


A Novel Approach to Vehicle Pose Estimation using Automotive Radar

Martijn Heller



A Novel Approach to Vehicle Pose Estimation using Automotive Radar

by

Martijn Heller

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday July 1, 2021 at 9:00 AM.

Student number: 4466675
Project duration: December 4, 2020 – July 1, 2021
Thesis committee: Prof. dr. A. Yarovoy, TU Delft MS3, Chairman & Thesis Advisor
Dr. ir. J.N. Driessen, TU Delft MS3
Dr. B. Shyrokau, TU Delft 3mE
Dr. N. Petrov, TU Delft MS3, Daily Supervisor

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

The accurate estimation of the pose, i.e. position and heading, of a vehicle while driving is of high importance in autonomous driving applications. Right now, the main tool to estimate the location of a vehicle is its GPS sensor. However, GPS data is known to be of very low accuracy, especially in urban environments, and is thus not ideal for this application.

In this report, the research into possible improvements to an existing vehicle pose estimation technique are presented, with the aim of making it applicable for automotive radar data. The existing technique is a scan-matching technique known as the Normal Distributions Transform (NDT), which was originally designed for LIDAR measurements. By adapting the technique to accommodate radar measurements, some of the drawbacks of LiDAR, such as the high cost and poor performance in certain weather conditions, can be overcome. Some of the main disadvantages of using radar compared to LiDAR, e.g. lower resolutions, are addressed in the presented techniques. The lower resolution results in significant spreading of the target response, this is especially prominent in the azimuth domain for a standard 3 Tx \times 4 Rx MIMO automotive radar system. By addressing the scan-matching problem in the polar domain, this spreading of the target response is better captured in the distribution used to perform the scan-matching. Further, the implementation in polar coordinates allows for incorporation of the Doppler measurements, which contain knowledge about the angles of arrival of targets and are generally measured at a much higher resolution than the angular measurements themselves. Moreover, the use of radar measurements introduces the availability of knowledge about the Radar Cross-Section of individual targets, this can be used to reduce the influence of false alarms.

The incorporation of Doppler additionally allows the exploitation of the relation between the Doppler measurements and the angle of arrival to estimate a bias in the angle measurements, this can be used for sensor calibration while driving.

The influence of the presented improvements to the NDT are examined through simulations and experiments. These results show significant reduction in the estimation errors. The sensor bias estimation technique also proves to provide extremely accurate estimates. Finally, a small extension is worked out to perform trajectory estimation using the individual poses by means of the Extended Kalman Filter.

Acknowledgements

I would like to thank my thesis advisor Alexander Yarovoy, he always took the time for our meetings and was very helpful when it came to deciding the further progress of the research.

I would also like to thank my daily supervisor Nikita Petrov, he always provided clear explanations when something was not clear to me. He kept up to date and steered me towards the most productive path. The feedback he provided each time significant progress was made was extremely valuable.

I would like to thank Yun Lu, who made time to work on a solution for the integration of GPS and radar data which would have served in the gathering of new experimental data had time permitted. And Hans Driessen, who helped work out the definition of the recursive filtering solution for construction of the trajectory.

Finally, I would like to thank my friends, my family and my girlfriend for the much needed distraction, support and motivation throughout my research.

Contents

List of Figures	ix
List of Tables	ix
1 Introduction	1
1.1 Background and Motivation	1
1.1.1 Background: Scan-matching Using Range Scans	1
1.1.2 Motivation: From LiDAR to Radar	2
1.2 Thesis Objective	3
1.3 Outline of the Report	3
2 State of the Art	5
2.1 Simultaneous Localization and Mapping	5
2.2 Scan-matching	7
2.2.1 Feature-based Methods	7
2.2.2 Distribution-based methods	8
2.3 On-line Sensor Calibration	9
3 The Conventional Normal Distributions Transform	11
3.1 The Piecewise Continuous Distribution	11
3.2 The Point Cloud Mapping Equations	12
3.3 The Optimization	13
3.4 Discussion	13
4 The Radar Normal Distributions Transforms	15
4.1 The Polar Normal Distributions Transform	15
4.1.1 The Piece-wise Continuous Distribution	15
4.1.2 The Point Cloud Mapping Equations	16
4.1.3 The Optimization	16
4.2 The Doppler Polar Normal Distributions Transform	16
4.2.1 The Piece-wise Continuous Distribution	17
4.2.2 The Point cloud Mapping Equations	17
4.2.3 The Optimization	18
4.2.4 Target Filtering using Doppler Measurements	18
4.3 The Range-Doppler Polar Normal Distributions Transform	18
4.4 Incorporation of Signal-to-Noise Ratio Measurements	18
4.4.1 Using Knowledge of the RCS in the Calculation of the Distribution	19
4.4.2 Using Knowledge of the RCS in the Calculation of the Score	19
4.5 Discussion	20
5 The Optimization Problem	21
5.1 The Objective function	21
5.2 The Hessian and Gradient of the Objective Function	28
5.3 The First- and Second-order Derivatives	28
5.3.1 NDT	28
5.3.2 PNDDT	29
5.3.3 DPNDT	30
5.4 Local Minima and not Positive Definite Hessian	31
6 Joint Pose Estimation and Calibration	33
6.1 Angular Bias as an Estimation Parameter	33
6.2 The Hessian and Gradient	33
6.3 Convergence to the Correct Minimum	34

7	Verification of the Performance of the NDT Techniques	37
7.1	Simulations	37
7.1.1	Set-up of the Simulations	37
7.1.2	Comparison of the Different Pose Estimation Techniques	39
7.2	Experiment	47
7.2.1	Set-up of the Experiment.	47
7.2.2	Comparison of the Different Pose Estimation Techniques	47
8	Calculation of the Trajectory	55
9	Conclusion	57

List of Figures

3.1	NDT Distribution	12
4.1	Comparison of the NDT and PNDT Distributions	16
4.2	Polar Mapping Equations	17
4.3	RCS-aware (P)NDT distributions	19
5.1	Convexity Analysis: NDT and PNDT	23
5.2	Convexity Analysis: DPNDT and RDPNDT	24
5.3	Convexity Analysis: DPNDT and RDPNDT, Velocity Related	25
5.4	Convexity Analysis: Interaction between t_y and ϕ	26
5.5	Convexity Analysis: Larger Cell Size	26
5.6	Convexity Analysis: Consideration of the RCS	27
5.7	Convexity Analysis: Saddle Point	32
6.1	Convexity Analysis: Offset in the Minimum Corresponding to the Bias	34
6.2	Convexity Analysis: Correct Minimum Corresponding to the Bias	35
7.1	Set-up of the Primitive Simulations	38
7.2	Set-up of the Realistic Simulations	40
7.3	Results of the Primitive Simulations: Pose Estimation	41
7.4	Results of the Primitive Simulations: Estimated Sensor Bias	43
7.5	Results of the Primitive Simulations: Estimated Sensor Bias for a Higher Doppler Resolution	44
7.6	Results of the Realistic Simulations: Pose Estimation	45
7.7	Set-up of the Experiment: Front of the Car	47
7.8	Set-up of the Experiment: View from the Car	48
7.9	Results of the Experiment: Reconstructed Trajectories Taking into Account RCS via the Received Power	50
7.10	Results of the Experiment: Reconstructed Trajectories Taking into Account RCS via the Received Power and the Range	51
7.11	Results of the Experiment: Reconstructed Trajectories for Varying Initialization Parameters	52
7.12	Results of the Experiment: Reconstructed Maps Made from the Reconstructed Trajectories	53

List of Tables

7.1	Initialization parameters of the NDT techniques	37
7.2	Chosen Initialization parameters of the NDT techniques	39
7.3	Results of the Primitive Simulations: Optimal Settings	42
7.4	Results of the Primitive Simulations: Reduction in RMSE compared to the NDT	42
7.5	Results of the Primitive Simulations: Normalized Execution Times of the Radar-appropriate NDTs	42
7.6	Results of the Primitive Simulations: Sensor Bias Estimation	43
7.7	Results of the Primitive Simulations: Sensor Bias Estimation for a Higher Doppler Resolution	43
7.8	Results of the Realistic Simulations: Optimal Settings	44

7.9 Results of the Realistic Simulations: Reduction in RMSE compared to the NDT	45
7.10 Results of the Realistic Simulations: Normalized Execution Times of the Radar-appropriate NDTs	46
7.11 Results of the Experiment: Normalized Execution Times of the Radar-appropriate NDTs for the Experimental Data	52
7.12 Results of the Experiment: Sensor Bias Estimation	54
7.13 Results of the Experiment: Sensor Bias Estimation	54

1

Introduction

Developments in the automotive industry show a trend towards a higher level of automation. Current estimates predict that in 2025 approximately 2% of car sales in the UK alone will comprise of vehicles capable of conditional automation. By 2035 it is projected that 40% of car sales will be Connected and Autonomous Vehicles (CAVs) [1]. CAV refers to a vehicle which accommodates autonomy levels of 3 (conditional automation) and above as defined by the SAE International Standard J3016 [2]. From this 40% around 65% is expected to refer to cars with autonomy levels 4 and 5, meaning high (level 4) to full (level 5) automation will be the standard in these types of vehicles. Examples of high automation are human activated systems such as automated driving on the freeway with presence of the driver or even automated parking in absence of the driver. Full automation refers to a fully automated driving system without intervention of a human driver [2]. These developments pose additional requirements on the on-board sensor systems, as they are expected to accommodate more complex tasks such as object recognition and localization in urban environments [3].

1.1. Background and Motivation

The most used approach to perform localization in automotive applications is through the Global Positioning System (GPS). However, especially in urban environments, GPS does not provide high enough accuracy for automated driving application due to multi-path interference and non-line-of-sight reception [4]. A technique known as differential GPS uses ground-based reference stations to improve on accuracy, however the accuracy of such systems is only around 1-3 meters [5]. In the applications of highly and fully automated driving, such low accuracy is unacceptable. Solutions to the localization problem exist that make use of range scanning sensors to estimate the vehicle location within a map of its surroundings. This map can be made a priori [6] or simultaneous with the estimation of the vehicle location, which is known as the Simultaneous Localization and Mapping (SLAM) problem [7]. The area of SLAM can be divided into two main groups, techniques that perform full SLAM and techniques that perform online SLAM. The methods for solving full SLAM are considered to be too complex to perform in real-time, whereas the online SLAM methods are well suited for real-time implementation. Over longer distances, SLAM techniques are known to suffer from drift in their location estimation, this is especially common in online SLAM [4]. A way to compensate for this drift is by using a scan-matching approach. Scan-matching techniques estimate the pose, i.e. position and heading, of the vehicle relative to a different point in time by maximizing the overlap between the respective range scans. Scan-matching does not require the creation and maintenance of a map of the surroundings [8].

1.1.1. Background: Scan-matching Using Range Scans

Scan-matching techniques use the point clouds that result from 2D or 3D range scans to estimate the relative pose. Range scans at two different time instances are considered, the so-called “reference scan” and the “current scan”. The objective is to find the relative pose of the vehicle between these two scans. The relative pose is estimated by finding the transformation of the current scan point cloud that results in maximum overlap with the reference scan and in turn relating this transformation of the point cloud to a transformation in the pose of the vehicle [9].

The scan-matching techniques can be divided into two main groups, the feature-based techniques [10], and the distribution-based approaches [11]. The feature-based approaches aim to find the maximum overlap between the scans by minimizing the Euclidean distance between individual features, such as points, lines or surfaces, explicitly. The distribution-based techniques overcome the required knowledge about the correspondence of individual points as they represent the reference scan as a piece-wise continuous distribution, after which the point of maximum overlap is found by maximizing the probability of the points of the current scan to be at a certain location within this distribution.

Pioneering the distribution-based approaches is the Normal Distributions Transform (NDT) [11]. The NDT converts the 2D range scans from the native polar coordinate frame to Cartesian coordinates, after which the reference scan is represented by a combination of bivariate Gaussian distributions related to the distribution of the points within the cells of a grid on the xy plane. The representation of the reference scan on a fixed Cartesian grid comes with drawbacks such as convergence to local minima which have been addressed in proposed techniques [12], [13]. Additional extensions have been proposed to perform scan-matching using three-dimensional scans [14], [15].

1.1.2. Motivation: From LiDAR to Radar

Due to the high resolution of laser scanners (LiDAR), especially in the angular domain [3], [15], [16], and due to the popularity of both laser sensors and scan-matching techniques within the robotics community, existing scan-matching techniques have been optimized to accommodate LiDAR technology. Laser scanners, however, come at a higher cost [16] and perform poorly in bad weather conditions such as heavy rain or fog [17]. Another line of research is devoted to the application of scan-matching techniques to (stereo) cameras [18], [19], which also become sensitive to weather and lighting conditions when mounted on a car.

These shortcomings of laser and optical sensors can be overcome using the mm-wave radar sensors available in the majority of cars with a high automation level. However, the existing SLAM and scan-matching techniques cannot be applied to radar measurements directly. Some work has been done on radar SLAM using scan-matching, using multiple previous scans [20], [21] or visual features [22]. The point-to-point scan-matching approaches are especially not well suited for radar measurements, due to the fluctuations in radar cross section (RCS) for changing observation angles. The RCS of a target is highly dependent on the observation angle, especially in the case of a complex shape [23, Ch. 6]. This in turn results in inconsistency in the reflected power of the same target after movement of the vehicle, which can result in missed detections in one of the scans used for scan matching, causing the point clouds to suffer from “floating points” – the primitive detections which do not have a counterpart in the other scan. Another problem that is encountered when using radar scans is the low angular resolution. In radar systems the angular resolution is determined by the antenna dimensions in terms of the wavelength. For uniform linear arrays (ULA) or MIMO radars with virtual arrays, it is determined by the number of spacial channels. The current generation of automotive radar performs distinction in the azimuth domain realized via 3 Tx \times 4 Rx MIMO arrays [24], [25]. This results in a 12-element virtual array capable of angle of arrival estimation with an angular resolution of the order of approximately $\Delta\theta_{radar} = 10^\circ$ at the broadside of the radar [26], degrading with deviation from the broadside. This angular resolution of radar is significantly worse than LiDAR measurements, which have an angular resolution of an order of $\Delta\theta_{LiDAR} = 1^\circ$ [15]. The low angular resolution of radar results in substantial spreading of the target response in the cross-range dimension, depending on the target’s range which results in inaccuracies in the estimation of the angle of arrival. The response looks even more complicated after transformation to the Cartesian grid for conventional NDT, which does not take this effect into account.

The standard output of radar (considering standard range-Doppler-angular processing followed by detection) offers additional information, which can be beneficial for localization. Of particular importance can be the RCS of detected objects and their relative radial velocity with respect to the radar due to Doppler processing. The access to Doppler measurements is foreseen to give an improvement in cross-range resolution for stationary targets. In application to forward-looking radar, this processing is referred to as Doppler Beam Sharpening [27], [28]. These particular aspects of radar data are however not exploited in the current state-of-the-art localization techniques.

1.2. Thesis Objective

This thesis details the steps towards a novel scan matching technique for vehicle localization using radar data. The scan matching solution of the NDT will be addressed in polar coordinates, which are native for the sensor, directly. The proposed technique, called polar normal distribution transform (PNDT), allows for the addition of Doppler measurements into the scan-matching algorithm for further improvement in localization accuracy. This extension will be referred to as Doppler polar NDT (DPNDT). Further, the ability to incorporate targets RCS into the localization problem is addressed along with the benefits of this approach. Finally, it is shown that the incorporation of Doppler measurements enforces a particular structure of the data, which gives the potential for the estimation of angular bias while the vehicle is driving.

1.3. Outline of the Report

First, the state of the art in vehicle localization is discussed in Chapter 2 through a literature review. Then, in Chapter 3, the idea behind the conventional Normal Distributions Transform (NDT), along with the changes needed to better suit radar scans are discussed briefly. The incorporation of solutions to the problems with the NDT when applied to radar measurements are addressed in Chapter 4. Chapter 5 describes the optimization that is performed in the NDT techniques. The extension of the proposed techniques to perform sensor bias estimation is discussed in Chapter 6. The newly introduced techniques are compared to each other and to the conventional NDT through simulations and experimental data, this is treated in Chapter 7. In Chapter 8, a way to use dynamic model-based filtering to improve the estimated trajectory is explored and finally, in Chapter 9, conclusions are drawn and suggestions for future work are made.

2

State of the Art

Using measurements of range scanning sensors to estimate the location of a vehicle within a map of the surroundings has been researched extensively. Initially, the map used for localization had to be made a priori, such as in the position estimation solution of [6]. Soon hereafter the Simultaneous Localization and Mapping (SLAM) problem was introduced, which aims to create a map of the surroundings while simultaneously estimating the position of the vehicle. The SLAM problem is paradoxical, in the sense that in order to accurately estimate the position of the measurement set-up within a certain map of the surroundings, this map should be extremely accurate, while in creating such a map from measurements taken from the measurement platform, the position of the platform has to be known accurately [7].

2.1. Simultaneous Localization and Mapping

In recent years, many solutions to the SLAM problem have been proposed. So-called full SLAM techniques aim to construct the whole trajectory and map using all measurements and control inputs, whereas online SLAM methods typically only use the current sensor inputs. The full SLAM problem is most often solved using the so-called optimization-based approaches. Examples of this are bundle adjustment and graph SLAM. The working principle of optimization-based SLAM solutions is a two-step approach. First, the measurements and control inputs are used to construct constraints after which both the vehicle location and the map are optimized to conform to these constraints. In bundle adjustment, this problem is tackled by minimizing a score function based on the Euclidean distance between the predicted projection of a set of points and the corresponding measured points. Such an optimization problem can be solved using linearized least-squares as in [29] or using non-linear least-squares optimization techniques such as Levenberg-Marquardt [30]. In graph-SLAM a set of nodes is used to represent vehicle poses or landmarks. These nodes are connected by edges representing a set of constraints between these nodes. The constraints describe the relation between two nodes based on measurements and are accompanied by an information matrix representing their uncertainty. Taking nodes i and j , a prediction can be made at node i of the observations from node j , the likelihood of all observations is in turn maximized taking into account the measurement uncertainty and the error between this prediction and the constraints [31]. The contents of the graph structure dictate the type of graph-SLAM that is considered. In general graph SLAM, the edges represent observations of landmarks, the nodes are both poses and landmarks. In pose-graph SLAM the nodes represent the vehicle poses and the edges represent relative poses. Solutions to the graph-based SLAM problem have been researched extensively, from solutions regarding the two-dimensional domain [32], [33] to solutions extended to three-dimensional movement [34], causing a much more involved derivation for the before-mentioned error. Since most of the full SLAM techniques consider the entire collection of information, they are considered too computationally complex for real-time applications [4].

Online SLAM is solved using filter-based approaches, making them well suited for real-time implementation. Filter-based approaches rely on a prediction step and an update step. The prediction step is based on a dynamic evolution model of the vehicle and the control inputs and produces a predicted estimate of the vehicle and map states. The update step in turn uses the sensor inputs to adjust these estimates. Assuming linearity, the models take the following form:

$$\begin{aligned}\mathbf{s}^{[t]} &= \mathbf{F}\mathbf{s}^{[t-1]} + \mathbf{w}^{[t-1]}, \\ \mathbf{z}^{[t]} &= \mathbf{G}\mathbf{s}^{[t]} + \mathbf{n}^{[t]}.\end{aligned}\tag{2.1}$$

Here, \mathbf{s} is the to be estimated state of the system and \mathbf{z} a vector containing observations. Superscript $(\cdot)^{[t]}$ denotes the time index t . The state at the current time index $\mathbf{s}^{[t]}$ is related to the state at the previous time index $\mathbf{s}^{[t-1]}$ through the dynamic model characterized by \mathbf{F} . The observation at time t , $\mathbf{z}^{[t]}$, is related to the state $\mathbf{s}^{[t]}$ through the observation model characterized by \mathbf{G} . \mathbf{w} and \mathbf{n} denote process noise and measurement noise, respectively.

So-called Bayesian filters use these models to produce a mean and covariance of the state estimate, using weights based on both the dynamic and observation model uncertainties through \mathbf{w} and \mathbf{n} . Over the years, a variety of implementations of these Bayesian filters have been proposed, some of which are the Kalman filter, the extended Kalman filter, the Unscented Kalman filter and the particle filter.

The Kalman filter utilizes the assumption made in the definition of (2.1) that both the dynamic and observation model are linear along with the assumption that all errors are Gaussian [35], [36]. In doing so, a recursive solution to the estimation problem can be derived, consisting of the following prediction step and update step:

$$\mathbf{s}^{[t+1|t]} = \mathbf{F}\mathbf{s}^{[t]},\tag{2.2}$$

$$\mathbf{Q}^{[t+1|t]} = \mathbf{F}\mathbf{Q}^{[t|t]}\mathbf{F}^T + \mathbf{W}^{[t]}.\tag{2.3}$$

$$\mathbf{s}^{[t|t]} = \mathbf{s}^{[t|t-1]} + \mathbf{K}^{[t]}(\mathbf{z}^{[t]} - \mathbf{G}\mathbf{s}^{[t|t-1]}),\tag{2.4}$$

$$\mathbf{Q}^{[t|t]} = (\mathbf{I} - \mathbf{K}^{[t]}\mathbf{G})\mathbf{Q}^{[t|t-1]},\tag{2.5}$$

$$\mathbf{K}^{[t]} = \mathbf{Q}^{[t|t-1]}\mathbf{G}^T(\mathbf{G}\mathbf{Q}^{[t|t-1]}\mathbf{G}^T + \mathbf{N}^{[t]})^{-1}.\tag{2.6}$$

Here, \mathbf{s} and \mathbf{Q} denote the estimated mean and covariance of the state and \mathbf{W} and \mathbf{N} are the covariance matrices characterizing the Gaussian noise vectors.

In order to apply this to SLAM, both the dynamic and the observation models have to be derived for the vehicle and the map features. A vehicle in motion and the observations of landmarks from this moving vehicle cannot be modelled linearly, causing the relations in (2.1) to describe nonlinear relations, i.e. $\mathbf{s}^{[t]} = \mathbf{f}(\mathbf{s}^{[t-1]}) + \mathbf{w}^{[t-1]}$ and $\mathbf{z}^{[t]} = \mathbf{g}(\mathbf{s}^{[t]}) + \mathbf{n}^{[t]}$. To solve this, SLAM uses the Extended Kalman Filter (EKF) instead of the standard Kalman Filter [37], [38]. In the EKF, a linearization of the non-linear models is performed through their Jacobians around the current estimates. This results in linear functions which can be used in the standard Kalman Filter algorithm. To accommodate models with very extreme non-linearities the Unscented Kalman Filter (UKF) can be used [39]. By taking samples according to a predefined algorithm that represents the estimated mean and covariance, the UKF propagates these points through the non-linear (and possibly non-Gaussian) system, using the points at the output to calculate the mean and covariance as subjected to the non-linear system. This concept can be applied to the state transition and observation models of the SLAM Kalman Filter which is presented in [40]. Another on-line SLAM technique that is robust to non-linearities and non-Gaussian noise in the models is Particle Filter SLAM. Similar to the UKF, the Particle Filter (PF) uses samples taken from the distributions characterized by the system noise (so-called ‘‘particles’’) to represent the estimated mean and covariance. Rather than using these particles to calculate the mean and covariance and continue with the standard KF approaches, the PF keeps the representation as particles. Additionally, in the PF, the points are taken from a random distribution defined by the system, rather than being chosen via a predefined algorithm [41]. A particle filter is employed in a technique known as FastSLAM. FastSLAM relies on the fact that for a known path, the posterior that is estimated can be decomposed for each of the individual landmarks. FastSLAM exploits this fact by using a particle filter to represent the possible paths and for each path apply the EKF to each individual landmark to estimate its position. With efficient implementation, this results in $\log(N)$ complexity for N landmarks, rather than N^2 as in the EKF [42].

In recent years some solutions to lower the computational complexity of using graph SLAM have been proposed. A combination of a filter-based approach with a graph-based representation of the map

is proposed in [43], known as the Sparse Extended Information Filter SLAM. By removing edges in the graph that can be deactivated, e.g. those connected to features that are far away from the vehicle, the information matrix is sparsified and the problem can be solved in real time. Similarly, by removing edges within the pose graph that are not relevant at the current pose, the pose graph can be converted into a pose-chain, significantly reducing the computational complexity, this is used in COP-SLAM [44].

2.2. Scan-matching

Scan-matching techniques can be used to compensate for errors in the estimations from SLAM that result from increasing uncertainty over time. The objective of scan-matching techniques is to find the transformation, i.e. translation and rotation, of a vehicle between two scans at different points in time. These two scans are referred to as the “reference scan” and the “current scan”. By iteratively transforming the point cloud of the current scan until the maximum overlap with the reference scan is found, the relative pose \mathbf{p} of the vehicle between the two scans can be found through its relation with the transformation of the point cloud. In most cases only the relative translation and rotation of the vehicle are of interest, leading to the relative pose being defined as $\mathbf{p} = [t_x, t_y, \phi]^T$. This vector contains the translation in the x -direction t_x , the translation in the y -direction t_y and the change in heading ϕ at the current scan, relative to the reference scan. t_x is defined in the direction of the heading of the car at the reference scan and t_y perpendicular to that. The change in heading ϕ is defined to be positive for rotation to the left and negative for rotation to the right. Since scan matching relies on only two scans per estimation, the creation and maintenance of a map of the surroundings is not necessary [8].

