

What Goes Around: Leveraging a Constant-curvature Motion Constraint in Radar Odometry

Roberto Aldera¹, Matthew Gadd^{1,†}, Daniele De Martini^{1,†}, Paul Newman¹

Abstract—This paper presents a method that leverages vehicle motion constraints to refine data associations in a point-based radar odometry system. By using the strong prior on how a non-holonomic robot is constrained to move smoothly through its environment, we develop the necessary framework to estimate ego-motion from a single landmark association rather than considering all of these correspondences at once. This allows for informed outlier detection of poor matches that are a dominant source of pose estimate error. By refining the subset of matched landmarks, we see an absolute decrease of 2.15 % (from 4.68 % to 2.53 %) in translational error, approximately halving the error in odometry (reducing by 45.94 %) than when using the full set of correspondences. This contribution is relevant to other point-based odometry implementations that rely on a range sensor and provides a lightweight and interpretable means of incorporating vehicle dynamics for ego-motion estimation.

Index Terms—radar, sensing, radar odometry, ego-motion estimation, motion constraints, field robotics

I. INTRODUCTION

EGO-MOTION estimation using radar-only methods has gained much attention recently in the robotics community. As a sensor that functions consistently in all weather and lighting conditions, radar has become an attractive alternative to traditional vision and laser-based approaches, which are of limited use in poor conditions like fog, rain, snow, direct sunlight and shadows, dust, or darkness. In addition to functioning reliably in diverse environmental conditions, radar has a significantly greater sensing range than its traditional counterparts. This can be of considerable benefit in enabling earlier vehicle and obstacle detection for improved safety or more robust autonomous navigation for mobile robots in environments where stable static features are sparse. For experimental evidence comparing vision and radar in the ego-motion estimation task, see [1, 2].

Building on the foundational Radar Odometry (RO) system developed in [1], this work seeks to leverage a strong prior that previously has not been exploited in this context. Namely, by acknowledging that the robot’s motion can be approximated by an underlying model with certain constraints, we aim to improve the accuracy of ego-motion estimates against some ground truth signal. This is achieved by developing a model

Manuscript received: February 22, 2022; Revised: May 24, 2022; Accepted: June 16, 2022.

This paper was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers’ comments.

We are grateful to our partners at Navtech Radar, the Assuring Autonomy International Programme – a partnership between Lloyd’s Register Foundation and the University of York – and EPSRC Programme Grant “From Sensing to Collaboration” (EP/V000748/1).

¹Authors are with the Mobile Robotics Group (MRG), University of Oxford {roberto,mattgadd,daniele,pnewman}@robots.ox.ac.uk

[†]Equal contribution.

Digital Object Identifier (DOI): see top of this page.

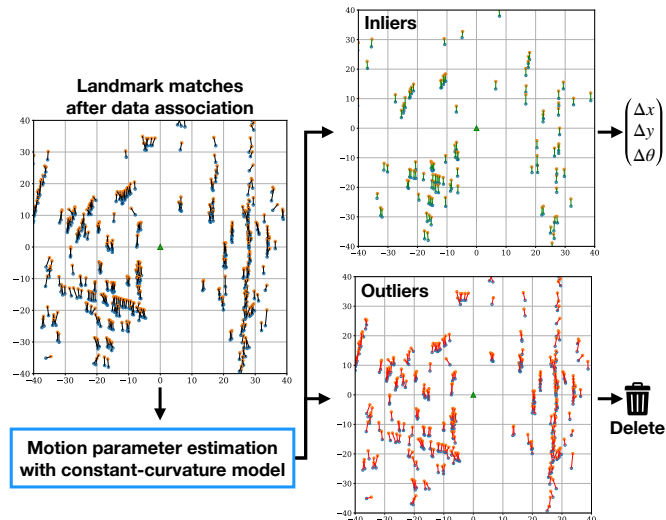


Figure 1. The system estimates motion parameters that result from each landmark match. These parameters provide an indication of match quality and can be classified as part of the inlier set to keep or the outlier set marked for deletion. Odometry accuracy improves when considering only matches corresponding to more plausible ego-motion, as defined by the parameters when under a constant-curvature constraint.

to estimate ego-motion from a single landmark match and expressing this result as a parameter pair. These parameters fully describe an arc of constant curvature traversed to generate this local observation, allowing each landmark association’s odometry estimate to be assessed directly. Once each match is in this form, it becomes easier to identify associations responsible for introducing significant error into the final estimate, which can naturally be refined based on the best available matches as shown in Fig. 1. Examples of these poor matches produced by the baseline RO system are seen in Fig. 2, which provides a clearer zoomed-in view.

Key contributions in this work are as follows:

- The derivation of a model to solve for relative motion in $\text{SE}(2)$ from a single landmark association, as observed with a range-bearing sensor
- Comparison of performance between the proposed method, a simpler baseline, and the original system without this additional refinement
- Validation of the theoretical contributions on real-world data in complex urban environments [3] alongside a comparison with other RO systems

We proceed by reviewing relevant literature in Sec. II and cover background material in Sec. III, before detailing the method in Sec. IV which illustrates how motion estimates are derived from a single landmark observed twice. Sec. V

presents the experimental setup used to validate the theoretical contribution of this work, with subsequent results and conclusions covered in Sec. VI and Sec. VII respectively.

