

Triangulation of 3D target points from radar range and bearing data

Magnus Oskarsson*

CVML, Centre for Mathematical Sciences, Lund University, Sweden

magnus.oskarsson@math.lth.se

Abstract

We propose a method for estimating the 3D position of a target point, given multiple measurements of it, using mm-wave radar data. Given azimuth headings and range estimates from posed radar positions, we find the 3D position, using an approximate, but geometrically and statistically meaningful cost. The 3D position is found in an optimal way, using this approximate cost. By deriving the Lagrangian of the corresponding maximum likelihood and maximum a posteriori estimates, we show that we can find all local minima by solving an eigenvalue problem. The global optimum can then easily and efficiently be extracted from these solutions. We validate the method on synthetic data and test it on several real world datasets, and release public code¹.

1. Introduction

Over the last years we have seen a renewed interest in radar applications, especially in combination with camera data. One of the attractive points of this combination is the sensors' complementary nature and failure modes. Cameras and lidars suffer from severe degradation in harsh environments (containing smoke, dust, fog, rain, and snow). With longer wavelengths, radars can penetrate through such matter to a larger extent. Radars are also installed on various mobile platforms in the desire to reduce the sensor payload, giving them larger roles in applications where cameras are traditionally used, such as positioning, localization, mapping, object detection and classification. These uses make radar sensors and networks of radar sensors prevalent in robotics [13, 15] and autonomous vehicles [34, 46, 48]. Of course, radar sensors also exhibit a number of limitations, e.g. low angular resolution, unique noise characteristics, from spurious returns throughout the sensor range, to complicated speckle noise and multi-path effects as well as problems with distinguishing between closely located ob-




Figure 1. Observations of a 3D point x , from two radar positions (y_1 and y_2) will constrain the point to lie in the intersection of two spheres and two planes. Without noise this will in general give a unique solution for x , but if we have noise in the measurements, we would like to find the optimal estimate of x .

jects, [46]. In this paper we work with imaging or scanning radars. A radar data point consists of a range, azimuth and doppler speed measurement. There are also radars that directly give the full 3D position, but most often the vertical angle resolution is very poor [34, 49]. The geometry of range and angle measurements from two sensors is illustrated in Figure 1.

The main objective in this paper is to present a method for estimating the 3D target point, given a number of such measurements, from known positions and orientations of the radar sensor, i.e. *the triangulation problem*. One of the main goals of this paper is to formulate this problem in a geometrically and statistically meaningful way. The triangulation problem is interesting for estimating 3D structure, given tracked measurements, in positioning applications such as SLAM, localization, path planning, and obstacle avoidance. It could also play a role in semantic tasks such as object recognition and tracking. As such, robust, accurate and efficient methods for triangulation represent important fundamental building blocks in many systems.

Our main contributions in the paper are:

- A geometrically motivated and statistically valid error model for the radar triangulation problem in 3D,
- derivation of the corresponding Maximum likelihood

*This work was supported by the ELLIIT strategic research projects RADICAL and Local Positioning Systems.

¹<https://github.com/hamburgerlady/radar-triangulation>

- (ML) and Maximum a posteriori (MAP) estimates,
- a fast and non-iterative method that finds the global optimizers to a close approximation of both the ML and the MAP triangulation problem.

1.1. Related work

Related to the radar triangulation problem, there has been much work on point cloud estimation [38], object detection [14, 47], and occupancy grid estimation [40]. These methods are typically data driven and learning based. For our approach we assume known positions of the sensors and feature correspondences. This can e.g. be obtained using SLAM systems [5, 19] or from some form of calibration process. Calibration of sensors is often done using specialized calibration targets [9, 37]. One can also assume planar scenes [42, 43] or other properties of the 3D scene [34]. Related to the radar triangulation problem that we investigate in this paper, is the triangulation problem given only range measurements. This problem is also known as single-source localization, and has broad application areas within e.g. communication, chemistry, and robotics [10, 12, 26, 32, 35]. The measurements can come from Time Of Arrival (TOA) or Received Signal Strength (RSS), where the estimation problem is known as *trilateration*. For Time Difference Of Arrival (TDOA) measurements the problem is known as *multilateration*. These problems are typically both nonlinear and nonconvex optimization problem, with multiple local minima. There is a rich body of work, but most approaches are iterative or based on relaxations of the problem [3, 4, 8, 20, 21, 29, 31, 39]. Most closely related to our approach is the optimal trilateration method presented in [25]. If we are given two angle measurements and a range measurement from each sensor, i.e. as in the case of lidar or time-of-flight cameras, we directly get an estimate for the 3D position from each sensor. The optimal triangulation then simply corresponds to doing averaging. If we do not have accurate estimates of the sensor positions, the problem can be solved using optimal registration methods [36, 45]. For triangulation using ordinary cameras, the 3D position can be found by simply intersecting the back-projected image rays. However, if we want to minimize some statistical error such as the reprojection error in many views, finding the global minimizer is inherently difficult, since we need to estimate a depth in every image. For two views it's possible, [17], but for three views it's numerically and theoretically much harder [7, 24, 41], and for more views even more difficult. Iterative methods, that do not guarantee any global optimum have also been proposed [2, 18, 23, 28]. Other methods include geometric methods that minimize the L_∞ -error [16, 33] and Branch-and-Bound methods [22, 30], with worst case exponential convergence.




