

SIGNAL PROCESSING BASED DISTORTION MITIGATION IN INTERFEROMETRIC
RADAR ANGULAR VELOCITY ESTIMATION

By

Eric Klinefelter

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Electrical Engineering – Doctor of Philosophy

2021

ABSTRACT

SIGNAL PROCESSING BASED DISTORTION MITIGATION IN INTERFEROMETRIC RADAR ANGULAR VELOCITY ESTIMATION

By

Eric Klinefelter

Interferometric angular velocity estimation is a relatively recent radar technique which uses a pair of widely spaced antenna elements and a correlation receiver to directly measure the angular velocity of a target. Traditional radar systems measure range, radial velocity (Doppler), and angle, while angular velocity is typically derived as the time-rate change of the angle measurements. The noise associated with the derived angular velocity estimate is statistically correlated with the angle measurements, and thus provides no additional information to traditional state space trackers. Interferometric angular velocity estimation, on the other hand, provides an independent measurement, thus forming a basis in \mathbb{R}^2 for both position and velocity.

While promising results have been presented for single target interferometric angular velocity estimation, there is a known issue which arises when multiple targets are present. The ideal interferometric response with multiple targets would contain only the mixing product between like targets across the antenna responses, yet instead, the mixing product between all targets is generated, resulting in unwanted ‘cross-terms’ or intermodulation distortion. To date, various hardware based methods have been presented, which are effective, though they tend to require an increased number of antenna elements, a larger physical system baseline, or signals with wide bandwidths. Presented here are novel methods for signal processing based interferometric angular velocity estimation distortion mitigation, which can be performed with only a single antenna pair and traditional continuous-wave or frequency-modulated continuous wave signals.

In this work, two classes of distortion mitigation methods are described: model-based and response decomposition. Model-based methods use a learned or analytic model

with traditional non-linear optimization techniques to arrive at angular velocity estimates based on the complete interferometric signal response. Response decomposition methods, on the other hand, aim to decompose the individual antenna responses into separate responses pertaining to each target, then associate like targets between antenna responses. By performing the correlation in this manner, the cross-terms, which typically corrupt the interferometric response, are mitigated. It was found that due to the quadratic scaling of distortion terms, model-based methods become exceedingly difficult as the number of targets grows large. Thus, the method of response decomposition is selected and results on measured radar signals are presented. For this, a custom single-board millimeter-wave interferometric radar was developed, and angular velocity measurements were performed in an enclosed environment consisting of two robotic targets. A set of experiments was designed to highlight easy, medium, and difficult cases for the response decomposition algorithm, and results are presented herein.

Copyright by
ERIC KLINEFELTER
2021

To my family for always believing in me and supporting my wild ideas.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Jeffrey A. Nanzer, for giving me the opportunity to join the Electromagnetics Research Group here at Michigan State. Your work ethic and subject matter expertise have been great sources of motivation for me personally, which I know I will carry on throughout my career. To my other committee members, I would like to thank Dr. Shanker Balasubramaniam for your always insightful comments and questions. Though computational electromagnetics is your expertise, it is your personal and career advice that I will remember the most. To Dr. Vishnu Boddeti, I thank you for the machine learning expertise I have gained from your course and our discussions, even when they were completely unrelated to my research, like teaching a neural net how to beat the stock market.

I am deeply indebted to Dr. Hayder Radha, Dr. Daniel Morris, and my advisor for writing the proposal and mentoring our team in the SAE Autodrive Challenge. This is where I found my true passion and where I will undoubtedly spend the rest of my career. Thank you to General Motors, Intel, and the other SAE Autodrive Challenge sponsors for providing vehicles, sensors, and other hardware, on which I gained invaluable experience developing software for our autonomous vehicle. Thank you to Jason Merlo for being our Autodrive project manager and for being a constant source of knowledge, ideas, and feedback. Also, thank you to the rest of my lab mates in both the Electromagnetics Research Group and the DELTA group. I have had so many great discussions with you all, which I will definitely miss. Finally, I'd like to thank my family, my friends, and my wife Rachel for always believing in me throughout my time in graduate school.

TABLE OF CONTENTS

LIST OF TABLES	ix
LIST OF FIGURES	x
CHAPTER 1 INTRODUCTION	1
1.1 Background	2
1.2 Research Problem and Objectives	8
1.3 Research Significance	10
1.4 Scope of This Research	11
1.5 Dissertation Organization	11
CHAPTER 2 GENERAL INTERFEROMETRIC RESPONSE TO AN N-POINT TARGET	13
2.1 Reference Geometry	13
2.2 Interferometric N-Point Target Model	15
2.2.1 Exact Model: Arbitrary Transmitter Location	16
2.2.2 Exact Model: Transmitter Co-located with Receiver 1	17
2.2.3 Approximate Model: Far-field and Small-angle Approximations	18
2.3 Simulated Model Response	19
CHAPTER 3 SIGNAL PROCESSING BASED DISTORTION MITIGATION FOR MULTITARGET ANGULAR VELOCITY ESTIMATION	24
3.1 Model-Based Methods	24
3.1.1 Learned Representations	24
3.1.2 Analytic Model	25
3.2 Response Decomposition Methods	29
3.2.1 Signal Decomposition	31
3.2.2 Data Association	31
CHAPTER 4 MEASUREMENT RESULTS FOR RESPONSE DECOMPOSITION INTERFEROMETRIC DISTORTION MITIGATION	33
4.1 Measurement System	33
4.1.1 Single-Board 40 GHz Interferometric Radar	33
4.1.2 V-Band Patch Antenna Design	37
4.1.3 Ground-Truth System and Robotic Targets	39
4.2 Measured Results	40
4.2.1 Angular Velocity Estimation via Response Decomposition	45
4.2.1.1 Known Target Frequencies and Data Associations	46
4.2.1.2 Unknown Target Frequencies and Data Associations	49
CHAPTER 5 PRACTICAL APPLICATIONS OF ANGULAR VELOCITY ESTI- MATION	51

5.1	Vehicle Speed Estimation	51
5.1.1	Signal Model	53
5.1.2	Velocity Estimation	55
5.1.2.1	Bayesian Estimator	55
5.1.2.2	Neural Net Estimator	57
5.1.3	Experimental Evaluation	59
5.1.3.1	Measurements	59
5.1.3.2	Estimator Results	62
5.2	Multiobject Tracking with Orthogonal Velocity Measurements	65
5.2.1	Object Tracking	66
5.2.1.1	Single Object Tracking	67
5.2.1.2	Multiobject Tracking	67
5.2.2	Experimental Results	68
5.2.2.1	Simulated Results	68
5.2.2.2	nuScenes Dataset Results	71
5.3	Hand Gesture Recognition	75
5.3.1	Experimental Measurements	76
5.3.2	CNN Classifier	77
5.3.3	Classification Results	79
CHAPTER 6 CONCLUSION		83
6.1	Contributions	83
6.2	Future Work	84
BIBLIOGRAPHY		87

LIST OF TABLES

Table 2.1: Simulation 1 Target Parameters	21
Table 2.2: Simulation 2 Target Parameters	22
Table 2.3: Simulation 3 Target Parameters	22
Table 4.1: Angular Velocity Estimation Statistics With Known Target Frequencies .	50
Table 4.2: Angular Velocity Estimation Statistics With Unknown Target Frequencies	50
Table 5.1: RMSE of Estimators	65
Table 5.2: Average RMSE over 10k realizations	70
Table 5.3: State RMSE for target vehicle in scene-0061	74
Table 5.4: Classification accuracy over all aspect angles for 0° and 45° radar rotation	81
Table 5.5: Classification accuracy over 0° and 45° radar rotation for all aspect angles	81

LIST OF FIGURES

<p>Figure 1.1: The fringe pattern of a two-element interferometric array has lobes of ambiguous absolute angle but increased angular resolution. The diagram shows the response of the two point sources as the interferometric fringe pattern scans the sky due to the rotation of the Earth. Image [1] © 2019 IET Press.</p>	3
<p>Figure 1.2: Block diagram of a two-element correlation interferometer with system baseline D, transmitted signal $s(t)$, and correlator output $r(t)$. The fringe pattern creates a series of peaks and nulls created by the constructive and destructive interference of the received signals at various angles. The geometric time-delay is shown as τ_g and assumes that the target is far-field to the array.</p>	4
<p>Figure 1.3: Simulated interferometric time-frequency response for two point scatterers. The two oscillating lower frequency terms are the ideal interferometric terms, while the higher frequency oscillations are the intermodulation distortion terms. Image [2] © 2021 IEEE.</p>	8
<p>Figure 2.1: The notation for the parameters used in this work. RX1 is placed at the origin, and point n contains both radial (blue) and tangential (orange) velocity components relative to the origin. In the case that TX1 is not co-located with RX1 we must project the velocities as shown, considering the second antenna to be the transmitter. Similarly, we must always project the velocities at RX2. In both cases the projected radial velocity will have two components, one from the original radial and one from the original tangential velocities. Image [2] © 2021 IEEE.</p>	14
<p>Figure 2.2: The target must be located in the 3 dB antenna overlap region of RX1 and RX2. This is usually not an issue for wide beam antennas, though can be a difficult requirement when using high gain (narrow beam) antennas with a large electrical baseline, shown here as $D_\lambda = \frac{D}{\lambda}$.</p>	15
<p>Figure 2.3: The percent error between the frequency estimates for target position (a) with fixed radial velocity of 1.0 m/s and angular velocity of 1.0 rad/s. The frequency percent error for radial and angular velocities (b) with target position fixed at a range of 20.0 m and angle of 0°. The error has been limited to 5% around the origin of plot (b) where the error increases rapidly as ω approaches zero. Image [2] © 2021 IEEE.</p>	19