II. RELATED WORK

The richness of radar data itself has meant a relatively diverse set of approaches to ego-motion estimation and localisation has been developed. Our system is based on the RO pipeline of [1] and successive related contributions [4, 5]. Learnt approaches to the odometry task have included an end-to-end method from Barnes et al. [6], which uses the cross-correlation between scans to predict pose after removing distractor features, as well as a point-based method that provides greater interpretability [7]. Burnett et al. looked at the effects of failing to correct errors in landmark positions resulting from motion distortion at higher speeds [8] and later presented a hybridised RO framework [9] that combines probabilistic trajectory estimation with learnt features to benefit from both classical and data-driven approaches.

Other work that has made use of radar as a navigation sensor includes that by Park et al. [10] which applies the Fourier-Mellin Transform to log-polar images computed from downsampled Cartesian images, and Kung et al. [11] which uses a normal distribution transform typically applied to 2D and 3D LiDAR. Adolfsson et al. [12] employ filtering to retain the strongest azimuthal returns and compute a sparse set of oriented surface points, while Hong et al. [2] use vision-based features and graph-matching in a radar context.

However, previous work on radar-only ego-motion estimation has directly attempted to leverage the dynamic constraints of the system in operation. This is a natural consideration, as exploiting the underlying smoothness in real trajectories traversed by sensors mounted on non-holonomic robots has proven useful in a wide range of contexts in the past. Using cubic B-splines, Lovegrove et al. [13] looked at leveraging the torque-minimal motion of their visual-inertial system approximated by a parameterised spline, built using relative poses as control points to describe a continuous pose estimate on an interval between discrete poses. This effectively constrains the underlying motion but is not directly applicable to online ego-motion estimation when future pose estimates are unavailable.

Tong et al. [14] developed Gaussian Process Gauss-Newton as a non-parametric state estimation framework where similarly to [13], the continuous-time reality of robot motion is acknowledged when processing measurements. Although our work does not directly impose a continuous-time framework, we do leverage the continuous nature of our robot’s trajectory when assuming that it moves smoothly between poses.

As a range-based sensor, radar is in some respects similar to 2D LiDAR, making the work in [15] of interest where authors used a non-holonomic extended Kalman Filter motion model to refine laser scan match estimates. Earlier, the authors in [16] proposed a Visual Odometry (VO) system that solved for Ackerman model parameters directly from feature pixel displacement using a ground-facing camera. However, these methods *enforce* a motion constraint to either directly estimate or smooth the odometry instead of *leveraging* the underlying model to isolate poorer matches as a refinement step.

Scaramuzza et al. [17] present 1-point RANSAC which draws on this idea of exploiting a vehicle’s motion constraints in an odometry system and is perhaps the closest work to our contribution, but is only applicable in the vision context. They show how outliers can be discarded in monocular odometry by using a restrictive motion model. However, the details between their work and that which is presented here differ significantly due to the sensor measurement model – radar is a range sensor that measures distances directly, unlike the vision context. We solve this particular problem for range-azimuth sensors, with the full derivation in the appendix. Furthermore, [17] do not use a baseline odometry system to propose their initial correspondences, as we do. Lastly, empirically we found the histogram voting used in [17] performs poorly in this domain and so present a new selection approach using the lower and upper quantiles of the estimated motion parameter. More details of these differences follow in Sec. IV.

III. PRELIMINARIES

We use a Navtech FMCW scanning radar which rotates about its vertical axis and continuously senses the environment using frequency-modulated radio waves. Measurements are captured along an azimuth (α) at one of M discrete angular positions and return a single power reading (a function of reflectivity, size, and orientation of objects) for N range *bins* corresponding to distances, r . We refer to one full rotation through all M azimuths as a *scan* \mathcal{S} . Furthermore, let $\mathbf{s}(k) \in \mathbb{R}^{N \times 1}$ be the power-range readings at time step k , where $t(k) = t_k$ is the time value at k and $\alpha(k) \in A$ is the azimuth associated with the measurement. The element $s_i(k) \in \mathbf{s}(k)$ is the power return at the i -th range bin, with $i \in \{1, \dots, N\}$; its measurement range is given by $r_i(k) = \beta(i - 0.5)$, where β is the range resolution.

Our contribution builds upon the RO pipeline first described in [1], which we summarise here. Firstly, landmark features are extracted from each azimuth’s power-range spectrum captured by the radar as it scans the environment. Although this process is not detailed here, it is well known to be difficult in practice due to the complexity of radar measurements themselves, which may often be obscured by multipath reflections, harmonics, speckle noise, and other hindering effects [18].

To compute odometry from the extracted features in two successive landmark sets \mathbb{L}_{t_k} and $\mathbb{L}_{t_{k-1}}$, landmarks in \mathbb{L}_{t_k} must be associated with a single corresponding landmark in $\mathbb{L}_{t_{k-1}}$. We pose this data association as a pairwise problem as in [19], postulating that the distance from each point to its neighbours as observed in one instance uniquely describes that point in any proximal observation, regardless of the relative translation and rotation between observations. This assumption allows landmarks to be matched to counterparts in any other landmark set containing sufficient overlap. Once unary candidates have created preliminary associations and final matches selected using the largest elements from the principal eigenvector of the pairwise compatibility matrix (see [1, 5] for more detail), a final pose estimate G_{t_{k-1}, t_k} is computed by a Singular Value Decomposition (SVD) algorithm [20]. For a graphical representation of this RO pipeline, please refer to Fig. 5 on page 5 of Cen and Newman [1].

