ego vehicle, and v_x and v_y are the velocity components of the tracked object in the world frame.

3) *Cluster associated to a track*: A cluster is associated to a track when the innovation of both position and velocity is less than 9 m^2 and $9 \text{ m}^2/\text{s}^2$ respectively, and the sum of their innovations is the least for that track. Because more than one cluster can be associated to the same track, the first step is to assimilate all points belonging to one track from the current time step, and recompute the parameters of this larger cluster. The state update is then performed using the parameters of this updated cluster (see Figure 3a).

4) *Track not associated to any cluster*: When an existing track is not associated to any cluster, no updates are performed and the track is flagged to denote that the track had no associations. If the track is flagged for 5 consecutive time steps, it is assumed to be stale and is deleted from memory.

III. VELOCITY ESTIMATION USING RLS

The instantaneous velocity of a cluster is essential for performing accurate track-cluster association and for assigning a good initial velocity estimate to new tracks so that fast moving vehicles can accurately be tracked by the tracker.

A. Problem Formulation

The velocity of a cluster can be explicitly computed using the Doppler velocity and the bearing w.r.t. the sensor of all points in the cluster. The Doppler velocity defines the rate at which the target is coming towards or moving away from the sensor, and bearing refers to the angle between the ray extending perpendicularly outwards from the sensor and a line pointing directly at the target as shown in Figure 3b. Assuming that all points on the target object have identical velocity vectors, i.e. all objects in the world are rigid, the velocity of a cluster can be computed by solving a linear system of equations as shown in (4).

$$\begin{aligned} \dot{r}_1 &= v_x \cos(\phi_1^w) + v_y \sin(\phi_1^w) \\ \dot{r}_2 &= v_x \cos(\phi_2^w) + v_y \sin(\phi_2^w) \\ &\vdots \\ \dot{r}_n &= v_x \cos(\phi_n^w) + v_y \sin(\phi_n^w) \end{aligned} \quad (4)$$

$$\begin{aligned} \phi^w &= \phi^s + \theta^s + \alpha \\ & \phi^s = \phi^s + \theta^s + \alpha \end{aligned} \quad (5)$$

Here, \dot{r}_n refers to the compensated range rate of the n^{th} point in the cluster, v_x and v_y correspond to the velocity components of the cluster in the world frame, and ϕ_n^w is the bearing of the n^{th} cluster point in the world frame computed using (5). In (5), ϕ^w and ϕ^s refer to the bearing of the point in the world and sensor frames respectively, θ^s refers to the angle of the radar w.r.t. the ego-vehicle, and α refers to the rotation angle between the sensor and world frames.

B. Compensated Range Rate

The compensated range rate, needed for estimating the cluster velocity, can be computed by first computing the range rate of ego-vehicle as experienced by the sensor and then adding the measured range rate as shown in (8). The range rate of the ego-vehicle is in-turn estimated by first computing the motion at the sensor using the known velocity of the ego-vehicle, and then projecting it onto the line joining the sensor and reflection point as shown in (6) and (7).

$$\begin{bmatrix} \omega^s \\ v_x^s \\ v_y^s \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -(y^s - y^e) & 1 & 0 \\ (x^s - x^e) & 0 & 1 \end{bmatrix} \begin{bmatrix} \omega^e \\ v_x^e \\ v_y^e \end{bmatrix} \quad (6)$$

$$\dot{r}_e = v_x^s \cos(\theta^s + \phi^s) + v_y^s \sin(\theta^s + \phi^s) \quad (7)$$

$$\dot{r}_{\text{comp}} = \dot{r}_{\text{meas}} + \dot{r}_e \quad (8)$$

Here v_x^s , v_y^s and ω^s represent the motion as experienced by the sensor, x^s , y^s and x^e , y^e denote the position of the sensor and ego-vehicle respectively, v_x^e , v_y^e and ω^e represent the known motion of the ego vehicle, \dot{r}_e , \dot{r}_{meas} , \dot{r}_{comp} denote the ego-vehicle's range rate, measured range rate and compensated range rate respectively, and θ^s and ϕ^s represent the angle of the sensor w.r.t. the car and the bearing of the point w.r.t. the sensor respectively.