Figure 2. A radar at y is swept in the xy -plane, giving a range estimate r and an azimuth direction θ for a target point x . The range estimate is assumed to have an error standard deviation of σ and the azimuth angle standard deviation is δ . At the target point x , the standard deviation relative the azimuth plane is approximated by $r\delta$ (assuming a small δ).

2. Problem formulation

A schematic of the geometry for two 2D radars is given in Figure 1. The radar measures the distance to a point in 3D (r), and also the angle to the point in the plane that the radar is swept in (θ). This means that in the coordinate system of the radar, the 3D point is constrained to lie on a great circle, on the intersection of a sphere and a plane in 3D.

An approximation that simplifies the geometry is to assume that all points lie in the plane that the radar is swept in. The benefit of this approximation is that one can directly compute the target 3D position from the radar measurement, simply as $x = (r \cos \theta, r \sin \theta, 0)$ if the radar is swept in the plane $z = 0$. Of course, if the true 3D position of the point is away from this plane, this is a bad approximation. This approximation is valid if the height (i.e. offset from the radar plane) is small compared to the distance from the radar.

If we would like to find the true 3D position of the target point, we need to look at the constraints that the measurement poses on this point. We know that the target point x should have a distance r to the radar position y , which means that we have

$$\|x - y\|_2 = r. \quad (1)$$

Furthermore, the angle measurement θ constrains the point to lie on a plane. We can write this constraint as

$$n^T(x - y) = 0, \quad (2)$$

where n is the normal to the plane. The normal is directly given by the angle θ and the plane in the local radar frame.

For instance, if the radar is swept in the xy -plane (with zero angle corresponding to the x -direction) then $\mathbf{n}$ is given by

$$\mathbf{n}(\theta) = \begin{pmatrix} \sin(\theta) \\ -\cos(\theta) \\ 0 \end{pmatrix}. \quad (3)$$

The downside of this model is of course that we now need more than one radar position to estimate a 3D position. This leads us to the following basic problem

Problem 1 *Given N observations (r_i, θ_i) , $i = 1, \dots, N$, of a 3D point, from radar positions $\mathbf{y}_i$, what is the best estimate of the 3D point position $\mathbf{x}$?*

Since each measurement puts two constraints on $\mathbf{x}$ we need at least two measurements to estimate $\mathbf{x}$.

2.1. Linear solution

Given two or more radar positions, we can linearly estimate the unknown target 3D position $\mathbf{x}$, using standard techniques. For example, if we have two radar measurements from $\mathbf{y}_1$ and $\mathbf{y}_2$ respectively, we have

$$\|\mathbf{x} - \mathbf{y}_1\|_2^2 = r_1^2, \quad (4)$$

$$\mathbf{n}_1^T(\mathbf{x} - \mathbf{y}_1) = 0, \quad (5)$$

$$\|\mathbf{x} - \mathbf{y}_2\|_2^2 = r_2^2, \quad (6)$$

$$\mathbf{n}_2^T(\mathbf{x} - \mathbf{y}_2) = 0. \quad (7)$$

Taking (4) – (6) will give a linear constraint on $\mathbf{x}$, giving a linear formulation for the solution of $\mathbf{x}$ according to

$$\begin{bmatrix} \mathbf{n}_1^T \\ \mathbf{n}_2^T \\ 2\mathbf{y}_2^T - 2\mathbf{y}_1^T \end{bmatrix} \mathbf{x} = \begin{bmatrix} \mathbf{n}_1^T \mathbf{y}_1 \\ \mathbf{n}_2^T \mathbf{y}_2 \\ r_1^2 - r_2^2 - \mathbf{y}_1^T \mathbf{y}_1 + \mathbf{y}_2^T \mathbf{y}_2 \end{bmatrix}. \quad (8)$$

If we have more measurements, we can simply stack the linear equations, giving an overdetermined system of $2N - 1$ equations given N radar positions.

2.2. Minimal solution

Given two radar observations we can also simply derive a minimal solution by dropping one of the quadratic constraints. Using (5) and (7) will constrain $\mathbf{x}$ to

$$\mathbf{x} = \mathbf{v}_o + \lambda \mathbf{v}_1, \quad (9)$$

where $\mathbf{v}_o$ and $\mathbf{v}_1$ are known vectors depending only on the data, and λ is an unknown parameter. Inserting this expression into (4) gives a second-degree polynomial in λ leading to two potential solutions for $\mathbf{x}$. These solutions can be checked in (6) for the final estimate. If we have more available radar measurements, we can directly use this simple minimal solver in a RANSAC framework [11] to get a robust estimate of $\mathbf{x}$.