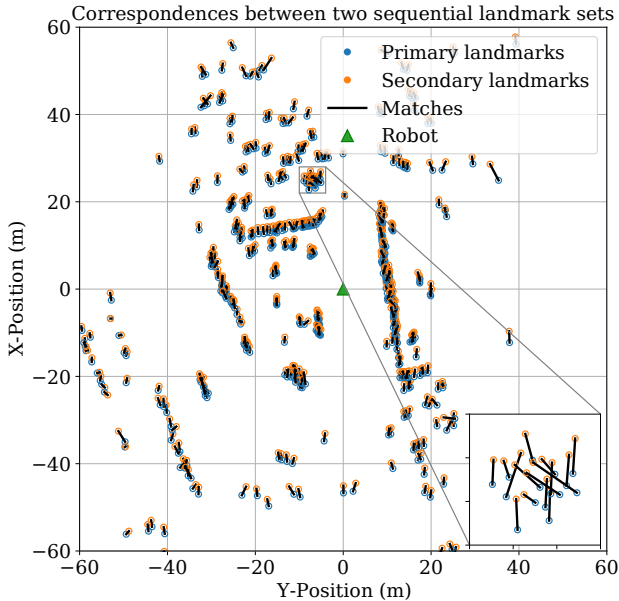


Figure 2. Two successive sets of associated radar landmarks are shown from a scan match, cropped from the full radar range of 165 m to 60 m for clearer visualisation. After observing the orange landmarks, the robot shown at the origin as a green triangle moved straight ahead by roughly 1.4 m and then observed the blue landmarks. Although most matches indicate this motion relatively clearly, there are few stark examples where associated points imply contradictory relative motions, which contribute disproportionately to ego-motion estimate error.

Our contribution interfaces with the system components described above, taking these associated landmarks as inputs. We seek to refine the association and pose estimation process by considering the physical constraints of the robot platform and leveraging this additional information to improve odometry accuracy.

Fig. 2 illustrates the problem we have identified where it is clear that despite the best efforts in landmark extraction and data association, poor matches will persist. Those shown here are from the baseline system after the data association step, which is tasked with selecting the best candidates. In this approach, a point is associated with the most similar available point based on distances and angles to other points (i.e the scene geometry), which does not consider vehicle kinematic constraints. At first it appears as if an outlier rejection scheme like RANSAC would explicitly remove these outliers, however we show in Sec. VI that a standard application of RANSAC only marginally improves ego-motion estimation accuracy, as it also does not exploit domain knowledge of how the vehicle moves. The effect of failing to remove these poor matches is shown in Fig. 3, where implausible lateral motion is estimated due to poor correspondences.

IV. METHOD

This section details approaches to mitigating problematic matches in Fig. 2, outlines the proposed model for leveraging a constant-curvature constraint, and describes pose estimation from a refined subset of candidates. We assume the motion

of our vehicle is described by a non-holonomic kinematic motion model, appropriate to two-wheel drive urban vehicles which traverse local arcs of constant curvature between two successive measurements, detailed in Sec. IV-B.

A. Outlier rejection baseline

Before leveraging a constant-curvature constraint, we first consider a sensible starting approach that provides a suitable baseline. One instinctive way to improve pose estimate accuracy is to identify and discard poor associations. Matches contributing disproportionately to the final pose error can be flagged as outliers by a simple RANSAC [21] implementation before the SVD calculation. By considering landmark set $\mathbb{L}_{t_{k-1}}$ in the reference frame of \mathbb{L}_{t_k} (which amounts to transforming $\mathbb{L}_{t_{k-1}}$ into the same frame using the computed pose estimate G_{t_{k-1},t_k}), associations can be assessed individually. Practically, RANSAC is run by randomly selecting two matched landmarks and computing a pose estimate with the SVD from this minimal subset. All other candidates are then classified as inliers or outliers based on their range-compensated distance away from their associated partner when under the proposed pose estimate. The sampled model that produces the most inliers after a fixed number of iterations yields the best proposal. These inliers are passed to the SVD method for a final pose estimate. This serves as a baseline to compare our method against, which takes a more informed approach that does not use RANSAC.

B. Constant-curvature motion constraints

Inspired by [17], we seek to describe ego-motion using only a single landmark association. However, in contrast to [17], our formulation does not use pixel coordinates and epipolar geometry required in vision, but instead develops the corresponding theory for range sensors that provide range and azimuth readings. To this end, we present a similar yet novel expression in eq. (A.7) in the appendix for the constraint parameter to be estimated, described below.