Figure 2.4:	Simulated interferometric time-frequency response for a single point scatterer (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.	20
Figure 2.5:	Simulated interferometric time-frequency response for two point scatterers (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.	21
Figure 2.6:	Simulated interferometric time-frequency response for four point scatterers (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.	23
Figure 3.1:	The gradient algorithm uses the measured and model response as inputs to the loss function. The gradient of the loss function is calculated and the update rule is used to calculate the new angular velocity estimates. This happens X times, or until the total error is below some threshold.	27
Figure 3.2:	The loss surface (a) for two estimates of angular velocity of an interferometric radar with true values of $\omega = [0.001; -0.189]$ rad/s. Notice that the entire loss space is non-convex and so the gradient method will only approach the true values when the initial estimate is close to the true parameter value. The contour plot (b) of the loss space with the initial estimate shown in as a blue 'x'. We run 50 iterations with each iterative estimate shown as a green '+'. The estimates approach the true value shown as a black 'o'.	28
Figure 4.1:	Schematic of millimeter-wave radar used for experimental measurements. System consists of two phase-quadrature, direct-downconversion receive channels with a continuous-wave transmitter operating at 40 GHz.	34
Figure 4.2:	Testing the 40 GHz grounded coplanar waveguide transmission lines in Ansys SIwave.	34
Figure 4.3:	The 40 GHz system on Rogers 4350 substrate has 3 IQ receive channels and a single transmit channel. The antennas are connectorized to allow for experimentation with various antenna types and arrays. Image [3] © 2020 IEEE.	35
Figure 4.4:	The measurement system used in this work consisted of a custom single board radar and L3-NARDA 15 dBi horn antennas with a 20λ interferometric baseline.	36

Figure 4.5: Model of the 1:2 Wilkinson power splitter in Ansys HFSS (a), and the entire 1:4 structure (b) for full-wave simulation.	37
Figure 4.6: The 40 GHz patch antenna with 2.4 mm connector. The design uses a quarter-wave transformer to match the transmission line impedance of 50Ω to that of the edge of the microstrip patch. The antenna has a dielectric superstrate included to shield the antenna from environmental conditions. Image [4] © 2020 IEEE.	38
Figure 4.7: Measurements of the S_{11} from two antennas each with a resonance frequency of 42.4 GHz. The overall decrease in S_{11} at higher frequencies is likely an effect of the losses incurred by the dielectric superstrate. Image [4] © 2020 IEEE.	39
Figure 4.8: Top-down view of the experimental setup. The six motion capture cameras form the capture volume shown as the dashed circle. The radar is placed at the edge of the volume and we set the origin to the center of the interferometric array. The targets travel in various trajectories as shown.	40
Figure 4.9: Robotis Turtlebot3 Waffle platforms used for measurements, the reflective targets consisted of polystyrene spheres coated in copper tape.	41
Figure 4.10: The time-frequency responses for a single target at various trajectory angles shown in each column. The plots are the Doppler response from antenna one (a), the Doppler response from antenna two with the expected Doppler shift overlaid in a white dashed line (b), and the interferometric response with the expected interferometric shift (c). The Doppler responses of columns two and six look identical, but we can notice a small negative interferometric shift in column two, and a small positive shift in column six, implying counter-clockwise and clockwise motion.	42
Figure 4.11: The time-frequency responses for two targets with the first target at a constant angle and the second target at various trajectory angles shown in each column. The plots from top to bottom are the Doppler response from antenna one (a), the Doppler response from antenna two with the expected Doppler shift overlaid in a white dashed line (b), and the interferometric response with the expected interferometric shift (c).	43
Figure 4.12: Robot trajectory directions for the easy (a), medium (b), and difficult case (c), with the antenna receiving pair shown at the bottom of each image. The large differential in target Doppler shifts will be easily separable in case (a), while the negligible Doppler shifts in case (c) will not be separable.	44

Figure 4.13: The measured time-frequency responses for the easy case include the antenna one Doppler response (a), antenna two Doppler response with ground-truth expected Doppler shift (b), and interferometric response with the expected interferometric shift (c). The same responses are shown for the medium case in (d), (e), (f), as well as the difficult case in (g), (h), and (i).	45
Figure 4.14: The reconstructed interferometric responses for the first case target one (a), and target two (d), the second case for target one (b), and target two (e), and the third case for target one (c), and target two (f). Here we can see good agreement for case one and two, while there is a visible bias in the responses for case three.	47
Figure 4.15: Angular velocity estimates for the first case with ground-truth angular velocity underlaid for target one (a) and target two (b), we see a low estimation error and variance for the easy case. The estimates for the second case are shown in (c) and (d) with a small increase in the variance as expected. The estimates for the third (difficult) case are shown in (e) and (f) where we see a larger increase in the estimation variance, as expected.	49
Figure 5.1: A single point scatterer, n , with linear velocity $v_{l,n}$ which can be decomposed into radial and tangential components. Here $\theta_{1,n}$ is the angle from receiver 1 to point n , and $d_{1,n}$ is the distance from receiver 1 to n . Image [5] © 2020 IEEE.	53
Figure 5.2: The backscattering pattern for a smooth surface (left) is primarily specular, a medium roughness surface (middle) has both a specular and a backscatter component, and the very rough surface (right) is primarily backscatter. Image [5] © 2020 IEEE.	55
Figure 5.3: A sample network with 256 input features and 16 hidden neurons. The input features represent the amplitudes of the frequency components from the spectrogram for a given time slice. Various networks were tested and a network with 3 layers of 64 neurons each was selected. Image [5] © 2020 IEEE.	58
Figure 5.4: The training and validation error shown over 20k iterations. The network is selected based on the iteration with the lowest overall validation error, after this point the validation error increases and the network starts to overfit the training data. Image [5] © 2020 IEEE.	58
Figure 5.5: The 40 GHz radar hardware suction mounted to the MSU CANVAS autonomous vehicle test platform. Approximately 20 minutes of driving data was captured at MSU’s Spartan Village. Image [5] © 2020 IEEE.	60

Figure 5.6:	The average measured spectrum over 472 samples where the test vehicle was traveling at 10 ± 0.05 m/s, along with the modeled response using 500 point scatterers. Notice that the model seems to capture the fringing effect that is present in the measurements, though the locations of the peaks are slightly misaligned. Image [5] © 2020 IEEE.	61
Figure 5.7:	The time-frequency response for the three individual interferometric data captures. The ground-truth velocity from the vehicle CAN bus is shown overlaid in blue. Notice the fringing effect predicted by the model is faintly observable in an instantaneous time slice of the response. Image [5] © 2020 IEEE.	62
Figure 5.8:	The estimates from the UKF tracker on the interferometric response shown in blue with the ground-truth velocity shown in cyan. Here we show performance over the same test data used with the neural net estimator. Image [5] © 2020 IEEE.	64
Figure 5.9:	The instantaneous interferometric spectral response shown in blue with the mean of the posterior state shown in magenta. Two sigma-point responses from the UKF are also shown in green and blue. Image [5] © 2020 IEEE.	65
Figure 5.10:	The output of the neural net estimator after smoothing with 3rd order local regression, the ground-truth velocity values are shown in cyan. Image [5] © 2020 IEEE.	66
Figure 5.11:	A single realization of a point-target with a constant-turn (CT) motion model with model noise of $\sigma_v = 0.1$ m/s and $\sigma_\Omega = \frac{\pi}{360}$ rad/s, and clutter intensity of $\lambda = 10$. The ground-truth position is shown in green along with the estimated position from the Gaussian sum filter (GSF). Image [6] © 2020 IEEE.	70
Figure 5.12:	The average RMSE of the state vector for each measurement model at various clutter rates. The value is taken as the average over 1k realizations, and a linear least-squares fit is shown to highlight the trend. Image [6] © 2020 IEEE.	71
Figure 5.13:	The ground-truth bounding box is shown in green along with its box velocity. The vehicle is slowing down and turning so in the local-frame the velocity is up and towards our vehicle which is shown with its baseline on the coordinate axes. The radar detections shown in blue are appended with the angular component of the ground-truth box velocity $v_{y,gt}$, plus Gaussian noise to emulate an angular velocity radar measurement. Image [6] © 2020 IEEE.	73

Figure 5.14: The tracked vehicle shown with green annotation and the magenta state estimates with covariance. The ground-truth position was chosen to be the center of the rear edge of the box to represent the rear bumper where most of the radar detections hit. The large magenta lines represent the birth model covariances where new objects are most likely to appear. Image [6] © 2020 IEEE.	74
Figure 5.15: The four hand gestures used in this work. The temporal sequence is shown top to bottom for each of the gestures. Notice that each gesture ends in the same position as it begins; this was done to simplify processing the data. The fifth gesture is the control in which no hand was present.	77
Figure 5.16: The radar antenna locations which form a grid separated by baseline $D_\lambda = 8$. The hand locations were measured at a height of $0.15 \pm 0.05\text{m}$ from the plane of the antennas, at angles $0 \pm 5^\circ$, $26 \pm 5^\circ$, $45 \pm 5^\circ$, and $56 \pm 5^\circ$	78
Figure 5.17: The time frequency responses for gestures one to four from top to bottom at the broadside location, the Doppler response is shown (left), and the interferometric response (right).	79
Figure 5.18: The topology of the CNN used in this work, each layer comprises a 2D convolution, rectified linear unit (ReLU), and max pooling. The outlined portion of the network remains the same for all of the cases being tested.	80
Figure 5.19: The confusion matrices for each of the four cases over all rotation and aspect angles, note that here the gestures are zero-indexed and thus gesture four is the control gesture.	81
Figure 5.20: The confusion matrices for what we consider the hardest case, a radar rotation angle of 45° and an aspect angle of 56° . Here we see the strength of the interferometric systems where case two perfectly classifies all gestures, and case four confuses only a single gesture.	82

CHAPTER 1

INTRODUCTION

Autonomous ground and aerial vehicles hold the potential to completely revolutionize how we move throughout this world. It is difficult to imagine a future in which seamless, completely autonomous transportation is not integrated into our everyday lives. Though immense strides have been made over the past decade, there is still much work to be done in this field. One of the most critical tasks is perception, that is, building a vision of the world that is digestible to advanced path-planning algorithms. This is absolutely imperative in order to make split second decisions in a safe and deterministic manner. Safety critical systems must have multiple redundant, individually engineered subsystems for measuring vehicle state, ideally with non-overlapping failure modes [7]. Sensors critical to this perception task include radar, lidar, and camera, all of which are able to provide a snapshot of the scene in which the autonomous vehicle is operating. All sensors are subject to noise, though luckily, many targets adhere to predictable motion models, which, when combined with measurements, provide a powerful method for tracking targets in the presence of noise [8]. Multitarget (MTT) tracking aims to solve the problem of estimating and tracking the states of multiple targets through time by observing measurements which contain noise [9]. MTT has a history spanning over 50 years, beginning in the aerospace industry, and finding applications in computer vision, remote sensing, robotics, and autonomous vehicles [10, 11, 12, 13, 14].