3. Optimal approximate triangulation

Using the linear solution given in Section 2.1 will often give a reasonable solution if we have many measurements, from radars that are spatially separated. However, if we have many measurements, and if we have some knowledge about the accuracy of our measurements, we would like to use this to find a statistically more valid solution. We would also like to find the target position where the cost is geometrically well founded. We will in this section describe how we can model our problem and derive a non-iterative method for finding the near optimal solution in a statistical sense. To this end, we assume that our measurements are corrupted by additive noise, that is also assumed to be independent between measurements. Our error model is illustrated in Figure 2. The range measurement is simply assumed to be offset with additive noise with zero mean and standard deviation σ (in e.g. meters). We will throughout the paper assume that we have a reasonable estimate of σ given. The azimuth heading is assumed to be disturbed with zero mean noise of standard deviation δ , given in angle units. Again, we assume this quantity to be known. Now, if we look at a measurement, the standard deviation of the orthogonal offset to the corresponding plane $\mathbf{n}$ (due to the error in θ) can then be well approximated with $r\delta$. If we assume that the errors follow normal distributions, we get for a measurement the following equations,

$$\|\mathbf{x} - \mathbf{y}\|_2 - r = \epsilon_r, \quad \epsilon_r \in \mathcal{N}(0, \sigma), \quad (10)$$

$$\frac{\|\mathbf{n}^T \mathbf{x} - \mathbf{n}^T \mathbf{y}\|_2}{\sqrt{\mathbf{n}^T \mathbf{n}}} = \epsilon_\theta, \quad \epsilon_\theta \in \mathcal{N}(0, r\delta). \quad (11)$$

We will in the following assume (w.l.o.g) that $\mathbf{n}^T \mathbf{n} = 1$.

3.1. Maximum likelihood formulation

Given N measurements, we would like to find the $\mathbf{x}$ that maximizes the likelihood of getting these measurements (the ML estimate). On the assumption that the errors in these measurements $\epsilon = (\epsilon_{r_1}, \epsilon_{r_2}, \dots, \epsilon_{\theta_N})$, are independent we get the following maximization problem

$$\arg \max_{\mathbf{x}} P(\epsilon|\mathbf{x}) = \arg \max_{\mathbf{x}} \prod_{i=1}^N P(\epsilon_{r_i}|\mathbf{x}) P(\epsilon_{\theta_i}|\mathbf{x}). \quad (12)$$

We will in this paper assume that the errors follow normal distributions, and hence the probabilities are given by

$$P(\epsilon_{r_i}|\mathbf{x}) = \frac{1}{\sqrt{2\pi}\sigma_i} e^{-\frac{(\|\mathbf{x} - \mathbf{y}_i\|_2 - r_i)^2}{2\sigma_i^2}}, \quad (13)$$

$$P(\epsilon_{\theta_i}|\mathbf{x}) = \frac{1}{\sqrt{2\pi}r_i\delta_i} e^{-\frac{(\mathbf{n}_i^T \mathbf{x} - \mathbf{n}_i^T \mathbf{y}_i)^2}{2r_i^2\delta_i^2}}. \quad (14)$$

Looking at the Negative Log Likelihood (NLL), and disregarding the normalization constants, will then give us the

following non-linear least squares problem,

$$\arg \max_{\mathbf{x}} P(\epsilon|\mathbf{x}) = \arg \min_{\mathbf{x}} -\log P(\epsilon|\mathbf{x}) = \quad (15)$$

$$\arg \min_{\mathbf{x}} \sum_{i=1}^N -\log P(\epsilon_{r_i}|\mathbf{x}) - \log P(\epsilon_{\theta_i}|\mathbf{x}) = \quad (16)$$

$$\arg \min_{\mathbf{x}} \sum_{i=1}^N \left(\frac{(\|\mathbf{x} - \mathbf{y}_i\|_2 - r_i)^2}{2\sigma_i^2} + \frac{(\mathbf{n}_i^T \mathbf{x} - \mathbf{n}_i^T \mathbf{y}_i)^2}{2r_i^2 \delta_i^2} \right). \quad (17)$$

This is a non-linear and non-convex problem, so it's difficult to derive a closed-form solution. If at least the cost was polynomial in $\mathbf{x}$ we could potentially differentiate the Lagrangian of (17), and solve for all possible local minima, and evaluate these to find the optimal $\mathbf{x}$. In order to do this, we linearize the square root in the first terms of the sum in (17) to get

$$(\|\mathbf{x} - \mathbf{y}_i\|_2 - r_i)^2 \approx \frac{1}{4r_i^2} (\|\mathbf{x} - \mathbf{y}_i\|_2^2 - r_i^2)^2. \quad (18)$$

Note that we choose the point that we linearize around to be the actual measurement r_i . Inserting this approximation into (17) will then give a cost on the form

$$L(\mathbf{x}) = \sum_{i=1}^N \omega_i (\|\mathbf{x} - \mathbf{y}_i\|_2^2 - r_i^2)^2 + \gamma_i (\mathbf{n}_i^T \mathbf{x} - \mathbf{n}_i^T \mathbf{y}_i)^2, \quad (19)$$

with $\omega_i^{-1} = 8r_i^2 \sigma_i^2$ and $\gamma_i^{-1} = 2r_i^2 \delta_i^2$. This is then the problem that we would like to solve,

Problem 2 Given N observations $(r_i, \mathbf{n}_i)$, with standard deviations σ_i and δ_i , $i = 1, \dots, N$, from radar positions $\mathbf{y}_i$, what is the approximate maximum likelihood estimate of the 3D point position $\mathbf{x}$, i.e. that minimizes (19)?