The iterative process to find maximum overlap takes an estimate $\hat{\mathbf{p}}$ at each iteration to map the point cloud of the current scan. $\hat{\mathbf{p}}$ can be an initialization vector or the estimate from a previous iteration in the estimation process. By then defining certain criteria related to the overlap, optimization techniques are used to converge to an estimate $\hat{\mathbf{p}}$ that results in maximum overlap between the point clouds. The scan-matching techniques can be divided into two main groups related to the criteria that are used for defining the overlap; feature-based and distribution-based techniques. The feature-based approaches try to minimize the distance between individual features within the scans while the distribution-based approaches represent one of the scans as a density which can be used to express a likelihood of overlap with the other scan.

2.2.1. Feature-based Methods

In feature-based scan-matching methods, the aim is to match certain features within the scans explicitly. The criterion for overlap is in this case the Euclidean distance between individual points, lines or surfaces. One of the first feature-based matching solutions is the Iterative Closest Point (ICP) algorithm [10]. It presents a technique to match range scans from a data set P to a free-form, complex three-dimensional model defined by a data set X by minimizing the distance between the range measurements and the surface of the model. The prerequisites for applying such an approach are derived first; formulas describing the distance of a point to different geometric shapes are derived and methods to compute the point within these shapes that is closest to an arbitrary point are discussed. A derivation of the least squares problem minimizing the distance between N_p points $p_i \in P$ to N_x corresponding points $x_i \in X$ with $N_p = N_x$ is presented. The least squares problem is formulated in terms of the spatial mapping of P to X , expressed as a translation vector and a rotation matrix:

$$\min_{\mathbf{q}} \frac{1}{N_p} \sum_{i=1}^{N_p} \|x_i - \mathbf{R}(\mathbf{q}_R)p_i - \mathbf{q}_T\|^2, \quad (2.7)$$

with $\mathbf{R}(\mathbf{q}_R)$ the rotation matrix based on rotation vector \mathbf{q}_R and \mathbf{q}_T the translation vector. Both the measurements and the model are discretized (curves become sets of lines, surfaces become sets of triangles) and in turn transformed to point sets. The result is two sets of points, P and X . For each point p_i , the point $y_i \in X$ that minimizes the Euclidean distance of p_i to X is found. The set of resulting correspondence points $Y = \mathcal{C}(P, X)$ is then used to solve the least squares problem in order to find the translation vector and rotation matrix. \mathcal{C} denotes the closest point operator. The spatial mapping is applied to P and the calculations are performed again until convergence. In essence, the distance between individual points is minimized which in the end results in the distance between the surfaces being minimized.

In [8] two techniques are presented to perform relative pose estimation for robots using two range scans, one of which is very similar to the ICP algorithm. The first technique uses tangent lines of individual points. An objective function is derived which relates the normals of the tangent lines to each other through a translation \mathbf{T} and a rotation $\mathbf{R}(\omega)$, where ω is the rotation angle. The objective function is solvable for the \mathbf{T} when ω is known. In order to perform the optimization, ω is estimated by a search on the objective function using the golden section method after which \mathbf{T} is optimized for this value of ω . This alternating approach is continued until convergence. The second technique is a point-to-point method, the criterion deciding correspondence between points in the respective scans extends on the ICP algorithm. Where ICP only uses the closeness of data points to model points the proposed algorithm also includes the assumption that corresponding points fall within a certain region when it comes to the measured bearing, this bound is determined by a predefined maximum rotation, along with the fact that their measured ranges are closest together. The technique is named the Iterative Dual Correspondence (IDC) algorithm and is found to perform better than the ICP, especially in case of curved shapes. The authors suggest to use the two presented techniques together, using the tangent line method, which can handle a large initial pose error but produces less accurate estimates, to obtain a crude estimation of the pose to be used in IDC, which results in more accurate estimates.

Apart from the additional criterion proposed in IDC, both of the aforementioned techniques operate in the Cartesian coordinate frame to perform scan-matching. A technique which utilizes the native polar coordinate frame of range scans is presented in [9]. The technique, known as Polar Scan Matching (PSM), uses a sparse set of laser measurements to perform the scan matching. At each iteration, the current scan is projected according to an estimate of the pose:

$$\begin{aligned} r'_m &= \sqrt{[t_x + r_m \cos(\phi + \theta_m)]^2 + [t_y + r_m \sin(\phi + \theta_m)]^2}, \\ \theta'_m &= \text{atan2}[t_y + r_m \sin(\phi + \theta_m), t_x + r_m \cos(\phi + \theta_m)]. \end{aligned} \quad (2.8)$$

The resulting points are mapped to the angle measurements of the reference scan using linear interpolation of the measurements of the current scan. In the case that multiple range values are mapped to a single bearing, only the closest one is maintained. The resulting set contains points with bearings equal to those measured at the reference scan and at most one interpolated range per angle. The residuals are then defined as the difference between the interpolated range and the measured range in the reference scan for each bearing angle. The sum of these residuals is to be minimized in order to estimate the translation. The rotation is estimated through a search approach shifting the set of ranges and bearings over the angle until a point of minimum average range residuals is found.

2.2.2. Distribution-based methods

In 2003 Peter Biber and Wolfgang Straßer proposed the Normal Distributions Transform (NDT) [11]. In the NDT, a score function is used as the criterion of overlap between the scans. This score function is then maximized in order to find the pose \mathbf{p} corresponding to maximum overlap. This score value is calculated by assessing the probability of each point in the point cloud of the current scan to be at a certain position within the reference scan after mapping. In order to do this, the reference scan is represented as a combination of distributions derived from the Normal distribution.

The distribution is calculated by transforming the range scans, which are received as sets of points with their respective range and angle values, to the Cartesian coordinate frame, on which a grid is constructed. For each cell C_k within this grid with a minimum of three points, a two-dimensional vector \mathbf{q}_k and matrix Σ_k are calculated representing the mean and covariance, respectively, of the positions (the x - and y -coordinates) of the points within the cell. Using this mean vector and covariance matrix, the positions of the points within a cell are approximated by a Gaussian distribution. This Gaussian distribution is used to calculate a score for each point x_m in the current scan after mapping of the current scan point cloud using an estimate of \mathbf{p} :

$$\text{score}(x'_m) = \exp \left[-\frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \Sigma_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right]. \quad (2.9)$$

This calculated value is related to the probability of the point m , characterized by its x and y coordinates \mathbf{x}'_m , to be within a certain cell of the grid on the reference scan. Formulas are derived to relate the pose change \mathbf{p} and the original position x_m to the resulting position \mathbf{x}'_m . These formulas

form the relation between the score value of (2.9) and \mathbf{p} , this relation is used to find the estimate \mathbf{p} that maximizes the sum of these score values over all points M . In the paper, a method to use the NDT for SLAM is also explained. The precise workings of the NDT are explained in detail in Chapter 3.

A technique using this distribution representation of the NDT for Monte-Carlo Localization (MCL) is presented in [45]. MCL makes use of a particle filter to estimate the posterior of the pose given a map, the controls and the observations. In traditional MCL an occupancy grid map is used, in the proposed technique both the map and the measurements are represented as the Normal Distributions Transform, using them for calculation of weights for the particles in the particle filter. Results show very accurate performance of the NDT-MCL, outperforming the traditional grid-based MCL in all tests.

The calculation of the distribution over a fixed Cartesian grid has been shown to lead to convergence to local minima. Solutions to these problems have been proposed in [12] and [13], which use k -means clustering to group points together after which the distribution is calculated for each cluster rather than for each cell. In [12] the clusters are formed on the xy -plane, iteratively decreasing the cluster size to guarantee convergence to the global minimum, whereas [13] makes use of the fact that range scans are often received ordered by the angle of incidence, which can be used to cluster the points of consecutive angle bins with a predefined maximum distance in range to each other. Experimental laser scan data shows that the polar clustering technique outperforms the Cartesian clustering technique considerably [13].

For scan-matching using three-dimensional range scans the relation between the measured point cloud and the relative pose of the vehicle changes significantly. These relations have been derived for the underlying application of underground mining by Magnusson et al. [14]. In the paper, special attention is paid to some of the problems with the conventional NDT. An in-depth analysis of different grid construction techniques such as additive, iterative and octree subdivision as well as the choice of the cell size and sampling method are presented. In this paper, a size constraint was applied to the descent step $\Delta\mathbf{p}$ to prevent divergence. In his dissertation Magnusson treats this 3D NDT technique in much more detail [15].

2.3. On-line Sensor Calibration

As part of the thesis, research was devoted to radar sensor calibration techniques suitable for calibration while driving.

A technique to estimate errors in the position of a sensor is explained in [46]. In this thesis, Direction of Arrival (DoA) estimation using a sparse representation of the array manifold is presented. A small portion of the thesis is devoted to calibration using this DoA approach. By including the explicit dependence of the measured bearings on the erroneous coefficient (in this case the position of the sensor), an error function can be constructed based on the error between the measurements and the array manifold, which now depends on both the DoA estimates and the position of the sensor:

$$\min_{\boldsymbol{\theta}, \mathbf{p}} \sum_{t=1}^T \|\mathbf{y}(t) - \mathbf{A}(\boldsymbol{\theta}, \mathbf{p})\mathbf{u}(t)\|_2^2. \quad (2.10)$$

In this expression, t represents a time index, \mathbf{y} represents the received signal, \mathbf{u} represents the baseband signal and \mathbf{A} the array manifold, now depending on the bearing vector $\boldsymbol{\theta}$ containing the bearings of the individual targets and the to be calibrated sensor position \mathbf{p} . Using block-coordinate descent both the position and the angle of arrival are estimated; alternating using each other's optimal values until convergence.

A combination of sensor calibration and SLAM is presented in [47]. The estimation of calibration coefficients is done by assuming one array element to be calibrated correctly and including the real and imaginary parts of the calibration coefficients of each of the other array elements in the state vector of the EKF SLAM algorithm. In order to do so, the observation model is extended to include the received signal power per element, normalized to the correctly calibrated element, which depends on the real and imaginary calibration coefficients. The dynamic model is extended to include additive Gaussian noise on the calibration coefficients. Both simulations and experiments show an accurate estimation of the calibration coefficients after only a few tens of measurements.

In robotics odometry can be used to initialize localization techniques. Odometry data often contains the velocity commands provided to the left and right wheel, which can be used, together with the wheel

distance and radii to calculate the traversed path. However, in the case that the wheels of the robot do not have equal radii, e.g. due to an uneven load distribution, upon control of the robot to drive in a straight trajectory, odometry data would indicate a straight path while the robot's path will have taken a slight turn. In [48] an approach to estimate the radii of the wheels of a robot is presented which makes use of pose-graph SLAM by incorporating the radii and wheel distance in the state vector. The resulting error function is linearized around the current initial guess through first order Taylor expansion which can be solved using techniques such as Gauss-Newton and Levenberg-Marquardt. Both simulations and experiments show that the wheel radius can be estimated with high accuracy. Such a technique could be used to estimate a sensor bias.

3

The Conventional Normal Distributions Transform

Despite its drawbacks, the NDT poses a very appropriate starting point for a radar-applicable scan matching technique due to its simplicity and its robustness against missed detections, which becomes a more prominent problem with its application to radar measurements.

The Normal Distributions Transform (NDT), as first proposed by Biber and Straßer in [11], aims to find the point of maximum overlap by representing the point cloud of the reference scan by a piece-wise continuous distribution that is closely related to the Gaussian or Normal probability distribution, hence the name of the NDT, and iteratively maximizing the overall probability of the points in the point cloud of the current scan inside this distribution. In doing so, the scan-matching problem can be expressed as maximizing the likelihood of overlap between the scans given the relative pose \mathbf{p} , this likelihood is expressed using the distribution resulting from the NDT.

The steps of the maximization are as follows: First a grid is constructed over the reference scan which is used to create the piecewise continuous distribution. Based on this distribution, a score function is proposed that aims to maximize the probability of overlap. Formulas are derived that relate the relative pose \mathbf{p} of the vehicle to a transformation of the points inside the point cloud and then by iteratively transforming the point cloud, the point of maximum overlap is found using Newton's method.

3.1. The Piecewise Continuous Distribution

The piece-wise continuous distribution is created by first converting the polar range scans to the Cartesian coordinate frame. On the, now Cartesian, reference scan four overlapping grids are constructed by creating a grid and shifting it by half a cell width in the x -direction, the y -direction and both the x - and y -direction. For simplicity, the following parts discuss the steps for only one grid, the steps are repeated for each grid and later averaged over the grids. Cells which contain at least three points will be used to construct the distributions, cells with less than three points will be represented by a probability of zero throughout the cell. In order to represent each cell as a Gaussian distribution, the points inside a cell C_k are used to calculate the mean and covariance in the x - and y -direction. The result is a two-dimensional mean vector $\mathbf{q}_k \in \mathbb{R}^{2 \times 1}$ and a two-dimensional covariance matrix $\mathbf{\Sigma}_k \in \mathbb{R}^{2 \times 2}$ for each cell C_k with three or more points. \mathbf{q}_k and $\mathbf{\Sigma}_k$ are used as the expected values and covariances in the 2-dimensional Gaussian-based representation of the reference scan. The probability of measuring a point m in the point cloud of the reference scan at a certain position $\mathbf{x}_m = [x_m, y_m]^T \in C_k$ is then related to this mean and covariance via [11]:

$$p(\mathbf{x}_m) \sim \exp \left[-\frac{(\mathbf{x}_m - \mathbf{q}_k)^T \mathbf{\Sigma}_k^{-1} (\mathbf{x}_m - \mathbf{q}_k)}{2} \right]. \quad (3.1)$$

This expression is closely related to the probability density function of a bivariate Gaussian distribution characterized by mean \mathbf{q}_k and covariance $\mathbf{\Sigma}_k$. The similarity symbol is used since the probability of point \mathbf{x}_m inside $\mathcal{N}(\mathbf{q}_k, \mathbf{\Sigma}_k)$ would require a scaling with $\frac{1}{2\pi\sqrt{|\mathbf{\Sigma}|}}$. For its utilization in the NDT this scaling

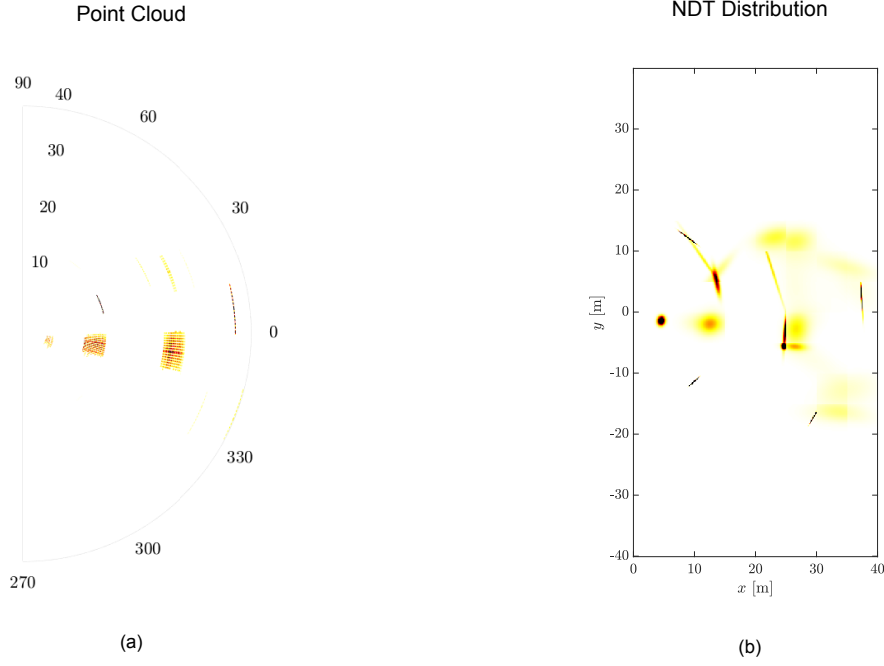


Figure 3.1: Representation of the point cloud as a piece-wise continuous distribution, (a) the original detections, (b) the Normal Distributions Transform

factor can be dropped. Figure 3.1 shows an example of a radar range scan with its NDT distribution. In the scan the colored dots represent the individual detections, their color represents the received signal power. From the figure, it is clear that the spreading in the azimuth direction results in an odd NDT distribution in the surrounding cells.

3.2. The Point Cloud Mapping Equations

In order to calculate the probability of overlap for a certain relative pose \mathbf{p} the, now Cartesian, point cloud of the current scan has to be transformed to represent the measurements were they taken from the position related to \mathbf{p} . This transformation is performed according to the equations in (3.2), where x'_m and y'_m are the mapped x and y -coordinates of point m originally located at $[x_m, y_m]^T$ and t_x, t_y and ϕ are the parameters in \mathbf{p} describing the translation in the x - and y - direction as well as the change in heading, both relative to the pose at the reference scan.

$$\begin{aligned} x'_m &= x_m \cos(\phi) - y_m \sin(\phi) + t_x \\ y'_m &= x_m \sin(\phi) + y_m \cos(\phi) + t_y \end{aligned} \quad (3.2)$$

For the actual pose \mathbf{p} , \mathbf{x}'_m contains the x and y coordinates in the reference scan corresponding to the point m with coordinates \mathbf{x}_m in the current scan. For each point, this new location $\mathbf{x}'_m = [x'_m, y'_m]^T$ is then used to calculate a score related to the probability of (3.1). The cell C_k in the grid of the reference scan is found in which \mathbf{x}'_m resides and the calculated mean and covariance of C_k are used to determine the score of point m :

$$\text{score}(\mathbf{x}'_m) = \exp \left[-\frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right] \quad (3.3)$$

The relation between $\text{score}(\mathbf{x}'_m)$ and \mathbf{p} is thus through (3.2), which are needed to calculate the mapped point \mathbf{x}'_m . This relation between the pose and the score will be used in the scan-matching algorithm to estimate \mathbf{p} by maximizing the sum of the scores for all points m through optimization techniques.

3.3. The Optimization

The scan-matching is performed iteratively. At each iteration the point cloud of the current scan is mapped according to an estimate \mathbf{p} . This estimate can be an initialization vector or the estimate of a previous iteration. At each iteration a descent step $\Delta\mathbf{p}$ is calculated using Newton's method. The descent step is added to the previous estimate of \mathbf{p} and the next iteration is initialized by mapping the current scan point cloud according to this new estimate. This process repeats until convergence. The objective is to maximize the sum of the scores in (3.3) of all points m in the current scan. Using a negative sign, the minimization problem of (3.4) can be constructed:

$$\min_{\mathbf{p}} - \sum_{m=1}^M \exp \left[- \frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right]. \quad (3.4)$$

In order to solve this minimization problem using Newton's method, calculation of the gradient and the Hessian of the objective function is required. By choosing the score to be based on the Gaussian distribution it is assured that the first and second-order derivative, needed for calculation of the gradient and Hessian, are well-defined. Detailed derivations of the gradient and Hessian, as well as the first and second-order derivatives of the mapping equations, can be found in Chapter 5.

3.4. Discussion

The NDT provides a powerful tool for scan matching of high-resolution images collected with laser scanners. In radar, the range and angular resolutions are much worse than those of the aforementioned sensors. This leads to spreading of target responses in these dimensions and significant uncertainty of range and angular measurements of the targets. The applicability of the NDT to radar measurements is limited by the targets' representation in the Cartesian grid, which does not take these uncertainties into account. Conventional automotive radars provide the detection point cloud along with additional information, such as Doppler velocity and information about the radar cross-section (RCS) of a target through its reflected power or SNR, which are related to the RCS via the radar equation. These aspects are not exploited by the standard NDT. Incorporation of this knowledge could be used to improve the performance, the following chapters provides details about the steps necessary to derive techniques that are more suited for scan-matching using radar data.

4

The Radar Normal Distributions Transforms

As mentioned previously and became evident from Figure 3.1b the standard NDT is not well suited for the representation of radar scans due to spreading of the targets in the azimuth direction. Additionally, in order to include knowledge about the relative Doppler velocity of targets into the optimization technique, a reference to their bearings is necessary. A target's bearing is available in the original range scans, as they are received in polar coordinates, however, a direct reference to this parameter is lost once the scans are converted to the Cartesian coordinate system. Representation of the scan as a distribution in the range-angle domain, which is the native domain for range scan measurements, allows for incorporation of Doppler as well as a better representation of the target spreading.

4.1. The Polar Normal Distributions Transform

Converting the NDT to polar coordinates requires changes to the algorithm: The grid for calculation of the piecewise continuous distribution will have to accommodate the polar coordinate system. Additionally, the mapping equations used to transform the point cloud of the current scan are changed in order to relate the movement of the vehicle to changes in the range and bearing of the targets, which in turn requires the derivation of their first and second-order derivatives for calculation of the gradient and Hessian for the optimization steps. The resulting technique is named the Polar Normal Distributions Transform (PNDT).

4.1.1. The Piece-wise Continuous Distribution

The steps in which the distribution is created are the same as for the NDT, the difference is the dimensions of the cells that constitute the grids. Again, four overlapping grids are constructed, each with cells \tilde{C}_k with fixed lengths in the range and angle direction. The overlap in the grids is again created by shifting the grid in each and both of the cell directions. The fixed cell width in degrees causes the surface of the cells to become bigger at larger ranges. By constructing the grid in this way, the resulting distribution shows a more authentic representation of the received radar measurements; in the resulting distribution the power reflected from the targets is represented as a Gaussian distribution around the targets, which is a reasonable approximation of the main beam of the target response in radar measurements [23]. Additionally, the widening of the cells with increasing range results in an appropriate representation of the increased spreading in the azimuth direction for larger ranges. Figure 4.1 shows the distribution representation for the NDT and the PNDT as a side-by-side comparison, for the same scan as in Figure 3.1. From the figure it is very clear that the PNDT depicts a more accurate representation of the original scene. For each cell \tilde{C}_k on the polar grid the mean vector $\tilde{\mathbf{q}}_k \in \mathbb{R}^{2 \times 1}$ and the covariance matrix $\tilde{\Sigma}_k \in \mathbb{R}^{2 \times 2}$ are calculated, now containing, respectively, the two-dimensional mean and covariance of the position of the points in the polar domain.

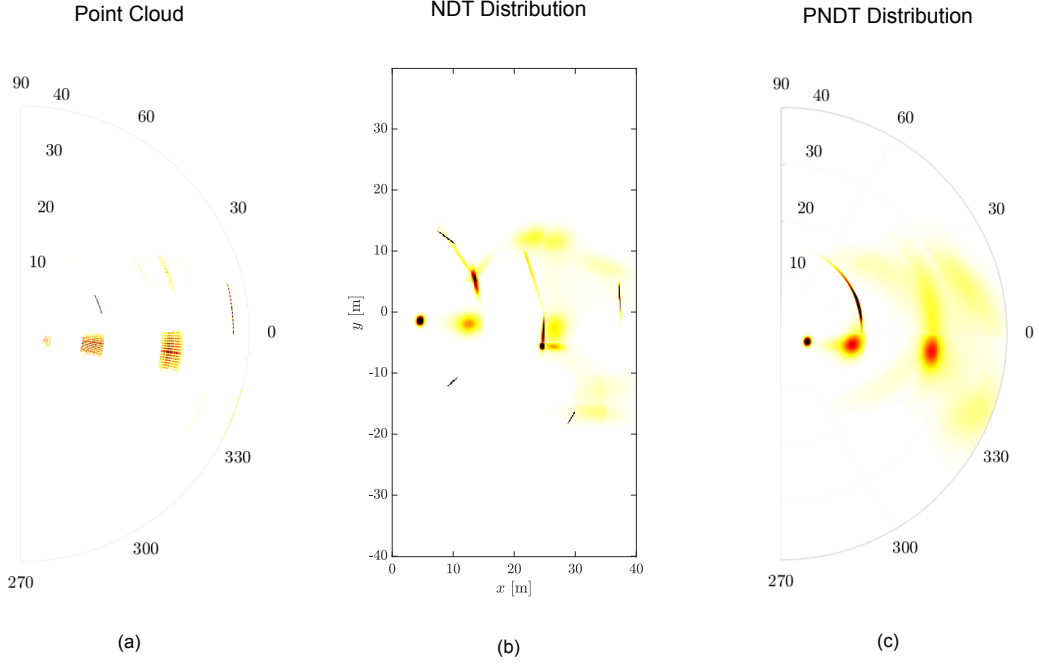


Figure 4.1: Different representations of the radar scan, (a) the original detections, (b) the Normal Distributions Transform, (c) the Polar Normal Distributions Transform

4.1.2. The Point Cloud Mapping Equations

The individual points in the point clouds are now represented by a range and a bearing measurement. These points have to be transformed based on the available estimate of \mathbf{p} . This results in slightly more involved mapping equations:

$$\begin{aligned} r'_m &= \sqrt{[t_x + r_m \cos(\phi + \theta_m)]^2 + [t_y + r_m \sin(\phi + \theta_m)]^2}, \\ \theta'_m &= \text{atan2}[t_y + r_m \sin(\phi + \theta_m), t_x + r_m \cos(\phi + \theta_m)]. \end{aligned} \quad (4.1)$$

Here, r_m and θ_m denote the original range and bearing measurements of point m , r'_m and θ'_m denote their mapped counterparts and t_x , t_y and ϕ again are the estimation parameters inside \mathbf{p} . Figure 4.2 shows how these mapping equations can be derived for a stationary target. These equations were also derived in [9] where they were used for a point-to-point scan matching technique in polar coordinates, as can be seen in (2.8).