Measurements may consist of kinematic states, like position, velocity, acceleration, jerk, etc., or alternatively, measured attribute quantities, which describe some attribute of the target, such as signal-to-noise ratio (SNR), radar cross-section (RCS), or target shape parameters, among others. We tend to assume that noise is uncorrelated between measured quantities, and thus, each state that we measure increases information, and thus improves tracking performance. Probably the greatest strength of radar-based sensing is that the longer wavelength signals (compared to optical frequencies) can penetrate dense rain, fog, and

snow. Thus, the radar sensor is able to provide a picture of the world when complementary sensors, such as camera and lidar, are unable to provide accurate measurements. Most standard radar systems measure range, radial velocity (Doppler), and angle. Some systems may compute the angular velocity of the target by calculating the time rate change of angle measurements. The noise of this derived value is statistically dependent on the noise of the angle measurements, and thus no information is gained by including the derived value in a state space tracker. In this work I explore the use of a fourth type of basic radar measurement, interferometric angular velocity [15]. This method provides a measure of angular velocity that is independent of angle measurements and requires only a single pair of widely spaced antenna elements. In this chapter I provide an introduction to interferometric angular velocity estimation, background into the field of study, and a discussion of the current research problem. Next, we discuss our specific goals with this work, the overall significance to the broader radar community, and finally limitations associated with this study.

1.1 Background

Interferometry is used to extract useful information from the interference pattern created by a distributed aperture, and was first used in optics over 100 years ago to discover the wave nature of light [16]. Fizeau and Michelson first proposed using optical interferometry to measure the angular dimensions of astronomical objects [17], which lead to Ryle and Vonberg's stellar measurements using an interferometric radio telescope [18]. An example of this type of measurement is shown in Fig. 1.1. The two-element correlation interferometer produces a fringe pattern which creates lobes of ambiguous angle, but increases angular resolution compared to a single antenna element. As the Earth rotates, a voltage oscillation is measured due to the phase differences between the two antenna elements from the radiation observed from the distant stellar sources. Radio interferometry is still used today to synthesize large antenna apertures from a series of smaller antennas, such as the Karl G. Jansky Very Large Array (VLA) in New Mexico [18], which is used to obtain high resolution images of distant

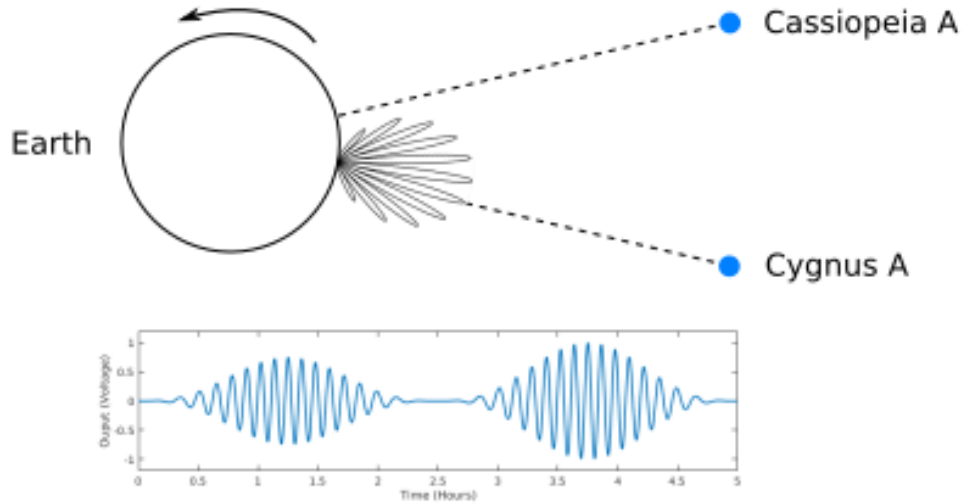


Figure 1.1: The fringe pattern of a two-element interferometric array has lobes of ambiguous absolute angle but increased angular resolution. The diagram shows the response of the two point sources as the interferometric fringe pattern scans the sky due to the rotation of the Earth. Image [1] © 2019 IET Press.

stellar sources. More recently, interferometric radar has found applications in radar imaging, where interferometric synthetic aperture radar (InSAR) has become a standard technique for generating topographic maps [19]. Also more recently, and the topic of this work, interferometric radar has been used to directly measure the angular velocity of moving targets [20, 21, 15]. If performed jointly with a traditional Doppler measurement then measurements of both radial and angular velocity may be obtained [22, 23, 24]. Similar to the micro-Doppler response, using interferometry, it is possible to generate time-frequency signatures which contain information relating to the angular motion of a target [25]. These methods have been used to detect walking people, rotating blades, unmanned aerial vehicles (UAV's), vehicle velocity, cooperative target tracking, and hand gesture recognition [22, 26, 27, 28, 3, 29, 30].

Interferometric angular velocity estimation works by correlating the received signals at two widely separated antenna elements from an angular moving target, which generates a signal whose frequency response is directly proportional to the angle rate of the target

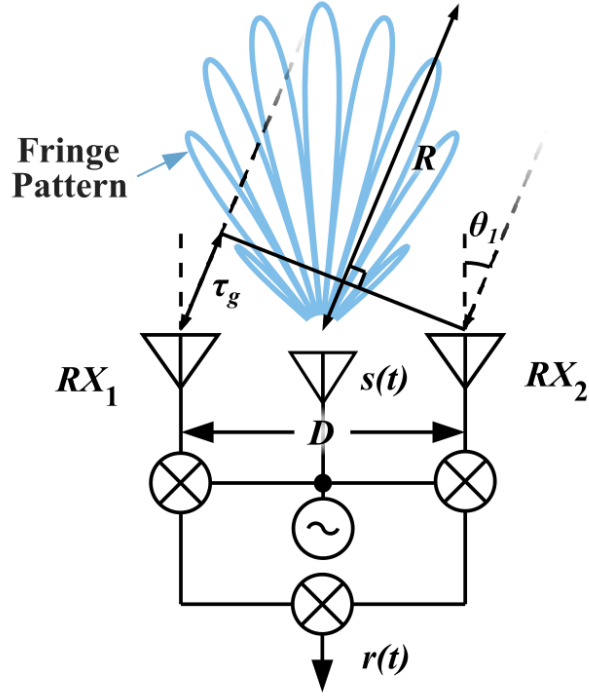


Figure 1.2: Block diagram of a two-element correlation interferometer with system baseline D , transmitted signal $s(t)$, and correlator output $r(t)$. The fringe pattern creates a series of peaks and nulls created by the constructive and destructive interference of the received signals at various angles. The geometric time-delay is shown as τ_g and assumes that the target is far-field to the array.

[15]. A diagram of a two-element correlation interferometer is shown in Fig. 1.2. Two antenna elements are separated by the system baseline D , and the received signals are downconverted and correlated to form the interferometric response, $r(t)$. In the image, the correlation is shown as being performed with a hardware mixer, though in practice, typically, the signals are individually sampled and correlated in software. The benefit of the latter method being that an individual antenna response may be used to obtain the traditional Doppler response. The fringe pattern is shown in blue, which represents the angular response of the correlation interferometer. The peaks and nulls in the fringe pattern appear at angles where the geometric time delay, and thus the phase difference, results in perfectly constructive or destructive interference between the signals received at each antenna element. As a target passes through the system pattern, which is a function of the individual

antenna patterns and the fringe pattern, the voltage output of the interferometer oscillates in conjunction with the peaks and nulls of the fringe pattern. An object which passes through, in angle, quickly results in an oscillation that is compressed in time, whereby a slow moving object results in an oscillation which appears stretched in time. Thus, it should be clear that the instantaneous frequency of this signal is directly related to the angular velocity of the target.

If a phase-quadrature, continuous-wave (CW) interferometer is viewing a point-target in the far-field to the array, the complex received signals at each antenna are

$$\begin{aligned} s_1(t) &= a_1 A_1(\theta) e^{j2\pi f t} + n_1(t) \\ s_2(t) &= a_2 A_2(\theta) e^{j2\pi f (t - \tau_g)} + n_2(t), \end{aligned} \quad (1.1)$$

where a_i is the complex signal amplitude, which is controlled by transmit power, receiver gain, range, RCS of the target, and multipath effects, A_i is the complex antenna voltage pattern, f is the system center frequency, τ_g is the geometrical time delay, and the n_i terms represent complex white Gaussian noise. The geometric time delay may be given in terms of the baseline D by

$$\tau_g = \frac{D}{c} \sin(\theta). \quad (1.2)$$

This expression assumes that the angles from the each antenna to the target are approximately equal which is not strictly true. In electromagnetics we typically use the Fraunhofer distance, $d_f = 2D^2/\lambda$, where D represents the maximum dimension of an antenna element, to represent the minimum range at which we can assume that the near-field spherical wavefronts may be represented as plane waves. Similarly, in a signal processing sense, we can define a distance which is ‘far-field’ to the array, in which we use the Fraunhofer distance, but now D is equal to the system baseline of the interferometer. If the range to the target is greater than this distance than the approximation for the geometrical time delay is valid. Many of these techniques still work in areas considered to be ‘near-field’ to the array, however there will be errors due to the exact expression for the geometric time delay.