3.2. Maximum a posteriori formulation

The previous section described how we can formulate the ML problem for triangulation. In some cases, we have some prior information on the position $\mathbf{x}$, such as e.g. that the height is close to some ground plane. We will now show how we can incorporate such priors in our model. Using Bayes formula, we can directly formulate this problem as

$$\arg \max_{\mathbf{x}} P(\mathbf{x}|\epsilon) = \arg \max_{\mathbf{x}} \frac{P(\mathbf{x})P(\epsilon|\mathbf{x})}{P(\epsilon)}. \quad (20)$$

The optimization problem will not depend on $P(\epsilon)$, so we can discard this factor. The second factor in the nominator is the likelihood described in the previous section. If we model the prior $P(\mathbf{x})$ using a Gaussian with mean $\mathbf{x}_0$ and

covariance matrix Φ and take the negative logarithm of (20) we will get the following problem

$$\mathbf{x} = \arg \min_{\mathbf{x}} L(\mathbf{x}) + (\mathbf{x} - \mathbf{x}_0)^T \Phi^{-1} (\mathbf{x} - \mathbf{x}_0), \quad (21)$$

with $L(\mathbf{x})$ given by (19). This gives us the Maximum a posteriori (MAP) problem formulation

Problem 3 Given N observations $(r_i, \mathbf{n}_i)$, with standard deviations σ_i and δ_i , $i = 1, \dots, N$, from radar positions $\mathbf{y}_i$, and the prior assumption that $\mathbf{x} \in \mathcal{N}(\mathbf{x}_0, \Phi)$, what is the approximate MAP estimate given by (21).

3.3. Solution by enumerating all local minima

We will now describe our method for finding the solution to Problem 2. We will do this by differentiating the cost in (19) w.r.t. $\mathbf{x}$, and finding all solutions. The global optimum of (19) will be located at one of these local minima. Here $L(\mathbf{x})$ is of fourth-degree in $\mathbf{x}$, so the derivatives will be of degree three,

$$L'(\mathbf{x}) = \sum_{i=1}^N 4\omega_i (\|\mathbf{x} - \mathbf{y}_i\|_2^2 - r_i^2) (\mathbf{x} - \mathbf{y}_i) + \quad (22)$$

$$2\gamma_i (\mathbf{n}_i^T \mathbf{x} - \mathbf{n}_i^T \mathbf{y}_i) \mathbf{n}_i. \quad (23)$$

We can write this in the following way, in terms of different degree terms in $\mathbf{x}$,

$$L'(\mathbf{x}) = a(\mathbf{x}^T \mathbf{x}) \mathbf{x} + (\mathbf{x}^T \mathbf{x} \mathbf{I} + \mathbf{x} \mathbf{x}^T) \mathbf{b} + C \mathbf{x} + \mathbf{d}, \quad (24)$$

with

$$\mathbf{a}_{1 \times 1} = \sum_i 4\omega_i, \quad \mathbf{b}_{3 \times 1} = \sum_i -4\omega_i \mathbf{y}_i,$$

$$C_{3 \times 3} = \sum_i 4\omega_i ((\mathbf{y}_i^T \mathbf{y}_i - r_i^2) \mathbf{I} + 2\mathbf{y}_i \mathbf{y}_i^T) + 2\gamma_i \mathbf{n}_i \mathbf{n}_i^T,$$

$$\mathbf{d}_{3 \times 1} = - \sum_i 4\omega_i (\mathbf{y}_i^T \mathbf{y}_i - r_i^2) \mathbf{y}_i + 2\gamma_i (\mathbf{n}_i^T \mathbf{y}_i) \mathbf{n}_i.$$

This problem has the same structure as the optimal trilateration problem described in [25]. We will follow the same solution strategy here, which will lead to an eigenvalue formulation. To see this, we first translate our coordinate system so that $\mathbf{b} = \sum_i -4\omega_i \mathbf{y}_i = 0$. This will eliminate second degree terms. We then do an orthogonal diagonalization of the symmetric matrix $C = U \hat{C} U^T$. Writing this in terms of $\hat{\mathbf{x}} = U \mathbf{x}$ and $\hat{\mathbf{d}} = U \mathbf{d}$ gives a separation of variables according to

$$(\hat{\mathbf{x}}^T \hat{\mathbf{x}}) \hat{x}_j + \hat{C}_{jj} \hat{x}_j + \hat{d}_j = 0, \quad j = 1, 2, 3, \quad (25)$$

where $\hat{\mathbf{x}} = [\hat{x}_1, \hat{x}_2, \hat{x}_3]^T$ and $\hat{\mathbf{d}} = [\hat{d}_1, \hat{d}_2, \hat{d}_3]^T$. Multiplying each equation with $\hat{x}_j$ then leads to the following

eigenvalue formulation of the problem

$$\underbrace{\begin{bmatrix} -\hat{C} & -\text{diag}(\hat{\mathbf{d}}) & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{3 \times 3} & -\hat{C} & -\hat{\mathbf{d}} \\ \mathbf{1}_{1 \times 3} & \mathbf{0}_{1 \times 3} & 0 \end{bmatrix}}_{Q_{7 \times 7}} \mathbf{w} = \lambda \mathbf{w}, \quad (26)$$

with $\lambda = \hat{\mathbf{x}}^T \hat{\mathbf{x}}$ and $\mathbf{w} = [\hat{x}_1^2, \hat{x}_2^2, \hat{x}_3^2, \hat{x}_1, \hat{x}_2, \hat{x}_3, 1]^T$. The matrix Q on the left-hand side is 7×7 , and there are hence seven eigenvectors. Each possible eigenvector will correspond to a local minimum of our problem, and we can find the global optimum by simply evaluating the cost for each solution, and choose the best one. The different steps are summarized in Algorithm 1. This approach gives us the