4.1.3. The Optimization

The objective function of the PNDD is formulated in the same manner as in (3.4). Accounting for the change of parameters, the following objective function is found:

$$\min_{\mathbf{p}} - \sum_{m=1}^M \exp \left[- \frac{(\mathbf{x}'_m - \tilde{\mathbf{q}}_k)^T \tilde{\Sigma}_k^{-1} (\mathbf{x}'_m - \tilde{\mathbf{q}}_k)}{2} \right]. \quad (4.2)$$

Here, $\mathbf{x}'_m = [r'_m, \theta'_m]^T \in \tilde{\mathcal{C}}_k$ contains the mapped range and bearing measurements. The optimization is again done using Newton's method, for which the first and second-order derivatives of (4.1) are needed. These will be derived in Chapter 5.

4.2. The Doppler Polar Normal Distributions Transform

Due to the high resolution and accuracy of Doppler measurements and their relation to the angle of incidence of a target, incorporation of Doppler into the scan-matching algorithm can serve as a solution to the low resolution and accuracy of angular measurements in radar. The relation between the angle




Figure 4.2: The range-bearing measurements of a target m mapped using the relative pose $\mathbf{p} = [t_x, t_y, \phi]^T$. ϕ is calculated as $\phi + \theta_m$. The circles denote the vehicle at two separate positions and the bars indicate the front of the car, in this case chosen to coincide with the broadside of the radar.

of incidence and the Doppler measurements depends on the vehicle velocity, for a stationary target this relation is given as follows:

$$v_m = v_{car} \cos(\theta_m). \quad (4.3)$$

Here, v_m is the measured Doppler velocity of target m , θ_m is the angle of incidence associated with target m and v_{car} is the velocity of the vehicle in the direction of the broadside of the radar, i.e. in the direction of t_x in \mathbf{p} .

As becomes evident from (4.3) in order to incorporate the Doppler measurements and create the Doppler Polar Normal Distributions Transform (DPNDT), an estimate of the vehicle velocity is required. This necessitates the extension of the estimation vector to include the vehicle velocity: $\check{\mathbf{p}} = [t_x, t_y, \phi, v_{car}]^T$.

4.2.1. The Piece-wise Continuous Distribution

The goal is to now optimize the overlap between the range, angle and Doppler measurements of the current scan and the reference scan. For this, the two-dimensional distribution of the PNDT has to be extended to a three-dimensional distribution spanning not only the range and angle, but also the Doppler measurements. In the range and angle plane the four overlapping grids are kept and a single additional grid is constructed over the Doppler measurements. The result is four overlapping grids with three-dimensional cells. For each cell \check{C}_k on this grid with at least three points again the, now three-dimensional, mean vector $\check{\mathbf{q}}_k$ and covariance matrix $\check{\Sigma}_k$ are calculated for the reference scan.

4.2.2. The Point cloud Mapping Equations

At each iteration, the point cloud of the current scan has to be transformed according to the estimate of the relative pose. In addition to the range and angle measurements as in (4.1), the Doppler measurements have to be mapped. Upon inspection of (4.3), this is easily done using the mapped angle θ'_m according to (4.1):

$$v'_m = v_{car} \cos(\theta'_m). \quad (4.4)$$

4.2.3. The Optimization

Similar to the PNDDT, the objective function keeps the same form of (3.4). The vectors and matrices now simply become three-dimensional:

$$\min_{\mathbf{p}} - \sum_{m=1}^M \exp \left[- \frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \check{\Sigma}_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right], \quad (4.5)$$

with $\mathbf{x}'_m = [r'_m, \theta'_m, v'_m]^T \in \check{C}_k$ and $\mathbf{q}_k \in \mathbb{R}^{3 \times 1}$, $\check{\Sigma}_k \in \mathbb{R}^{3 \times 3}$. For the optimization of the score function, the first and second order partial derivatives of the new mapping equation (4.4) have to be calculated, along with the first and second order partial derivatives of (4.1) with respect to the vehicle velocity v_{car} . This is done in Chapter 5.

4.2.4. Target Filtering using Doppler Measurements

The proposed incorporation of Doppler explained in this section results in rewriting of the problem, however, the availability of Doppler measurements can also be used to improve the performance of scan-matching techniques in much simpler ways. Since the overlap between two scans is tested for estimation of the pose, non-stationary targets are undesirable. These targets will appear in different locations in the consecutive scans which leads to lower overlap. With the availability of Doppler non-stationary targets, as well as artifacts due to thermal noise, can be filtered out once an estimate of the velocity is available. Such an estimate \hat{v}_{car} can for example be calculated by rewriting (4.3) into the following form:

$$\hat{v}_{car,m} = \frac{v_m}{\cos(\theta_m)}, \quad (4.6)$$

and taking the average or median value over all points. Each detection is then tested to see if their measured Doppler velocity conforms to a constraint related to (4.3) such as:

$$|v_m - \hat{v}_{car} \cos(\theta_m)| \leq \epsilon, \quad (4.7)$$

where ϵ is a pre-defined bound. Additionally, all targets outside of the measurable azimuth region can be filtered out by removing any target m with a measured Doppler velocity lower than $\hat{v}_{car} \cos(\theta_{max})$, where θ_{max} is the maximum measurable angle.

4.3. The Range-Doppler Polar Normal Distributions Transform

In an effort to reduce the computational load, the possibility of using only the range and Doppler measurements for the optimization is investigated. Both the distribution and the score are calculated for only the range and Doppler measurements. This results in $\mathbf{x}_m = [r_m, v_m]^T \in \check{C}_k$ with \check{C}_k denoting a cell inside the range-Doppler grid and $\mathbf{q}_k \in \mathbb{R}^{2 \times 1}$, $\check{\Sigma}_k \in \mathbb{R}^{2 \times 2}$ being the mean and covariance in the range and Doppler direction. The pose vector $\check{\mathbf{p}}$ is the same as in Section 4.2. The resulting scan-matching technique is called the Range-Doppler PNDDT (RDPNDDT).

4.4. Incorporation of Signal-to-Noise Ratio Measurements

Using radar rather than LiDAR brings the availability of the received power of a target in the measurements through its Signal-to-Noise ratio (SNR). This received power is related to the radar cross section (RCS) of a target via the radar equation [23, p. 64-66], as a function of $\frac{1}{R^4}$. Knowledge about the received power can be used to reduce the influence of missed detections resulting in floating points between the scans, i.e. a point in one of the scans that has no corresponding counterpart in the other scan, since these missed detections are most common for weaker targets. The incorporation of this knowledge into the scan-matching algorithm can be done by addressing the received power directly or by constructing a reference to the RCS of the individual targets through a multiplication of the received power with R^4 . However, since the thermal noise level is constant over the range, this multiplication

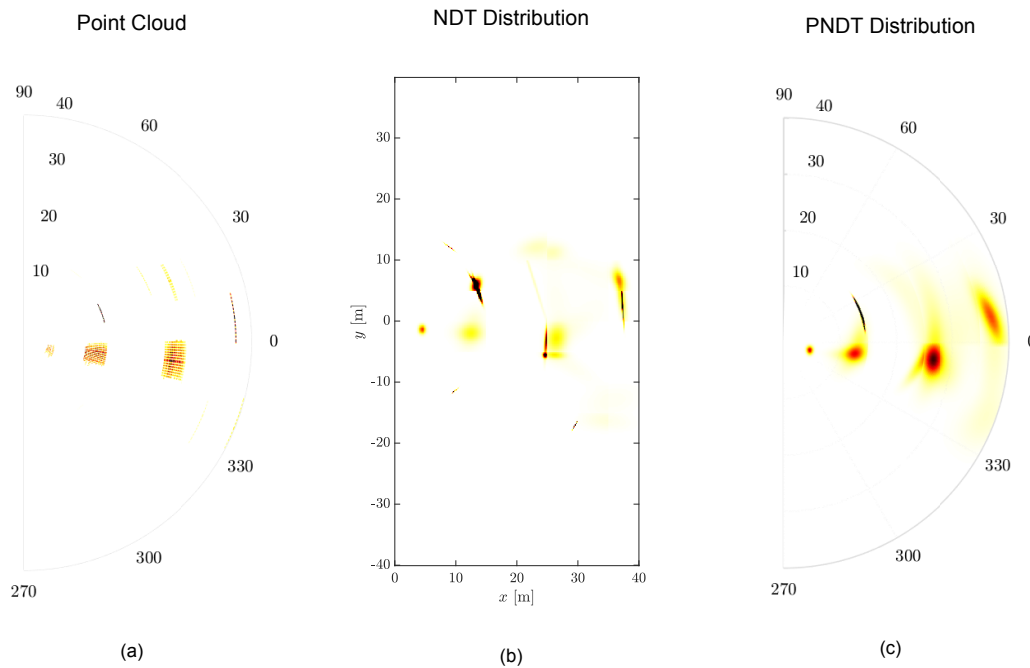


Figure 4.3: Different representations of the radar scan, (a) the original detections, (b) the SNR-aware Normal Distributions Transform, (c) the SNR-aware Polar Normal Distributions Transform

with R^4 does increase the influence of false alarms at further ranges. Both of these approaches can be implemented in two ways, by considering them in the calculation of the grid or by using it in the calculation of the score values of each point according to (3.3) in calculation of the objective functions (3.4), (4.2), (4.5).

4.4.1. Using Knowledge of the RCS in the Calculation of the Distribution

In order to account for the received power in the calculation of the distribution, the weighted mean and covariance are calculated. This results in the following relations [49]:

$$\mathbf{q}_k^i = \frac{1}{W_k} \sum_{m=1}^M w_m \mathbf{x}_m^i, \quad (4.8)$$

$$\boldsymbol{\Sigma}_k^{i,j} = \frac{1}{W_k} \sum_{m=1}^M w_m (\mathbf{x}_m^i - \mathbf{q}_k^i)(\mathbf{x}_m^j - \mathbf{q}_k^j).$$

In these definitions, i and j denote the entries of the vector and matrix, m denotes the measured target whose associated weight w_m is defined as the received power if it were addressed directly or the received power multiplied by R^4 to obtain a reference to the RCS directly, M is the total number of points inside C_k and $W_k = \sum_{m=1}^M w_m$. The distributions in Figures 4.1 have been constructed without taking the SNR into account. Figure 4.3 shows the resulting distributions for the case where the SNR is taken into account according to (4.8).

Again, the original scan of Figure 3.1a is used. Clearly, through incorporation of the SNR into the calculation of the distribution, the scan is represented as a more authentic representation compared to Figure 4.1. The SNR-aware distributions show clear regions of high density around the areas of large received power in the original scan and lower concentrations in the weaker regions.

4.4.2. Using Knowledge of the RCS in the Calculation of the Score

The same weights can be used in the calculation of the score, the adapted score function takes the following form:

$$\text{score}(\mathbf{x}'_m) = w_m \exp \left[-\frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right]. \quad (4.9)$$

With the weight factors w_m defined in the same way as in (4.8). This definition results in the following adjusted minimization problem:

$$\min_{\mathbf{p}} - \sum_{m=1}^M w_m \exp \left[-\frac{(\mathbf{x}'_m - \mathbf{q}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}'_m - \mathbf{q}_k)}{2} \right]. \quad (4.10)$$

This scaling with w_m is a linear operation, it is thus trivial to include it in the Hessian and gradient for the optimization steps.

4.5. Discussion

All of the adjustments to the conventional NDT as discussed in this chapter could improve the performance of the scan-matching. The PNNT aims to do so by creating a more authentic representation of the scan, with a distribution that more closely resembles the extended targets that are present in radar detections. Addressing the problem in polar coordinates additionally allows for the incorporation of the Doppler measurements, resulting in the DPNT. The Doppler measurements are related to the bearings of the targets and provide higher resolution than detection in the azimuth direction. The RDPNT is a mere simplification of the DPNT in an effort to reduce the computational load. Finally, incorporation of the received signal power can increase the robustness of the algorithm against missed detections. Missed detections are most prominent for targets that reflect less power (so-called “weak targets”). By incorporating a weighting factor related to this received power, the scan-matching relies more on stronger targets.

5

The Optimization Problem

In order to perform the optimization of (3.4)¹ the gradient and Hessian of the objective function with respect to \mathbf{p} have to be derived. Since the structure of the score function itself does not change between the techniques, the gradient and Hessian of (3.4) are derived in terms of the first and second order derivatives of the mapped points towards \mathbf{p} . These derived equations can also be found in [11].

5.1. The Objective function

Recall (3.4), the objective function aims to maximize the sum of all scores of each point m , in turn maximizing the overlap. The objective function is formulated as a minimization as is conventional in optimization problems [50, p. 1]. Applying Newton's method requires the objective function to be convex, which introduces the condition that the Hessian \mathbf{H} of the objective function is always positive semidefinite [50, p. 71]. This requirement implies that any local minimum is the global minimum. Determining convexity analytically for (3.4) is non-trivial, as it relies on the positions of individual points within the reference scan. However, an objective function can be visualized by calculating its value over an interval on the optimization parameters in \mathbf{p} in order to investigate convexity. Figures 5.1, 5.2 and 5.3 show the shape of the objective functions corresponding to the scan in 3.1a in the matching with its consecutive scan for some of these intervals, for a cell size of 1 meter. The true minima are located at $t_x = 0.2047$ m, $t_y = -2 \times 10^{-5}$ m, $\phi = -0.005^\circ$ and $v = 4.0942$ m/s.

From these figures it immediately becomes clear that the objective functions are not convex and suffer from many local minima, which leads to problems in the absence of an accurate initial estimate of the pose. Such an estimate can for example be made by assuming both the translation in the y -direction t_y and the change in heading ϕ are small, which is a plausible assumption given the fact that the velocity of the car is defined in the x -direction and the difference in time frames taken for scan-matching is small. Additionally, t_x can be initialized through an estimate of the vehicle velocity along with knowledge of the time between the scans. Such an estimate of the velocity can be made using (4.6). Noteworthy is the unusual structure in any of the figures showing the objective values for varying t_y and ϕ . The figures show the minimum as an oblique line, this is seen more clearly in Figure 5.4. This can be explained by the fact that for small angles a change in heading has approximately the same effect as a translation in the y -direction.

When it comes to the incorporation of Doppler measurements, some observations can be made. First of all, when looking at the values around the optimal t_x and t_y in Figures 5.2a and 5.2b, it can be seen that a somewhat more well-defined minimum in the direction of t_y is present, especially the minimum in Figure 5.1a has a seemingly flat shape in comparison. Furthermore, it seems like the objective function of the DPNDT presents a less smooth path towards the minimum of t_x and ϕ in Figure 5.2c, which could result in additional local minima within the main valley of the global minimum.

Figures 5.2d and 5.2f show that the minima corresponding to t_x and ϕ and t_y and ϕ for the RDPNDT are found at an offset to the correct minimum of ϕ . This problem can be explained by the fact that by

¹and its counterparts for the adjusted techniques

addressing the problem in the range-Doppler domain, any information to distinguish between left and right is lost since the only relation to the angle is through the cosine, which does not maintain the sign. Because of this, objects with approximately equal but opposite angles end up very close together in the range-Doppler representation. This causes the two individual targets to overlap, resulting in an area of high density in the calculated distribution at the center of this overlap. The probability that is maximized by the objective function now has a maximum value for these two particular targets in the case where they would perfectly overlap, which is the case for a mapping which would result in their angles having exactly equal values with opposite signs. In a realistic situation detected objects, such as lampposts, parked cars or buildings, are positioned at approximately equal distance from the center of the road. Since the measurements of Figure 3.1a represent measurements of a vehicle driving on the right side of the road this results in objects to the right of the road appearing at a closer angle than objects to the left. This causes the objective function to favor a positive heading, which is also seen in 5.2d.

A benefit of the introduction of the Doppler measurements is the well-defined and narrow minima in Figure 5.3. This, however, leads to the requirement of an accurate initialization of the vehicle velocity to prevent the optimization process from starting in a zero score region. Such a velocity estimate can again be calculated using (4.6).

An additional way to reduce the local minima is by increasing the cell size of the grids, this however leads to minima that are less well defined and appear more flat, sometimes even leading to bias in the minimum due to crude discretization in the grid, leading to errors in estimation of the true minimum. The counterparts of Figures 5.1a, 5.1b, 5.2a and 5.2b for a cell size of 10 meters are presented in 5.5, the shapes corresponding to the other combinations of parameters show similar smoothing.

Utilizing the availability of the SNR in the calculation of either the distribution, the score or both also has significant influence on the distinction of the global minimum. Figure 5.6 shows the shape of the original objective function of 5.1a along with the shapes resulting from the incorporation of the SNR. It can be seen that the local minima reduce in depth in the direction of t_x and the global minimum can be found through a smoother path.

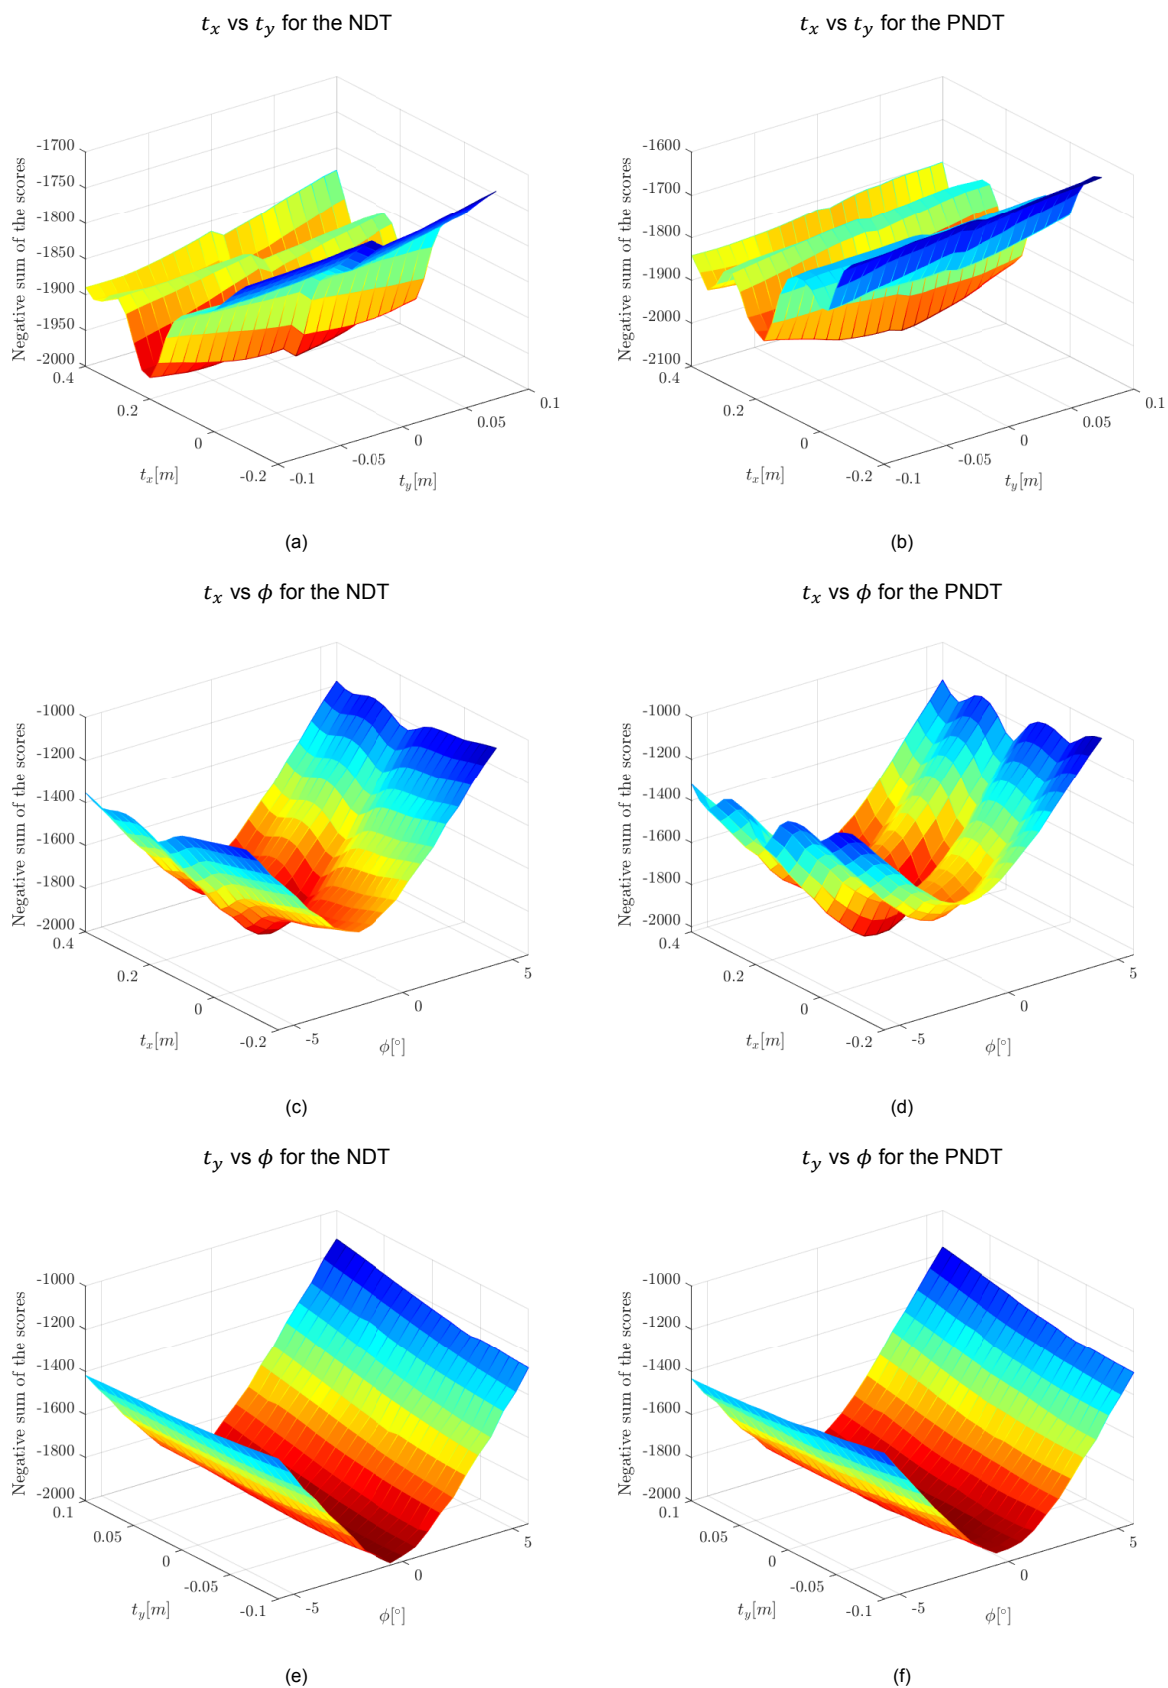


Figure 5.1: The minima of the objective functions visualized by calculating their values for combinations of the optimization parameters for the NDT and the PNDT for a cell size of 1 meter, (a) the NDT scanning through t_x and t_y , (b) the PNDT scanning through t_x and t_y , (c) the NDT scanning through t_x and ϕ , (d) the PNDT scanning through t_x and ϕ , (e) the NDT scanning through t_y and ϕ , (f) the PNDT scanning through t_y and ϕ