The complex correlation interferometer performs conjugate multiplication, resulting in

$$\begin{aligned}
r(t) &= \langle s_1(t)s_2^*(t) \rangle \\
&= \langle (a_1A_1(\theta)e^{j2\pi ft} + n_1(t)) (a_2^*A_2^*(\theta)e^{-j2\pi f(t-\tau_g)} + n_2^*(t)) \rangle \\
&= \langle a_1A_1(\theta)e^{j2\pi ft}a_2^*A_2^*(\theta)e^{-j2\pi f(t-\tau_g)} \rangle \\
&\quad + \langle a_1A_1(\theta)e^{j2\pi ft}n_2^*(t) \rangle \\
&\quad + \langle n_1(t)a_2^*A_2^*(\theta)e^{-j2\pi f(t-\tau_g)} \rangle \\
&\quad + \langle n_1(t)n_2^*(t) \rangle, \tag{1.3}
\end{aligned}$$

where $\langle \cdot \rangle$ represents time-averaging which can be approximated with a low-pass filter. Since the noise is uncorrelated its time average will tend to zero. Assuming the receiver antenna patterns are identical the response is

$$r(t) = \langle a_1a_2^*|A(\theta)|^2e^{j2\pi f\tau_g} \rangle. \tag{1.4}$$

The instantaneous frequency is the time derivative of the phase term, thus,

$$f_i = \frac{1}{2\pi} \frac{d\Phi}{dt} = \frac{1}{2\pi} \frac{d2\pi f\tau_g}{dt} = \frac{d\frac{D}{\lambda} \sin(\theta)}{dt}, \tag{1.5}$$

where λ is the wavelength of the received signal. The angular velocity is given as the time rate change of angle as

$$\omega = \frac{d\theta}{dt} \implies \theta = \omega t + \theta_0 \tag{1.6}$$

where θ_0 represents the initial angle which is assumed to be zero without loss of generality.

Thus the instantaneous frequency of the complex correlation interferometer is

$$f_i = \frac{\omega D}{\lambda} \cos(\theta), \tag{1.7}$$

and we see that in the small angle regime, where $\cos(\theta) \approx 1$, the instantaneous frequency is directly proportional to the target angular velocity. Measurements conducted experimentally have also confirmed this relationship [15, 31].

There is a known issue which arises when multiple targets are present, that is, the cross-terms from the correlation processing appear as distortion in the interferometric signal response. We can easily observe this issue by examining the response of a system observing two targets. Ignoring noise, the received signals at each antenna due to both targets are

$$\begin{aligned} s_1(t) &= a_{11}A(\theta_1)e^{j2\pi f_1 t_1} + a_{12}A(\theta_2)e^{j2\pi f_2 t_2} \\ s_2(t) &= a_{21}A(\theta_1)e^{j2\pi f_1(t_1 - \tau_{g,1})} + a_{22}A(\theta_2)e^{j2\pi f_2(t_2 - \tau_{g,2})} \end{aligned} \quad (1.8)$$

where $f_n = f_c + f_{D,n}$, and $t_n = t + \tau_n$, which represent the independent Doppler shifts and time-delays. Thus, the correlator output is

$$\begin{aligned} r(t) &= \langle s_1(t)s_2^*(t) \rangle \\ &= \langle a_{11}a_{21}^* |A(\theta_1)|^2 e^{j2\pi f_1 \tau_{g,1}} \rangle + \langle a_{12}a_{22}^* |A(\theta_2)|^2 e^{j2\pi f_2 \tau_{g,2}} \rangle \\ &\quad + \langle a_{11}a_{22}^* A(\theta_1)A^*(\theta_2) e^{j2\pi(f_1 t_1 - f_2 t_2 + f_2 \tau_{g,2})} \rangle \\ &\quad + \langle a_{12}a_{21}^* A(\theta_1)A^*(\theta_2) e^{j2\pi(f_2 t_2 - f_1 t_1 + f_1 \tau_{g,1})} \rangle. \end{aligned} \quad (1.9)$$

Here we notice that the first two terms are similar to the single target response in (1.4), but for each individual target. The third and fourth terms in the response are the cross-terms, which depend both on the geometric time delay and the differential Doppler between the two targets. Thus, we are able to correlate target one with itself (from the other antenna), and target two with itself, but unfortunately we must also correlate target one with target two and vice versa. Shown in Fig. 1.3 is the interferometric time-frequency response of two targets with sinusoidally time-varying angular and radial velocities. The two lower frequency terms represent the ideal interferometric terms which are directly proportional to the targets' angular velocities, while the higher frequency oscillating terms are the intermodulation distortion. The greatest challenge associated with multitarget interferometric angular velocity estimation is the intermodulation distortion terms which arise when multiple target responses are present. Specifically, when N targets are present, the interferometric response will contain N^2 terms. N of these terms will be directly proportional to the angular velocity of each target, and $N^2 - N$ terms will be intermodulation distortion.




Figure 1.3: Simulated interferometric time-frequency response for two point scatterers. The two oscillating lower frequency terms are the ideal interferometric terms, while the higher frequency oscillations are the intermodulation distortion terms. Image [2] © 2021 IEEE.

1.2 Research Problem and Objectives

The issue of interferometric angular velocity estimation distortion mitigation was first discussed in [32]. In that work, hardware based methods were presented to mitigate the intermodulation distortion. One method was to use a pulsed system to temporally separate the responses from targets at different ranges. With this method, responses from like targets are only correlated with themselves and thus the mixing products of unlike targets are not generated. Later in this work we propose a signal processing based distortion mitigation technique that works in this same regard, though instead of separating the responses in the time domain, we work to separate them in the frequency domain. Another hardware method was to use long-wavelength signals, thus driving the differential Doppler terms to zero. With this method the distortion terms collapse onto the ideal interferometric terms. One problem with this technique is the need for a large physical baseline (in order to keep the electrical baseline of a given system constant), and also the inability to measure radial velocity (Doppler) with the system. In [33] the authors propose a dual frequency system to mitigate the latter issue. In this work the authors use a lower center frequency for angular velocity estimation, thus driving the distortion terms to zero, and a higher center frequency

for measuring the Doppler frequencies of the targets. Other hardware-based methods exist as well: in [34] the authors propose using the sum of the responses of multiple interferometric baselines, while in [35] the authors similarly use the product of multiple interferometric baselines. All of these methods help to mitigate the issue of intermodulation distortion, yet come at the cost of additional hardware resources, larger physical baselines, or large signal bandwidths.

The aim of this work is to provide methods for distortion mitigation which are purely signal processing based. In this manner, they operate either on the complete interferometric response, or on a pair of individual antenna responses themselves to perform angular velocity estimation of multiple targets. These approaches require no additional antenna elements or RF hardware. Among the questions this research seeks to answer are: how practical are either of these methods at jointly estimating the angular velocities of multiple targets, and are these methods scalable as the number of targets grows large? The major contributions of this work are as follows:

1. We present a complete signal model which describes the response of a correlation interferometer to an arbitrary target comprising N independent point targets. We first describe the exact response, then present simplified versions taking far-field and small angle approximations into account. This model is critical for designing and assessing signal processing based distortion mitigation techniques.
2. Two signal processing based distortion mitigation strategies are presented which we refer to as model-based methods, and response decomposition methods. Model-based methods use a learned or known signal model response and measured response to iteratively arrive at angular velocity estimates based on minimizing a loss function. When an analytic model is used it is framed as a non-linear optimization problem, which may be solved using common optimization techniques, such using gradient descent. We find that due to the N^2 nature of the frequency components in the interferometric

response this method will not scale as the number of targets grows large. The response decomposition methods operate on the individually received antenna signals to form angular velocity estimates. This method works by decomposing each antenna response into individual target responses, then performing data association between the target responses in the two antenna element signals. With this we are able to correlate only like targets and thus mitigate the cross-terms found in the traditional correlated response.

3. Experimental demonstration of the response decomposition distortion mitigation technique is presented with measurements conducted with a custom, single-board 40 GHz interferometric radar. Three specific multitarget cases are presented, which are designed to demonstrate easy, medium, and difficult cases for the proposed distortion mitigation algorithm. As expected, the accuracy of the angular velocity estimates decreases as we approach the difficult case.

1.3 Research Significance

This research contributes to the body of work surrounding the relatively recent method of interferometric angular velocity estimation. Specifically, it addresses the known issue of non-linear distortion which occurs in the interferometric signal response when multiple targets are present. By mitigating this distortion, and formalizing the process for jointly estimating the angular velocity of multiple targets, this work provides techniques to obtain an additional independent kinematic measurement. Since this measurement is independent of angle measurements, it may be used to improve state tracking performance. This measurement, along with the other three traditional radar measurements, form a complete set in \mathbb{R}^2 for both position and velocity, thus allowing measurements of both position and the full 2D velocity vector.

1.4 Scope of This Research

The scope of this research is focused on the use of signal processing techniques for distortion mitigation, and thus constraints have been applied in various aspects to maintain this focus. First, we provide results in the case where only two targets are present. This is the simplest case for interferometric distortion mitigation, which becomes exceedingly more difficult as more targets are present. Specifically, the number of distortion terms grows as N^2 , thus we will see that any methods which act on the interferometric response alone may be impractical as the number of targets grows large. The measurements performed in this work were obtained in a controlled environment, and thus may not be representative of real-world measurements, such as in the case of automotive radar in a dense urban environment. The trajectories of the targets were chosen to highlight cases which were both easy (targets easily separable in the Doppler domain), and difficult (targets moving at similar speeds with similar Doppler shifts) to extract and estimate angular velocity. All algorithms are performed in post-processing and thus little effort has been made to optimize run-time speed for real-time processing. However, these limitations do not impact the focus of this work, which is to demonstrate the efficacy of signal processing based distortion mitigation techniques, and thus, nonetheless provides a significant contribution to the field.