When matching two landmark sets to compute a relative pose in $\mathbb{SE}(2)$, at least two pairs of associated landmarks need to be available to constrain the solution. With only a single pair, the relative motion cannot be fully resolved as there would be infinitely many poses on a circle around that landmark that satisfy the single constraint. However, by leveraging a constant-curvature constraint to describe the local robot motion between poses, it is possible to resolve relative motion using only one landmark association pair. This model serves as a suitable approximation for the motion of the vehicle in the experimental dataset (see Sec. VI) which was captured in an urban environment where wheel slippage is minimal. Additionally, the short 0.25 s duration between successive scan measurements means that the curvature of the traversed path can be assumed to be constant. Fig. 4 shows two successive robot poses where an arc has been traversed along the circumference of a circle centred at the Instantaneous Centre of Rotation (ICR) of radius R_{ICR} , through an angle of θ . The constant curvature constraint here is applied only between these two poses at which consecutive scans are

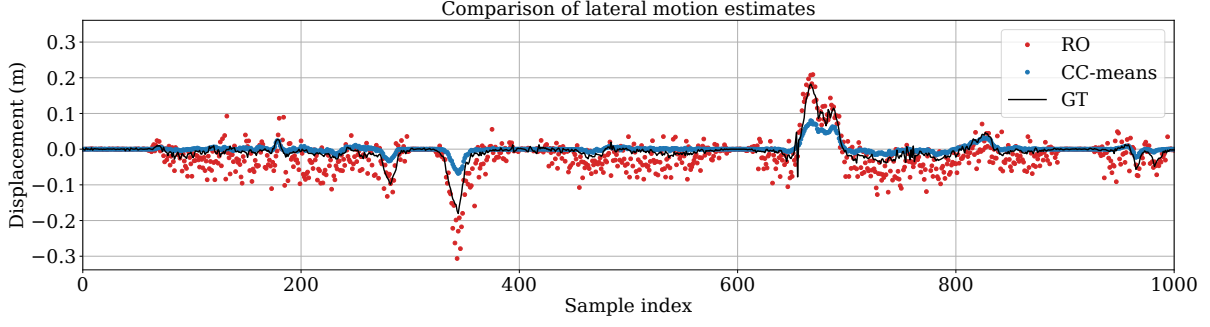


Figure 3. Without leveraging any knowledge of platform constraints, pose estimates from RO in red tend to produce relatively poor performance in lateral motion estimation, particularly evident when driving straight, where estimates deviate by up to as much as 10 cm. Our contribution shown in blue produces smoother estimates in these Δy -values which are a nearer approximation of the ground truth signal in black.




Figure 4. Associated observations of a landmark L from the robot at pose G_{t-1} and then at G_t can be used to describe the local relative motion as a circular arc.

collected, and does not consider a window of scans or poses before this current estimate. The final pose is a distance ρ from the starting position, with a yaw offset of this same θ . Relative to the starting pose, the final pose is at an angle of $\theta/2$ in this case, observing that both poses are at a distance R_{ICR} from the ICR, forming an isosceles triangle containing θ and two other equal angles. For straight trajectories, the robot can be described as moving along a circle of infinite radius.

C. Pose estimation from a single landmark association

A single landmark observed from the starting pose is at a range r_0 and bearing β_0 . At the next instance, once the robot has reached the final pose, that same landmark is observed again, now at a range r_1 and bearing β_1 . These four parameters allow us to express the robot's motion along the arc that subtends θ from the ICR where θ and R_{ICR} are expressed as our constant-curvature motion parameters:

$$\theta = 2 \arctan \left(\frac{r_0/r_1 \sin \beta_0 - \sin \beta_1}{r_0/r_1 \cos \beta_0 + \cos \beta_1} \right) \quad (1)$$

$$R_{ICR} = \frac{r_1 \sin(\beta_0 - \beta_1 - \theta)}{2 \sin(\theta/2) \sin(\theta/2 - \beta_0)} \quad (2)$$




Figure 5. Once a landmark L has been observed initially from pose G_{t-1} to be at a certain range and bearing angle, the task is now to find where G_t is relative to G_{t-1} given a new range and bearing to that same associated landmark from the new pose. This diagram illustrates that although there are infinitely many possible G_t poses that satisfy the second range and angle measurement around a locus, which is a circle centred on L of radius r_1 , only one of them will satisfy the constant-curvature motion constraint too. Notice how for the pose slightly below the correct G_t where a faded proposed circle has been drawn, there is an error ϵ between the point a at which their extended axes meet and the centre of that circle b . These points must coincide if this G_t guess were to fully satisfy both the landmark observation and the constant-curvature motion constraints.

And from these parameters, Δx , and Δy are simply

$$\Delta x = \rho \cos(\theta/2) \quad (3)$$

$$\Delta y = \rho \sin(\theta/2) \quad (4)$$

where $\rho = 2R_{ICR} \sin \theta/2$ and $\Delta \theta$ is the same as its constant-curvature parameter counterpart. The reader is directed to the appendix for a full derivation. Note that, for the edge case when driving in a perfectly straight line, $\theta = 0$ and R_{ICR} is infinite, as mentioned in Sec. IV-B. Although theoretically possible, this tends only to happen in practice when $\beta_0 = \beta_1$ and $r_0 = r_1$ which occurs when the robot is stationary, where we set $\Delta x = \Delta y = \theta = 0$. Fig. 5 attempts to illustrate intuitively how the combination of a single associated landmark and this constant-curvature motion constraint is all that is required to solve for a relative pose estimate.