Algorithm 1 ML Eigenvalue solver

Require: $\mathbf{y}_i, r_i, \mathbf{n}_i, \sigma_i, \delta_i, \quad i = 1, \dots, N$

Ensure: $\mathbf{x}$

- 1: Change coordinate system
 - 2: Build the matrix Q in (26)
 - 3: Find the eigenvectors $\mathbf{w}_k$ of $Q, \quad k = 1, \dots, 7$
 - 4: Change back to original coordinate system
 - 5: For all eigenvectors extract local minimizers $\mathbf{x}_k$
 - 6: $\mathbf{x} = \arg \min_{\mathbf{x}_k} L(\mathbf{x}_k), \quad k = 1, \dots, 7$
-

solution to Problem 2, but we can use the exact same approach to solve Problem 3 using some minor changes. We get the derivative of the cost in (21) by simply adding the term $S(\mathbf{x} - \mathbf{x}_0)$ to (24), where $S^T S = \Phi^{-1}$. The only change that this leads to is that we will use the modified constants

$$C' = C + S, \quad \mathbf{d}' = \mathbf{d} - S\mathbf{x}_0, \quad (27)$$

in the algorithm.

4. Experimental evaluation

We will now evaluate our proposed solver in a number of ways, on synthetic and real data. The Matlab implementation of our solver runs in $60\mu s$ on an Apple M3 Pro. For real radar datasets, it's difficult to annotate and find ground truth for 3D target point data. For this reason, we have constructed semi-synthetic data, based on publicly available large structure from motion reconstruction datasets.

4.1. Synthetic data

We will start by evaluating the solver on synthetic data with no noise added. We generated problems with $N = 15$ radar positions, and a randomly positioned target $\mathbf{x}_t$. From this we calculated corresponding measurements r_i and $\mathbf{n}_i$. Running Algorithm 1 gives an estimate position $\mathbf{x}$. We can then calculate the norm between the ground truth and the estimate. We repeat this a large number of times (100000).




Figure 3. Histogram of logarithmic errors on synthetic data with no noise, given 100000 random instances using 15 radars. The largest error was in this case $3.2e-12$

A histogram of the logarithm of the errors is given in Figure 3. One can see that we get close to machine precision errors in this case. Next, we would like to test how the method performs in the presence of noise. To this end we again generated synthetic sequences of data but added different amounts of noise to the measurements. We do the tests by varying the standard deviations of the range measurements σ and of the angle measurements δ , individually and in combination. We compare our proposed solver with the linear approach described in Section 2.1. The distance between the estimates and the ground truth as a function of standard deviation of the added noise is shown in Figure 4. From left to right we vary σ , δ and both (σ, δ) in combination, respectively. One can see that the proposed solver consistently gives more accurate solutions than the linear solver. One can also see that when only one of the standard deviations increase, the error in the estimate becomes bounded, even for very large standard deviations. When both standard deviations are increased the error increases linearly in the standard deviations. In the graphs we also show the oracle maximum likelihood estimate (in dashed yellow). This is obtained by minimizing the real (not approximate) cost given in (17). We find the minimum using a Levenberg-Marquardt solver with the ground truth position as starting point. One can see that our proposed optimal solver consistently finds a very close estimate to the true ML estimate. All metrics are based on a large number of runs (10000) for each standard deviation, and the mean of the runs is reported in the graphs. The previous graphs showed the errors compared to the ground truth positions. We can also look at the errors that we optimize, i.e. the cost in (19). Figure 5 shows these errors for the exact same experiment. Since we normalize the cost with the given standard deviations, we should get a constant cost, for varying




Figure 4. Results on synthetic data, for varying amounts of added noise. We compare our optimal solution with the linear solver and the oracle maximum likelihood solver. On the y -axis the distance between the ground truth and the estimates are shown. From left to right the graphs show reconstruction errors as functions of increasing noise in range individually, angle individually and range and angle in combination, respectively. Due to the complementary nature of range and angle measurements, we get bounded errors if one of the errors is fixed. Note that the linear solver is very sensitive to range errors due to (8).




Figure 5. Results on synthetic data, as in Figure 4, but showing the maximum likelihood cost (17) for the estimates on the y -axis.




Figure 6. Comparison between the proposed optimal solver and the linear solver when we have varying standard deviations for the 15 different radar positions. The graph shows histograms of logarithms of distances between estimates and the ground truth 3D positions.

degrees of added noise. One can see that when we vary both the standard deviations, we overestimate the standard

deviation for very large added errors. For all these synthetic experiments we added gaussian noise with the same standard deviation to all the radar measurements. In Figure 6 we show the result of an experiment where we randomly chose standard deviations for the different radar measurements (but these are still assumed to be known). The graph shows the histogram of the logarithms of the distances to the ground truth target position. One can see that the proposed optimal solver gives much lower errors than the linear solution in this case also.