1.5 Dissertation Organization

In this chapter we have introduced the background of our study, described the research problem, our objectives, the significance to the overall radar community, and described the scope of work. In chapter 2 we derive the exact interferometric response to an arbitrary set of targets containing N independent point scatterers. In chapter 3 we present two frameworks for signal processing based interferometric distortion mitigation. In chapter 4 we present the radar hardware used in this work, along with multitarget measurements and angular velocity estimation results for the two target case. In chapter 5 we describe practical implementations of interferometric angular velocity estimation, and finally, in chapter 6 we conclude with a

discussion of contributions and topics for future work.

CHAPTER 2

GENERAL INTERFEROMETRIC RESPONSE TO AN N-POINT TARGET

In this chapter we derive the response of a complex correlation interferometer to an arbitrary target containing N independent point scatterers [2]. We first derive the response for arbitrary transmitter and receiver locations. Next, we make the approximation that the transmitter is co-located with one of the receiving elements, which helps to simplify the response substantially. This may be the case if a single antenna is used for both transmit and receive, or if the transmit and receive antennas are placed sufficiently close together. In practice, often times the transmitter is placed in the middle of the two receiving elements, and we have found that this approximation works well even in that case. Finally, we make the far-field and small-angle approximations to arrive at the simplest form of the interferometric response, and describe the regions in which this approximation holds. We end this chapter with simulations demonstrating the computed and model responses for arbitrary target trajectories.

2.1 Reference Geometry

The reference geometry for the parameters used in this work are shown in Fig. 2.1 for a single target n . Here the system baseline is shown exaggerated to better visualize the vector projections, but in practice this value is typically on the order of $10 - 30\lambda$. The baseline must be large enough to ensure multiple cycles of the phase at the output of the correlation interferometer and, in general, increasing the baseline improves target resolution while also increasing the physical size of the system [36]. The location of antenna one (RX1 = index 1) is set arbitrarily at the origin and we define all radial and tangential velocities with respect to the origin. The radial and tangential velocities of the target are shown in the figure as the vectors $\mathbf{v}_{r,1,n}$ (blue) and $\mathbf{v}_{t,1,n}$ (orange), respectively. The first number in the subscript denotes the antenna element and the second value denotes the point scatterer we

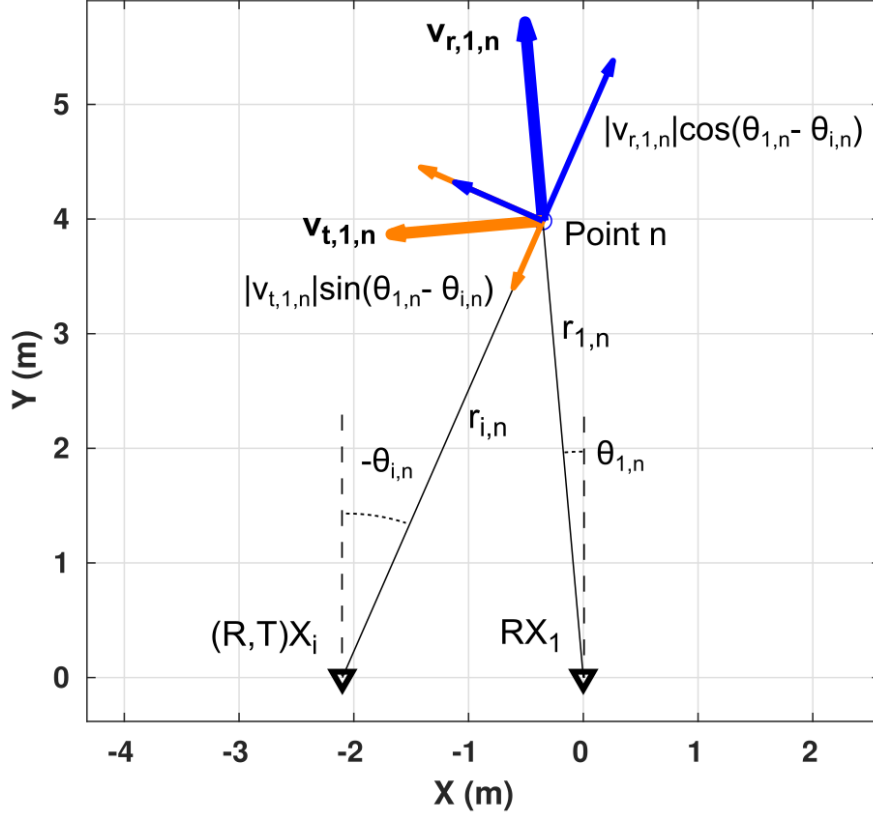


Figure 2.1: The notation for the parameters used in this work. RX1 is placed at the origin, and point n contains both radial (blue) and tangential (orange) velocity components relative to the origin. In the case that TX1 is not co-located with RX1 we must project the velocities as shown, considering the second antenna to be the transmitter. Similarly, we must always project the velocities at RX2. In both cases the projected radial velocity will have two components, one from the original radial and one from the original tangential velocities. Image [2] © 2021 IEEE.

are referencing. Here we show a second antenna which could either be the transmitter (TX = index 0) (if it is not co-located with RX1 as in one of the approximations), or the second antenna (RX2 = index 2). In practice the transmitter will be located somewhere in between RX1 and RX2 to ensure the radiated signal has equal coverage in the RX1 and RX2 antenna overlap regions. We should note that it is critical that the target is observed in this overlap region as shown in Fig. 2.2. Both antennas must have sufficient gain in order to receive the reflected signal, which may be of concern when using high-gain (narrow beam) antennas and large system baselines. We have found that approximating TX co-located with RX1 is often sufficient and greatly reduces the complexity of the model response. The geometric

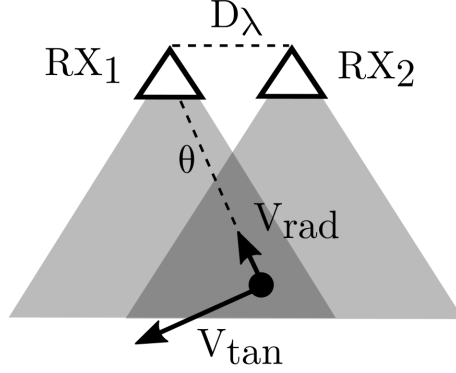


Figure 2.2: The target must be located in the 3 dB antenna overlap region of RX1 and RX2. This is usually not an issue for wide beam antennas, though can be a difficult requirement when using high gain (narrow beam) antennas with a large electrical baseline, shown here as $D_\lambda = \frac{D}{\lambda}$.

distances and angles are labeled as $r_{i,n}$ and $\theta_{i,n}$ respectively. It is important to note that the radial velocity at the second antenna has two velocity components. One component of the point n radial velocity with respect to RX1 (shown in blue as $|v_{r,1,n}| \cos(\theta_{1,n} - \theta_{i,n})$) and one component due to the tangential velocity of point n with respect to RX1 (shown in orange as $|v_{t,1,n}| \sin(\theta_{1,n} - \theta_{i,n})$).

2.2 Interferometric N-Point Target Model

In each model we represent a target or scene as a summation of individual, independent point scatterers. The received voltage at antenna i , due to the target, is

$$s_i(t) = \sum_{n=1}^N A(\theta_{i,n}) e^{j2\pi f_c(t - \tau_{i,n}(t))}, \quad (2.1)$$

where N is the number of scatterers, $A(\theta_{i,n})$ is the complex antenna voltage pattern, $\theta_{i,n}$ is the angle from broadside to point n at antenna i , and $\tau_{i,n}$ is the time-delay due to the distance the wave propagates from the transmitter to point n and back to antenna i . The time delay for a target n at antenna i may be defined as

$$\tau_{i,n}(t) = \frac{r_{0,n}(t) + r_{i,n}(t)}{c}, \quad (2.2)$$

where $r_{0,n}$ is the distance, which is time-varying, from the transmitter to point n , and $r_{i,n}$ is the time-varying distance from receiver i to point n .

2.2.1 Exact Model: Arbitrary Transmitter Location

A complex correlation interferometer performs conjugate multiplication, this results in

$$\begin{aligned}
r(t) &= \langle (s_1(t) + n_1(t))(s_2^*(t) + n_2^*(t)) \rangle \\
&= \langle s_1(t)s_2^*(t) \rangle + \langle s_1(t)n_2^*(t) \rangle \\
&\quad + \langle s_2(t)n_1(t) \rangle + \langle n_1(t)n_2^*(t) \rangle
\end{aligned} \tag{2.3}$$

where $n_i(t)$ are complex Gaussian noise terms. Since the noise is uncorrelated, their time average will tend to zero, resulting in

$$\begin{aligned}
r(t) &= \langle s_1(t)s_2^*(t) \rangle \\
&= \left\langle \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n}) e^{j2\pi f_c(t-\tau_{1,n}(t))} A^*(\theta_{2,k}) e^{-j2\pi f_c(t-\tau_{2,k}(t))} \right\rangle \\
&= \left\langle \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n}) A^*(\theta_{2,k}) e^{-j\frac{2\pi}{\lambda_c}(r_{0,n}(t)+r_{1,n}(t)-r_{0,k}(t)-r_{2,k}(t))} \right\rangle.
\end{aligned} \tag{2.4}$$