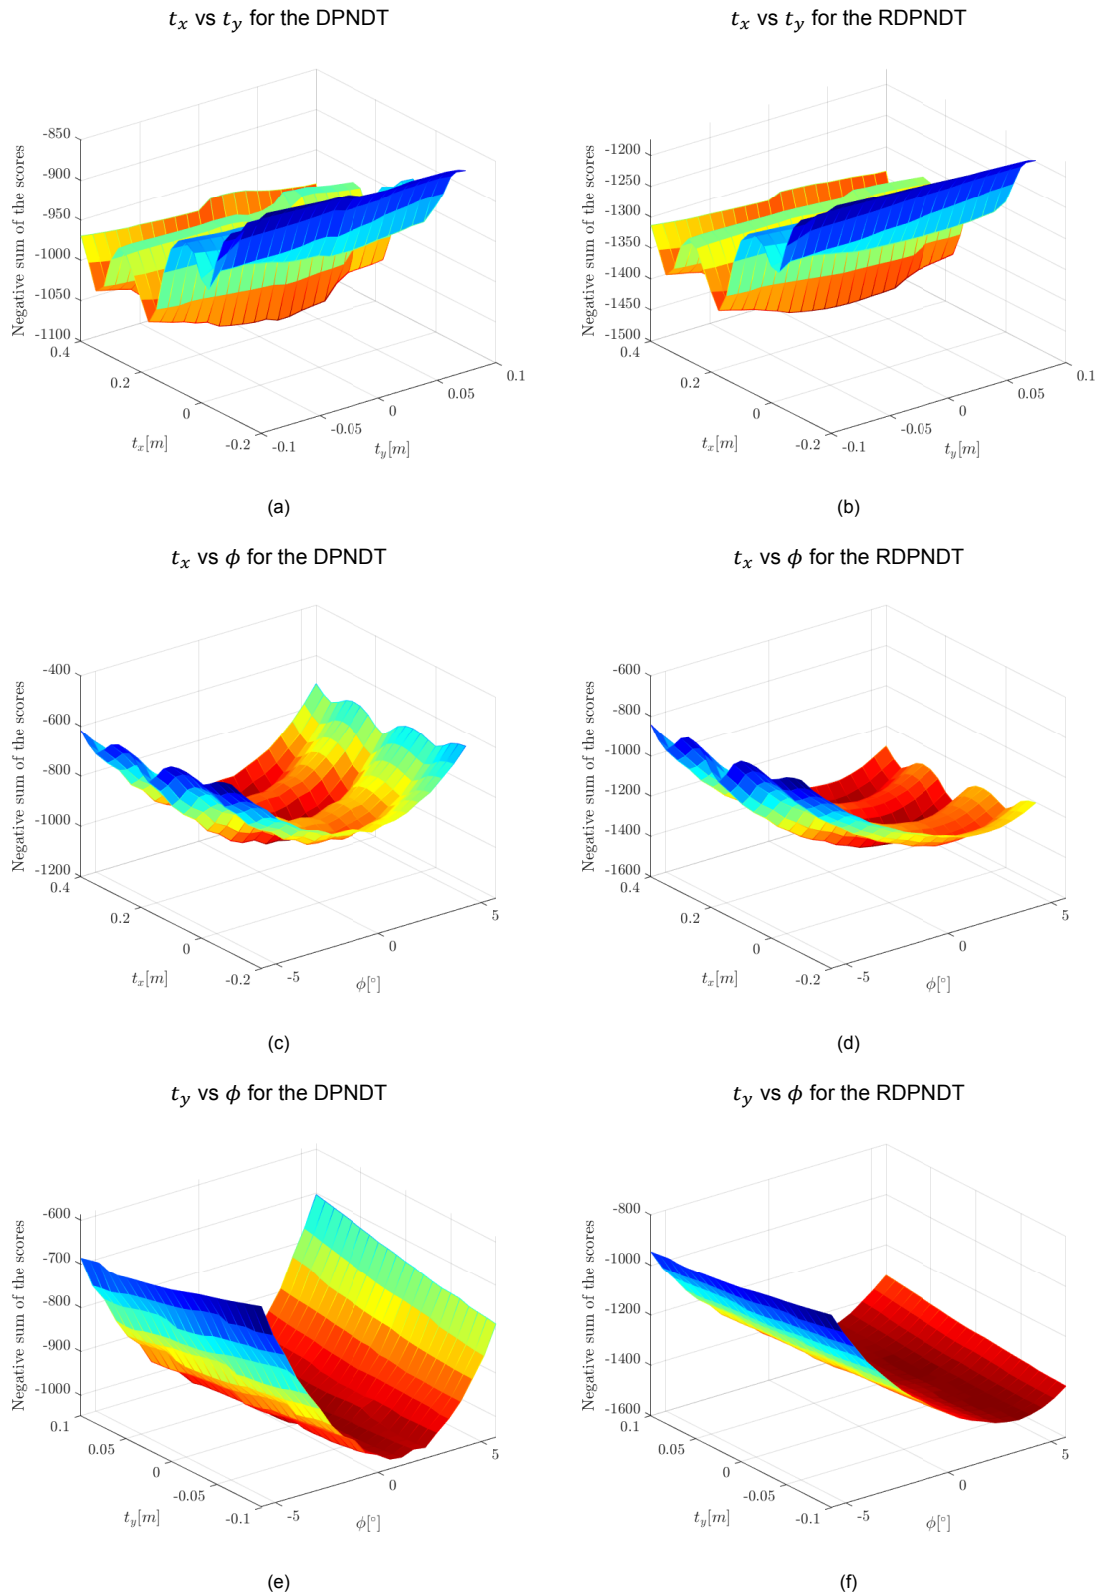


Figure 5.2: The minima of the objective functions visualized by calculating their values for combinations of the optimization parameters for the DPNDT and the RDPNDT for a cell size of 1 meter, (a) the DPNDT scanning through t_x and t_y , (b) the RDPNDT scanning through t_x and t_y , (c) the DPNDT scanning through t_x and ϕ , (d) the RDPNDT scanning through t_x and ϕ , (e) the DPNDT scanning through t_y and ϕ , (f) the RDPNDT scanning through t_y and ϕ

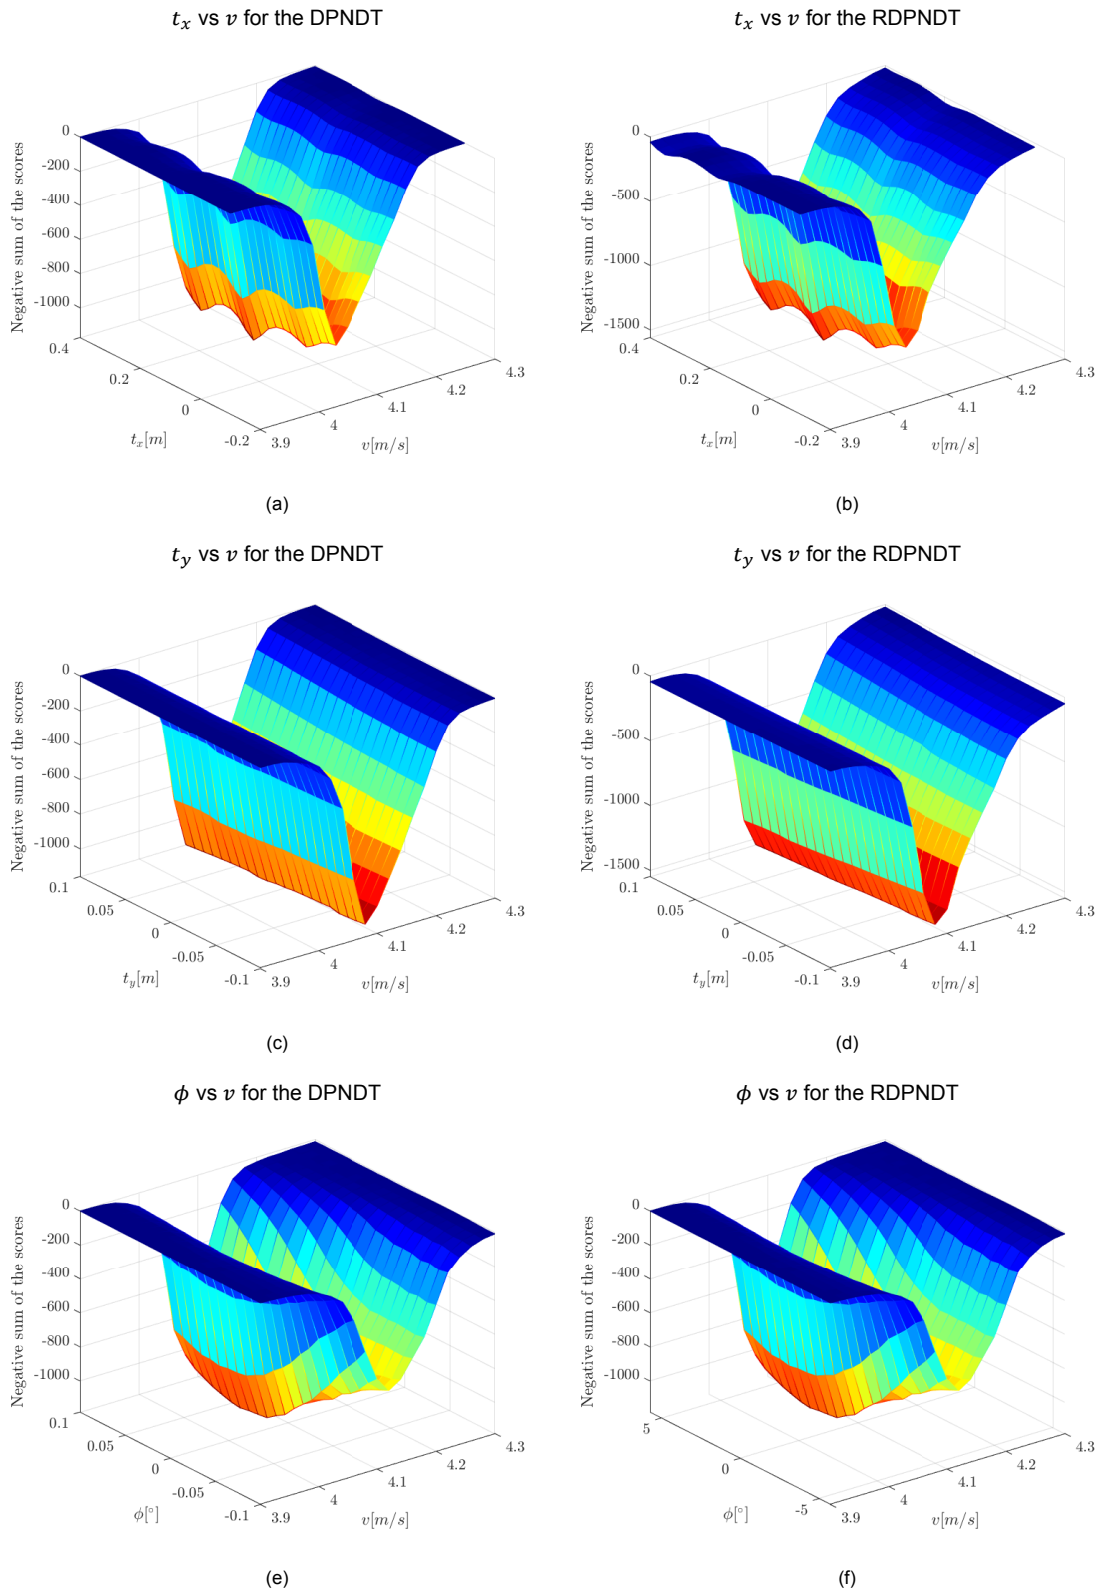


Figure 5.3: The minima of the objective functions related to the velocity estimate visualized by calculating their values for combinations of the optimization parameters for the DPNDT and the RDPNDT for a cell size of 1 meter, (a) the DPNDT scanning through t_x and v , (b) the RDPNDT scanning through t_x and v , (c) the DPNDT scanning through t_y and v , (d) the RDPNDT scanning through t_y and v , (e) the DPNDT scanning through ϕ and v , (f) the RDPNDT scanning through ϕ and v

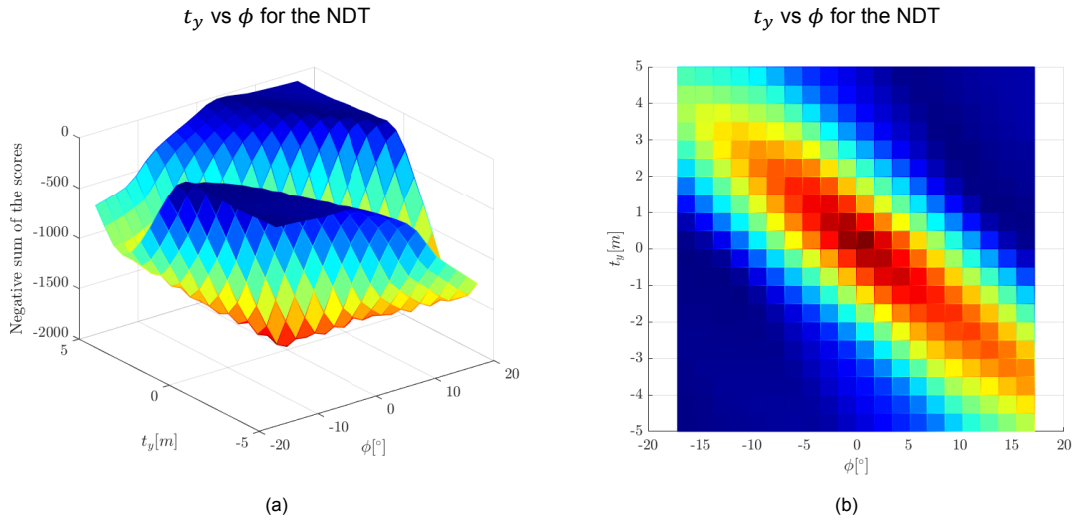


Figure 5.4: The oblique minimum inside Figure 5.1e zoomed out, (a) default view, (b) top-down view

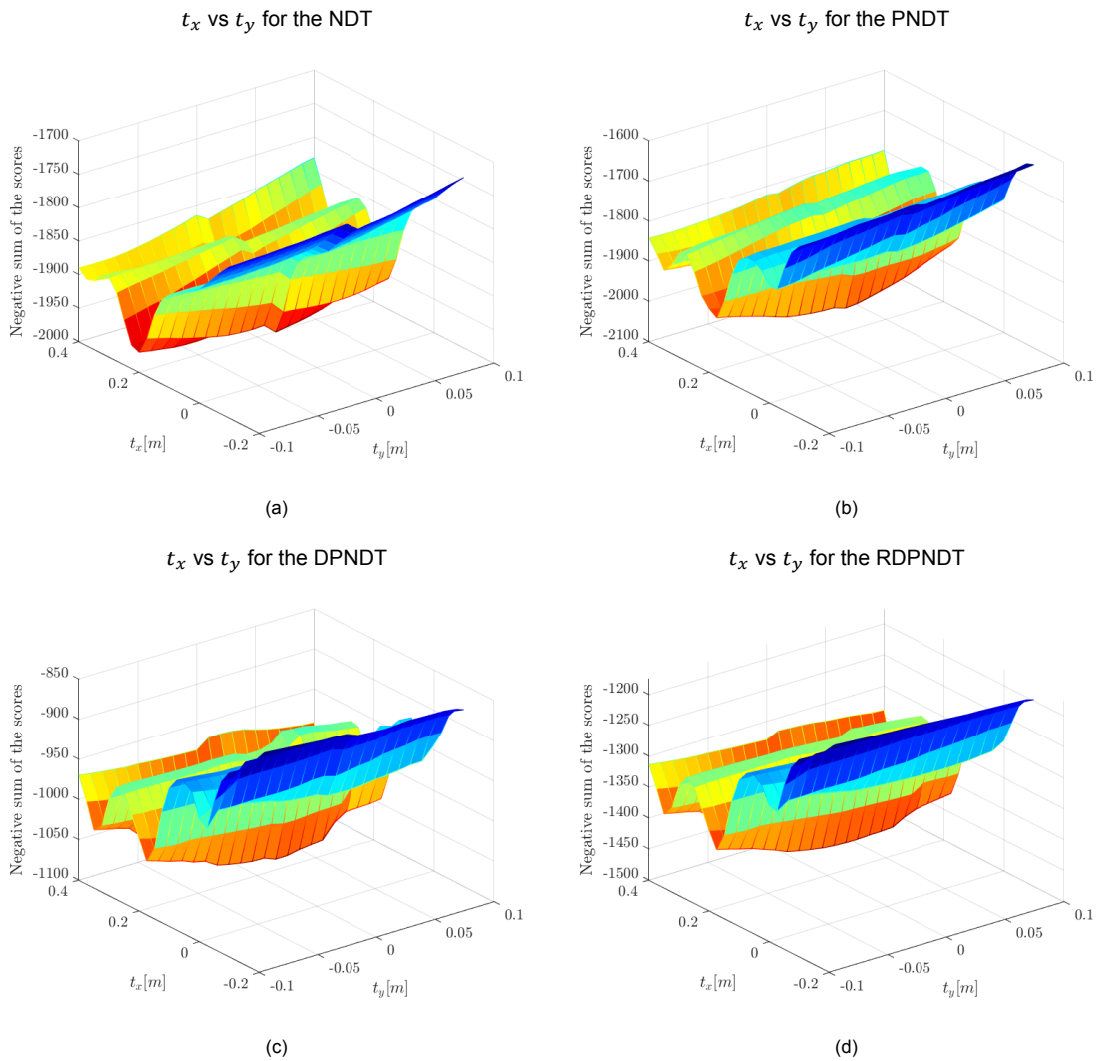


Figure 5.5: The visualization of the objective functions for a cell size of 10 meters, (a) the NDT scanning through t_x and t_y , (b) the PNDDT scanning through t_x and t_y , (c) the DPNDT scanning through t_x and t_y , (d) the RDPNDT scanning through t_x and t_y

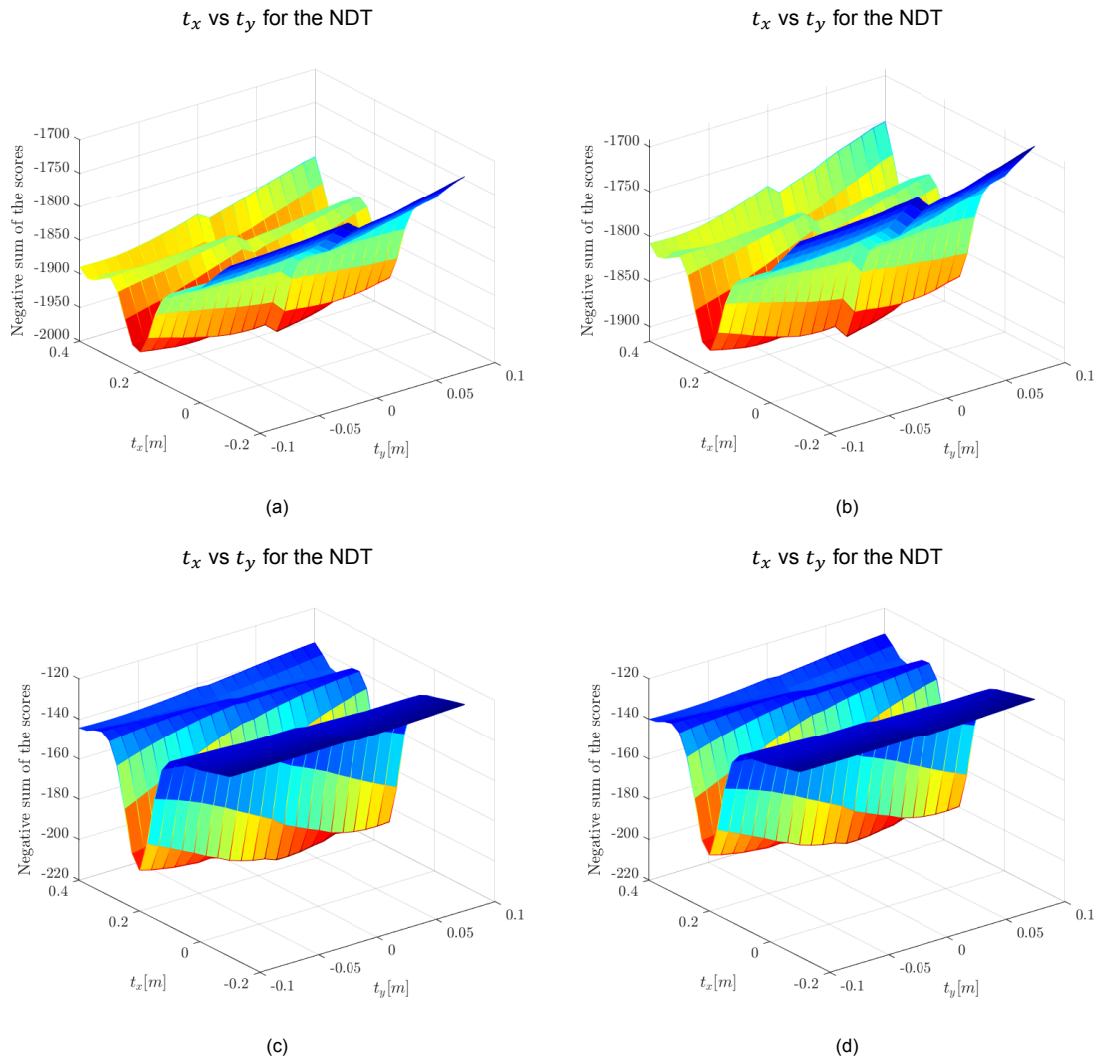


Figure 5.6: The minima of the objective function of the NDT with consideration of the RCS, visualized by calculating their values for combinations of the optimization parameters for a cell size of 1 meters, (a) using no knowledge about the RCS, (b) using knowledge about the RCS in the calculation of the distribution, (c) using knowledge about the RCS in the calculation of the score, (d) using knowledge about the RCS in the calculation of both the distribution and the score

5.2. The Hessian and Gradient of the Objective Function

The fact that a Gaussian-based distribution is used ensures twice differentiability, which is needed in order to calculate the Hessian. The gradient and Hessian of the objective function in terms of the first and second-order derivatives of the mapped points can be found in (5.1) and (5.2) [11]. Here, $\hat{\mathbf{q}}_m = \mathbf{x}'_m - \mathbf{q}_m$, \mathbf{g}_m^i is the i -th entry of the gradient vector, corresponding to entry i of \mathbf{p} and $\mathbf{H}_m^{i,j}$ is the entry on the i -th row and j -th column of \mathbf{H}_m , corresponding to entries i and j of \mathbf{p} . As is obvious from the subscripts, the gradient and Hessian are calculated for each point m after which they are averaged over the number of points to find \mathbf{g} and \mathbf{H} which are used to calculate the increment step. The size of \mathbf{g}_m and \mathbf{H}_m is determined by the size of \mathbf{p} . Thus, for the NDT and the PNNT we find $\mathbf{g}_m, \tilde{\mathbf{g}}_m \in \mathbb{R}^{3 \times 1}$, $\mathbf{H}_m, \tilde{\mathbf{H}}_m \in \mathbb{R}^{3 \times 3}$ and for the DPNT and the RDPNT we find $\check{\mathbf{g}}_m, \bar{\mathbf{g}}_m \in \mathbb{R}^{4 \times 1}$ and $\check{\mathbf{H}}_m, \bar{\mathbf{H}}_m \in \mathbb{R}^{4 \times 4}$.

$$\mathbf{g}_m^i = \hat{\mathbf{q}}_m^T \Sigma_m^{-1} \frac{\partial \hat{\mathbf{q}}_m}{\partial \mathbf{p}^i} \exp\left(-\frac{\hat{\mathbf{q}}_m^T \Sigma_m^{-1} \hat{\mathbf{q}}_m}{2}\right) \quad (5.1)$$

$$\mathbf{H}_m^{i,j} = -\exp\left(-\frac{\hat{\mathbf{q}}_m^T \Sigma_m^{-1} \hat{\mathbf{q}}_m}{2}\right) \left[\left(-\hat{\mathbf{q}}_m^T \Sigma_m^{-1} \frac{\partial \hat{\mathbf{q}}_m}{\partial \mathbf{p}^i}\right) \left(-\hat{\mathbf{q}}_m^T \Sigma_m^{-1} \frac{\partial \hat{\mathbf{q}}_m}{\partial \mathbf{p}^j}\right) + \left(-\hat{\mathbf{q}}_m^T \Sigma_m^{-1} \frac{\partial^2 \hat{\mathbf{q}}_m}{\partial \mathbf{p}^i \partial \mathbf{p}^j}\right) + \left(-\frac{\partial \hat{\mathbf{q}}_m^T}{\partial \mathbf{p}^i} \Sigma_m^{-1} \frac{\partial \hat{\mathbf{q}}_m}{\partial \mathbf{p}^j}\right) \right] \quad (5.2)$$

At each iteration, the current scan point cloud is first transformed according to the mapping equations, then a single Newton step is taken to calculate the increment $\Delta \mathbf{p}$ by solving the following equation:

$$\mathbf{H} \Delta \mathbf{p} = -\mathbf{g} \quad (5.3)$$

$\Delta \mathbf{p}$ is added to the previous estimate of \mathbf{p} to obtain a new pose estimate, to be used in the next iteration for mapping the point cloud. This process repeats until convergence.

5.3. The First- and Second-order Derivatives

5.3.1. NDT

The first and second-order derivatives of (3.2) have been calculated in [11], here they will be summarized in short:

The First Order Partial Derivatives

x-direction:

$$\begin{aligned} x'_m &= x_m \cos(\phi) - y_m \sin(\phi) + t_x \\ \frac{\partial x'_m}{\partial t_x} &= 1 \\ \frac{\partial x'_m}{\partial t_y} &= 0 \\ \frac{\partial x'_m}{\partial \phi} &= -x_m \sin(\phi) - y_m \cos(\phi) \end{aligned}$$

y-direction:

$$\begin{aligned} y'_m &= x_m \sin(\phi) + y_m \cos(\phi) + t_y \\ \frac{\partial y'_m}{\partial t_x} &= 0 \\ \frac{\partial y'_m}{\partial t_y} &= 1 \\ \frac{\partial y'_m}{\partial \phi} &= x_m \cos(\phi) - y_m \sin(\phi) \end{aligned}$$

The Second Order Partial Derivatives

Any second order derivative towards t_x or t_y is equal to zero, the only remaining expressions are:

$$\begin{aligned} \frac{\partial^2 x'_m}{\partial \phi^2} &= -x_m \cos(\phi) + y_m \sin(\phi) \\ \frac{\partial^2 y'_m}{\partial \phi^2} &= -x_m \sin(\phi) - y_m \cos(\phi) \end{aligned}$$

5.3.2. PNDT

As mentioned previously, the objective function in terms of the mapped points does not change between the techniques. The gradient and Hessian of the objective function in terms of the first and second-order derivatives of the mapping equations is thus unaltered. However, due to the change in coordinate frame, the mapping equations have changed to the ones in (4.1). The first and second order derivatives of these equations are derived in the following sections.