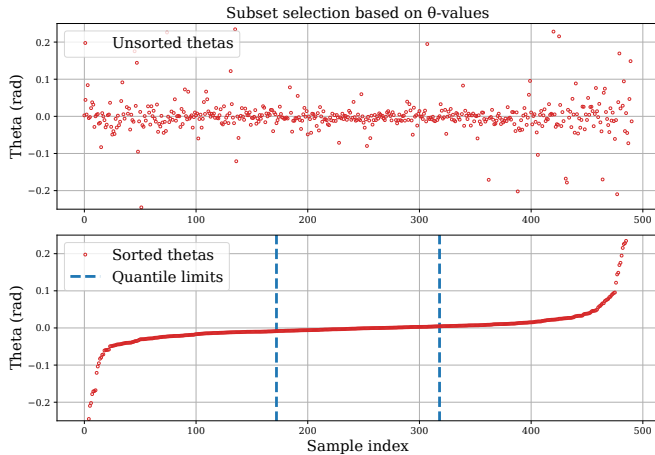


Figure 6. After calculating a θ -value for each individual match between landmark sets from adjacent scans, we can sort them to get an indication on what the majority of matches are estimating. Poor matches tend to produce more extreme θ -values that do not align with the consensus. By taking a subset between two quantiles, we are able to exclude these less reliable (and in some instances more obviously erroneous) matches when computing the final pose estimate.

D. Final pose estimation using a conforming subset

Once each associated landmark match is expressed in terms of θ and R_{ICR} , those that yield θ -values near the median generally correspond to matches that have captured the underlying ego-motion. Associations that produce more extreme θ -values usually correspond to dynamic objects or spurious radar power returns that were similar enough between two observations to pass through the initial data association algorithm. Unlike in [17], we found in this context that simply using the median θ association to calculate the final pose estimate led to inconsistent estimation performance. It appears in [17] that θ -values are discretised, with the middle of the most populated bin taken as the solution. In practice, we found in early experiments that the histogram approach performed poorly in this context. We instead choose to retain a conforming subset and avoid any discretisation errors which may arise. We also experimented with considering the R_{ICR} values when evaluating candidates but found they did not give a strong indication of a candidate’s credibility. Based on these observations, we select a subset of matches when calculating the final pose estimate, which can be done by sorting all θ -values in ascending order and taking an inter-quantile set, where specific quantiles to use are left as a hyperparameter. The final pose estimate is produced by passing every match in the selected subset to the SVD algorithm as before.

Fig. 6 shows how when these θ -values are sorted they tend to form a general consensus, with outliers being pushed to the extremities. However, we found that estimates improve if we instead produced Δx , Δy , and $\Delta \theta$ estimates for individual matches in the subset and selected the mean of each of these parameters as our final pose estimate. Note that although we rely on our motion model to assess the reliability of each landmark association, the final estimate is not forced to conform to these constraints and instead is free to describe more complex and realistic robot motions.

V. EXPERIMENTAL SETUP

To evaluate our proposed system, we utilise the *Oxford Radar RobotCar Dataset* [3] as its optimised ground truth signal is required to assess if RO accuracy has been improved. Our experimental data consists of the first five sequences of the dataset, totalling approximately 50 km of driving. Landmarks are extracted and associated as previously described in [1] prior to our proposed refinement process. The use of this RO framework is convenient as it allows us to build on a familiar system, but we emphasise that our contribution is transferable to other point-based odometry implementations that use measurements from a range sensor.

A. RANSAC baseline

As discussed in Sec. IV-A, we implement a rudimentary RANSAC-based outlier rejection system as a baseline. This provides a reference point in assessing performance when matches have undergone some form of refinement but without leveraging a constant-curvature constraint.

B. Selecting a subset using constant-curvature parameters

When determining which matches to consider in the final pose estimation step described in Sec. IV-D, we found using $\{0.35, 0.65\}$ as the quantile limits yields satisfactory performance for this dataset, slightly more so than using the first and third quartiles. These may need to be adjusted when operating in settings with a different distribution of misleading associations, like in open rural environments where there may be fewer reflections or scenarios with an increased number of dynamic objects in the scene.

C. Processing the selected subset to produce pose estimates

After landmark associations have been reduced to a conforming subset, Sec. IV-D described two methods to compute a final pose estimate. Final poses are estimated using both the SVD method as well as the simpler approach of taking the means of Δx , Δy , and $\Delta \theta$ to illustrate relative performance.

D. Performance metrics

For evaluation we follow the KITTI odometry benchmark [22] which evaluates the distance-normalised translational and rotational error over all possible fixed segment lengths ranging from $\{100 \text{ m}, 200 \text{ m}, \dots, 800 \text{ m}\}$.

This approach is commonly used when assessing odometry performance, as it provides a balanced metric that unlike end-point errors is not overly sensitive to *when* in the sequence an error occurs, i.e. an error in a rotational estimate near the start would tend to produce higher end-point errors than if occurs near the end. Our evaluation is based on the implementation provided by the STARS laboratory¹.

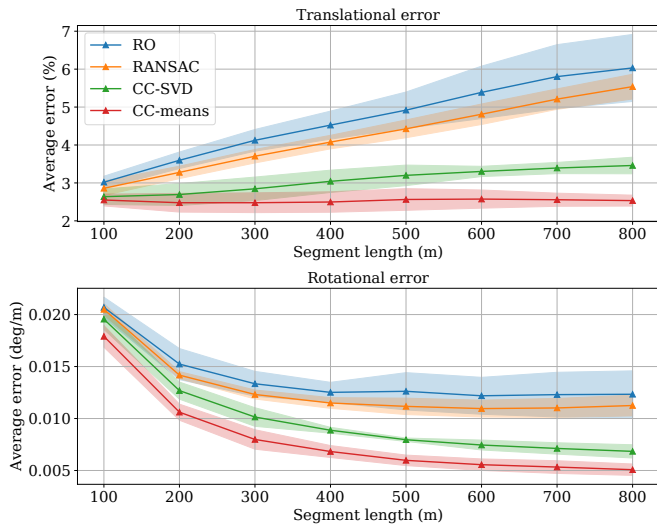


Figure 7. Mean segment errors with corresponding standard deviation envelopes over the course of the first five sequences of the Oxford 10 km route showing translational and rotational error performance for all four methods.