The instantaneous frequency is the time derivative of the phase term, thus, for a single component $\{n, k\}$, its instantaneous frequency is

$$f_{n,k} = \frac{v_{r,0,n}}{\lambda_c} + \frac{v_{r,1,n}}{\lambda_c} - \frac{v_{r,0,k}}{\lambda_c} - \frac{v_{r,2,k}}{\lambda_c}, \tag{2.5}$$

where we have defined the radial velocity for an approaching target n with respect to antenna i as the time derivative of the position as $\frac{\partial r_{i,n}(t)}{\partial t} = -v_{r,i,n}$. The radial velocities at the transmitter ($i = 0$) and antenna two ($i = 2$) due to point n are found by projecting the radial and angular velocities referenced from antenna one as

$$v_{r,i,n} = v_{r,1,n} \cos(\theta_{1,n} - \theta_{i,n}) - \omega_{1,n} r_{1,n} \sin(\theta_{1,n} - \theta_{i,n}) \tag{2.6}$$

where the relationship between tangential and angular velocity is $v_{t,1,n} = \omega_{1,n}r_{1,n}$. Thus the instantaneous frequency components become

$$\begin{aligned}
f_{n,k} = & \frac{v_{r,1,n}}{\lambda_c} \\
& + \frac{v_{r,1,n} \cos(\theta_{1,n} - \theta_{0,n}) - \omega_{1,n}r_{1,n} \sin(\theta_{1,n} - \theta_{0,n})}{\lambda_c} \\
& - \frac{v_{r,1,k} \cos(\theta_{1,k} - \theta_{0,k}) - \omega_{1,k}r_{1,k} \sin(\theta_{1,k} - \theta_{0,k})}{\lambda_c} \\
& - \frac{v_{r,1,k} \cos(\theta_{1,k} - \theta_{2,k}) - \omega_{1,k}r_{1,k} \sin(\theta_{1,k} - \theta_{2,k})}{\lambda_c}.
\end{aligned} \tag{2.7}$$

When we reference the velocities in (2.5) to antenna one, each component breaks out into two components, one due to the radial velocity, and the other due to the angular velocity of that point with respect to antenna one. The exception is $v_{r,1,n}$ since it is already referenced to antenna one and thus contains the full radial motion. The spectral response is the superposition of the individual frequency components, resulting in

$$R(f) = \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n})A^*(\theta_{2,k})\delta(f - f_{n,k}), \tag{2.8}$$

where δ is the Dirac delta function.

2.2.2 Exact Model: Transmitter Co-located with Receiver 1

The expressions may be simplified in the case where the transmitter and antenna one (RX1) are co-located. This could be the case if a single antenna is used for both transmitting and receiving, or more commonly, if the two antennas are close together, as was the case in our experimental results shown later. With this assumption, the frequency components become

$$f_{n,k}^{(1)} = \frac{2v_{r,1,n}}{\lambda_c} - \frac{v_{r,1,k}}{\lambda_c} - \frac{v_{r,2,k}}{\lambda_c}, \tag{2.9}$$

where we use the superscript (1) to differentiate this approximated response from the exact model with arbitrary TX location. Thus, the values with respect to antenna one are

$$f_{n,k}^{(1)} = \frac{2v_{r,1,n}}{\lambda_c} - \frac{v_{r,1,k}}{\lambda_c} - \frac{v_{r,1,k} \cos(\theta_{1,k} - \theta_{2,k})}{\lambda_c} + \frac{\omega_{1,k}r_{1,k} \sin(\theta_{1,k} - \theta_{2,k})}{\lambda_c}. \tag{2.10}$$

2.2.3 Approximate Model: Far-field and Small-angle Approximations

The exact models, shown above, require the measurement of the range and angle of each scatterer. However, under some additional simplifying assumptions, we can reduce these expressions further. Under the small-angle approximation, ($\sin(\theta) \approx \theta$ and $\cos(\theta) \approx 1$), then (2.10) becomes

$$f_{n,k}^{(2)} = \frac{2v_{r,1,n}}{\lambda_c} - \frac{2v_{r,1,k}}{\lambda_c} + \frac{\omega_{1,k}r_{1,k}(\theta_{1,k} - \theta_{2,k})}{\lambda_c} \quad (2.11)$$

where we use the superscript (2) to signify that this representation makes further approximations. From the geometry of the problem, the baseline D may be defined as

$$D = r_{1,k} \sin(\theta_{1,k}) - r_{2,k} \sin(\theta_{2,k}). \quad (2.12)$$

If the scatterers are sufficiently far from the radar such that $r_{1,k} \approx r_{2,k}$, we can re-write eq. (2.11) as

$$f_{n,k}^{(2)} = \frac{2v_{r,1,n}}{\lambda_c} - \frac{2v_{r,1,k}}{\lambda_c} + \frac{\omega_{1,k}D}{\lambda_c}. \quad (2.13)$$

As we can see from this response, each frequency component contains three terms, two which are defined by the differential Doppler shift between the two points n and k , and one which is directly proportional to the angular velocity of point k . Ideally, the response would only contain the self-terms ($n = k$ in the double summation of (2.8)). When $n = k$ the differential Doppler terms will cancel and we arrive the expected small-angle interferometric response in (1.7) for that specific target. This would allow us to directly measure the angular velocity of each target. Unfortunately, when $n \neq k$, the cross-terms still contain information related to the angular velocity, but are also corrupted by an amount which is proportional to the difference in the corresponding radial velocities and thus Doppler shifts between the two targets.

It is important to be aware of the regimes in which we are able to use this approximate model because it is much simpler and requires no specific range and angle measurement

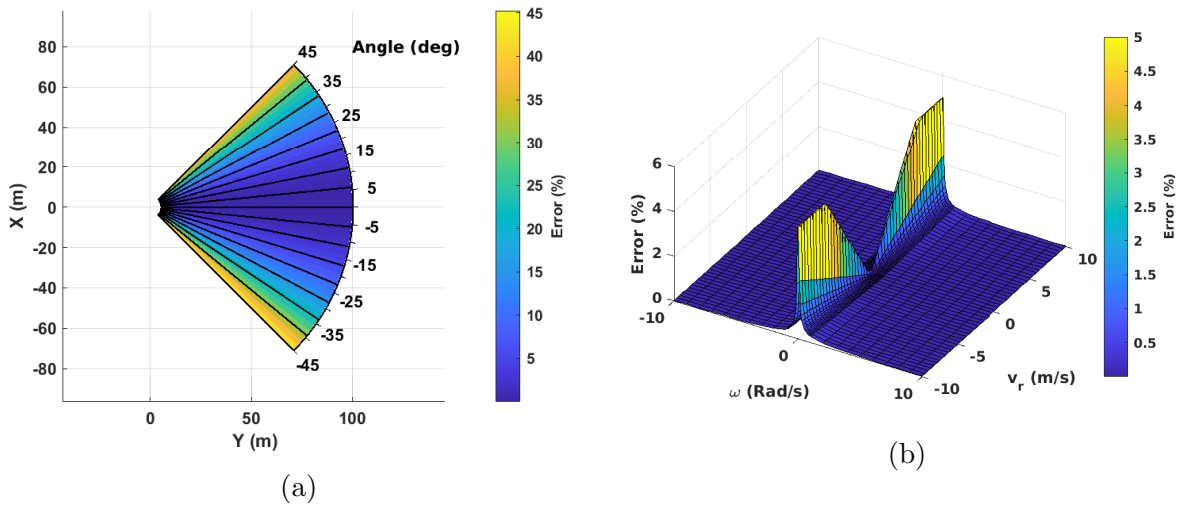



Figure 2.3: The percent error between the frequency estimates for target position (a) with fixed radial velocity of 1.0 m/s and angular velocity of 1.0 rad/s. The frequency percent error for radial and angular velocities (b) with target position fixed at a range of 20.0 m and angle of 0° . The error has been limited to 5% around the origin of plot (b) where the error increases rapidly as ω approaches zero. Image [2] © 2021 IEEE.


of the targets. This region will depend upon the acceptable error between the approximate and actual response. Shown in Fig. 2.3 (a) is the percent error of the frequency components between the two models for varying target position, with the angular and radial velocities fixed at 1.0 rad/s and 1.0 m/s, respectively. We see that at angles close to broadside where the small-angle approximation holds, the approximate model frequency values very closely resemble the actual model. Shown in Fig. 2.3 (b) is the frequency component percent error for various angular and radial velocity values, with the range and angle fixed at 20 m and 0° . These values have little effect on the accuracy of the approximate model, except in the extreme case of a target containing purely radial motion.

2.3 Simulated Model Response

To verify the accuracy of the model, a simulation was developed using MATLAB based on a similar example in [37]. The radar parameters were chosen to match the experimental setup (described in a later section), with a center frequency of 40 GHz and a system baseline of 20λ . The number of targets was parameterized, and for each scattering center, random



(a)




(b)


Figure 2.4: Simulated interferometric time-frequency response for a single point scatterer (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.

sinusoidal angular and radial velocity trajectories were realized. Each target also was given a random starting range and angle. RX1 and the transmitter were placed at the origin, while RX2 was placed at $[-D, 0]$. The phase at each antenna was calculated as in (3.6) at each time instant, while the amplitude was set equal to one in order to visualize the ideal response without antenna affects. A total of 5 s was simulated with a sampling frequency of 1 kHz. The interferometric response was calculated by conjugate multiplication of the complex time-domain signals, and the time-frequency response was computed via the short-time Fourier transform (STFT) [38] with window size and bin size of 128 samples.

The interferometric time-frequency response for a single point scatterer is shown in Fig. 2.4 (a). The parameters were selected randomly for the target and shown in Table 2.1. At each time instant in the simulation we compute the frequency components from



(a)



(b)

Figure 2.5: Simulated interferometric time-frequency response for two point scatterers (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.

Table 2.1: Simulation 1 Target Parameters

Ang. Vel. (Rad/s):	Rad.Vel. (m/s):	Range (m):	Ang. (Deg.)
$0.60 \sin(-1.40t)$	$0.10 \sin(-2.0t)$	10.03	27.66

the exact model response given in (2.10), and the approximate response as given in (2.13). The calculated frequency components for the single point scatterer are shown in Fig. 2.4 (b) in blue for the exact response, and in orange for the approximate response. In simulation the agreement between the exact response and the simulated response is very high due to availability of perfect knowledge of the target parameters. For the approximate response, the largest source of error comes from the angle of the target as shown in Fig. 2.3 (a). In this realization the minimum and maximum angle of the target was -21.5° and 27.7° , which

Table 2.2: Simulation 2 Target Parameters

Ang. Vel. (Rad/s):	Rad.Vel. (m/s):	Range (m):	Ang. (Deg.)
$0.93 \sin(-1.23t)$	$0.45 \sin(2.27t)$	8.45	-13.1
$0.46 \sin(-0.33t)$	$0.14 \sin(-0.05t)$	10.44	0.7


Table 2.3: Simulation 3 Target Parameters

Ang. Vel. (Rad/s):	Rad.Vel. (m/s):	Range (m):	Ang. (Deg.)
$-0.46 \sin(0.21t)$	$-0.02 \sin(2.12t)$	10.05	3.2
$0.56 \sin(-2.20t)$	$0.50 \sin(1.31t)$	9.76	23.8
$-0.27 \sin(1.25t)$	$-0.38 \sin(-0.70t)$	9.94	-0.6
$0.14 \sin(2.02t)$	$0.30 \sin(-1.02t)$	11.27	-1.9


are still within an acceptable region of the small angle approximation.