4.2. Large-scale semi-synthetic data

In order to perform a controlled, but large-scale experiment with realistic data, we used real structure from motion data, and generated simulated radar data from the given camera positions. We show results from two reconstructions, the Notre Dame data from [44] and the Coliseum part of the Rome16K dataset [27]. We use the estimated 3D structure as ground truth. For each camera position, we calculate the range and heading angle to all visible 3D points in that camera. We use the visibility given by the feature point associations in each camera. We then add random noise to the measurements with some realistic standard deviation ($\sigma = 0.024m$ and $\delta = 0.45^\circ$). The resulting reconstruc-




Figure 7. Example 3D reconstructions from the semi-synthetic experiment. Left shows the original 3D reconstruction from the datasets. Also shown in white are the camera positions. Second shows the reconstruction from the simulated radar data, using the proposed optimal method. The simulated radar positions are the same as the camera positions in white. Third shows the reconstruction using the linear method. To the right the reconstruction error distributions are shown (on the x -axis we have the logarithmic distance from the estimates to the ground truth 3D positions).




Figure 8. Figure shows 3D reconstruction from real radar data from [6]. The blue path shows the ground truth path of the vehicle that was used in the reconstruction. In green, the 3D points estimated using the proposed optimal method are shown.

tions from the simulated radar data is given in Figure 7. To the left the original SfM reconstruction is shown. In the middle the proposed optimal solver is shown, and to the right is the reconstruction using the linear solver. The 3D points are colored with the colors from the reconstruction for visualization purposes. Also shown in each image are the camera/radar positions in white. In the rightmost panel the error distributions for the two datasets are shown. The

graphs show histograms of the logarithms of the distances to the ground truth positions. One can see that we get small errors and better estimates using the proposed solver compared to the linear solver.

4.3. Real data use-case

In order to test our method on real data we have used the large-scale dataset Boreas (A Multi-Season Autonomous




Figure 9. Reprojections for eight radar positions for the experiment based on the real data from [6]. The grayscale images show the raw radar measurements in the cartesian frame centered and aligned with the radar. The blue dots are the extracted and tracked feature points using the SLAM system of [1, 19]. The red circles are the corresponding reprojected 3D points, estimated by the proposed optimal method. The rms values for these reprojected points are $0.41m$, $1.1m$, $8.6m$, $0.97m$, $0.53m$, $0.92m$, $0.72m$, and $6.1m$ respectively.

Driving Dataset) [6]. There is unfortunately not any ground truth available for the radar detections. We take the raw radar measurement and run a radar SLAM pipeline to extract features and track them throughout a sequence. We have used the open source implementation given in [1], which is an implementation of the SLAM odometry system described in [19]. We then use the ground truth positions and orientations of the radar positions when we triangulate the 3D points. We estimate points that are seen in at least five radar positions, and where the baseline of the radars is at least five meters in total. As standard deviations we used the half-width of the measurement resolution ($\sigma = 0.024m$ and $\delta = 0.45^\circ$). The resulting reconstruction is shown in Figure 8. In blue the ground truth path of the car (and the radar) is shown. The estimated 3D points are shown in green, with a number of close-ups given. We see that we get reasonable estimates of the 3D structure that follows the path of the vehicle. In Figure 9 the reprojections for a number of radar positions are shown. The grayscale image shows the raw radar data in the cartesian frame centered around the vehicle (at pixel (320,320) in the image). The blue dots indicate the feature points that were extracted and tracked by the SLAM system. The red circles show the reprojections of the estimated 3D points in these frames. One can see that we get small reprojection errors for most points. There seem to be some outliers in the data, but we haven't done any additional filtering before we ran our solver. It's also clear that the tracking system fails to detect and track a large portion of the potential data points in the radar images. The SLAM system in [19] is working purely in 2D. We believe that having stronger 3D representations could benefit radar SLAM and reconstruction systems. Here the use of our solver could potentially make the mapping both more robust and accurate.

5. Conclusion

We have in this paper presented a geometrically based framework for triangulation of 3D target points, given radar measurements on range and azimuth angle. By formulating a tractable and statistically meaningful error cost, we can estimate the 3D point by simply solving a small eigenvalue problem. We have shown that this gives an estimate that is very close to the true maximum likelihood estimate, and which degrades nicely with increasing degrees of noise in the measurements. We have also shown how to incorporate priors on our model, in order to find MAP estimates. We believe that geometric algorithms, such as the one presented in this paper, could be important building blocks in many radar SLAM system, leading to increase in robustness and accuracy. There is also a great potential in using such geometric methods in bootstrapping learning based feature extraction and tracking algorithm. This is especially true, since there is a lack of good ground truth for 3D point data in radar datasets. Finally, our proposed methods could also serve other downstream perception tasks, such as detection, recognition and scene understanding.