C. Solving the System of Linear Equations using RLS

The system of linear equations shown in (4) is trivial to solve and a solution can be obtained when there are at least two distinct points in the cluster. The challenge, however, arises from the fact that the radar measurements are very noisy and using solutions from trivial methods results in incorrect velocity estimates. A similar set of linear equations was solved by the authors in [20] wherein the outliers are rejected using RANSAC and the velocity is computed by using the Levenberg-Marquardt algorithm. This approach, however, is ineffective in our case due to the sparsity of the point cloud, and because the outlier count is comparable to the inlier count in many cases, rendering RANSAC unusable.

We instead propose a novel RLS-based approach to compute the velocity of the cluster, the pipeline of which is shown in Figure 4. The RLS algorithm is an adaptive filter algorithm that recursively updates the parameters to minimise the weighted least squares cost function. The algorithm starts with an initial estimate of the true velocity which is updated every time a new data point enters the filter. It also exhibits quick convergence and is computationally very fast because it updates the parameters using only the state vector and

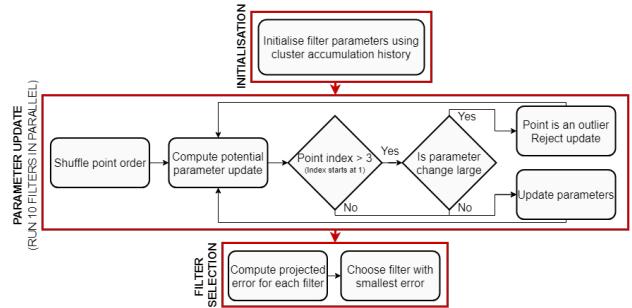


Fig. 4: Pipeline of the proposed RLS-based algorithm to compute the instantaneous velocity of a cluster. Image best viewed on a computer.

the new data point. The update equations are shown in (9) and (10).

$$\mathbf{v}_n = \mathbf{v}_{n-1} + \frac{\mathbf{P}_{n-1}\phi_n}{1 + \phi_n^T} (\dot{r}_n - \phi_n^T \mathbf{v}_{n-1}) \quad (9)$$

$$\mathbf{P}_n = \mathbf{P}_{n-1} - \frac{\mathbf{P}_{n-1}\phi_n\phi_n^T\mathbf{P}_{n-1}}{1 + \phi_n^T\mathbf{P}_{n-1}\phi_n} \quad (10)$$

Here, $\mathbf{v}_n = [\hat{v}_x^n \ \hat{v}_y^n]$ where \hat{v}_x^n and \hat{v}_y^n are the estimated velocity components after using the n^{th} point of the cluster, \mathbf{P} is a positive-definite matrix that can be thought of as the strength of the parameter update, $\phi_n = [\cos(\phi_n^w) \ \sin(\phi_n^w)]$, where ϕ_n^w is the bearing of the n^{th} point in the world frame, and \dot{r}_n is the measured compensated range rate value of that data point.

Literature review suggests that the RLS filter is usually initialised by setting \mathbf{v}_0 to $\mathbf{0}$ and \mathbf{P}_0 to some random positive definite matrix, usually an identity matrix multiplied by a small positive constant [21, 22, 23]. It was, however, observed that this initialisation did not converge to the true value in many cases because 1) the number of points was too small, 2) one component of the velocity vector dominated the other, and 3) large outliers in the data pulled the filter away from the true value. One scheme that works very well is the computation of the initial velocity estimate from the track-history. Since the position estimates from the radar are accurate, the initial velocity can be estimated by computing the change in position and dividing it by the time between the two frames as shown in (11).

$$\mathbf{v}_0^{t_2} = \frac{\mathbf{p}^{t_2} - \mathbf{p}^{t_1}}{t_2 - t_1} \quad (11)$$

Here, $\mathbf{v}_0^{t_2}$ refers to the initial velocity estimate at time step t_2 , \mathbf{p}^{t_2} and \mathbf{p}^{t_1} are the positions of the cluster at time steps t_2 and t_1 respectively. As will be shown in section IV, this estimate is only good for initialisation and is not accurate enough to be directly used for the track-cluster association and track initialisation.