Table I
ERROR METRICS FOR EGO-MOTION
ESTIMATION METHODS AVERAGED OVER
ALL TESTED SEQUENCES

Method	Tr. [%]	Rot. [$^{\circ}$ /m]
RO (full matches)	4.6755	0.0139
RO w/ RANSAC	4.2379	0.0129
RO w/ CC*, SVD	3.0717	0.0101
RO w/ CC*, means	2.5287	0.0082

*Constant-curvature method

VI. RESULTS

Over the five 10 km test sequences, trajectory metrics are summarised in Tab. I for each method’s performance. Tab. II details the performance of the four different systems on each of these five sequences to provide a more in-depth set of results.

As expected, because of its widespread application to outlier removal, the naïve RANSAC method performs better than when using the full set of matches. However, this improvement is marginal. Therefore, the implausible associations in Fig. 2 are evidently not dealt with effectively by standard RANSAC. We expect that this is because the initial match selection performed by RO has already selected matches based on their plausibility from a geometric standpoint, which is similar to what this baseline method is doing. However, when leveraging the constant-curvature constraint to select matches, we see a much larger drop in both translational and rotational error. For a direct comparison between methods, the SVD algorithm is used to compute poses on the refined subset and demonstrates the quantitative benefit of removing spurious matches.

Further to this, estimating pose from a single landmark association allows the constant-curvature motion parameters to be converted into their $\mathbb{SE}(2)$ form for every inlier, where we can find the mean of Δx , Δy , and $\Delta \theta$. This outperforms

the SVD approach, as it reduces the impact of relatively poor associations more so than the SVD algorithm, which assigns all members of the selected subset an equal weighting.

Taking a closer look at the estimates, we see that the lateral error that is present when we originally use the full set of matches (i.e. *RO (full matches)* in Tabs. I and II) is erratic while our proposed system is steady in its estimates, particularly in areas when there should be smooth forward motion i.e. no lateral displacement, shown in Fig. 3. During sharp turns, there is some underestimation of the lateral motion likely due to the applied quantile limits selecting a narrower set of associations as compared to more common gradual turns or straight trajectories. While being a limitation of our system, this method still tracks the ground truth signal *far* closer than the baseline system (see Tabs. I and II and Fig. 7), which *itself* sometimes over- or underestimates the true motion and presents highly deviant estimates when the vehicle is travelling straight forwards. Stated plainly, the ego-motion estimates that our proposed system produces are more plausible given the real constraints that the vehicle is under – i.e. cars do not suddenly move sideways.

Aggregating over all segment lengths for each sequence, we see from Tab. II translational error in the range of 2.25 % to 2.84 % as compared to 4.26 % to 5.55 % for the baseline RO system, a reduction in error in the range 51 % to 53 %. Similar gains are made for rotational errors.

When looking across all five sequences, our contribution (*CC-means*) has been able to reduce absolute error in translation by approximately 2.15 %, almost halving the original error in RO. Rotational error was just $0.0082^{\circ}\text{m}^{-1}$, down from $0.0139^{\circ}\text{m}^{-1}$ when using all matches.

Fig. 8 (best viewed in colour) shows the overall trajectories from each method compared with ground truth, which provides more of a qualitative result. Here, the reduction in drift over kilometres of driving can be exemplified by paying attention to the square loop located at approximate coordinates (200 m; 400 m) in the ground truth trace shown in black. This section of the route occurs near to the end of the 10 km sequence, where the majority of any accumulated error is observable. At this point, our proposed method has drifted approximately 140 m with a few degrees of yaw error, while the original system is more than 450 m off-course and has a yaw estimate more or less perpendicular to the ground truth.

Making direct comparisons with other RO systems is difficult as they publish their results on different sets of either 7 or 8 of the 32 sequences from the *Oxford Radar Robotcar Dataset* [3], as there is no universally recognised “test set”. However, the translational error can provide a loose comparison metric, despite being reported on different sequences from the dataset. In achieving 2.53 %, our proposed system brings the baseline RO [1] on which it is applied (having been relegated with 4.68 %) back into competition with more recently developed top-performing RO systems, which include: RadarSLAM [2] (3.11 % to 2.21 %), Barnes’ Under The Radar [7] (2.06 %), HERO [9] (1.99 %), CFEAR [12] (1.76 %), and Barnes’ Masking by Moving [6] (2.78 % to 1.16 %). This is achieved without any use of learning techniques that introduce additional complications,

¹Evaluation metrics implementation available here: <https://github.com/tiasSTARS/pyslam>