References

- [1] Raw-roam: Really adverse weather-radar odometry and mapping (python reimplementation of radarslam). <https://github.com/Samleo8/RadarSLAMPy/>, 2025. 8
- [2] Chris Aholt, Sameer Agarwal, and Rekha Thomas. A qcqp approach to triangulation. In *European Conference on Computer Vision (ECCV)*, 2012. 2
- [3] A. Beck, P. Stoica, and J. Li. Exact and Approximate Solutions of Source Localization Problems. *IEEE Transactions on Signal Processing*, 56(5):1770–1778, 2008. Conference Name: IEEE Transactions on Signal Processing. 2
- [4] Amir Beck, Marc Teboulle, and Zahar Chikishev. Iterative Minimization Schemes for Solving the Single Source Lo-

- calization Problem. *SIAM Journal on Optimization*, 19(3): 1397–1416, 2008. 2
- [5] Keenan Burnett, Yuchen Wu, David J Yoon, Angela P Schoellig, and Timothy D Barfoot. Are we ready for radar to replace lidar in all-weather mapping and localization? *IEEE Robotics and Automation Letters*, 7(4):10328–10335, 2022. 2
- [6] Keenan Burnett, David J Yoon, Yuchen Wu, Andrew Z Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith YK Leung, Angela P Schoellig, and Timothy D Barfoot. Boreas: A multi-season autonomous driving dataset. *The International Journal of Robotics Research*, 42(1-2):33–42, 2023. 7, 8
- [7] Martin Byröd, Klas Josephson, and Kalle Åström. Fast optimal three view triangulation. In *Asian Conference on Computer Vision (ACCV)*, 2007. 2
- [8] Yiu-Tong Chan, H. Yau Chin Hang, and Pak-chung Ching. Exact and approximate maximum likelihood localization algorithms. *IEEE Transactions on Vehicular Technology*, 55(1):10–16, 2006. Conference Name: IEEE Transactions on Vehicular Technology. 2
- [9] Joris Domhof, Julian FP Kooij, and Darius M Gavrilă. A joint extrinsic calibration tool for radar, camera and lidar. *IEEE Trans. on Intelligent Vehicles*, 6(3):571–582, 2021. 2
- [10] Jack D Dunitz. Distance geometry and molecular conformation, 1990. 2
- [11] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 3
- [12] Z Jason Geng and Leonard S Haynes. A “3-2-1” kinematic configuration of a stewart platform and its application to six degree of freedom pose measurements. *Robotics and computer-integrated manufacturing*, 11(1):23–34, 1994. 2
- [13] Timo Grebner, Vinzenz Janoudi, Pirmin Schoeder, and Christian Waldschmidt. Self-calibration of a network of radar sensors for autonomous robots. *IEEE Trans. on Aerospace and Electronic Systems*, 59(5):6771–6781, 2023. 1
- [14] Yiduo Hao, Sohrab Madani, Junfeng Guan, Mohammed Aloulah, Saurabh Gupta, and Haitham Hassanieh. Bootstrapping autonomous driving radars with self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15012–15023, 2024. 2
- [15] Kyle Harlow, Hyesu Jang, Timothy D Barfoot, Ayoung Kim, and Christoffer Heckman. A new wave in robotics: Survey on recent mmwave radar applications in robotics. *arXiv:2305.01135*, 2023. 1
- [16] Richard Hartley and Frederik Schaffalitzky. L_{inf} minimization in geometric reconstruction problems. In *Computer Vision and Pattern Recognition (CVPR)*, 2004. 2
- [17] Richard I Hartley and Peter Sturm. Triangulation. *Computer Vision and Image Understanding (CVIU)*, 68(2):146–157, 1997. 2
- [18] Johan Hedborg, Andreas Robinson, and Michael Felsberg. Robust three-view triangulation done fast. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 152–157, 2014. 2
- [19] Ziyang Hong, Yvan Petillot, Andrew Wallace, and Sen Wang. Radar slam: A robust slam system for all weather conditions. *arXiv preprint arXiv:2104.05347*, 2021. 2, 8
- [20] Darya Ismailova and Wu-Sheng Lu. Penalty convex-concave procedure for source localization problem. In *2016 IEEE Canadian conference on electrical and computer engineering (CCECE)*, pages 1–4. IEEE, 2016. 2
- [21] R. Jyothi and P. Babu. SOLVIT: A Reference-Free Source Localization Technique Using Majorization Minimization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2661–2673, 2020. Conference Name: IEEE/ACM Transactions on Audio, Speech, and Language Processing. 2
- [22] Fredrik Kahl, Sameer Agarwal, Manmohan Krishna Chandraker, David Kriegman, and Serge Belongie. Practical global optimization for multiview geometry. *International Journal of Computer Vision (IJCV)*, 2008. 2
- [23] Kenichi Kanatani, Yasuyuki Sugaya, and Hirotaka Niitsuma. Triangulation from two views revisited: Hartley-sturm vs. optimal correction. In *British Machine Vision Conference (BMVC)*, 2008. 2
- [24] Zuzana Kukelova, Tomas Pajdla, and Martin Bujnak. Fast and stable algebraic solution to l2 three-view triangulation. In *International Conference on 3D Vision (3DV)*, 2013. 2
- [25] Martin Larsson, Viktor Larsson, Kalle Astrom, and Magnus Oskarsson. Optimal trilateration is an eigenvalue problem. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5586–5590. IEEE, 2019. 2, 4
- [26] Mao Li, Feng Jiang, and Cong Pei. Review on positioning technology of wireless sensor networks. *Wireless Personal Communications*, 115(3):2023–2046, 2020. 2
- [27] Yunpeng Li, Noah Snavely, and Daniel P Huttenlocher. Location recognition using prioritized feature matching. In *European conference on computer vision*, pages 791–804. Springer, 2010. 6
- [28] Peter Lindstrom. Triangulation made easy. In *Computer Vision and Pattern Recognition (CVPR)*, 2010. 2
- [29] Thomas Lipp and Stephen Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17:263–287, 2016. 2
- [30] Fangfang Lu and Richard Hartley. A fast optimal algorithm for l2 triangulation. In *Asian Conference on Computer Vision (ACCV)*, 2007. 2
- [31] D. Russell Luke, Shoham Sabach, Marc Teboulle, and Kobi Zatlaway. A simple globally convergent algorithm for the nonsmooth nonconvex single source localization problem. *Journal of Global Optimization*, 69(4):889–909, 2017. 2
- [32] Chengqi Ma, Bang Wu, Stefan Poslad, and David R Selviah. Wi-fi rtt ranging performance characterization and positioning system design. *IEEE Transactions on Mobile Computing*, 21(2):740–756, 2020. 2
- [33] Yang Min. L-infinity norm minimization in the multiview triangulation. In *International Conference on Artificial Intelligence and Computational Intelligence*, pages 488–494. Springer, 2010. 2