The First Order Partial Derivatives

Range:

$$r'_m = \sqrt{\frac{[t_x + r_m \cos(\phi + \theta_m)]^2 + [t_y + r_m \sin(\phi + \theta_m)]^2}{\alpha}}$$

$$\frac{\partial r'_m}{\partial t_x} = \frac{[t_x + r_m \cos(\phi + \theta_m)]}{\sqrt{\alpha}}$$

$$\frac{\partial r'_m}{\partial t_y} = \frac{[t_y + r_m \sin(\phi + \theta_m)]}{\sqrt{\alpha}}$$

$$\frac{\partial r'_m}{\partial \phi} = \frac{r_m [t_y \cos(\phi + \theta_m) - t_x \sin(\phi + \theta_m)]}{\sqrt{\alpha}}$$

Angle:

$$\theta'_m = \text{atan2}[t_y + r_m \sin(\phi + \theta_m), t_x + r_m \cos(\phi + \theta_m)]$$

$$= \arctan \left[\frac{t_y + r_m \sin(\phi + \theta_m)}{t_x + r_m \cos(\phi + \theta_m)} \right] + C,$$

with C a constant depending on the quadrant. Using \mathbf{p}^i to represent the i -th element of $\mathbf{p} = [t_x, t_y, \phi]$, we can write:

$$\frac{\partial \theta'_m}{\partial \mathbf{p}^i} = \frac{1}{1 + \beta^2} \frac{\partial \beta}{\partial \mathbf{p}^i}$$

$$\frac{\partial \beta}{\partial t_x} = \frac{-t_y - r_m \sin(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^2}$$

$$\frac{\partial \beta}{\partial t_y} = \frac{1}{t_x + r_m \cos(\phi + \theta_m)}$$

$$\frac{\partial \beta}{\partial \phi} = \frac{r_m^2 + t_x r_m \cos(\phi + \theta_m) + t_y r_m \sin(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^2}$$

The Second Order Partial derivatives

Range:

$$\frac{\partial^2 r'_m}{\partial t_x^2} = \frac{\sqrt{\alpha} - [t_x + r_m \cos(\phi + \theta_m)] \frac{\partial r'_m}{\partial t_x}}{\alpha}$$

$$\frac{\partial^2 r'_m}{\partial t_x \partial t_y} = -\frac{[t_x + r_m \cos(\phi + \theta_m)]}{\alpha} \frac{\partial r'_m}{\partial t_y}$$

$$\frac{\partial^2 r'_m}{\partial t_x \partial \phi} = -\frac{r_m \sqrt{\alpha} \sin(\phi + \theta_m) + [t_x + r_m \cos(\phi + \theta_m)] \frac{\partial r'_m}{\partial \phi}}{\alpha}$$

$$\frac{\partial^2 r'_m}{\partial t_y^2} = \frac{\sqrt{a} - [t_y + r_m \sin(\phi + \theta_m)] \frac{\partial r'_m}{\partial t_y}}{\alpha}$$

$$\frac{\partial^2 r'_m}{\partial t_y \partial \phi} = \frac{r_m \sqrt{a} \cos(\phi + \theta_m) - [t_y + r_m \sin(\phi + \theta_m)] \frac{\partial r'_m}{\partial \phi}}{\alpha}$$

$$\frac{\partial^2 r'_m}{\partial \phi^2} = \frac{-r_m \sqrt{a} [t_x \cos(\phi + \theta_m) + t_y \sin(\phi + \theta_m)]}{\alpha}$$

$$+ \frac{r_m [t_x \sin(\phi + \theta_m) - t_y \cos(\phi + \theta_m)] \frac{\partial r'_m}{\partial \phi}}{\alpha}$$

Angle: Using the variable β as defined above and using the same definition for the elements of \mathbf{p} we can write:

$$\frac{\partial^2 \theta'_m}{\partial \mathbf{p}^i \partial \mathbf{p}^j} = \left[\frac{1}{1 + \beta^2} \frac{\partial^2 \beta}{\partial \mathbf{p}^i \partial \mathbf{p}^j} - \frac{2\beta}{(1 + \beta^2)^2} \frac{\partial \beta}{\partial \mathbf{p}^i} \frac{\partial \beta}{\partial \mathbf{p}^j} \right].$$

The second-order partial derivatives of β are then given by:

$$\frac{\partial^2 \beta}{\partial t_x^2} = 2 \frac{t_y + r_m \sin(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^3}$$

$$\frac{\partial^2 \beta}{\partial t_x \partial t_y} = \frac{-1}{[t_x + r_m \cos(\phi + \theta_m)]^2}$$

$$\frac{\partial^2 \beta}{\partial t_x \partial \phi} = \frac{-r_m^2 [\sin^2(\phi + \theta_m) + 1] - 2t_y r_m \sin(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^3} - \frac{t_x r_m \cos(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^3}$$

$$\frac{\partial^2 \beta}{\partial t_y^2} = 0$$

$$\frac{\partial^2 \beta}{\partial t_y \partial \phi} = \frac{r_m \sin(\phi + \theta_m)}{[t_x + r_m \cos(\phi + \theta_m)]^2}$$

$$\frac{\partial^2 \beta}{\partial \phi^2} = \frac{[t_y r_m \cos(\phi + \theta_m) - t_x r_m \sin(\phi + \theta_m)] [t_x + r_m \cos(\phi + \theta_m)]}{[t_x + r_m \cos(\phi + \theta_m)]^3}$$

$$+ \frac{2r_m \sin(\phi + \theta_m) [r_m^2 + t_x r_m \cos(\phi + \theta_m) + t_y r_m \sin(\phi + \theta_m)]}{[t_x + r_m \cos(\phi + \theta_m)]^3}$$

These expressions are much more involved than the ones in [11]. However, during the verification process it was found that the increase in execution time per iteration of the PNDDT compared to the NDT is very low.

5.3.3. DPNDDT

For the DPNDDT, a new mapping equation has been introduced and the estimation vector has been extended to include the velocity of the car. The range and angle mapping equations do not depend on the vehicle velocity v_{car} , so the first and second order partial derivatives of Equations (4.1) towards v_{car} are equal to zero. Therefore, only the first and second order partial derivatives of Equation (4.4) are treated in this section. For simplicity, the partial derivatives derived in this section are expressed in terms of the first and second order partial derivatives of the mapped angle as derived in 5.3.2. Naturally, the derivatives calculated in this section will also be used in the RPNDDT.

The First Order Partial Derivatives

$$\begin{aligned}\frac{\partial v'_m}{\partial t_x} &= -v_{car} \sin(\theta'_m) \frac{\partial \theta'_m}{\partial t_x} \\ \frac{\partial v'_m}{\partial t_y} &= -v_{car} \sin(\theta'_m) \frac{\partial \theta'_m}{\partial t_y} \\ \frac{\partial v'_m}{\partial \phi} &= -v_{car} \sin(\theta'_m) \frac{\partial \theta'_m}{\partial \phi} \\ \frac{\partial v'_m}{\partial v_{car}} &= \cos(\theta'_m)\end{aligned}$$

The Second Order Partial derivatives

$$\begin{aligned}\frac{\partial^2 v'_m}{\partial t_x^2} &= -v_{car} \cos(\theta'_m) \left(\frac{\partial \theta'_m}{\partial t_x} \right)^2 - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial t_x^2} \\ \frac{\partial^2 v'_m}{\partial t_x \partial t_y} &= -v_{car} \cos(\theta'_m) \frac{\partial \theta'_m}{\partial t_x} \frac{\partial \theta'_m}{\partial t_y} - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial t_x \partial t_y} \\ \frac{\partial^2 v'_m}{\partial t_x \partial \phi} &= -v_{car} \cos(\theta'_m) \frac{\partial \theta'_m}{\partial t_x} \frac{\partial \theta'_m}{\partial \phi} - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial t_x \partial \phi} \\ \frac{\partial^2 v'_m}{\partial t_x \partial v_{car}} &= -\sin(\theta'_m) \frac{\partial \theta'_m}{\partial t_x}\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 v'_m}{\partial t_y^2} &= -v_{car} \cos(\theta'_m) \left(\frac{\partial \theta'_m}{\partial t_y} \right)^2 - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial t_y^2} \\ \frac{\partial^2 v'_m}{\partial t_y \partial \phi} &= -v_{car} \cos(\theta'_m) \frac{\partial \theta'_m}{\partial t_y} \frac{\partial \theta'_m}{\partial \phi} - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial t_y \partial \phi} \\ \frac{\partial^2 v'_m}{\partial t_x \partial v_{car}} &= -\sin(\theta'_m) \frac{\partial \theta'_m}{\partial t_y}\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 v'_m}{\partial \phi^2} &= -v_{car} \cos(\theta'_m) \left(\frac{\partial \theta'_m}{\partial \phi} \right)^2 - v_{car} \sin(\theta'_m) \frac{\partial^2 \theta'_m}{\partial \phi^2} \\ \frac{\partial^2 v'_m}{\partial \phi \partial v_{car}} &= -\sin(\theta'_m) \frac{\partial \theta'_m}{\partial \phi}\end{aligned}$$

$$\frac{\partial^2 v'_m}{\partial v_{car}^2} = 0$$

5.4. Local Minima and not Positive Definite Hessian

As became apparent in Section 5.1, the objective function often suffers from local minima, especially when a smaller cell size is used. This is something that has also been addressed in previous research [12], [13], [15]. Solutions include techniques such as iterative subdivision of the grid cells and representation via k -means clustering instead of a fixed grid. As can be seen from (5.3), Newton's method ensures the descent step to be in the direction of the negative gradient, thus moving towards the minimum, for convex functions (the Hessian being positive (semi-)definite). The Newton step even requires the Hessian to be positive definite for the inverse to be defined.

The local minima in the objective function cause so-called saddle points. Saddle-points are points within the objective function, corresponding to a combination of optimization variables \hat{p} that denote the minimum of the function in one dimension and the maximum of the function in another dimension [50, p. 238], an example of such a point as found within Figure 5.3a is shown in Figure 5.7. At such a point,

the calculated Hessian \mathbf{H} is not positive definite and thus taking the Newton step using the calculated Hessian will result in divergence. In order to solve this problem, a predefined value is chosen to act as the minimum eigenvalue λ_{min} . After every calculation of the Hessian \mathbf{H} its eigenvalues are checked using the eigenvalue decomposition, any eigenvalue that is below λ_{min} is set equal to it and the Hessian is reconstructed using its new eigenvalues.




Figure 5.7: A saddle point inside the objective function of the DPNDT for the matching of the scan in 3.1a with its consecutive scan

6

Joint Pose Estimation and Calibration

Some research was devoted to the possibility of using the DPNDT to perform sensor bias estimation. Accurate localization using automotive radar requires high-quality radar measurements. Many factors, e.g. temperature drift, calibration errors, presence of a water or dirt layer on the bumper or its deformation, can dramatically distort radar performance. Angular measurements, obtained via digital beam-forming, are the most sensitive to these types of distortions, leading to the widening and distortion of the main beam, as well as bias from its nominal direction [51]. This section describes the extension of the DPNDT to perform sensor bias estimation.

6.1. Angular Bias as an Estimation Parameter

In order to estimate the angular sensor bias, the properties of Doppler measurements can be utilized. Upon close inspection of Equation (4.3) it is seen that the measured Doppler velocity of a target is related to its actual, i.e. unbiased, angle of incidence θ_m . The angles in (4.1) (Fig. 4.2) are unbiased angles of the targets. By making the bias in the error explicit in the mapping equations and including it in the estimation vector \mathbf{p} , it can be estimated using the DPNDT through the standard steps as explained in 4.2.

The bias in angular measurements can be modeled by $\theta_{m,meas} = \theta_{m,act} + \epsilon_\theta$, independently on the scanning angle. The mapping equations of (4.4) are based on the actual angle, through the definition of the measured angles, the modified mapping equations then become:

$$\begin{aligned} r'_m &= \sqrt{[t_x + r_m \cos(\phi + \theta_{m,meas} - \epsilon_\theta)]^2 + [t_y + r_m \sin(\phi + \theta_{m,meas} - \epsilon_\theta)]^2}; \\ \theta'_{m,act} &= \text{atan2}[t_y + r_m \sin(\phi + \theta_{m,meas} - \epsilon_\theta), t_x + r_m \cos(\phi + \theta_{m,meas} - \epsilon_\theta)]; \\ \theta'_{m,meas} &= \theta'_{m,act} + \epsilon_\theta; \\ v'_m &= v_{car} \cos(\theta'_{m,act}). \end{aligned} \quad (6.1)$$

By now including the sensor bias ϵ_θ in the estimation vector \mathbf{p} and defining the gradient and Hessian for this new parameter, the sensor bias can be found through the standard optimization steps.

6.2. The Hessian and Gradient

The gradient and Hessian of the objective function in terms of the first and second order derivatives towards the entries of \mathbf{p} remain the same as in (5.1) and (5.2). However, \mathbf{p} has been extended to contain the sensor bias ϵ_θ , for which the first and second order derivatives still have to be derived. Upon closer inspection of the mapping equations in (6.1) it can be seen that many of the partial derivatives will be closely related to the partial derivatives of the heading ϕ , due to the similarity in occurrence of the two variables. Keeping this in mind, calculation of the derivatives becomes a trivial task, they can be expressed in terms of the partial derivatives towards ϕ as follows:

The First Order Partial Derivatives

$$\begin{aligned}\frac{\partial r'_m}{\partial \epsilon_\theta} &= -\frac{\partial r'_m}{\partial \phi} \\ \frac{\partial \theta'_m}{\partial \epsilon_\theta} &= -\frac{\partial \theta'_m}{\partial \phi} + 1 \\ \frac{\partial v'_m}{\partial \epsilon_\theta} &= -\frac{\partial v'_m}{\partial \phi}\end{aligned}$$

The Second Order Partial Derivatives

Range:

$$\begin{aligned}\frac{\partial^2 r'_m}{\partial t_x \partial \epsilon_\theta} &= -\frac{\partial^2 r'_m}{\partial t_x \partial \phi} \\ \frac{\partial^2 r'_m}{\partial t_y \partial \epsilon_\theta} &= -\frac{\partial^2 r'_m}{\partial t_y \partial \phi} \\ \frac{\partial^2 r'_m}{\partial \phi \partial \epsilon_\theta} &= -\frac{\partial^2 r'_m}{\partial \phi^2} \\ \frac{\partial^2 r'_m}{\partial \epsilon_\theta^2} &= \frac{\partial^2 r'_m}{\partial \phi^2}\end{aligned}$$

Angle:

$$\begin{aligned}\frac{\partial^2 \theta'_m}{\partial t_x \partial \epsilon_\theta} &= -\frac{\partial^2 \theta'_m}{\partial t_x \partial \phi} \\ \frac{\partial^2 \theta'_m}{\partial t_y \partial \epsilon_\theta} &= -\frac{\partial^2 \theta'_m}{\partial t_y \partial \phi} \\ \frac{\partial^2 \theta'_m}{\partial \phi \partial \epsilon_\theta} &= -\frac{\partial^2 \theta'_m}{\partial \phi^2} \\ \frac{\partial^2 \theta'_m}{\partial \epsilon_\theta^2} &= \frac{\partial^2 \theta'_m}{\partial \phi^2}\end{aligned}$$

Velocity:

$$\begin{aligned}\frac{\partial^2 v'_m}{\partial t_x \partial \epsilon_\theta} &= -\frac{\partial^2 v'_m}{\partial t_x \partial \phi} \\ \frac{\partial^2 v'_m}{\partial t_y \partial \epsilon_\theta} &= -\frac{\partial^2 v'_m}{\partial t_y \partial \phi} \\ \frac{\partial^2 v'_m}{\partial \phi \partial \epsilon_\theta} &= -\frac{\partial^2 v'_m}{\partial \phi^2} \\ \frac{\partial^2 v'_m}{\partial \epsilon_\theta^2} &= \frac{\partial^2 v'_m}{\partial \phi^2}\end{aligned}$$

6.3. Convergence to the Correct Minimum

During the verification process it was found that a bias could be estimated, however this bias suffered from a fixed offset. This fixed offset was the same for any arbitrarily chosen bias value. The reason for this can be seen in Figure 6.1, which shows similar plots to the ones in Chapter 5, now regarding the minima for varying values of the bias ϵ_θ . From this plot, it can be seen that the minimum of the objective function is actually situated at an offset to the true value, which is $\epsilon_\theta = -0.5^\circ$, denoted by the black cross.

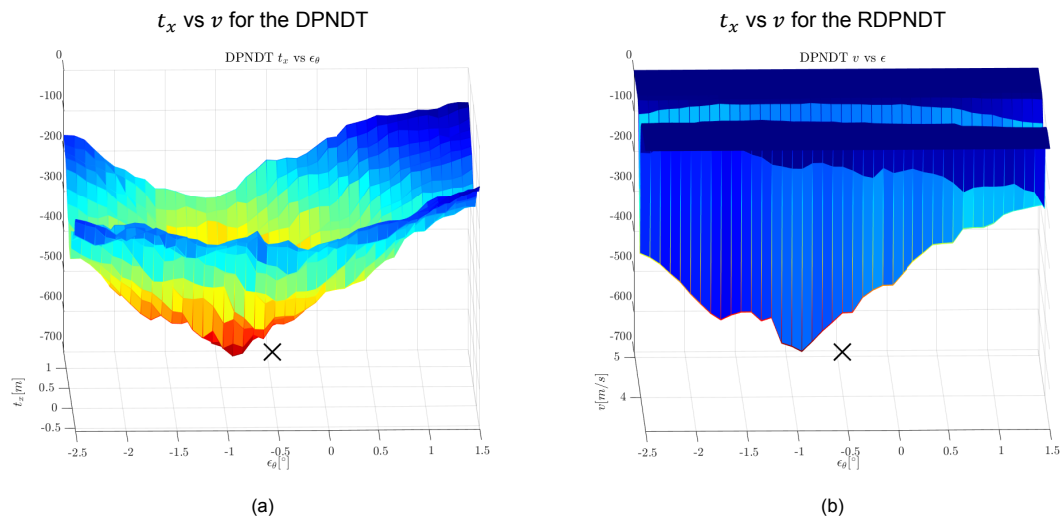


Figure 6.1: The minimum of the objective function of the DPNDT with sensor bias estimation, (a) scanning through t_x and ϵ_θ , (b) scanning through v and ϵ_θ

This problem can be explained by taking a look at the relation between velocity and angle as seen in (4.3). If we want to derive any information about the angle from the Doppler measurements, the

resolution of these Doppler measurements determines the resolution that our derived knowledge of the angle has. This can be seen more clearly through manipulation of (4.3), if we want to distinguish between any angles θ and $\theta + \Delta\theta$, the Doppler measurements should be able to distinguish between $v_1 = v_{car} \cos(\theta)$ and $v_2 = v_{car} \cos(\theta + \Delta\theta)$, in other words, the Doppler resolution should at least be:

$$\Delta v = v_{car} \cos(\theta) - v_{car} \cos(\theta + \Delta\theta). \quad (6.2)$$

For a vehicle velocity of 15 km/h and a Doppler resolution of $\Delta v \approx 7.81$ cm/s, as was the case during the verification process in Section 7.1, the resolution of the angle information that can be derived becomes $\Delta\theta \approx 1.074^\circ$, not enough to estimate a bias of -0.5° . Due to discretization this causes the minimum to be in the incorrect spot. If now the Doppler resolution were increased, either through lowering of the unambiguous velocity or by increasing the number of sweeps (either directly or indirectly, by concatenating consecutive frames in the slow time), knowledge about the angle can be derived at a finer resolution. This solution to the problem is confirmed by Figure 6.2, which shows the counterpart of Figure 6.1 for a Doppler resolution of $\Delta v \approx 1.95$ cm/s.




Figure 6.2: The minimum of the objective function of the DPNDT with sensor bias estimation for a higher Doppler resolution, (a) scanning through t_x and ϵ_θ , (b) scanning through v and ϵ_θ

Verification of the Performance of the NDT Techniques

In this section the influence of the suggested improvements are examined. All four approaches for incorporation of the SNR are compared; no weighting, weighting in the calculation of the distribution, weighting in the calculation of the score values, and both weighting in the distribution and the score. These comparisons are made for all of the four discussed techniques; the NDT, the PNNT, the DP-NDT and the RDPNT. In this way all of the proposed novelties are examined; transformation to polar coordinates, inclusion of Doppler and incorporation of the SNR.

The performance has been tested and compared using both simulated and real data. In order to do so, some fixed parameters have to be chosen on which the techniques depend, these parameters are the following:

Table 7.1: Initialization parameters of the NDT techniques

R_{cell}	The size of a cell in the range direction ¹
$[R_{min}, R_{max}]$	The minimum and maximum range for which the distribution is calculated
$[\theta_{min}, \theta_{max}]$	The minimum and maximum angle for which the distribution is calculated ²
$[v_{min}, v_{max}]$	The minimum and maximum velocity for which the distribution is calculated
$\Delta scan$	The position in discrete time of the current scan relative to the reference scan
$\ \Delta \mathbf{p}\ _{min}$	The minimum 2-norm of the Newton step, used to determine convergence ³
$ \Delta \mathbf{p}^i _{max}$	The maximum Newton step size for estimation parameter i , this is constrained to prevent divergence, similar to in [14]
λ_{min}	The minimum eigenvalue for the Hessian, as discussed in 5.4

¹Using this initialized value the sizes in the other dimensions are chosen such that all techniques have an equal number of cells. The only dimension for which this does not hold true is for the velocity dimension, which is divided into 100 cells in all cases

²Together the minimum and maximum range this determines the size of the Cartesian grid in the x - and y -direction

³If this constraint is not met within a certain number of iterations, the optimization process for the scans of interest is exited as well

During the verification process it was found that the performance of the different techniques depends very highly on the choice of most of these variables. Unfortunately there was not enough time to investigate their influences individually and values were chosen that were considered to result in decent outcomes most of the time.

7.1. Simulations

7.1.1. Set-up of the Simulations

For the simulation of measured radar data, a highly idealized environment is created. Shapes are created within the vicinity of the car using gray-scale values. These gray-scale values will be used as the RCS of individual targets to determine the received power. Within this created map a starting pose is

chosen whose measurements will be used as the reference scan. 100 poses relative to the starting pose are randomly generated. The magnitude of the individual relative pose parameters is in line with a vehicle driving along a straight trajectory at a randomly generated velocity with mean 15 km/h and standard deviation 0.1 km/h. At a framerate of 10 Hz, each timeframe will have a randomly generated velocity which, in combination with randomly generated steering irregularities (which represent the change in heading) determines the translation in the x - and y -direction. The randomly generated relative poses are stored to be used as a ground truth. Figure 7.1 shows the surroundings of the vehicle along with the generated relative poses, the circles representing the position and the bars indicating the heading.

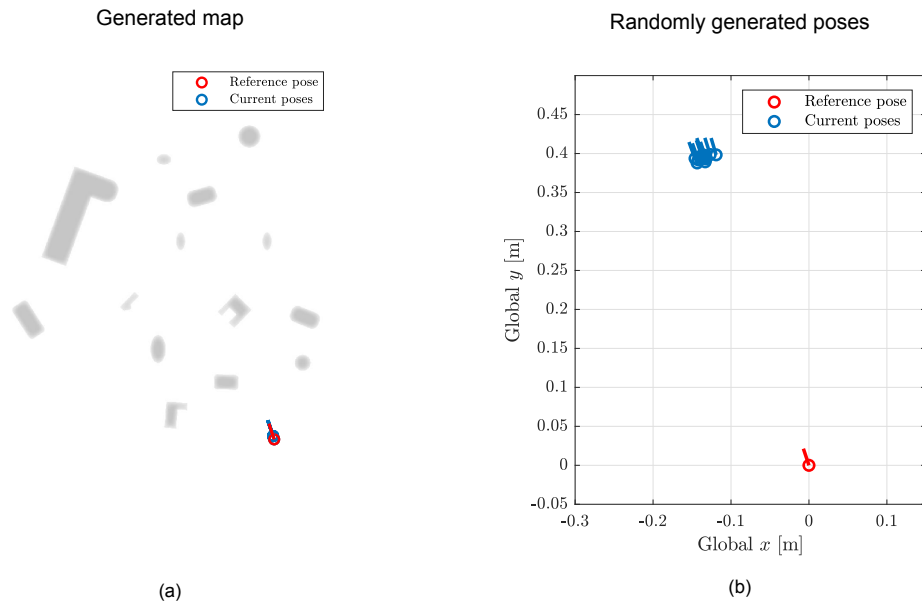


Figure 7.1: The simulated surroundings of the vehicles and the relative poses, (a) the map surrounding the vehicle, created using shapes that represent different objects, (b) the randomly generated poses along with the starting pose

For each timeframe, radar measurements are created of the surroundings. A radar scan is generated according to the following steps:

1. Extract the part of the image in the immediate surroundings of the vehicle, corresponding to a predefined maximum range
2. Add RCS fluctuation to the gray-scale values resulting in the distribution of the received power to be non-central chi-squared with two degrees of freedom [23]¹
3. Interpolate the Cartesian image to polar coordinates over a fine grid, restricting the field of view to $\pm 60^\circ$ ²
4. Extend the data into a radar data cube by calculating the Doppler velocity for each range-angle cell using (4.3)
5. Apply the inverse Fast Fourier Transform in the angular dimension to obtain a range-Doppler-antenna element cube
6. Discard any information outside of the first 12 (3 Tx \times 4 Rx) elements
7. Add target spreading in the Doppler and Range domain
8. Add white Gaussian noise to simulate thermal noise

¹Modeling the gray-scale values representing received signal magnitudes as a Rician distributed variable similar to [52]

²This grid is constructed such that the cell size in the range dimension coincides with typical radar range resolution and the cell size in the angle dimension is chosen such as to accommodate the chosen Doppler resolution

9. Perform median detection in the range-Doppler domain followed by digital beamforming

The result is a collection of scans, one for each time frame, each containing detected targets with corresponding range, angle, Doppler velocity and SNR values.