D. Handling Noise in the Data

Each cluster usually contains a large number of outlying compensated range rate measurements which can throw off and cause an incorrect convergence of the filter. Such outliers, which are usually a result of reflections from wheel wells and other moving parts of the vehicle, cannot be easily mapped to different parts of a vehicle and thus cannot be

removed using trivial geometric methods. We, instead, use the update step of the RLS filter to determine whether a point is an outlier and if its filter update should be used. A point is said to be an outlier if its update results in a change greater than a pre-defined threshold - 0.4 m/s in our case - to the v_x and v_y values. This check for outliers is performed only after the first 3 updates of the filter to allow for some inaccuracies in the initialisation and because it has been noted that most of the filters converge to a value that is close enough to the final estimate in the first 3 steps. Furthermore, since the order of points entering the filter affects its final output, especially in cases where the first few updates are performed only by outliers which can result in inliers being classified as noise, 10 filters are initialised for each cluster with each filter being given a randomised order of points as input. To choose the filter with the best parameters, we introduce an error metric called *re-projection error* which computes the sum of errors between the measured and the computed range rates for all points in the cluster marked as inliers. Mathematically,

$$\hat{r}_n = \hat{v}_x \cos(\phi_n^w) + \hat{v}_y \sin(\phi_n^w) \quad (12)$$

$$e_n^{\text{reproj}} = |\hat{r}_n - \dot{r}| \quad (13)$$

where \hat{r}_n is the computed range rate for the n^{th} inlier point using the estimated \hat{v}_x and \hat{v}_y from the filter and ϕ_n^w from the measurement, $|\cdot|$ computes the absolute value, and e_n^{reproj} is the re-projection error for that point. The filter having the least re-projection error amongst the 10 filters wins and its velocity estimate is chosen to be the velocity of the cluster.

IV. EXPERIMENTS

A. Experimental Setup

All the experiments described in this section have been performed using Renault Zoes equipped with multiple 77 GHz long-range Continental ARS 408-21 radars in the front and rear bumpers as shown in Figure 5a. The radar sensor has a positional resolution of 0.4 m with an ability to distinguish targets when they are at least 0.6 m apart, and a velocity resolution of 0.12 m/s. Henceforth, the vehicle performing the tracking is referred to as the *tracker*, and the vehicle being tracked is referred to as the *trackee*. The true velocity of the trackee is obtained by using an IMU in conjunction with wheel speed sensors with an accuracy of 0.1 m/s, and its true position is obtained by using a LiDAR coupled with pre-annotated HD maps and GPS receivers with centimetre accuracy. The performance of various parts of the algorithm is evaluated using three distinct scenarios as shown in Figure 5b. In scenario A, the tracker follows the trackee around a ~ 750 m long test track, and in scenarios B and C the tracker is stationary and the trackee drives towards the tracker, and perpendicular to the tracker respectively.

B. Velocity Estimation

In this section, we evaluate the performance of our RLS-based velocity estimation algorithm using the scenarios described above. Figures 5(c-e), 6(a-c), 6(d-f) qualitatively compare the performance of our algorithm to the velocity

Scenario		CAH	RLS	L+C	OLS	RANSAC
A	μ	0.269	0.038	-0.102	0.301	0.552
	\tilde{x}	0.061	0.014	-0.071	0.053	0.054
	σ^2	4.950	0.834	0.044	1.418	4.010
B	μ	0.458	0.171	-0.075	1.461	0.239
	\tilde{x}	0.601	-0.013	-0.035	0.315	0.025
	σ^2	4.825	0.456	0.101	10.108	2.369
C	μ	-	-0.221	-0.137	-	-
	\tilde{x}	-	-0.137	-0.131	-	-
	σ^2	-	0.780	0.088	-	-

TABLE I: Mean (μ), median (\tilde{x}), and variance (σ^2) of the difference in speed between various algorithms and the ground truth. The smallest and second smallest value in terms of absolute error value are coloured using red and blue respectively. It can be seen that our RLS-based algorithm performs very well and is comparable to the information-rich L+C tracker across all metrics for all scenarios.

from the cluster accumulation history (CAH) and ground truth (GT), and to widely used algorithms like Ordinary Least Squares (OLS), RANSAC followed by OLS, and also to the velocity estimated from LiDAR + Camera (L+C) fusion for scenarios A, B, and C² respectively. Table I quantitatively compares the deviation of all algorithms to GT using metrics like mean, median and variance.