Table II
ERROR METRICS FOR EGO-MOTION ESTIMATION METHODS OVER ALL EXPERIENCES FROM 2019-01-10

Sequence →	11-46-21		12-32-52		14-02-34		14-50-05		15-19-41	
Method ↓	Tr. [%]	Rot. [°/m]	Tr. [%]	Rot. [°/m]	Tr. [%]	Rot. [°/m]	Tr. [%]	Rot. [°/m]	Tr. [%]	Rot. [°/m]
RO (full matches)	4.54	0.013	5.55	0.017	4.26	0.013	4.44	0.014	4.61	0.013
RO w/ RANSAC	4.50	0.013	4.20	0.012	4.09	0.013	4.34	0.013	4.05	0.012
RO w/ CC*, SVD	2.96	0.010	3.15	0.010	3.32	0.011	2.92	0.010	3.32	0.011
RO w/ CC*, means	2.35	0.008	2.57	0.009	2.63	0.008	2.25	0.007	2.84	0.009

*Constant-curvature method




Figure 8. The ground traces for the four methods from the 2019-01-10-15-19-41 sequence show accumulated odometry performance over time as compared to optimised ground truth. This figure is best viewed in colour.

like those required in Under The Radar [7], HERO [9], and Masking by Moving [6]. Furthermore, our new method is applicable to any point-based RO pipeline (e.g. CFEAR [12]), where it should improve match refinement while introducing relatively negligible additional processing overhead.

VII. CONCLUSION

This work shows that, by considering the physical constraints on how a robot is able to move through its environment, we can drastically reduce both translational and rotational error in RO. Crudely filtering out matches based on the overall consensus of the group is potentially helpful, but ultimately a naïve approach to this problem. Instead, if we consider matched landmarks based on how an appropriately-constrained robot would have had to move in order to generate those associated candidates, the match refinement task becomes far more natural and yields compelling improvements in ego-motion estimation at a low cost.

Specifically, this approach has the advantage of being extremely light-weight: finding the two parameters that describe ego-motion for each match requires two simple equations, before a sorting and subset-selection step refines the set of matches. Solving for the final pose estimate then just involves taking a mean for Δx , Δy , and $\Delta \theta$ which is also inexpensive.

Our contribution makes use of a strong prior on how our vehicle is constrained to move. By folding this information into our system at the match-selection level, we see reduced lateral errors introduced by poor matches that adversely affect odometry estimates. This is achieved without the need to collect hundreds of kilometres of training data which is usually required to learn how a vehicle is likely to move.

Lastly, our method is portable: any point-based odometry system that uses a range sensor (radar, LiDAR) to measure ranges and bearings to landmarks in the environment, captured from a non-holonomic vehicle, should benefit from our contribution presented here.

REFERENCES

- [1] S. Cen and P. Newman, "Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2018.
- [2] Z. Hong, Y. Petillot, and S. Wang, "Radarslam: Radar based large-scale SLAM in all weathers," in *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, 2020, pp. 5164–5170.
- [3] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The Oxford radar robotcar dataset: A radar extension to the Oxford robotcar dataset," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [4] R. Aldera, D. De Martini, M. Gadd, and P. Newman, "Fast radar motion estimation with a learnt focus of attention using weak supervision," in *International Conference on Robotics and Automation (ICRA)*, 2019.
- [5] R. Aldera, D. De Martini, M. Gadd, and P. Newman, "What could go wrong? Introspective radar odometry in challenging environments," in *IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019.
- [6] D. Barnes, R. Weston, and I. Posner, "Masking by moving: Learning distraction-free radar odometry from pose information," in *Conference on Robot Learning*, 2020.
- [7] D. Barnes and I. Posner, "Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [8] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?" *Robotics and Automation Letters*, vol. 6, no. 2, pp. 771–778, 2021.
- [9] K. Burnett, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Radar odometry combining probabilistic estimation and unsupervised feature learning," in *Robotics: Science and Systems*, 2021.
- [10] Y. S. Park, Y. S. Shin, and A. Kim, "Pharao: Direct radar odometry using phase correlation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [11] P. C. Kung, C. C. Wang, and W. C. Lin, "A normal distribution transform-based radar odometry designed for scanning and automotive radars," in *International Conference on Robotics and Automation*. IEEE, 2021, pp. 14 417–14 423.
- [12] D. Adolfsson, M. Magnusson, A. Alhashimi, A. J. Lilienthal, and H. Andreasson, "CFEAR radarodometry — conservative filtering for efficient and accurate radar odometry," in *International Conference on Intelligent Robots and Systems*. IEEE/RSJ, 2021, pp. 5462–5469.
- [13] S. Lovegrove, A. Patron-Perez, and G. Sibley, "Spline fusion: A continuous-time representation for visual-inertial fusion with application to rolling shutter cameras." *British Machine Vision Conference*, pp. 93.1–93.12, 2013.

- [14] C. H. Tong, P. Furgale, and T. D. Barfoot, "Gaussian process gauss-newton for non-parametric simultaneous localization and mapping," *The International Journal of Robotics Research*, vol. 32, no. 5, 2013.
- [15] J. Almeida and V. M. Santos, "Real time egomotion of a nonholonomic vehicle using lidar measurements," *Journal of Field Robotics*, vol. 30, no. 1, 2013.
- [16] N. Nourani-Vatani, J. Roberts, and M. V. Srinivasan, "Practical visual odometry for car-like vehicles," in *IEEE International Conference on Robotics and Automation*, 2009.
- [17] D. Scaramuzza, F. Fraundorfer, and R. Siegwart, "Real-time monocular visual odometry for on-road vehicles with 1-point RANSAC," in *IEEE International conference on robotics and automation*, 2009.
- [18] E. J. M. Adams, J. Mullane and B. Vo, *Robot Navigation and Mapping with Radar*. Artech House, 2012.
- [19] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints," in *IEEE International Conference on Computer Vision*, vol. 2, 2005.
- [20] J. H. Challis, "A procedure for determining rigid body transformation parameters," *Journal of biomechanics*, vol. 28, no. 6, 1995.
- [21] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, 1981.
- [22] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.