For generating the measurements, the thermal SNR is chosen to be 30 dB with a detection threshold of 15 dB. In this simplistic case, the RCS fluctuation is not taken into account. Furthermore, a maximum measurable range of $R_{max} = 40$ m is chosen and the range and Doppler resolution are chosen to be $\Delta R = 17.38$ cm and $\Delta v \approx 7.81$ cm/s, respectively. These values are in a similar order of magnitude as the ones achieved in the experiment detailed in Section 7.2.

In the simulations, the influence of the different methods to incorporate knowledge about the RCS is examined for each of the four NDT techniques. The RCS is taken into account directly by using the received power. The comparison criteria are the error in the position estimate, expressed as the Euclidean distance to the correct position, and the error in heading, where the randomly generated poses are used as the ground truth.

All estimation errors are expressed as the root mean squared values of the aforementioned errors over the 100 time frames, along with their standard deviation.

Additionally, some more realistic simulations were performed based on a map of the Leeghwaterstraat near the EEMCS building, where the experiment of Section 7.2 was also conducted. The map is created by tracing a section of the Leeghwaterstraat in a vector-based design and drawing program using grayscale values with varying intensities for different targets. Again the grayscale values will be used to simulate the RCS of the different targets, this time being modeled as $\text{Rice}(A_0, \sigma_A)$ distributed random variables where A_0 is the grayscale value and $\frac{\sigma_A}{A_0} = 0.1$.

Many point-like targets (roughly the size of a lamppost) are added which, along with the buildings on the map and some added objects such as cars and bicycles, will be used as reflective targets. Along the road that is present on the map a perfectly straight trajectory is created where the only uncertainty is generated through random fluctuation of the vehicle velocity, which is generated by taking samples from a Gaussian distribution characterized by a mean of 15 km/h and a standard deviation of 1 km/h. The trajectory consists of vehicle poses taken at intervals of 0.1 seconds, allowing for a framerate of 10 Hz, and the resulting trajectory is stored for verification purposes. The trajectory, along with the map can be seen in Figure 7.2. The black line denotes the trajectory at each time frame, the circles and the bars indicate the vehicle's position and heading at intervals of 10 time frames. In order to complete the trajectory, 100 frames were needed.

For each timeframe along the trajectory, i.e. at a framerate of 10 Hz, radar measurements are created of the surroundings. A radar scan is generated using the same method and parameters as before.

7.1.2. Comparison of the Different Pose Estimation Techniques

For examination of the results, the parameters of Table 7.1 have to be chosen. As mentioned earlier, these parameters, especially the values determining the cell size and λ_{min} , highly influence the end result. Unfortunately time did not permit to investigate the direct influences of all of these parameters in order to obtain the optimal values, only the cell size R_{cell} is investigated for some values. The chosen values can be found in Table 7.2.

Table 7.2: Chosen Initialization parameters of the NDT techniques

R_{cell}	0.5, 1 and 2 meter
$[R_{min}, R_{max}]$	[0, 40] meter
$[\theta_{min}, \theta_{max}]$	[-90, 90] degrees
$[v_{min}, v_{max}]$	[0, 5] meter per second
$\Delta scan$	+1 scan
$\ \Delta \mathbf{p}\ _{min}$	1×10^{-5} ¹
$ \Delta \mathbf{p}^t _{max}$	0.05 meter or radians
λ_{min}	1

¹With a maximum iteration count of 50




Figure 7.2: The created map and generated trajectory

Primitive simulations

The Pose Estimation: Figure 7.3 shows the results of the pose estimation techniques for the different RCS incorporation techniques and different cell sizes. Shown is the average Euclidean distance to the correct position and the root mean square error (RMSE) of the heading estimate, along with the standard deviation of these errors denoted by the errorbars. These values are found by averaging over the 100 realisations. In each case the PNDT outperforms all of the other techniques considerably. For each of the techniques, the dependence on incorporation of RCS information is not very prominent, which can be expected in an idealized case as was used here. The PNDT seems to either outperform the NDT or show similar performance. The RDPNDT shows very poor performance overall, it shows no improvements over the NDT apart from the case where the cell size is chosen to be $R_{cell} = 0.5$ m.

A clear relation between the cell size and the overall performance of all the techniques cannot be found, the PNDT shows considerable improvements in performance for the larger cell size of $R_{cell} = 2$ m, especially when it comes to the standard deviation in the heading estimate, however when comparing the performance of the heading estimate in the case of $R_{cell} = 1$ m to that in the case of $R_{cell} = 0.5$ m it seems like both the error and the standard deviation increase for the larger cell size of the two. With this same increase in cell size from $R_{cell} = 0.5$ m to $R_{cell} = 1$ m for the NDT both the error and the standard deviation reduce, but they increase again as the cell size is increased further to $R_{cell} = 2$ m. The PNDT shows slightly reduced errors and standard deviations for cell sizes $R_{cell} = 1$ and $R_{cell} = 2$ m, both showing a slight reduction in both error and standard deviation compared to the case where $R_{cell} = 0.5$ m. The achieved errors for each of the techniques for their optimal implementation out of the investigated ones can be found in Table 7.3.

The improvements can be expressed more quantitatively by looking at the average RMSE over all these realisations, i.e. all three values of R_{cell} and all four of the SNR incorporation techniques, for both the position and the heading and compare them to the average RMSEs of the NDT. The reduction in average RMSE expressed as a percentage can be calculated. The results of this are shown in Table 7.4. Here, a minus sign denotes an increase in the RMSE and thus a reduction in performance.

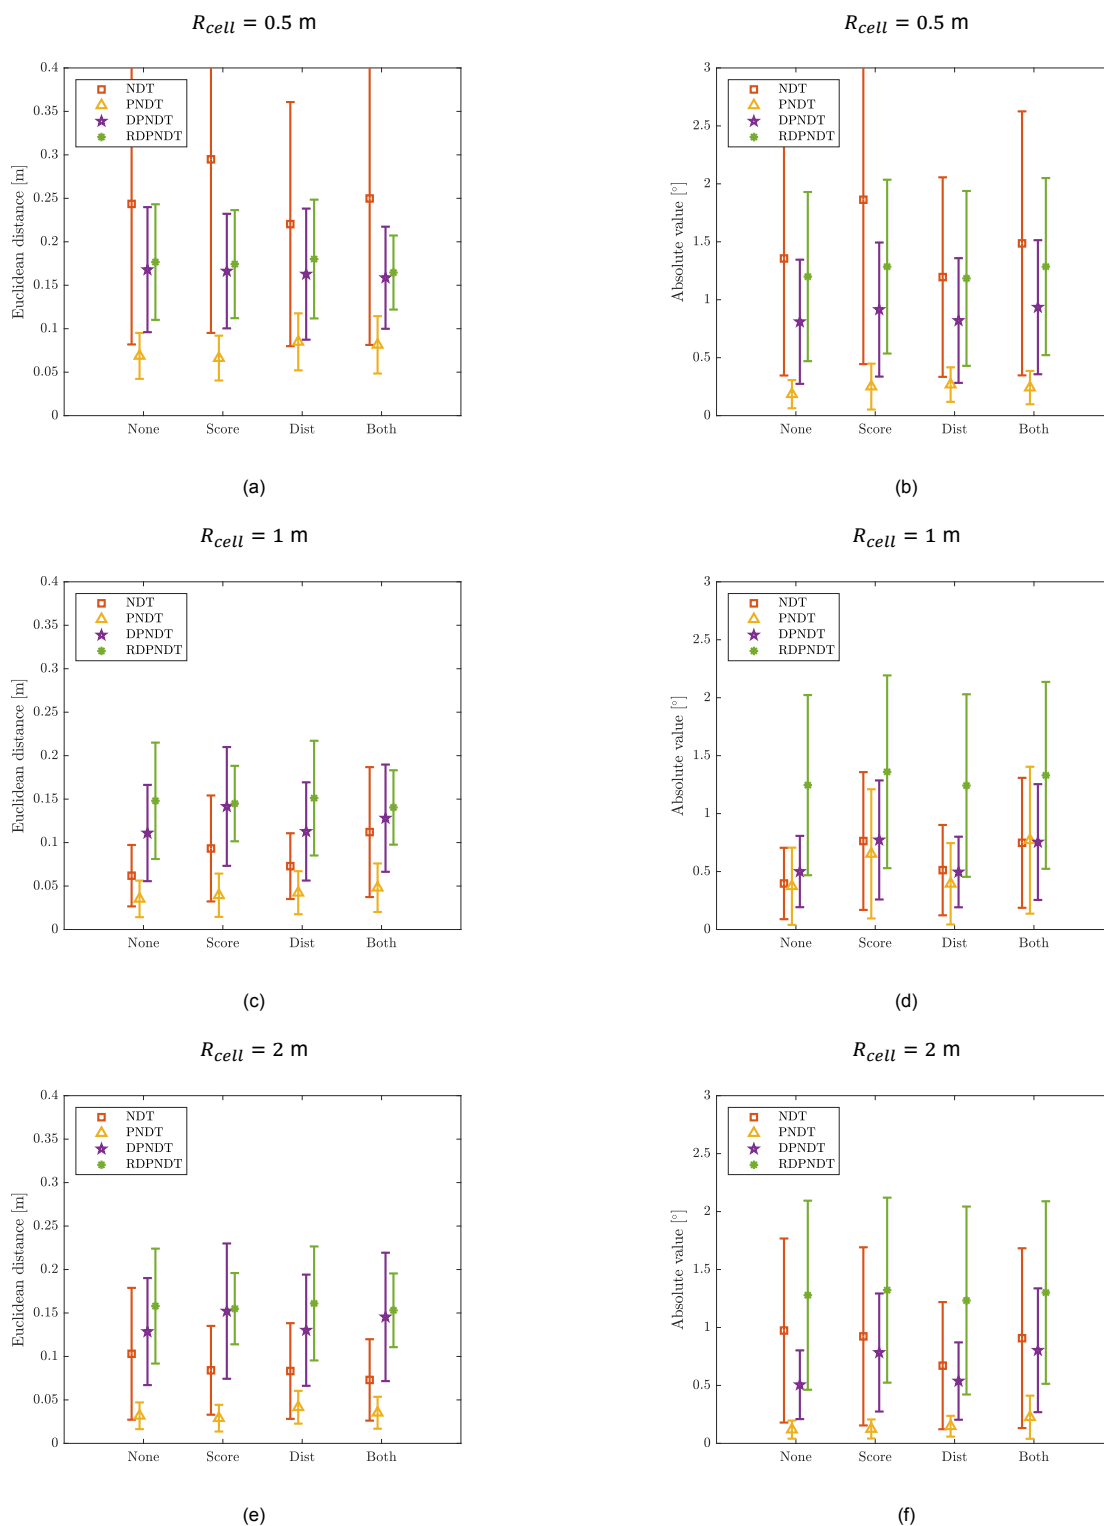


Figure 7.3: The errors, denoted by the symbols, along with the standard deviation, denoted by the error bar, of the pose estimates, (a) the Euclidean distance to the actual position using a cell size of $R_{cell} = 0.5$ m, (b) the RMSE of the heading using a cell size of $R_{cell} = 0.5$ m, (c) the Euclidean distance to the actual position using a cell size of $R_{cell} = 1$ m, (d) the error in the heading using a cell size of $R_{cell} = 1$ m, (e) the Euclidean distance to the actual position using a cell size of $R_{cell} = 2$ m, (f) the RMSE of the heading using a cell size of $R_{cell} = 2$ m

Table 7.3: The achieved errors for each technique, using the optimal settings

	Optimal set-up	Euclidean Distance [cm]	RMSE Heading [°]
NDT	$R_{cell} = 1$, RCS none	6.19	0.397
PNDT	$R_{cell} = 2$, RCS score	2.90	0.124
DPNDT	$R_{cell} = 1$, RCS none	11.1	0.500
RDPNDT	$R_{cell} = 1$, RCS none	14.81	1.245

Table 7.4: The reduction of the RMSE for each of the proposed techniques as compared to the conventional NDT, found by comparing the average RMSEs over the varying grid sizes and methods to implement RCS knowledge

	Position [%]	Heading [%]
PNDT	59.54	58.50
DPNDT	-27.03	19.09
RDPNDT	-45.60	-54.74

The Execution Times: The average execution times of the presented approaches are shown in Table 7.5, normalized to the execution time of the NDT³. It is clear that the increase in performance of the PNDT comes with only a very slight increase of the computational time. And while the computation time of the RDPNDT is lower than that of the NDT, the results are worse in most cases.

Table 7.5: Execution times of the NDT techniques, normalized to the execution time of the conventional NDT

	$R = 0.5$ m				$R = 1$ m				$R = 2$ m			
	None	Score	Dist	Both	None	Score	Dist	Both	None	Score	Dist	Both
PNDT	1.03	1.05	1.03	1.06	1.09	1.09	1.09	1.09	1.09	1.10	1.09	1.10
DPNDT	6.23	7.50	6.13	7.22	2.39	2.60	2.47	2.66	1.39	1.40	1.45	1.50
RDPNDT	0.87	0.90	0.86	0.88	0.98	1.00	0.99	0.99	0.97	0.97	0.97	0.98

The Sensor Bias Estimation: The sensor bias estimation technique was examined using an artificial bias added to the original angular measurements. The artificial bias was chosen to be $\epsilon_{\theta} = -0.5^\circ$ and the algorithms were ran using the same settings as before, with a cell size of $R_{cell} = 0.5$ m. Figure 7.4 shows the estimated sensor bias per frame. It shows that there are few but very severe outliers present among the estimates. The mean of the estimated bias over 100 frames can be seen in Table 7.6, along with their standard deviations, for the different types of RCS incorporation. The notation is read as “mean (standard deviation)”. Due to the small number of very severe outliers, most of them to one side, the median value is presented as well. It is clear that the estimation technique has detected a bias, it seems, however, that this bias is estimated with an offset of approximately -0.55° in each case when looking at the median. This offset was discovered during this verification process, but its cause has been explained in 6.3.

As mentioned before, the offset is due to discretization of the Doppler measurements and can be solved using a higher resolution. Figure 7.5 shows the estimated sensor bias per frame using a Doppler resolution of $\Delta v \approx 1.95$ cm/s for each of the SNR incorporation techniques. Calculating the mean, standard deviation and median over these 100 frames results in the values in Table 7.7, the notation is the same as in Table 7.6. Clearly, taking the median here is preferable and results in an

³The techniques were not yet optimized for computational efficiency, so at this moment they were not implemented in a way suitable for pose estimation in real-time




Figure 7.4: The estimated sensor bias per frame for the different RCS incorporation methods

Table 7.6: Average value, median and standard deviation of the estimated sensor bias for the different RCS incorporation techniques

	Estimated bias [°]	median [°]
No RCS incorporation	-0.9099 (0.4325)	-1.0615
RCS incorporation in the distribution	-0.8807 (0.6182)	-1.0635
RCS incorporation in the score	-0.9100 (0.4433)	-1.0540
RCS incorporation in both	-0.8610 (0.5951)	-1.0824

accurate estimate of the sensor bias, with an error of only 0.005° in the case of RCS incorporation in both the score and the distribution.

Table 7.7: Average value, median and standard deviation of the estimated sensor bias for the different RCS incorporation techniques

	Estimated bias [°]	median [°]
No RCS incorporation	-0.7995 (0.7022)	-0.5072
RCS incorporation in the distribution	-0.7001 (0.6733)	-0.4911
RCS incorporation in the score	-0.5881 (0.6631)	-0.5196
RCS incorporation in both	-0.6952 (0.6526)	-0.5050




Figure 7.5: The estimated sensor bias per frame for the different RCS incorporation methods for a higher Doppler resolution

More Realistic Leeghwaterstraat simulations

The Pose Estimation: Figure 7.6 shows again the Euclidean distance for the position and the RMSE for the heading estimate, along with their standard deviations, for each of the NDT techniques and each of the RCS incorporation techniques in the case of the straight trajectory. Comparisons are made using different cell sizes. It becomes immediately clear that the NDT without any adjustments is generally not suited for the application to radar data. Again, almost every time the PNDDT seems to outperform all of the other techniques and no clear influence of the RCS incorporation techniques is seen, which might be due to the high SNR along with a low fluctuation level. The DPNDT performs best when the cell size is chosen to be $R_{cell} = 2$ m. What stands out is that in the case of the RDPNDT, the incorporation of RCS knowledge actually reduces the performance quite considerably, especially regarding the position estimate. A clear relation between the cell size and the performance is again not found. Each technique has their own optimal settings, using these optimal settings, the obtained errors can be found in Table 7.8.

Table 7.8: The achieved errors for each technique, using the optimal settings

	Optimal set-up	Euclidean Distance [cm]	RMSE Heading [°]
NDT	$R_{cell} = 0.5$, RCS score	18.76	1.734
PNDDT	$R_{cell} = 0.5$, RCS score	12.62	0.9087
DPNDT	$R_{cell} = 0.5$, RCS dist	15.6	1.711
RDPNDT	$R_{cell} = 1$, RCS score	15.49	1.86

Again, the improvements are looked at from a more quantitative perspective by looking at the reduction in average RMSE compared to the NDT expressed as a percentage. The results of this are shown in Table 7.9, a minus sign denoting an increase in the RMSE and thus a reduction in performance. Apart from the position estimates using the RDPNDT, the overall reduction in error is very significant for each of the techniques.

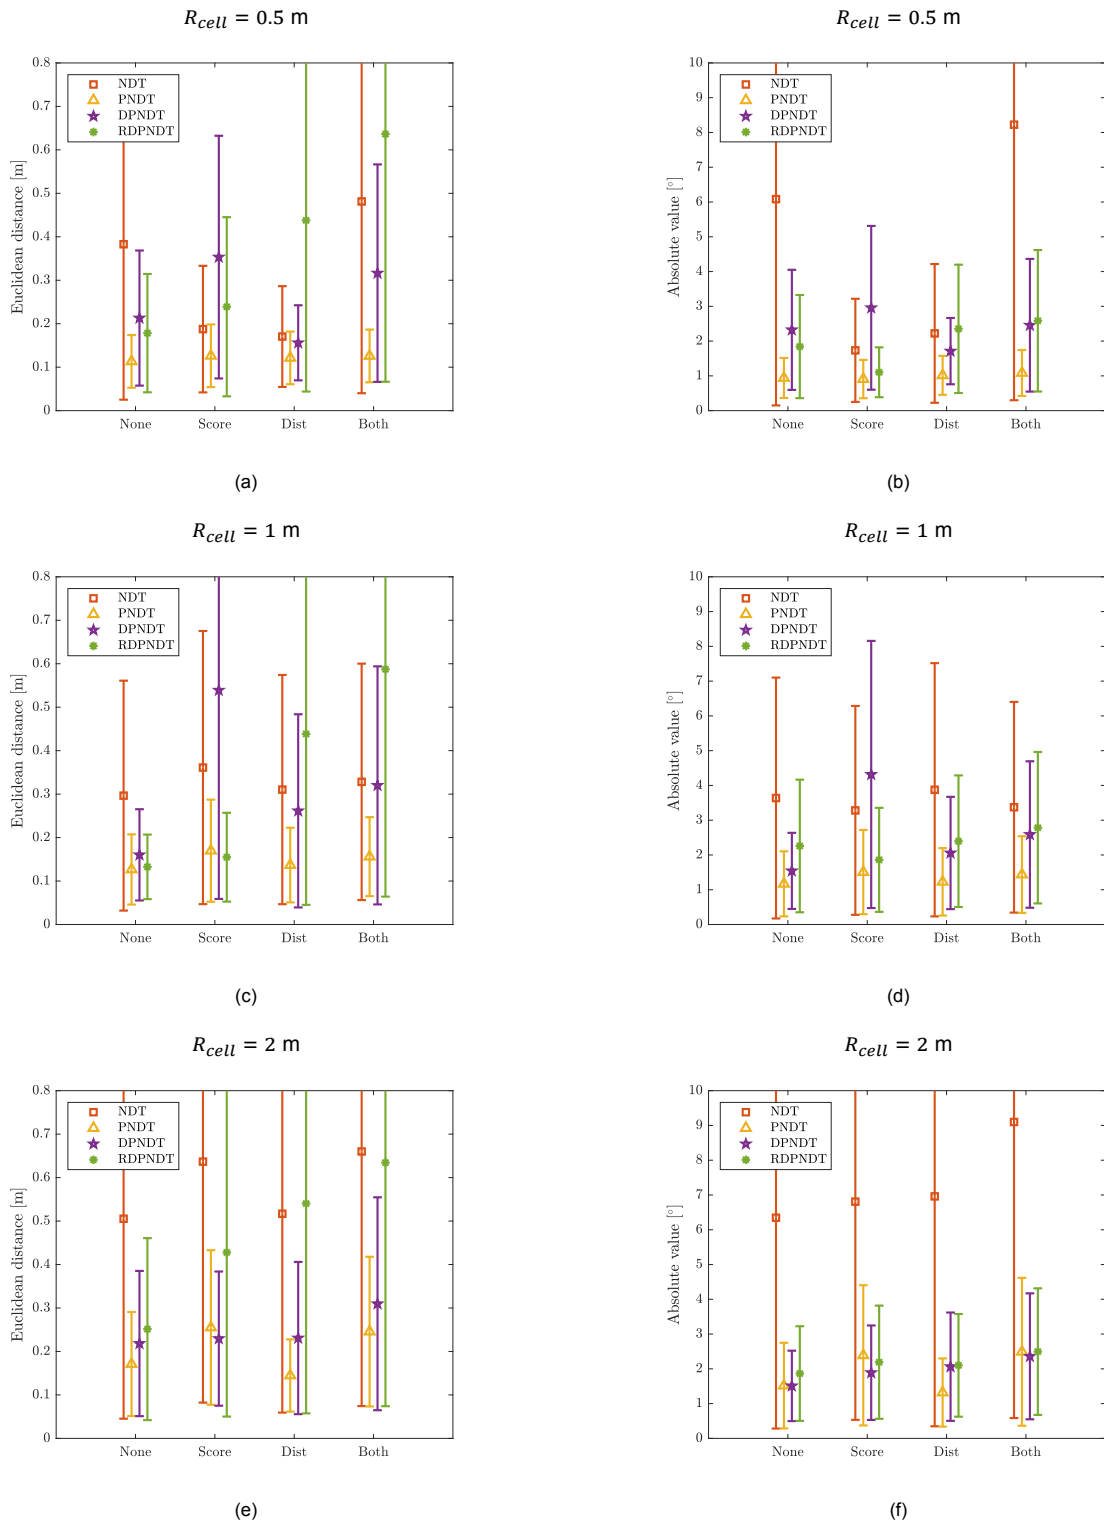


Figure 7.6: The errors, denoted by the symbols, along with the standard deviation, denoted by the error bar, of the pose estimates, (a) the Euclidean distance to the actual position using a cell size of $R_{cell} = 0.5$ m, (b) the RMSE of the heading using a cell size of $R_{cell} = 0.5$ m, (c) the Euclidean distance to the actual position using a cell size of $R_{cell} = 1$ m, (d) the error in the heading using a cell size of $R_{cell} = 1$ m, (e) the Euclidean distance to the actual position using a cell size of $R_{cell} = 2$ m, (f) the RMSE of the heading using a cell size of $R_{cell} = 2$ m

Table 7.9: The reduction of the RMSE for each of the proposed techniques as compared to the conventional NDT, found by comparing the average RMSEs over the varying grid sizes and methods to implement RCS knowledge

	Position [%]	Heading [%]
PNDT	57.11	67.97
DPNDT	20.25	39.47
RDPNDT	-7.36	48.88

The Execution Times: The average execution time of the presented approaches is shown in Table 7.10, again normalized to the execution time of the NDT. Again, the very significant increase in performance of the PNNDT comes at only a very slight cost of increased computational complexity.

Table 7.10: Execution times of the NDT techniques, normalized to the execution time of the conventional NDT

	$R = 0.5 \text{ m}$				$R = 1 \text{ m}$				$R = 2 \text{ m}$			
	None	Score	Dist	Both	None	Score	Dist	Both	None	Score	Dist	Both
PNNDT	1.02	1.03	1.03	1.03	1.08	1.07	1.07	1.06	1.10	1.09	1.08	1.06
DPNDT	2.23	2.13	2.29	2.18	1.83	1.78	1.81	1.80	1.76	1.69	1.74	1.66
RDPNDT	1.33	1.33	1.36	1.36	1.39	1.39	1.38	1.38	1.43	1.38	1.36	1.36

7.2. Experiment

For experimental verification, measurements were made using a commercially available automotive radar attached to the front of a car. During the experiment, raw radar data with synchronized GPS locations were recorded, all processing was performed offline. This section describes the experiment set-up along with the results.

7.2.1. Set-up of the Experiment

For collection of the data an NXP Dolphin automotive radar operating at a center frequency of 79 GHz was used. It was attached to the front of a car, slightly offset from the front of the vehicle. Figure 7.7 shows the radar system attached to the front of the car.



Figure 7.7: The front of the car with the MIMO radar system attached

The radar system contains a $3 \text{ Tx} \times 4 \text{ Rx}$ MIMO array operating at a bandwidth of 860 MHz. The angular resolution is approximately $\Delta\theta = 8^\circ$ at the broadside, the range resolution is $\Delta R \approx 0.1738$ m and the Doppler resolution is approximately $\Delta v \approx 3.81$ cm/s. The maximum measurable range is $R_{max} = 94$ m and the field of view of the radar is $\pm 60^\circ$. The data was captured at a frequency of 10 Hz. During the experiment, the car was driving at approximately $v_{car} = 15$ km/h on the quiet road the Leeghwaterstraat on the TU Delft campus. Figure 7.8 shows the view from the car during the experiment.