It can be inferred from Figure 5c that the RLS algorithm is able to accurately estimate the velocity of the trackee over extended periods of time. Some deviation can be observed around 32 s when the algorithm underestimates the velocity by around 0.5 m/s which can be attributed to an error in the ego-vehicle motion compensation causing the compensated range rate to be lower than the true value. Additionally, some outlier values observed around 52 s are caused by measurements being included from other tracks due the absence of track IDs in the velocity estimation phase. In addition, the extreme values that appear in the plot can be attributed to the few cases wherein the algorithm is unable to reject all noise points and is said to have failed. Lastly, it is worth noting that the velocity obtained from CAH is extremely noisy, but using it as a seed for the RLS filter results in very accurate velocity estimates. Furthermore, it is evident from Figure 5d that our algorithm outperforms both OLS and RANSAC followed by OLS, which can be verified from Figure 5e and Table I wherein our algorithm has a smaller deviation from GT, and also a lower inter-quantile-distance and variance as compared to OLS and RANSAC followed by OLS. These approaches fail because of the presence of large outliers that are not effectively handled by either of them due to the inherent lack of outlier rejection, sparsity of the point cloud, and a large outlier count. It is also noted that RANSAC performs worse than OLS due to a large outlier count causing RANSAC to choose noise points as inliers which results in multiple incorrect velocity estimates. OLS suffers to a lesser extent because the true inliers pull the velocity estimate closer to GT resulting in lower error metrics. Figures 5d, 5e and Table I show that the velocity computed by our algorithm is comparable to that estimated by the L+C tracker which works by first associating objects over frames using both LiDAR and cameras and then computing the velocity using

²Due to issues with data in scenario C, we compare the performance of our RLS-based algorithm only to the velocities obtained from GT and L+C tracker, both of which were computed online during data collection.

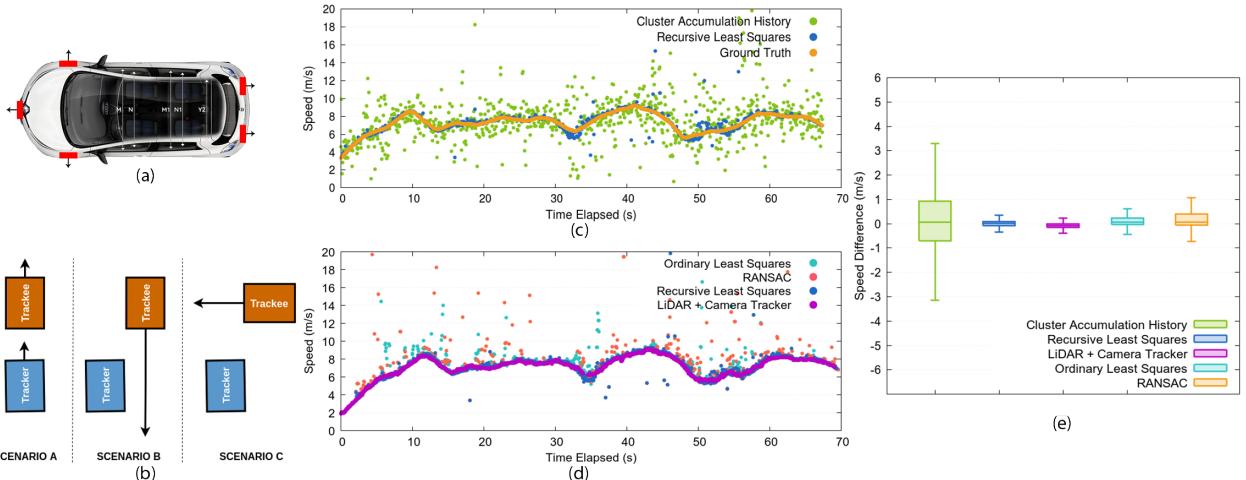


Fig. 5: (a) The 77 GHz long-range Continental ARS 408-21 radars in the front and rear bumpers of the Zoe are shown using red rectangles and the arrows denote the perpendicularly outwards direction w.r.t. the sensor. (b) The three scenarios A, B, C used for the experiments. (c, d) qualitatively compare the velocity estimated by our RLS-based algorithm to that of CAH and GT, and that estimated by OLS, RANSAC, and L+C tracker for scenario A respectively. The deviation from GT for each approach is depicted using box-plots in (e). It is evident that our algorithm is consistent with GT, is comparable in performance to the information-rich L+C tracker and notably outperforms other algorithms like OLS and RANSAC followed by OLS.