- [34] Ghina El Natour, Omar Ait-Aider, Raphael Rouveure, François Berry, and Patrice Faure. Toward 3d reconstruction of outdoor scenes using an mmw radar and a monocular vision sensor. *Sensors*, 15(10):25937–25967, 2015. [1](#), [2](#)
- [35] Evgeny Ochinnikov. Fundamentals of structural and functional organization of gnss. In *GPS and GNSS technology in geosciences*, pages 21–49. Elsevier, 2021. [2](#)
- [36] Carl Olsson, Fredrik Kahl, and Magnus Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, 2008. [2](#)
- [37] Juraj Peršić, Ivan Marković, and Ivan Petrović. Extrinsic 6dof calibration of a radar–lidar–camera system enhanced by radar cross section estimates evaluation. *Robotics and Autonom. Syst.*, 114:217–230, 2019. [2](#)
- [38] Akarsh Prabhakara, Tao Jin, Arnav Das, Gantavya Bhatt, Lilly Kumari, Elahe Soltanaghai, Jeff Bilmes, Swarun Kumar, and Anthony Rowe. High resolution point clouds from mmwave radar. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4135–4142. IEEE, 2023. [2](#)
- [39] N. Sirola. Closed-form algorithms in mobile positioning: Myths and misconceptions. In *Navigation and Communication 2010 7th Workshop on Positioning*, pages 38–44, 2010. [2](#)
- [40] Liat Sless, Bat El Shlomo, Gilad Cohen, and Shaul Oron. Road scene understanding by occupancy grid learning from sparse radar clusters using semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#)
- [41] Henrik Stewenius, Frederik Schaffalitzky, and David Nister. How hard is 3-view triangulation really? In *International Conference on Computer Vision (ICCV)*, 2005. [2](#)
- [42] Shigeki Sugimoto, Hayato Tateda, Hidekazu Takahashi, and Masatoshi Okutomi. Obstacle detection using millimeter-wave radar and its visualization on image sequence. In *Intl. Conf. on Pattern Recognition, ICPR*, pages 342–345, 2004. [2](#)
- [43] Tao Wang, Nanning Zheng, Jingmin Xin, and Zheng Ma. Integrating millimeter wave radar with a monocular vision sensor for on-road obstacle detection applications. *Sensors*, 11(9):8992–9008, 2011. [2](#)
- [44] Kyle Wilson and Noah Snavely. Robust global translations with 1dsfm. In *European conference on computer vision*, pages 61–75. Springer, 2014. [6](#)
- [45] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-icp: A globally optimal solution to 3d icp point-set registration. *IEEE transactions on pattern analysis and machine intelligence*, 38(11):2241–2254, 2015. [2](#)
- [46] Shanliang Yao, Runwei Guan, Xiaoyu Huang, Zhuoxiao Li, Xiangyu Sha, Yong Yue, Eng Gee Lim, Hyungjoon Seo, Ka Lok Man, Xiaohui Zhu, et al. Radar-camera fusion for object detection and semantic segmentation in autonomous driving: A comprehensive review. *IEEE Trans. on Intelligent Vehicles*, 2023. [1](#)
- [47] Ryoma Yataka, Pu Wang, Petros Boufounos, and Ryuhei Takahashi. Sira: scalable inter-frame relation and association for radar perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15024–15034, 2024. [2](#)
- [48] De Jong Yeong, Gustavo Velasco-Hernandez, John Barry, and Joseph Walsh. Sensor and sensor fusion technology in autonomous vehicles: A review. *Sensors*, 21(6):2140, 2021. [1](#)
- [49] Ao Zhang, Farzan Erlik Nowruzi, and Robert Laganriere. Raddet: Range-azimuth-doppler based radar object detection for dynamic road users. In *18th Conf. on Robots and Vision (CRV)*, pages 95–102, 2021. [1](#)