The collected data was pre-processed into a range-angle-Doppler data cube after which median detection was performed in the Range-Doppler domain followed by classical beamforming. In order to filter out any non-stationary targets, the vehicle velocity was estimated according to the process in 4.2.4, taking the median value over the number of targets. Using this estimate of the velocity, non-stationary targets were filtered out by discarding any target that did not meet the criterion set by (4.7) for $\epsilon = 0.5$ m/s.

7.2.2. Comparison of the Different Pose Estimation Techniques

From the pre-processed data 150 consecutive frames were taken to perform the scan-matching. The performance of each of the four scan-matching techniques is compared and the influence of the incorporation of knowledge about the target RCS is examined. The accuracy of the pose estimation is compared visually, since the only available ground truth is the GPS data which, as mentioned earlier, is generally not of very high quality. To visualize the performance, the trajectories as a result of the combination of the individual pose estimates have to be calculated according to (7.1):



Figure 7.8: A snapshot of the view from the car during the experiment

$$\begin{aligned}\psi[t] &= \psi[t-1] + \phi[t] \\ x[t] &= x[t-1] + \sqrt{(t_x[t])^2 + (t_y[t])^2} \cos(\psi[t-1] + \text{atan2}(t_y[t], t_x[t])), \\ y[t] &= y[t-1] + \sqrt{(t_x[t])^2 + (t_y[t])^2} \sin(\psi[t-1] + \text{atan2}(t_y[t], t_x[t])),\end{aligned}\quad (7.1)$$

where x and y denote the position in the global Cartesian coordinate system at discrete time instance t and ψ denotes the heading in the global coordinate system. The starting point (the very first reference scan) is taken to be $x[0] = 0$, $y[0] = 0$, $\psi[0] = 0$. Extreme outliers in the pose estimates are filtered out before application of (7.1). This is done using a simple rule, if one of the estimation parameters is larger than a certain threshold, the estimate is removed and replaced by the estimated x and y translations of the previous timeframe along with a heading ϕ of 0. These trajectories are then plotted on top of a map of the Leeghwaterstraat. The GPS is used for an approximation of the relevant part of the road corresponding to the measurements. Additionally, the trajectory according to the GPS data is plotted. It should be noted that while conducting the verification, it was found that the GPS data is even less accurate than expected. When looking at the distance traversed according to the GPS within any given time frame, the velocity needed to accomplish this is much higher than the velocity that was driven during the experiment according to the people that were conducting the experiment. This observation is corroborated by the fact that the maximum measured Doppler velocities are consistently approximately 0.5-0.8 m/s lower than needed to travel the distance provided by the GPS data. This overestimation in the GPS data will also be seen in the results. Moreover, it should be clear that this representation of the estimates is not ideal. Since the individual poses are estimated, and not the total trajectory, any error early on, especially in heading, results in a trajectory that might not look like it follows the traversed path.

Additionally, the reconstructed trajectories are used to create a map of the surroundings. These maps are created by stitching together the individual range scans using the estimated poses between them. These maps are created only for the approaches that show good performance.

Finally, the performance of the sensor bias estimation is investigated. From Figure 7.7 it can be seen that the radar system was attached at a slight offset from the center of the car. Due to the curvature of the bumper this results in a bias in angle. An effort to estimate this bias will be made, this time by taking into account 300 consecutive frames.

The Reconstructed Trajectories

Figure 7.9 shows the results of the estimated trajectory for the different techniques, utilizing the different ways to take into account the knowledge about the RCS through the received power directly. The black lines indicate the full reconstructed trajectory for each time frame, the shapes together with the bars indicate the estimated global pose of the vehicle at intervals of ten time frames. The initialization

parameters are the same as in the verification of the simulations, the cell size is fixed at $R_{cell} = 0.5$ m. Any extreme outliers in the pose estimates are first filtered by removing estimated poses of which one of the variables lies outside a certain bound. For t_x this bound is chosen to be 75 cm, for t_y 20 cm and for ϕ 10 degrees. These are very reasonable bounds, even for a trajectory that is not straight, considering the velocity of the vehicle of approximately 15 km/h (≈ 4.2 m/s) at a time difference between frames of 0.1 seconds. The removed estimates are replaced by a vector containing the estimates of the previous time frame with a heading $\phi = 0$. The offset from the road that can be seen in the image is caused by the initialization of the trajectory using the GPS data.

The first thing that stands out is the extremely poor performance of the RDPNDT. The trajectory shows a curve, which can be explained by the problem in the minimum of the objective function due to loss of distinction between left and right, as was explained in Chapter 5. The curve is towards the left as expected. Further, the results show that the standard NDT does not perform well when applied to radar data if no changes are made to the algorithm. It is clear from the figures that each technique is affected differently by the incorporation of the RCS-awareness, very poor performance for each technique is seen when the RCS is not considered at all, whereas the PNDT shows its best performance with consideration of the RCS in only the score and the DPNDT shows it benefits very heavily from the consideration in both the score and the distribution. The best performing techniques (PNDT in 7.9c and DPNDT in 7.9d) underestimate the traveled distance compared to the GPS data by approximately 8 and 9 meters, respectively, this is in line with the higher velocity according to the GPS data of approximately 0.5-0.8 m/s over 15 seconds (150 consecutive frames at 10 Hz).

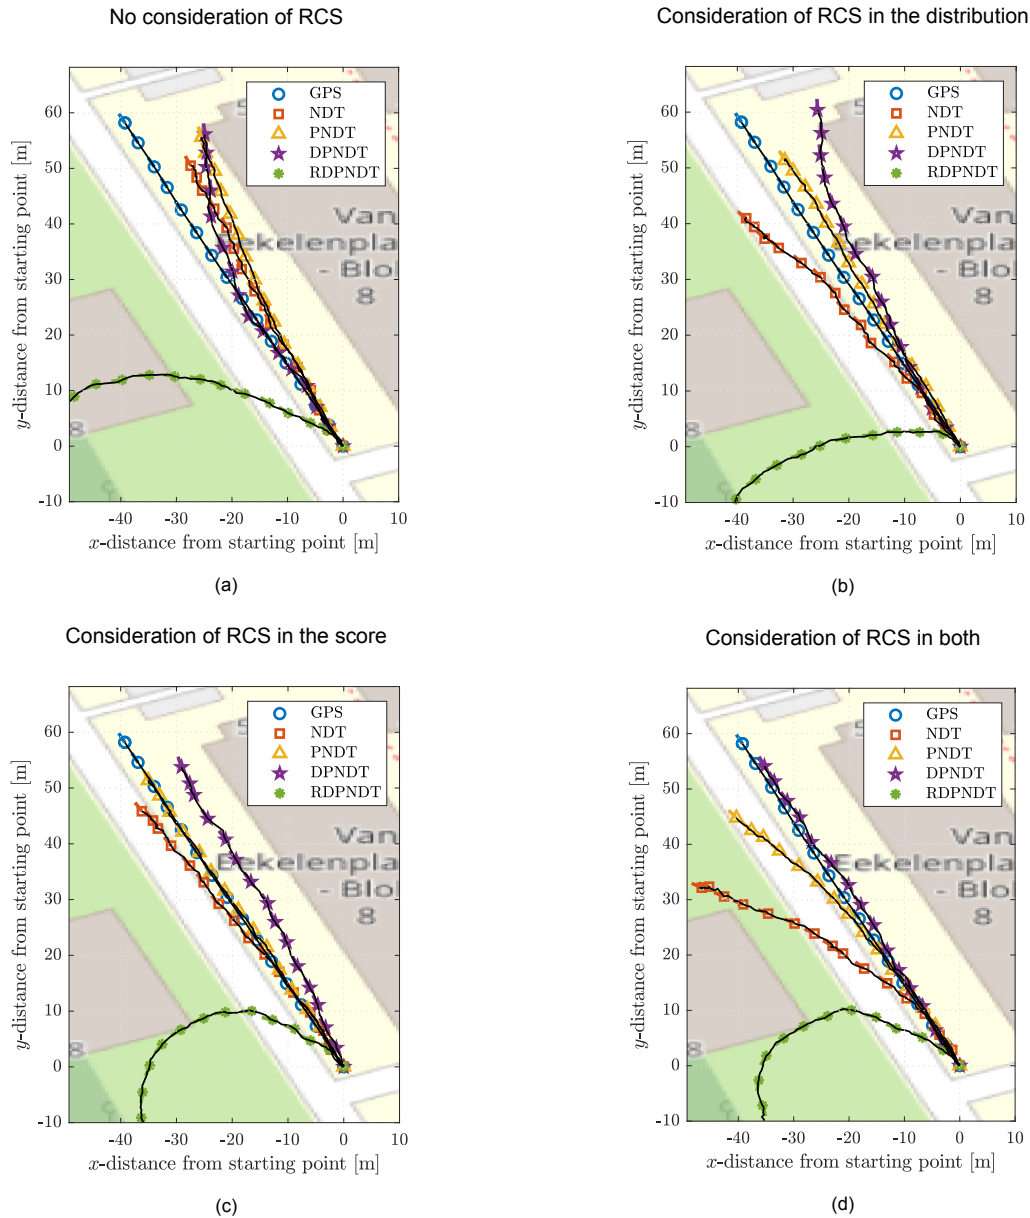


Figure 7.9: The trajectories calculated using the individual pose estimates of the NDT, PNNT, DPNT and RDPNT with the GPS data as a reference. Knowledge about the RCS is taken into account using the received power directly; (a) without incorporation of knowledge about the RCS, (b) with incorporation of knowledge about the RCS via the calculation of the score function, (c) with incorporation of knowledge about the RCS via the calculation of the distribution, (d) with incorporation of knowledge about the RCS via the calculation of both the score function and the distribution

Figure 7.10 shows results that were obtained in the same way, now considering the RCS by its relation with the received power through $\frac{1}{R^4}$. As explained in 4.4 this is done by first multiplying the received power associated with a detected target by its range to the fourth power. This results in a direct relation to the target's RCS. Overall, the DPNT benefits more from the incorporation of this direct reference to the RCS than simply the received power, as can be seen when comparing Figures 7.9 and 7.10. Strangely enough however, the combination of considering the RCS in both the score and the distribution leads to a curved shape of the trajectory, even more so than without consideration of the RCS in 7.10a.

The dependency of the results on the initialization parameters from Table 7.1 is illustrated by Figure 7.11. The images clearly show that the performance of each technique is very highly dependent on the choice of these parameters, so it can be expected that the presented results can be improved with

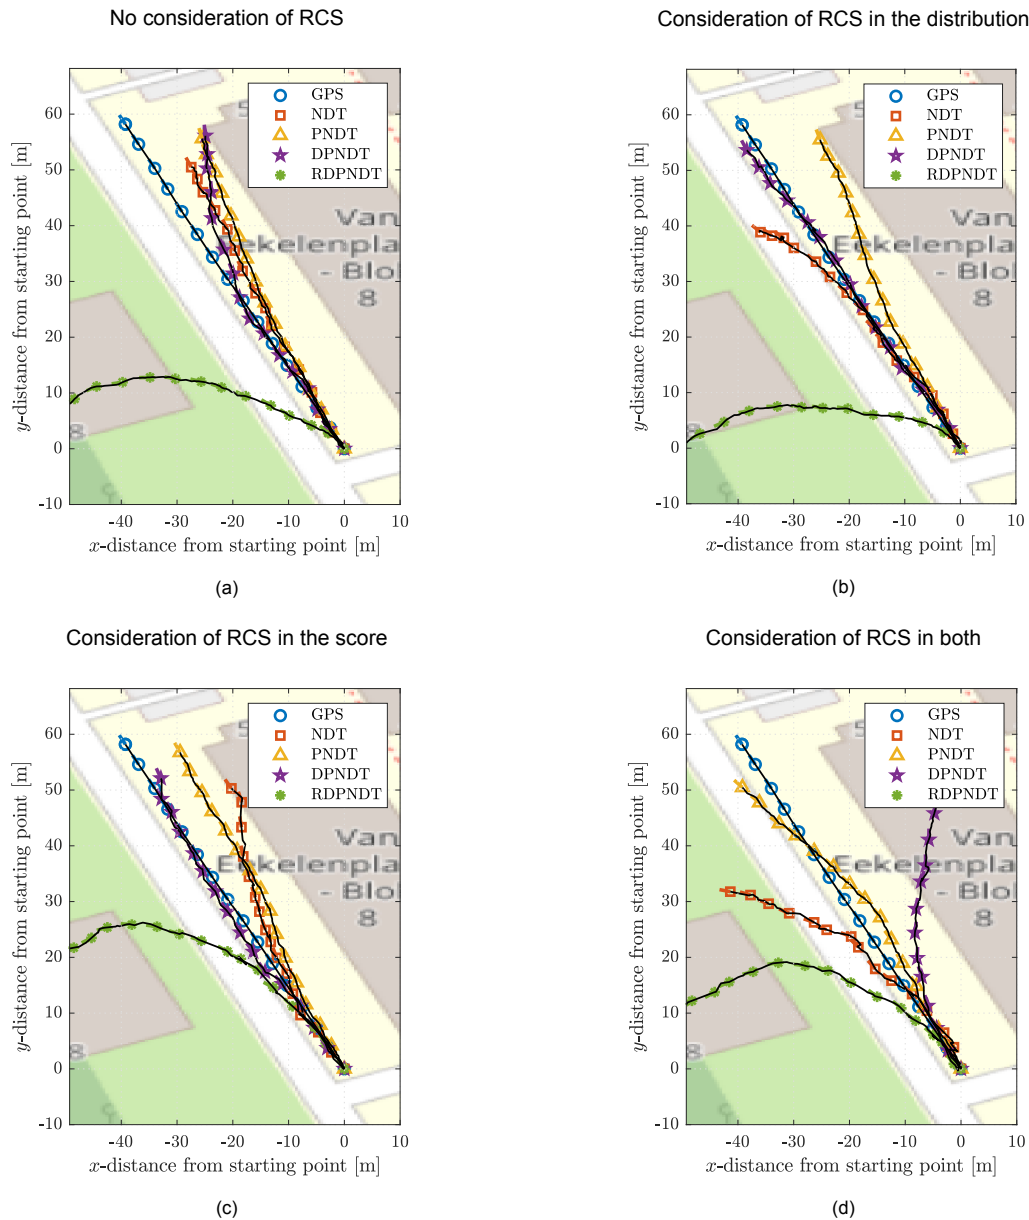


Figure 7.10: The trajectories calculated using the individual pose estimates of the NDT, PNDT, DPNDT and RDPNDT with the GPS data as a reference. Knowledge about the RCS is taken into account using the received power multiplied by R^4 , R being the target's measured range; (a) without incorporation of knowledge about the RCS, (b) with incorporation of knowledge about the RCS via the calculation of the score function, (c) with incorporation of knowledge about the RCS via the calculation of the distribution, (d) with incorporation of knowledge about the RCS via the calculation of both the score function and the distribution

further investigation into the optimality of these parameters. In the literature, problems regarding the cell size are often solved using techniques such as iterative subdivision [12], [15].

Building a Map of the Surroundings

Maps of the surroundings are constructed for the PNDT in case of RCS incorporation via the received power in the score, corresponding to Figure 7.9c, and for the DPNDT in the cases of RCS incorporation via the received power in both the score and the distribution, as in Figure 7.9d and in the case of RCS incorporation via the relation between received power and range in the distribution and in the score separately, corresponding to the trajectories in Figures 7.10b and 7.10c. The construction of these maps is performed by transforming the radar scans at each time frame according to the global translations \hat{x} and \hat{y} and rotation $\hat{\psi}$. By then transforming each range scan to the Cartesian coordinate frame a global map can be created, which can be visualized using the SNR values of the individual

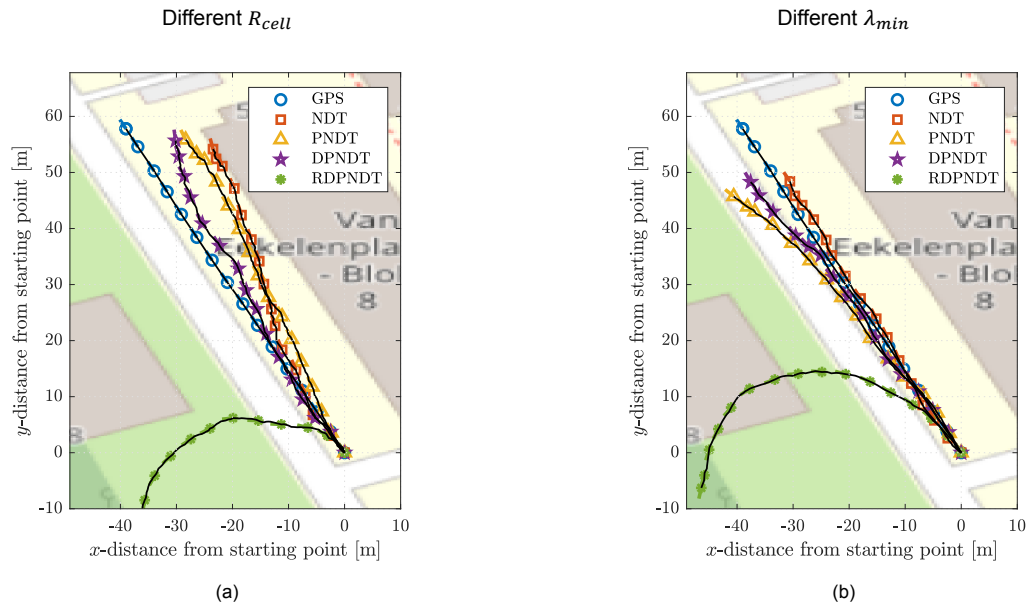


Figure 7.11: The influence of the choice of the initialization parameters of Table 7.1 visualized through the changes in the estimated trajectories of Figure 7.9d; (a) the constructed trajectory for $R_{cell} = 2$ m, (b) the constructed trajectory for $\lambda_{min} = 2$

targets. The resulting maps corresponding to the aforementioned trajectories are found in Figure 7.12.

Although from Figures 7.9d and 7.10c it seems like the DPNDT estimates the trajectory accurately, the maps in Figures 7.12b and 7.12d show many irregularities in what are assumed to be straight lines. The straightest maps are created by the trajectories estimated by the PNDT using knowledge about the SNR via the received power in the score and by the DPNDT using knowledge about the RCS directly in the distribution.

Execution Times

The average execution times for the estimates of Figure 7.9 are shown in Table 7.11. Again, the execution times have been normalized to the execution time of the NDT. It is debatable whether the immense increase in computational time of the DPNDT is worth the occasional increase in accuracy.

Table 7.11: Execution times of the NDT techniques, normalized to the execution time of the conventional NDT

	None	Score	Dist	Both
PNDT	1.2730	1.3387	1.2532	1.2853
DPNDT	15.8910	15.8661	14.4365	14.1560
RDPNDT	1.2109	1.2363	1.0399	1.0475

Sensor Bias Estimation

As mentioned before, the radar system was mounted to the vehicle at an angle to preserve vision of the license plate. Additionally, the radar system had not been calibrated for quite some time so the presence of a bias in angle was highly likely. The precise value of the bias is not known, so in an attempt to validate the sensor bias extension of the DPNDT, an artificial bias is added to the measurements, simply through addition of the same fixed value to the angle measurements of each frame, and the estimated bias before and after introduction of this artificial bias is compared. An estimate of the sensor bias can be found by taking the mean or median of the estimated biases $\hat{\epsilon}_\theta$ over the 300 time frames (30 seconds), its standard deviation is used as a measure of accuracy. The different ways to incorporate measurements of the RCS are compared. They are implemented by taking into account the received power directly, i.e. without multiplication with R^4 .

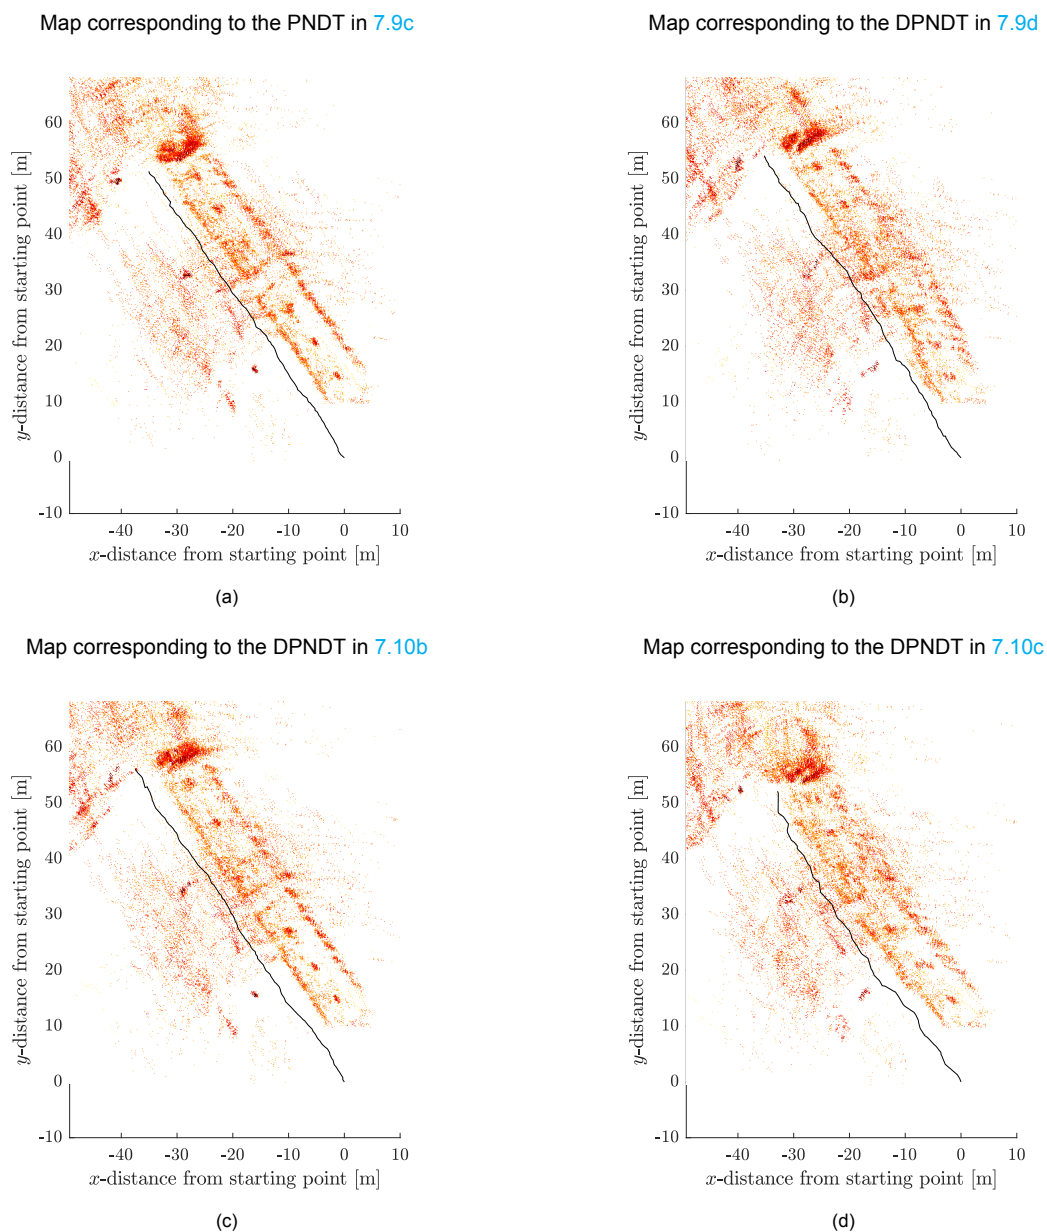


Figure 7.12: The resulting maps as created by transforming the individual range scans according to the estimated trajectories; (a) the map corresponding to the trajectory as calculated by the PNDT in 7.9c, (b) The map corresponding to the trajectory as calculated by the DPNDT in 7.9d, (c) The map corresponding to the trajectory as calculated by the DPNDT in 7.10b, (d) The map corresponding to the trajectory as calculated by the DPNDT in 7.10c distribution


The average estimated biases with their standard deviation, for each of the four RCS incorporation techniques, before and after introduction of the artificial bias are shown in Table 7.12, along with the difference of these averages. The notation of the mean and standard deviation is again “mean (standard deviation)”, the artificial bias was chosen to have a value of $\epsilon_{\theta} = 1^{\circ}$, so the difference should be as close as possible to 1. The Doppler resolution that was used during the experiment should be sufficient to estimate this artificial bias. Clearly, by taking into consideration measurements of the RCS in both the score and the distribution, an accurate estimate of the artificial sensor bias can be made. To take into account any outliers, the median value over the 100 frames is presented in Table 7.13. Here it can be concluded that by taking into account knowledge of the RCS in both the score and the distribution, the difference in bias can be estimated down to less than 0.02° of error.