an Extended Kalman Filter on the dense LiDAR point cloud. The variance, however, is much larger for our algorithm due to the few instances when all the noise points are not effectively filtered out resulting in the RLS filter computing very large and incorrect velocity estimates.

It can be inferred from Figures 6a and 6d that our algorithm is able to accurately estimate the velocity of the trackee during both the acceleration and deceleration phases. The gap in the output from 9 s to 11 s in Figure 6a is due to the trackee being closer to the tracker than the minimum allowed distance resulting in no measurements from any sensor. The extreme values, like the two points at 7 s, are due to the filter not rejecting all the outliers and can be classified as a failure of our algorithm. Furthermore, the few erroneous velocity estimates between 9 s and 11 s in Figure 6d are due to both the inclusion of the velocity estimates from other clusters, and the trackee being exactly perpendicular to the tracker resulting in valid range rate estimates being filtered out by our pre-processing step due to them being very close to 0 m/s. Lastly, the superior performance of our algorithm is also evident in Figures 6b and 6c wherein OLS fails due to the lack of noise handling capabilities. RANSAC, on the other hand, is able to reject most of the outliers resulting in its performance being comparable to that of ours in scenario B. Furthermore, it can be observed that in scenarios B and C, the velocity from the L+C tracker is better than that estimated by our algorithm because of the improved position estimates between frames due to the presence of a larger “width” view of objects in the point cloud. The radar, however, does not benefit from this because the velocity is computed directly using the measured Doppler velocity of the points on the object. Nevertheless, as seen in Table I, the performance of our RLS-based algorithm is still very accurate and is numerically comparable to the fusion approach.

C. Tracking Performance

This section evaluates the tracking performance of our EOT using two metrics, i.e., positional tracking accuracy,

and extents tracking accuracy, using scenario A. Figure 7a compares the trackee’s position computed by the tracker to that of GT. It is noted that the tracker is very accurate along straights but overestimates the radius during turns. This deviation can be attributed to the high measurement noise we use to account for the inaccuracies in the radar which results in the update step favouring the propagated state over the new measurements. Because the velocity vector is tangential to the curve, the prediction step pushes the predicted position out of the circle which is only partially corrected by the measurement update resulting in an overestimation of the radius of the curve. Such an overestimation can be addressed by using more sophisticated models like Constant Turn Rate Velocity (CTRV) and so on. Furthermore, the deviations on straights like those observed at (45, -15) and (-65, -45) are due to several erroneous points appearing on the road surface. Because of the proximity of these points to true ones, they are clustered together resulting in an incorrect measurement update and subsequently an incorrect track position. The L+C tracker, in contrast, tracks the trackee very precisely and its prediction almost overlaps with that of GT, which can be attributed to the presence of a very dense and noise-free 3D point cloud from the LiDAR.

Figure 7b compares the extents estimate of the trackee to GT. As previously stated, the length and width of the track are initialised to 4 m and 2 m respectively when the track is created and are updated as part of the measurement update step. The initial sharp drop can be explained by the lack of measurements with range rate greater than 0.5 m/s resulting in the extents being severely underestimated. Once the cars start moving more point pass through the filter which immediately improves the dimension estimates. Lastly, due to the large angular resolution of the radar which limits the number of returns from an object width-wise, the width is always underestimated by our tracker. This can be solved by either adding a bias, using radars with a smaller angular resolution, or using a model that accounts for such a behaviour. The L+C tracker predicts the extents of the