Shown in Fig. 2.5 is the interferometric time-frequency and model responses for a scene containing two targets with parameters as shown in Table 2.2. Due to the double summation in (2.8) there will be a total of N^2 frequency components, thus we have four components. The two lower frequency components pertain to the ideal interferometric shifts while the two higher frequency terms are the intermodulation terms. Notice after 1.5 s we see some disagreement between the approximate and exact responses. The minimum and maximum angle subtended by either target was -99.9° and 0.7° , violating the small-angle approximation, and thus accounting for the larger error in the approximate response. Nonetheless, the results are still reasonably accurate, since the trajectory of the targets sometimes places them within the small-angle approximation window and sometimes outside. Shown in Fig. 2.6 is the interferometric time-frequency and model responses for a scene containing four targets with parameters as shown in Table 2.3, resulting in 16 frequency components. In this instance the minimum and maximum angle by any target was -61.0° and 23.8° , resulting in relatively close agreement between the exact and approximate model responses.

In this chapter we provided the complete interferometric response to an arbitrary



(a)



(b)

Figure 2.6: Simulated interferometric time-frequency response for four point scatterers (a). The model response is shown (b) in blue for the exact response, and in dashed orange for the approximate response. Image [2] © 2021 IEEE.

target containing N independent point scatterers. Prior to this work, the response had been derived for a scene containing two targets in the far-field [32], but this is the first time the complete response has been derived without any approximations. Now that we have the exact signal response, we are able to know where the distortion terms will exist spectrally, and how to best mitigate them to recover angular velocity information. Armed with this knowledge we are ready to discuss methods for signal processing based distortion mitigation.

CHAPTER 3

SIGNAL PROCESSING BASED DISTORTION MITIGATION FOR MULTITARGET ANGULAR VELOCITY ESTIMATION

Now that we have described the complete interferometric model response, we can analyze different distortion mitigation approaches. The material in this chapter was first described in our work here [39]. We can see from the model that because of the double summation in (2.8), the interferometric response of a scene containing N scatterers will contain N^2 frequency components. N of these frequency components will be directly proportional to the target angular velocity, which occurs when a target response is correlated with itself from each of the antenna responses. There will also be $N^2 - N$ distortion terms, which occur when each target is correlated with every other target across the antenna responses. The purpose of any distortion mitigation method is to recover the angular velocity information of each target which is contained in the self-terms. The difficulty with this problem is that the closed-form expression is not one-to-one. This means that two scenes with different targets having different angular and radial velocities can generate the same spectral response. Thus, given an interferometric spectrum we can not directly solve for a single solution. To help solve this problem, we have identified two main classes of techniques which have guided our approach. The first uses the model response, described above, with common non-linear optimization techniques to estimate the underlying angular velocity parameters. These approaches we refer to as model-based methods. The second technique involves methods aimed at not generating the $N^2 - N$ distortion terms at all, which we call response decomposition methods.

3.1 Model-Based Methods

3.1.1 Learned Representations

Even without the interferometric model in (2.8), we could still use model-based methods by learning the model from data. This may be achieved using common machine learning

methods such as feedforward or recurrent neural networks [40]. The most difficult problem with this approach is that which is inherent for all learning-based approaches- a large amount of training data is required, especially for use with deep neural networks in which the amount of data required for training is proportional to the size of the network. This is easy to do in simulation because we have access to the individual target responses. We can perform the pairwise correlations only between like targets and arrive at the ideal interferometric response which contains no distortion. This is the ideal response which may be used as the target for the neural networks. The actual inputs to the network may be the amplitude and phase values of the complete interferometric response with distortion, and even the individual Doppler responses if so desired. For learning using gradient based methods the spectrum should not be represented with delta functions, which will have appreciable loss function gradients only at discrete intervals. The real challenge with this approach is that, with physical measurements, we never have access to the ideal interferometric response. To train on physical measurements the distortion would either need to be hand filtered or filtered using some other distortion mitigation method. We can address this issue by using the analytic model which was derived in the previous chapter.

3.1.2 Analytic Model

The availability of a closed-form representation, as given in (2.8), simplifies the problem of distortion mitigation substantially. The response is simply a parameterized non-linear function for which we would like to solve for the unknown parameter values, which in our case is the angular velocity values. We can accomplish this by finding parameter values which minimize some loss function, or the error between the observed response and the calculated model response with the estimated parameters. A brute-force approach would be to try all possible combinations of angular velocity values to find some model response which matches the measurement within some error tolerance. In practice, more efficient optimization techniques exist, such as gradient descent, which is guaranteed to approach the

true parameter values when the loss space is convex.

For gradient-based methods we represent our signal in a form that is differentiable, here we choose to represent the spectrum of the interferometric response as a weighted Gaussian sum of real signals. The spectral response is then

$$R(f) = \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n}) A(\theta_{2,k}) e^{\frac{-(f-f_{n,k})^2}{2\sigma^2}}, \quad (3.1)$$

where σ is the standard deviation which defines the width of the frequency component. Each component is centered around $f_{n,k}$ as given in (2.13). We define the loss function which measures the similarity between the measured, or simulated spectrum, and the spectrum we generate using our angular velocity estimates. We define this loss function as the sum of the squared differences as

$$L = \frac{1}{2} \sum_i^{N_f} (M(f_i) - R(f_i))^2, \quad (3.2)$$

where N_f is the number of FFT bins in the measured spectrum, and M is the measured (or simulated) spectrum. The derivative of the loss function with respect to a single target angular velocity component ω_β is

$$\frac{\partial L}{\partial \omega_\beta} = \sum_{i=1}^{N_f} (M(f_i) - R(f_i)) \frac{-\partial R(f_i)}{\partial \omega_\beta}. \quad (3.3)$$

The derivative of the model response with respect to a single target angular velocity parameter is

$$\frac{\partial R(f)}{\partial \omega_\beta} = \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n}) A(\theta_{2,k}) e^{\frac{-(f-f_{n,k})^2}{2\sigma^2}} \left(\frac{f_{n,k} - f}{\sigma^2} \right) \left(\frac{-1_{k=\beta} D}{\lambda_c} \right), \quad (3.4)$$

where $1_{k=\beta}$ is the indicator function. This function is equal to one when $k = \beta$, and zero when $k \neq \beta$. In the case where the loss space is convex, we will approach the true value of the parameter if we take small steps in the opposite direction of the gradient as

$$\hat{\omega}^{(x+1)} = \hat{\omega}^{(x)} - \alpha \frac{dL}{d\omega}, \quad (3.5)$$

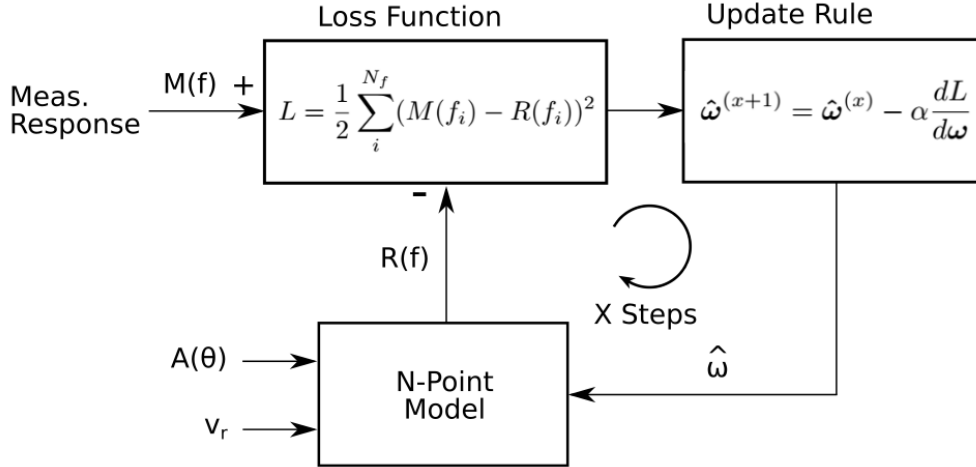


Figure 3.1: The gradient algorithm uses the measured and model response as inputs to the loss function. The gradient of the loss function is calculated and the update rule is used to calculate the new angular velocity estimates. This happens X times, or until the total error is below some threshold.

where $\hat{\omega}$ is a vector containing the estimated target angular velocities, n is the optimization index number, and α is the learning rate that controls the size of the step. If the loss space is non-convex, as is typically the case, there is a likelihood of arriving at estimates in local minima. The steps of the algorithm are shown in Fig. 5.13. Here we start with the measured response and the model response as inputs to the loss function, then the gradient is calculated and used in the update rule to update the parameter estimates. The new estimates are used to generate the new model response and the process is repeated X times or until the final error between the measured and model response is below some pre-defined error threshold.