Table 7.12: Average values and standard deviations of the estimated sensor bias over 100 frames, along with the difference after addition of an artificial bias

	w/o art. bias [°]	w/ art. bias [°]	Difference [°]
No RCS incorporation	0.4100 (0.6242)	1.1355 (1.0992)	0.7255
RCS incorporation in the distribution	0.3929 (1.8346)	1.1159 (2.2529)	0.7230
RCS incorporation in the score	0.5054 (1.2484)	1.3145 (1.6980)	0.8092
RCS incorporation in both	0.3663 (2.0340)	1.3108 (2.3662)	0.9444

Table 7.13: Median value of the estimated sensor bias over 100 frames, along with the difference after addition of an artificial bias

	w/o art. bias [°]	w/ art. bias [°]	Difference [°]
No RCS incorporation	0.3846	1.1073	0.7227
RCS incorporation in the distribution	0.3043	1.1575	0.8532
RCS incorporation in the score	0.4694	1.3784	0.9090
RCS incorporation in both	0.3464	1.3641	1.0177



Calculation of the Trajectory

In the results presented in Chapter 7 the trajectories were constructed using (7.1), filtering out any extreme outliers by re-using the previous estimates for t_x and t_y and setting $\phi = 0$ each time a calculated estimate was outside a certain bound. A more appropriate way to generate the trajectory is by means of a Bayesian filter, so that a dynamic model is used to compensate for uncertainties in the pose estimation.

In order to do so, a state evolution model and an observation model should be derived. Take state vector $\mathbf{s}^{[t]} = [x^{[t]}, y^{[t]}, \psi^{[t]}, v^{[t]}, x^{[t-1]}, y^{[t-1]}, \psi^{[t-1]}, v^{[t-1]}]^T$ containing the global position $[x, y]^T$, the global heading ψ and the vehicle velocity v at the current and previous time samples t and $t - 1$; using the nearly constant velocity model as used in [47], which assumes a close to straight trajectory, the following dynamic state evolution model is derived:

$$\underbrace{\begin{bmatrix} x^{[t+1]} \\ y^{[t+1]} \\ \psi^{[t+1]} \\ v^{[t+1]} \\ x^{[t]} \\ y^{[t]} \\ \psi^{[t]} \\ v^{[t]} \end{bmatrix}}_{\mathbf{s}^{[t+1]}} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & T \cos(\psi^{[t]}) \\ 0 & 1 & 0 & T \sin(\psi^{[t]}) \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\mathbf{f}(\mathbf{s}^{[t]})} \underbrace{\begin{bmatrix} x^{[t]} \\ y^{[t]} \\ \psi^{[t]} \\ v^{[t]} \\ x^{[t-1]} \\ y^{[t-1]} \\ \psi^{[t-1]} \\ v^{[t-1]} \end{bmatrix}}_{\mathbf{s}^{[t]}} + \underbrace{\begin{bmatrix} 0 \\ 0 \\ w_\psi^{[t]} \\ w_v^{[t]} \\ \mathbf{0}_{4 \times 1} \end{bmatrix}}_{\mathbf{w}^{[t]}}, \quad (8.1)$$

where $w_\psi^{[t]} \sim \mathcal{N}(0, \sigma_\psi)$ and $w_v^{[t]} \sim \mathcal{N}(0, \sigma_v)$ are the process noise related to fluctuations in velocity and steering respectively and T is the elapsed time between two frames. The global pose of the previous timeframe is introduced to accommodate the observation model, which is defined as follows: Take the observation vector $\mathbf{z}^{[t]} = [\hat{t}_x^{[t]}, \hat{t}_y^{[t]}, \hat{\phi}^{[t]}]^T$ containing the estimated relative vehicle pose. Now, the observation model can be derived:

$$\underbrace{\begin{bmatrix} \hat{t}_x^{[t]} \\ \hat{t}_y^{[t]} \\ \hat{\phi}^{[t]} \end{bmatrix}}_{\mathbf{z}^{[t]}} = \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 1} & -\mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 1} \end{bmatrix}}_{\mathbf{g}^{[t]}} \underbrace{\begin{bmatrix} x^{[t]} \\ y^{[t]} \\ \psi^{[t]} \\ v^{[t]} \\ x^{[t-1]} \\ y^{[t-1]} \\ \psi^{[t-1]} \\ v^{[t-1]} \end{bmatrix}}_{\mathbf{s}^{[t]}} + \underbrace{\begin{bmatrix} n_x^{[t]} \\ n_y^{[t]} \\ n_\psi^{[t]} \end{bmatrix}}_{\mathbf{n}^{[t]}}, \quad (8.2)$$

where $n_x^{[t]}$, $n_y^{[t]}$ and $n_\psi^{[t]}$ are terms related to the uncertainty of the local pose estimate. The Hessian \mathbf{H} corresponding to the pose estimate $\mathbf{p}^{[t]}$ provides a measure of the certainty of the estimate, its inverse

can be seen as an estimate of the covariance matrix [15]. $\mathbf{n}^{[t]}$ will thus be described as a zero mean multivariate Gaussian with covariance matrix \mathbf{H}^{-1} ; $\mathbf{n}^{[t]} \sim \mathcal{N}(\mathbf{0}, \mathbf{H}^{-1})$.

Due to the non-linearity of the evolution model, the traditional Kalman filter cannot be applied directly. Instead the Extended Kalman Filter (EKF) or the Iterated Extended Kalman Filter (IEKF) can be used.

The dynamic evolution model and the observation model have the following forms:

$$\begin{aligned}\mathbf{s}^{[t+1]} &= \mathbf{f}(\mathbf{s}^{[t]}) + \mathbf{w}^{[t]}, \\ \mathbf{z}^{[t]} &= \mathbf{G}^{[t]}\mathbf{s}^{[t]} + \mathbf{n}^{[t]},\end{aligned}\quad (8.3)$$

with $\mathbf{f}(\mathbf{s}^{[t]})$ a non-linear function of $\mathbf{s}^{[t]}$. The Extended Kalman filter estimates the state and its covariance through a prediction step, using the dynamic model, and an update step, using the observation model. The predicted state $\mathbf{s}^{[t+1|t]}$ and its predicted covariance are found according to:

$$\begin{aligned}\mathbf{s}^{[t+1|t]} &= \mathbf{f}(\mathbf{s}^{[t|t]}), \\ \mathbf{Q}^{[t+1|t]} &= \mathbf{F}\mathbf{Q}^{[t|t]}\mathbf{F}^T + \mathbf{W}^{[t]},\end{aligned}\quad (8.4)$$

where \mathbf{F} is the Jacobian of $\mathbf{f}(\mathbf{s}^{[t]})$ and $\mathbf{W}^{[t]} \in \mathcal{R}^{8 \times 8}$ is defined by σ_ψ and σ_v on the corresponding entries. This Jacobian is straightforward to calculate and has the following form:

$$\mathbf{F} = \left[\begin{array}{cccc|cc} 1 & 0 & -v^{[t]}T \sin(\psi^{[t]}) & T \cos(\psi^{[t]}) & & \\ 0 & 1 & v^{[t]}T \cos(\psi^{[t]}) & T \sin(\psi^{[t]}) & & \\ 0 & 0 & 1 & 0 & \mathbf{0}^{4 \times 4} & \\ 0 & 0 & 0 & 1 & & \\ \hline & & & & \mathbf{I}^{4 \times 4} & \mathbf{0}^{4 \times 4} \end{array} \right] \quad (8.5)$$

The update step has the same form as in the standard Kalman filter, since the observation model is linear:

$$\begin{aligned}\mathbf{s}^{[t|t]} &= \mathbf{s}^{[t|t-1]} + \mathbf{K}^{[t]}(\mathbf{z}^{[t]} - \mathbf{G}\mathbf{s}^{[t|t-1]}), \\ \mathbf{Q}^{[t|t]} &= (\mathbf{I} - \mathbf{K}^{[t]}\mathbf{G})\mathbf{Q}^{[t|t-1]}, \\ \mathbf{K}^{[t]} &= \mathbf{Q}^{[t|t-1]}\mathbf{G}^T (\mathbf{G}\mathbf{Q}^{[t|t-1]}\mathbf{G}^T + (\mathbf{H}^{[t]})^{-1})^{-1}\end{aligned}\quad (8.6)$$

Here, $\mathbf{s}^{[t|t-1]}$ and $\mathbf{Q}^{[t|t-1]}$ are the predicted estimates according to (8.4).

9

Conclusion

In this thesis, the steps taken during the research into a relative pose estimation technique for automotive radar were presented. During this research, several techniques were investigated for this purpose. The techniques are based on the Normal Distributions Transform (NDT), a distribution-based scan-matching technique originally developed for the robotics field with applications to LiDAR measurements. Scan-matching techniques aim to find the relative pose between two timestamps using range scans taken at those timestamps. They do so by trying to maximize the overlap between the scan at the first timestamp, the so-called “reference scan”, and the second timestamp, the so-called “current scan”. The NDT aims to solve this problem by addressing the scans in the Cartesian coordinate system and using one of the scans to construct a piecewise continuous distribution and use the probability of the points of the other scan within this distribution as a measure of overlap. This measure of overlap is then maximized in an iterative process using Newton’s method.

Some of the drawbacks associated with the use of laser scanners are their high price and poor performance in adverse weather conditions, making them less suitable for automotive applications. In an effort to develop techniques applicable in the automotive field, scan-matching techniques were explored that utilize radar measurements. In order to ensure applicability of the existing NDT to radar scans, a drawback that results from using radar measurements has to be addressed, which is the low resolution of the range scans compared to laser scans. Radar measurements contain properties that can be exploited to compensate for this. The aspects that were explored during the thesis research were the incorporation of Doppler measurements, which are known to contain higher resolution, as well as knowledge about the received signal power, which can be used to improve the robustness of the measure of overlap between the scans.

By reformulating the problem in polar coordinates through the Polar Normal Distributions Transform (PNDT), the representation of the data takes into account the spreading of the target responses and allows for the incorporation of Doppler measurements. The use of Doppler measurements was introduced in the Doppler Polar Normal Distributions Transform (DPNDT), which extends the PNDT to a third dimension, aiming to maximize overlap in not only the range and angle, but also the Doppler measurements. In an effort to reduce the computational complexity, the DPNDT was rewritten into a form only using the Range-Doppler domain to perform the optimization resulting in the Range-Doppler PNDT (RDPNDT). Furthermore, an extension to the DPNDT was investigated that allows for the estimation of a sensor bias which would affect the angle measurements but not the Doppler measurements. Additionally, for the construction of the traversed trajectory of the vehicle from the individual pose estimates, expressions to perform recursive filtering using the Extended Kalman Filter have been presented.

Results of applications of the pose estimation techniques to both simulation data and experimental data prove that the standard NDT has limited applicability to radar data due to its unnatural representation of the radar data in Cartesian coordinates. Improvements were made by the reformulation of the problem in the polar coordinate system, leading to the distribution used for maximization of the overlap to be a more authentic representation of the scanned scene. Additionally, it was found that the incorporation of knowledge about the Signal-to-Noise Ratio (SNR) of the detected targets leads to

further improvements in some cases for the experiments. It was found that the performance of each of the techniques is heavily depended on the initialization parameters, the optimality of which was not researched. Validation using simulation data suggests that the PNNT yields the most consistent improvements over the traditional NNT at the cost of a small increase in the computational complexity, showing a reduction in error of around 60% at an increase in computation time of only 5-10%. The experimental results show that under the right conditions the DPNT outperforms all other techniques in the pose estimation, however this comes at a computational cost that is more than tenfold the cost of the conventional NNT.

The estimation of a sensor bias in the angular measurements was examined using both experimental and simulated data. Simulated data showed the bias can be estimated with an error of less than 0.005° . In the case of the experiment, the initially present bias was not known, however, the difference in estimated biases before and after addition of an artificial bias corresponded to this artificial bias with an error of less than 0.02° .

Recommendations for future research include investigation of the optimality of the techniques for different initialization parameters. Moreover, the techniques should be validated for practical data with a more reliable ground truth. During the validation of the experimental data, it was found that the GPS data suffered some problems, leading to the combination of the pose estimates to underestimate the traversed trajectory relative to the trajectory according to the GPS. Ideally a ground truth would be available that provided accurate knowledge about the relative poses directly, rather than about the global poses as is the case with GPS data. Further verification of the sensor bias estimation should be conducted using experimental data with a known ground truth. Finally, the construction of the trajectories can be improved significantly by considering a dynamic model, the formulas needed to do so were derived but not yet implemented.

Bibliography

- [1] Connected Places Catapult, Cambridge Econometrics, Element Energy Limited., and Great Britain. Department for Transport, *Market Forecast For Connected and Autonomous Vehicles*, English, 2021. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/919260/connected-places-catapult-market-forecast-for-connected-and-autonomous-vehicles.pdf.
- [2] On-Road Automated Driving (ORAD) committee, *Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems*, 2021. [Online]. Available: https://www.sae.org/standards/content/j3016%5C_202104/.
- [3] Z. Chai, T. Nie, and J. Becker, "Technologies for autonomous driving," in *Autonomous Driving Changes the Future*. Singapore: Springer Singapore, 2021, pp. 17–61.
- [4] G. Bresson, Z. Alsayed, L. Yu, and S. Glaser, "Simultaneous Localization and Mapping: A Survey of Current Trends in Autonomous Driving," *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 3, pp. 194–220, 2017. DOI: [10.1109/TIV.2017.2749181](https://doi.org/10.1109/TIV.2017.2749181).
- [5] D. Odijk and L. Wanninger, "Differential positioning," in *Springer Handbook of Global Navigation Satellite Systems*, P. J. Teunissen and O. Montenbruck, Eds. Cham: Springer International Publishing, 2017, pp. 753–780, ISBN: 978-3-319-42928-1.
- [6] I. Cox, "Blanche – An Experiment in Guidance and Navigation of an Autonomous Robot Vehicle," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 2, pp. 193–204, 1991. DOI: [10.1109/70.75902](https://doi.org/10.1109/70.75902).
- [7] J. Leonard and H. Durrant-Whyte, "Simultaneous Map Building and Localization for an Autonomous Mobile Robot," in *Proceedings IROS '91: IEEE/RSJ International Workshop on Intelligent Robots and Systems '91*, 1991, 1442–1447 vol.3. DOI: [10.1109/IROS.1991.174711](https://doi.org/10.1109/IROS.1991.174711).
- [8] F. Lu and Milios, "Robot Pose Estimation in Unknown Environments by Matching 2D Range Scans," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Los Alamitos, CA, USA: IEEE Computer Society, 1994, pp. 935–938.
- [9] A. Diosi and L. Kleeman, "Laser Scan Matching in Polar Coordinates with Application to SLAM," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 3317–3322. DOI: [10.1109/IROS.2005.1545181](https://doi.org/10.1109/IROS.2005.1545181).
- [10] P. J. Besl and N. D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992. DOI: [10.1109/34.121791](https://doi.org/10.1109/34.121791).
- [11] P. Biber and W. Strasser, "The Normal Distributions Transform: a New Approach to Laser Scan Matching," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 3, 2003, 2743–2748 vol.3. DOI: [10.1109/IROS.2003.1249285](https://doi.org/10.1109/IROS.2003.1249285).
- [12] A. Das and S. L. Waslander, "Scan Registration with Multi-Scale k-Means Normal Distributions Transform," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 2705–2710. DOI: [10.1109/IROS.2012.6386185](https://doi.org/10.1109/IROS.2012.6386185).
- [13] P. Saranrittichai, N. Niparnan, and A. Sudsang, "2D-Laser Scan Registration Using Multi-Scale NDT with Polar Scan Clustering," in *Advanced Research in Material Science and Mechanical Engineering*, ser. Applied Mechanics and Materials, vol. 446, Trans Tech Publications Ltd, 2014, pp. 981–985. DOI: [10.4028/www.scientific.net/AMM.446-447.981](https://doi.org/10.4028/www.scientific.net/AMM.446-447.981).
- [14] M. Magnusson, A. Lilienthal, and T. Duckett, "Scan Registration for Autonomous Mining Vehicles using 3D-NDT," *Journal of Field Robotics*, vol. 24, no. 10, pp. 803–827, 2007. DOI: <https://doi.org/10.1002/rob.20204>.

- [15] M. Magnusson, "The Three-Dimensional Normal-Distributions Transform: an Efficient Representation for Registration, Surface Analysis, and Loop Detection," Ph.D. dissertation, Örebro University, School of Science and Technology, 2009, p. 201, ISBN: 978-91-7668-696-6.
- [16] S. Campbell, N. O'Mahony, L. Krpalcova, D. Riordan, J. Walsh, A. Murphy, and C. Ryan, "Sensor Technology in Autonomous Vehicles: A review," in *2018 29th Irish Signals and Systems Conference (ISSC)*, 2018, pp. 1–4. DOI: [10.1109/ISSC.2018.8585340](https://doi.org/10.1109/ISSC.2018.8585340).
- [17] M. Kutilla, P. Pyykönen, H. Holzhüter, M. Colomb, and P. Duthon, "Automotive LiDAR Performance Verification in Fog and Rain," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 1695–1701. DOI: [10.1109/ITSC.2018.8569624](https://doi.org/10.1109/ITSC.2018.8569624).
- [18] F. Bertolli, P. Jensfelt, and H. I. Christensen, "SLAM using Visual Scan-Matching with Distinguishable 3D Points," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2006, pp. 4042–4047.
- [19] N. Naikal, J. Kua, G. Chen, and A. Zakhor, "Image Augmented Laser Scan Matching for Indoor Dead Reckoning," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2009, pp. 4134–4141.
- [20] M. Holder, S. Hellwig, and H. Winner, "Real-Time Pose Graph SLAM based on Radar," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1145–1151. DOI: [10.1109/IVS.2019.8813841](https://doi.org/10.1109/IVS.2019.8813841).
- [21] Y. Li, Y. Liu, Y. Wang, Y. Lin, and W. Shen, "The Millimeter-Wave Radar SLAM Assisted by the RCS Feature of the Target and IMU," eng, *Sensors (Basel, Switzerland)*, vol. 20, no. 18, p. 5421, 2020.
- [22] J. Callmer, D. Törnqvist, F. Gustafsson, H. Svensson, and P. Carlbom, "Radar SLAM using Visual Features," *EURASIP Journal on Advances in Signal Processing*, vol. 2011, no. 1, pp. 1–11, 2011. DOI: [10.1186/1687-6180-2011-71](https://doi.org/10.1186/1687-6180-2011-71).
- [23] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar, Volume I - Basic Principles*. SciTech Publishing, 2010, ISBN: 978-1-891121-52-4.
- [24] Infineon Technologies AG, *XENSIV™ - Sensor solutions for automotive, industrial and consumer applications*, Munich, Germany, 2021. [Online]. Available: https://www.infineon.com/dgdl/Infineon-SensorSelectionGuide-ProductSelectionGuide-v01%5C_00-EN.pdf?fileId=5546d462636cc8fb0164229c09f51bbe.
- [25] NXP B.V, *TEF82xx 77 GHz Automotive Radar Transceiver Fact Sheet*, 2020. [Online]. Available: <https://www.nxp.com/docs/en/fact-sheet/TEF82XX-FS.pdf>.
- [26] J. Gamba, "Fundamentals of radar systems," in *Radar Signal Processing for Autonomous Driving*. Springer Singapore, 2020, pp. 1–14, ISBN: 978-981-13-9193-4.
- [27] P. Tait, *Introduction to Radar Target Recognition*. London: Institution of Engineering and Technology, 2009.
- [28] L. Daniel, A. Stove, E. Hoare, D. Phippen, M. Cherniakov, B. Mulgrew, and M. Gashinova, "Application of Doppler Beam Sharpening for Azimuth Refinement in Prospective Low-THz Automotive Radars," *IET Radar, Sonar & Navigation*, vol. 12, no. 10, pp. 1121–1130, 2018.
- [29] H.-Y. Shum, Q. Ke, and Z. Zhang, "Efficient Bundle Adjustment with Virtual Key Frames: a Hierarchical Approach to Multi-Frame Structure from Motion," in *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, vol. 2, 1999, 538–543 Vol. 2. DOI: [10.1109/CVPR.1999.784733](https://doi.org/10.1109/CVPR.1999.784733).
- [30] E. Royer, M. Lhuillier, M. Dhome, and T. Chateau, "Localization in Urban Environments: Monocular Vision Compared to a Differential GPS Sensor," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, 2005, 114–121 vol. 2. DOI: [10.1109/CVPR.2005.217](https://doi.org/10.1109/CVPR.2005.217).
- [31] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based slam," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 4, pp. 31–43, 2010. DOI: [10.1109/MITS.2010.939925](https://doi.org/10.1109/MITS.2010.939925).

- [32] T. Duckett, S. Marsland, and J. Shapiro, "Fast, on-line learning of globally consistent maps," *Autonomous Robots*, vol. 12, no. 3, pp. 287–300, May 2002, ISSN: 1573-7527. DOI: [10.1023/A:1015269615729](https://doi.org/10.1023/A:1015269615729).
- [33] E. Olson, J. Leonard, and S. Teller, "Fast iterative alignment of pose graphs with poor initial estimates," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, 2006, pp. 2262–2269. DOI: [10.1109/ROBOT.2006.1642040](https://doi.org/10.1109/ROBOT.2006.1642040).
- [34] G. Grisetti, S. Grzonka, C. Stachniss, P. Pfaff, and W. Burgard, "Efficient Estimation of Accurate Maximum Likelihood Maps in 3D," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 3472–3478. DOI: [10.1109/IROS.2007.4399030](https://doi.org/10.1109/IROS.2007.4399030).
- [35] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960. DOI: [10.1115/1.3662552](https://doi.org/10.1115/1.3662552).
- [36] S. M. Kay, "Chapter 13: Kalman filters," in *Fundamentals of Statistical Signal Processing Vol. I: Estimation Theory*. Prentice-Hall, 1993.
- [37] M. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, "A Solution to the Simultaneous Localization and Map Building (SLAM) Problem," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, 2001. DOI: [10.1109/70.938381](https://doi.org/10.1109/70.938381).
- [38] J. Guivant and E. Nebot, "Optimization of the Simultaneous Localization and Map-Building Algorithm for Real-Time Implementation," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, pp. 242–257, 2001. DOI: [10.1109/70.938382](https://doi.org/10.1109/70.938382).
- [39] S. J. Julier and J. K. Uhlmann, "New Extension of the Kalman Filter to Nonlinear Systems," in *SPIE*, 1997, pp. 182–193, ISBN: 9780819424839. DOI: [10.1117/12.280797](https://doi.org/10.1117/12.280797).
- [40] D. Chekhlov, M. Pupilli, W. Mayol-Cuevas, and A. Calway, "Real-Time and Robust Monocular SLAM Using Predictive Multi-resolution Descriptors," in, ser. *Advances in Visual Computing*. Springer Berlin, Heidelberg, 2006, pp. 276–285.
- [41] A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," English, *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, 107–113(6), 2 Apr. 1993, ISSN: 0956-375X. [Online]. Available: <https://digital-library.theiet.org/content/journals/10.1049/ip-f-2.1993.0015>.
- [42] M. Montemerlo, S. Thrun, and B. Siciliano, *FastSLAM: a Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics*, English. Berlin: Springer, 2007. DOI: [10.1007/978-3-540-46402-0](https://doi.org/10.1007/978-3-540-46402-0).
- [43] S. Thrun, Y. Liu, D. Koller, A. Y. Ng, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous localization and mapping with sparse extended information filters," *The International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, Aug. 2004. DOI: [10.1177/0278364904045479](https://doi.org/10.1177/0278364904045479).
- [44] G. Dubbelman and B. Browning, "COP-SLAM: Closed-Form Online Pose-Chain Optimization for Visual SLAM," *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1194–1213, 2015. DOI: [10.1109/TRO.2015.2473455](https://doi.org/10.1109/TRO.2015.2473455).
- [45] J. Saarinen, H. Andreasson, T. Stoyanov, and A. J. Lilienthal, "Normal distributions transform monte-carlo localization (ndt-mcl)," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 382–389. DOI: [10.1109/IROS.2013.6696380](https://doi.org/10.1109/IROS.2013.6696380).
- [46] D. M. Malioutov, *A Sparse Signal Reconstruction Perspective for Source Localization With Sensor Arrays*, Mass. Inst. Technol., Cambridge, MA, 2003. [Online]. Available: http://ssg.mit.edu/~dmm/publications/malioutov_MS_thesis.pdf.
- [47] N. Petrov, O. Krasnov, and A. G. Yarovoy, "Auto-calibration of automotive radars in operational mode using simultaneous localisation and mapping," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 3, pp. 2062–2075, 2021. DOI: [10.1109/TVT.2021.3058778](https://doi.org/10.1109/TVT.2021.3058778).
- [48] R. Kümmerle, G. Grisetti, and W. Burgard, "Simultaneous Parameter Calibration, Localization, and Mapping," *Advanced Robotics*, vol. 26, no. 17, pp. 2021–2041, 2012. DOI: [10.1080/01691864.2012.728694](https://doi.org/10.1080/01691864.2012.728694).
- [49] G. R. PRICE, "Extension of covariance selection mathematics," *Annals of Human Genetics*, vol. 35, no. 4, pp. 485–490, Apr. 1972. DOI: [10.1111/j.1469-1809.1957.tb01874.x](https://doi.org/10.1111/j.1469-1809.1957.tb01874.x).

-
- [50] S. Boyd, *Convex optimization*. Cambridge, UK New York: Cambridge University Press, 2004, ISBN: 978-0-521-83378-3.
 - [51] C. M. Schmid, S. Schuster, R. Feger, and A. Stelzer, "On the Effects of Calibration Errors and Mutual Coupling on the Beam Pattern of an Antenna Array," *IEEE Transactions on Antennas and Propagation*, vol. 61, no. 8, pp. 4063–4072, 2013.
 - [52] N. Petrov, O. Krasnov, and A. Yarovoy, "Statistical Approach for Automotive Radar Self-Diagnostics," in *2019 16th European Radar Conference (EuRAD)*, 2019, pp. 117–120.