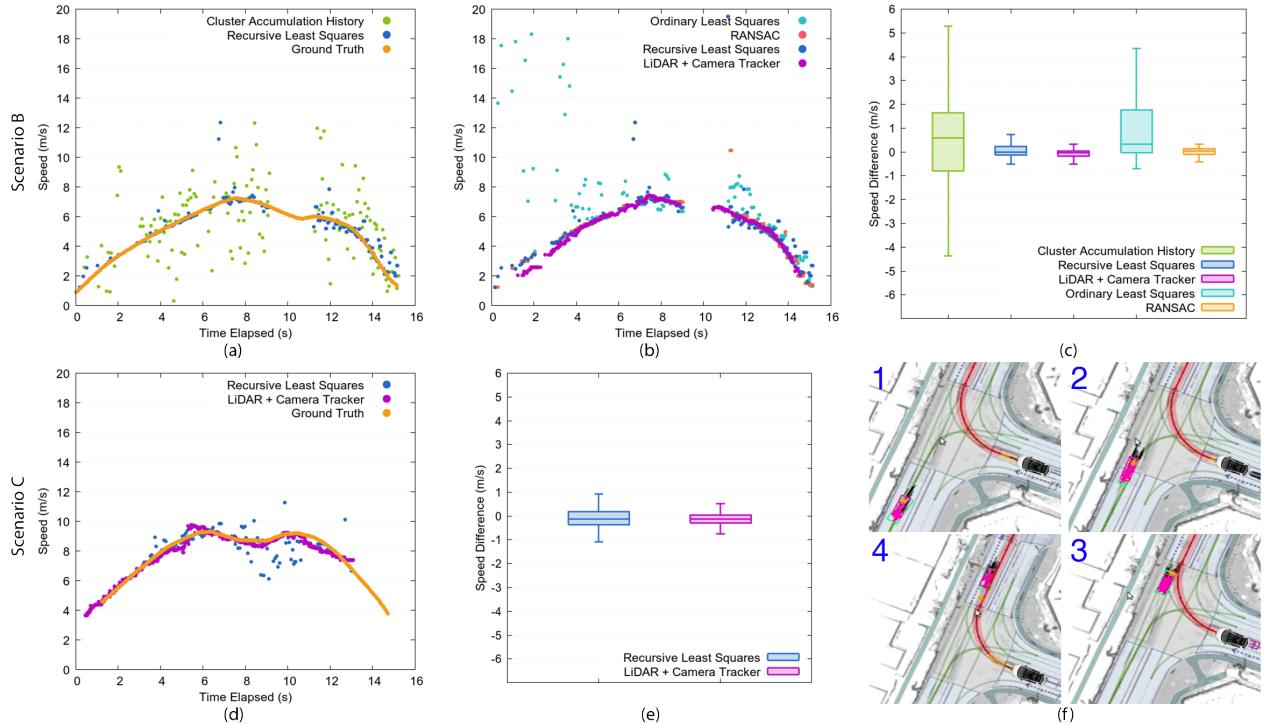


Fig. 6: Comparison of the velocity estimated by our RLS-based algorithm to that of (a) the velocity from CAH and GT and (b) the velocity estimated by OLS, RANSAC followed by OLS, and L+C tracker for scenario B. (d) compares the output of our RLS-based algorithm to that from GT and L+C tracker for scenario C. The deviation from GT for each of the approaches in scenarios B and C is depicted using box-plots in (c, e) respectively. Lastly, (f), in the clockwise direction, qualitatively shows the successful tracking of a dynamic object (pink) moving perpendicular to the ego-vehicle using 4 time-steps of a real-world scenario. Please refer to our accompanying video for more information. Our algorithm agrees with GT in both the scenarios and is significantly better than OLS in Scenario B. RANSAC is able to discard most of the noise points and its performance in scenario B is comparable to that of ours. Lastly, even in the presence of sparse and noisy data, the performance of our algorithm is comparable to that of its information-rich L+C counterpart.

car very accurately due to the presence of an information-rich and noise-free point cloud. The extents are constant throughout because the algorithm does not update the extents once a “good-enough” view of the trackee is obtained.