APPENDIX

As referenced in Sec. IV-C, given the ranges r_0 , r_1 and bearings β_0 , β_1 we can find θ and R_{ICR} . The diagram has the landmark on the left of both relative poses to avoid overcrowding of symbols on the right-hand side, but this derivation is clearer if we avoid reflex angles. As such, we define angles $\phi_0 = \beta_0 - 2\pi$ and $\phi_1 = \beta_1 - 2\pi$ before getting started. Firstly, we express γ in terms of ϕ_1 and θ as

$$\gamma = \phi_1 - \theta \quad (\text{A.1})$$

Then find α by summing angles in the constructed triangle $\Delta PG_{t-1}G_t$ and rearranging:

$$\begin{aligned} \pi &= \alpha + \gamma + \theta/2 \\ \therefore \alpha &= \pi - \phi_1 + \theta/2 \end{aligned} \quad (\text{A.2})$$

Then, by the sine rule, we can express

$$\begin{aligned} \frac{r_0}{\sin \alpha} &= \frac{r_1}{\sin(\phi_0 + \theta/2)} \\ \frac{r_0}{\sin(\pi - (\phi_1 - \theta/2))} &= \frac{r_1}{\sin(\phi_0 + \theta/2)} \\ \frac{r_0}{r_1} &= \frac{\sin(\phi_1 - \theta/2)}{\sin(\phi_0 + \theta/2)} \end{aligned}$$

Expanding using the trigonometric identity $\sin(A + B) = \sin A \cos B + \cos A \sin B$ gives

$$\frac{r_0}{r_1} = \frac{\sin(\phi_1) \cos(-\theta/2) + \cos(\phi_1) \sin(-\theta/2)}{\sin(\theta/2) \cos(\phi_0) + \cos(\theta/2) \sin(\phi_0)}$$

and dividing all terms through by $\cos(\theta/2)$

$$\begin{aligned} r_0/r_1 (\tan(\theta/2) \cos(\phi_0) + \sin(\phi_0)) \\ = \sin(\phi_1) - \cos(\phi_1) \tan(\theta/2) \end{aligned}$$

then moving θ -terms to the left-hand side

$$\begin{aligned} r_0/r_1 \tan(\theta/2) \cos(\phi_0) + \cos(\phi_1) \tan(\theta/2) \\ = \sin(\phi_1) - r_0/r_1 \sin(\phi_0) \end{aligned}$$




Figure 9. As in Fig. 4, the robot observes a landmark L from two successive poses. Included here are additional labels required to step through the derivation of the constant-curvature motion equations.

and finally rearranging to express θ as

$$\theta = 2 \arctan \left(\frac{\sin \phi_1 - r_0/r_1 \sin \phi_0}{r_0/r_1 \cos \phi_0 + \cos \phi_1} \right) \quad (\text{A.3})$$

The diagram in Fig. 9 shows that we can express Δx and Δy in terms of the trigonometric components of ρ , which is computable given ranges r_0 , r_1 and angles $\phi_0 = -\beta_0$ and $\phi_1 = -\beta_1$ as before. Firstly, we find angle ψ between the start and end pose from the landmark's position at L by summing angles in the triangle $\Delta LG_{t-1}G_t$ and using Eq. (A.2) to give

$$\begin{aligned} \pi &= \alpha + \psi + \phi_0 + \theta/2 \\ \therefore \psi &= \phi_1 - \phi_0 - \theta \end{aligned} \quad (\text{A.4})$$

and relating ρ to R_{ICR} as

$$\rho = 2R_{ICR} \sin(\theta/2) \quad (\text{A.5})$$

Then using the sine rule, we have

$$\begin{aligned} \frac{\rho}{\sin \psi} &= \frac{r_1}{\sin(\phi_0 + \theta/2)} \\ \rho &= \frac{r_1 \sin \psi}{\sin(\phi_0 + \theta/2)} \end{aligned}$$

and substituting in Eqs. (A.4) and (A.5) to solve for R_{ICR}

$$\begin{aligned} 2R_{ICR} \sin(\theta/2) &= \frac{r_1 \sin \psi}{\sin(\phi_0 + \theta/2)} \\ \therefore R_{ICR} &= \frac{r_1 \sin(\phi_1 - \phi_0 - \theta)}{2 \sin(\theta/2) \sin(\phi_0 + \theta/2)} \end{aligned} \quad (\text{A.6})$$

Lastly, expressing θ and R_{ICR} in terms of our bearings β_0 and β_1 after the above derivation yields:

$$\theta = 2 \arctan \left(\frac{r_0/r_1 \sin \beta_0 - \sin \beta_1}{r_0/r_1 \cos \beta_0 + \cos \beta_1} \right) \quad (\text{A.7})$$

$$R_{ICR} = \frac{r_1 \sin(\beta_0 - \beta_1 - \theta)}{2 \sin(\theta/2) \sin(\theta/2 - \beta_0)} \quad (\text{A.8})$$