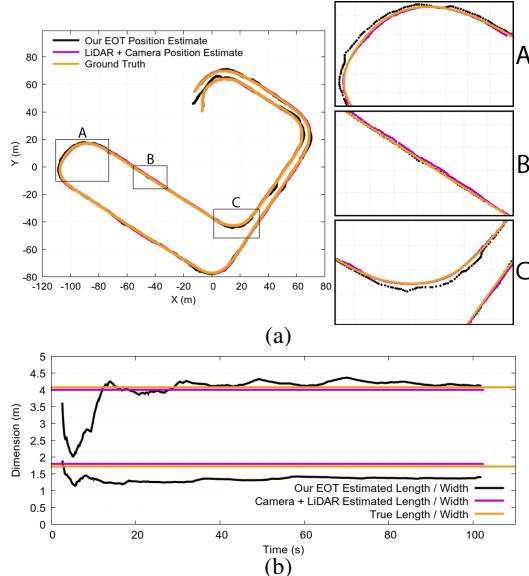


Fig. 7: Tracking is evaluated using two metrics, namely, position and extents. (a) overlays GT of trackee over the tracked estimates from our EOT. Very few deviations from GT are observed proving the positional accuracy of our tracker. (b) compares the estimated extents with the ground truth, from which we infer that although the length is quite accurate, the width is always underestimated due to the dearth of points in the width-wise direction.

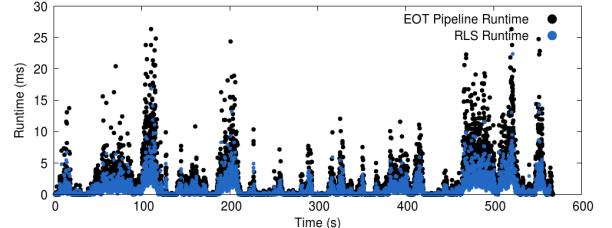


Fig. 8: Runtime of both the EOT pipeline and the RLS-based velocity estimation module. It can be noted that both of them run in real-time, thus allowing for their direct implementation on autonomous vehicles.

D. Runtime

Figure 8 plots the runtime of both the extended tracking pipeline, and the RLS-based velocity estimation module in real-world traffic with the number of objects ranging from 5 to 30 per frame. It is noted that the RLS module runs in real-time averaging 0.89 ms for all clusters in a frame. Similarly, the entire pipeline also runs in real-time averaging 2.1 ms per frame with all frames being computed in less than 30 ms. We can thus safely conclude that our EOT pipeline can run in real-time even in dense traffic scenarios.

V. CONCLUSION

In this paper we presented a novel RLS-based approach to estimate the instantaneous velocity of a vehicle in real-time. The RLS filter uses the compensated range rate and the bearing w.r.t. the sensor of all points in the cluster to solve a linear system of equations to estimate the velocity. Because the range rate estimates are very noisy, an efficient outlier rejection module was built into the filter to separate

the inliers from the outliers. It was shown that our approach performs significantly better than the other widely used approaches such as OLS and RANSAC followed by OLS, and is comparable to the performance of the information-rich LiDAR + Camera tracker. Furthermore, we also presented a pipeline to track extended objects in real-time which uses the computed velocity estimates to initialise new tracks and perform track-cluster association. This tracker runs in real-time and was also shown to be very precise enabling its direct implementation on autonomous vehicles. Some shortcomings of our approach like the overestimation of radius during turns and the underestimation of the width of objects can be resolved either by using more sophisticated models that account for these intricacies, or by using a higher resolution and more accurate radar sensor, and we intend to address them in a future work.

REFERENCES

- [1] J. Wei et al. “Towards a viable autonomous driving research platform”. In: *2013 IEEE Intelligent Vehicles Symposium (IV)*. June 2013, pp. 763–770.
- [2] *Autonomous Vehicle Radar Perception in 360 Degrees*. <https://devblogs.nvidia.com/autonomous-vehicle-radar-perception-in-360-degrees/>. Accessed: 2020-02-01.
- [3] J. Dickmann et al. “Automotive radar the key technology for autonomous driving: From detection and ranging to environmental understanding”. In: *2016 IEEE Radar Conference (RadarConf)*. May 2016, pp. 1–6.
- [4] Ekim Yurtsever et al. *A Survey of Autonomous Driving: Common Practices and Emerging Technologies*. 2019.
- [5] Karl Granstrom, Marcus Baum, and Stephan Reuter. “Extended Object Tracking: Introduction, Overview, and Applications”. In: *Journal of Advances in Information Fusion* 12 (Dec. 2017).
- [6] Alexander Scheel and Klaus Dietmayer. *Tracking Multiple Vehicles Using a Variational Radar Model*. 2018. arXiv: 1711.03799v2 [eess.SP]. URL: <https://arxiv.org/abs/1711.03799>.
- [7] J. Dickmann et al. “Making Bertha See Even More: Radar Contribution”. In: *IEEE Access* 3 (2015), pp. 1233–1247. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2015.2454533.
- [8] Erik Henriksson and Viktor Kardell. “Radar based target tracking for 360-degree environmental perception”. In: 2016.
- [9] J. Elfring, R. Appeldoorn, and M. Kwakkernaak. “Multisensor simultaneous vehicle tracking and shape estimation”. In: *2016 IEEE Intelligent Vehicles Symposium (IV)*. June 2016, pp. 630–635. DOI: 10.1109/IVS.2016.7535453.
- [10] F. Roos et al. “Reliable Orientation Estimation of Vehicles in High-Resolution Radar Images”. In: *IEEE Transactions on Microwave Theory and Techniques* 64.9 (Sept. 2016), pp. 2986–2993. ISSN: 0018-9480. DOI: 10.1109/TMTT.2016.2586476.
- [11] J. Schlichenmaier et al. “Template matching for radar-based orientation and position estimation in automotive scenarios”. In: *2017 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*. Mar. 2017, pp. 95–98. DOI: 10.1109/ICMIM.2017.7918865.
- [12] K. Gilholm and D. Salmond. “Spatial distribution model for tracking extended objects”. In: *IEE Proceedings - Radar, Sonar and Navigation* 152.5 (Oct. 2005), pp. 364–371. ISSN: 1350-2395. DOI: 10.1049/ip-rsn:20045114.
- [13] K. Granstrom and C. Lundquist. “On the use of multiple measurement models for extended target tracking”. In: *Proceedings of the 16th International Conference on Information Fusion*. July 2013, pp. 1534–1541.
- [14] Alexander Scheel et al. “Multi-sensor multi-object tracking of vehicles using high-resolution radars”. In: *2016 IEEE Intelligent Vehicles Symposium (IV)* (2016), pp. 558–565.
- [15] Christina Knill, Alexander Scheel, and Klaus Dietmayer. “A direct scattering model for tracking vehicles with high-resolution radars”. In: June 2016, pp. 298–303. DOI: 10.1109/IVS.2016.7535401.
- [16] L. Hammarstrand, M. Lundgren, and L. Svensson. “Adaptive Radar Sensor Model for Tracking Structured Extended Objects”. In: *IEEE Transactions on Aerospace and Electronic Systems* 48.3 (July 2012), pp. 1975–1995. ISSN: 0018-9251. DOI: 10.1109/TAES.2012.6237574.
- [17] L. Hammarstrand et al. “Extended Object Tracking using a Radar Resolution Model”. In: *IEEE Transactions on Aerospace and Electronic Systems* 48.3 (July 2012), pp. 2371–2386. ISSN: 0018-9251. DOI: 10.1109/TAES.2012.6237597.
- [18] J. Gunnarsson et al. “Tracking vehicles using radar detections”. In: *2007 IEEE Intelligent Vehicles Symposium*. June 2007, pp. 296–302. DOI: 10.1109/IVS.2007.4290130.
- [19] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In:
- [20] D. Kellner et al. “Instantaneous lateral velocity estimation of a vehicle using Doppler radar”. In: *Proceedings of the 16th International Conference on Information Fusion*. July 2013, pp. 877–884.
- [21] Monson H. Hayes. *Statistical Digital Signal Processing and Modeling*. 1st. USA: John Wiley Sons, Inc., 1996. Chap. 9.4, pp. 541–546. ISBN: 0471594318.
- [22] Weifeng Liu, Jose C. Principe, and Simon Haykin. *Kernel Adaptive Filtering: A Comprehensive Introduction*. 1st. Wiley Publishing, 2010. Chap. 4.1, pp. 94–96. ISBN: 0470447532.
- [23] Simon Haykin. *Adaptive Filter Theory (3rd Ed.)* USA: Prentice-Hall, Inc., 1996. Chap. 13, pp. 569–570. ISBN: 013322760X.