To demonstrate this we calculate the loss space for a theoretical interferometric radar detecting two targets, with $\lambda_c = 0.007$ m, and $D = 30\lambda_c$. The true parameters of the targets are $v_r = [0.479; 0.213]$ m/s, $\omega = [0.001; -0.189]$ rad/s, $\theta = [-0.461; 0.191]$ rad, and unit amplitude. The error between the model generated using the known parameters and the model using angular velocity estimates ranging from -2.0 to 2.0 rad/s is shown in Fig. 3.2. As mentioned above, the loss space is non-convex, thus if we start with random initial estimates we have no guarantee of reaching the true parameter values. To mitigate this, we perform a rough grid-search to arrive at the best guess of the initial estimates. Our

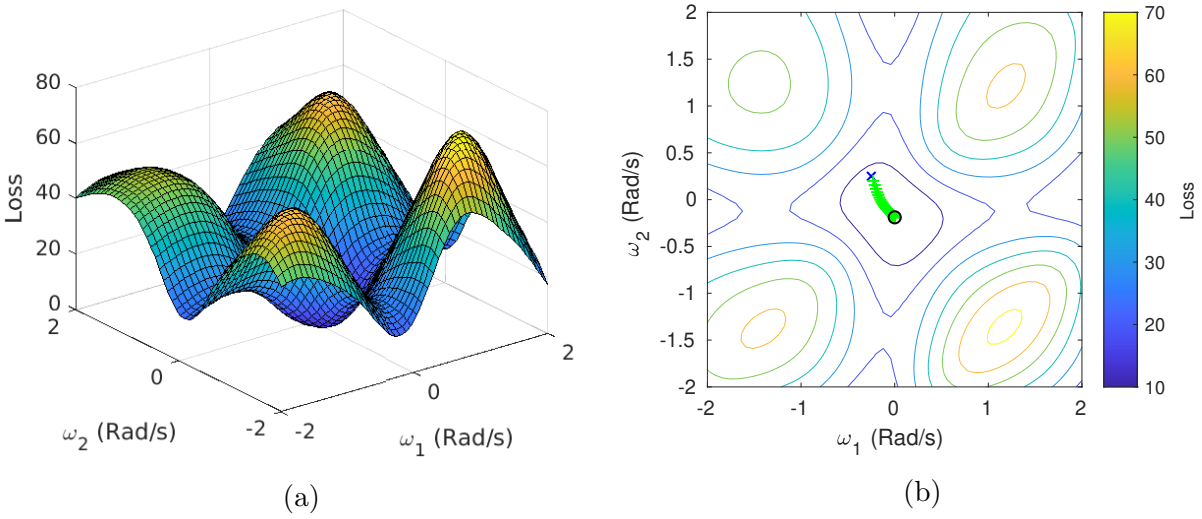


Figure 3.2: The loss surface (a) for two estimates of angular velocity of an interferometric radar with true values of $\omega = [0.001; -0.189]$ rad/s. Notice that the entire loss space is non-convex and so the gradient method will only approach the true values when the initial estimate is close to the true parameter value. The contour plot (b) of the loss space with the initial estimate shown in as a blue 'x'. We run 50 iterations with each iterative estimate shown as a green '+'. The estimates approach the true value shown as a black 'o'.

initial estimates are chosen to be $\hat{\omega} = [-0.25; 0.25]$ rad/s and we implement the gradient descent algorithm as shown in Fig. 5.13. The initial estimates are shown as a blue 'x' in (b), with each iterative estimate shown as a green '+'. After 50 iterations the estimates approach the true value shown as a black 'o'. The final estimates are $\hat{\omega} = [-0.001; -0.189]$ rad/s. The average error between the true and estimated parameters is 0.001 rad/s.

The difficulty with this method is that, in practice, the loss space is non-convex. We can mitigate this by first performing a coarse grid search of parameter values, within some bounds, to find initial parameter seed values that gets the estimates onto a convex part of the loss space that contains the global minima. Another difficulty with this approach, and any method that operates on the complete interferometric response, is that the number of frequency components of the response grows as N^2 , thus as the number of targets grows the number of components grows quadratically. This makes for a highly non-convex loss space with many local minima. This challenge has motivated the next class of distortion mitigation techniques which aim to not generate the $N^2 - N$ distortion terms at all.

3.2 Response Decomposition Methods

In the case of a single target there is no need for distortion mitigation because the target is only correlated with itself between antenna elements and thus there is only a single frequency component as described by (1.7). Earlier we mentioned a hardware-based distortion mitigation method in which a pulsed system was used to temporally separate the responses from individual targets [32]. This could be considered a hardware based version of signal response decomposition. In a similar manner, with the method described here, our goal is to isolate the individual responses from each target in each antenna response, but to do so in signal processing in the frequency domain.

The method of distortion mitigation using response decomposition ensures that the number of targets N in (2.8) will be equal to one for each correlation operation, and thus represents the self-term case when $n = k$. Thus, the response at antenna i due to a scene with N independent point targets is

$$s_i(t) = \sum_{n=1}^N s_{i,n}(t), \quad (3.6)$$

and the individual response due to a single target n is

$$s_{i,n}(t) = A(\theta_{i,n})e^{j2\pi f_c(t-\tau_{i,n}(t))}, \quad (3.7)$$

where f_c is the transmitted center frequency, and $\tau_{i,n}$ is the time-delay due to the geometric distance from the transmitter to point n and back to antenna i . A complex correlation interferometer performs the conjugate multiplication as

$$r(t) = \langle s_1(t)s_2^*(t) \rangle. \quad (3.8)$$

This is the general interferometric response, and will result in the spectral output given in (2.8). The response decomposition method seeks to decompose the measured signals at antenna one, $s_1(t)$, and antenna two, $s_2(t)$, into their individual N terms as described in (3.6). It assumes that after decomposing each signal response we are also able to associate

the terms with their correct counterpart in the other antenna response. For example, if each antenna response contains two terms we must know which terms were generated by the same target, and match those two. Thus, using this method, the output of the correlation interferometer using is

$$r_D(t) = \sum_{n=1}^N \langle s_{1,n}(t) s_{2,n}^*(t) \rangle. \quad (3.9)$$

We see now that the total response is a summation of the correlations of the responses at each antenna only due to the individual points n . Thus the total spectral response is

$$R_D(f) = \sum_{n=1}^N |A(\theta_{1,n})|^2 \delta \left(f - \frac{\omega_{1,n} D}{\lambda_c} \right). \quad (3.10)$$

From this we can see that the response contains frequency shifts which are only directly proportional to each target's angular velocity. In practice, we measure each antenna signal $R_1(f)$ and $R_2(f)$, which are the spectral responses of the received signals at antenna one and two, though, the response required to measure angular velocity is $R_D(f)$. Thus, we decompose each antenna's spectral response into components due to each target as

$$S_i(f) = \sum_{n=1}^N S_{i,n}(f), \quad (3.11)$$

where i is the antenna element, and N is the number of targets. Conjugate multiplication is equivalent to cross-correlation in the frequency domain, thus we generate the ideal response for a specific target n , by cross-correlating its spectral response from each antenna element as

$$R_{D,n}(f) = S_{1,n}(f) \star S_{2,n}(f). \quad (3.12)$$

The total spectral response using response decomposition is then the sum of the individual ideal responses as

$$R_D(f) = \sum_{n=1}^N R_{D,n}(f). \quad (3.13)$$

The ability to generate this response depends on how well we are able to isolate individual target responses in the each antenna response, and then associate them between the two responses.

3.2.1 Signal Decomposition

For ideal point scatterers, detection of the response of an individual antenna signal is equivalent to sinusoidal parameter estimation in the presence of noise. An approximation to the maximum likelihood estimator (MLE), for the parameter of frequency, is the maximum of the periodogram [41]. Thus, a simple method for signal decomposition may consist of detecting peaks above a threshold in each antenna spectral response. If the number of targets is known beforehand, super-resolution techniques may be used, such as MUSIC or ESPRIT to obtain higher resolution frequency estimates [42, 43]. In general, the number of targets may be unknown and time-varying, and thus some method must be used to estimate the number of targets present in each antenna signal response. Due to noise and limited frequency resolution it is possible that the number of targets is not equivalent between the antenna responses, and so this must be handled appropriately. In some cases, as in the measurements shown later, a target response is not sufficiently modeled as an ideal sinusoid with a single peak in the frequency domain. Instead of peak detection, we use a metric based on maximum integrated power to find target window locations in order to detect the target locations in the frequency domain. Regardless of which method is used, it is critical to detect, locate, then isolate the portion of the frequency spectrum that is attributable to a each individual target response. With this, we assume that the each antenna response has been separated into individual target responses either through masking, bandpass filtering, or some other signal decomposition method.

3.2.2 Data Association

The data association problem is well described in the multitarget tracking (MTT) literature [11]. Typically the likelihood is computed that two responses are associated with each other based on some a priori knowledge about the motion of the target. For example, if we are tracking range and have an expected motion model, we compute the most likely location at the next time step based on the target's velocity and current position. We can

define a Gaussian probability density function (PDF) centered at the expected location and would evaluate that PDF at the value of the new measurement to find the probability that the new measurement belongs to the target. The process is quite similar for associating detections between the responses at each antenna element of a correlation interferometer. In the case of response decomposition distortion mitigation, the signal response between the two antenna elements will be shifted, not by range as in the MTT example, but in frequency. The difference in frequency will be proportional to the target's angular velocity and the interferometric baseline. There are two common PDF's to use to represent the distribution of this frequency shift. One assumes that the angular velocity of a target can fall anywhere between some minimum and maximum values with equal likelihood, in this case it would be a uniform distribution. We may also assume, specifically in the case that the targets are far from the radar system, that the angular velocity is most likely small or close to zero. In this example we would most likely use a Gaussian PDF where the maximum expected angular velocity value would be set to three standard deviations or similar. The simplest of all data association methods in the interferometric case would be similar to global nearest neighbor (GNN) which is used in MTT. With this method we associate the assignments which minimize the total distance (in frequency) between the targets.

The problems of signal decomposition and data association are not trivial, but they are well defined in the MTT literature [44, 45, 46]. Since these problems are well defined and response decomposition solves the issue associated with model-based methods which have the quadratic nature of the growing distortion terms we propose this method over the other model-based techniques described earlier. In the next chapter we describe the radar hardware used, and evaluate this method on actual measurements.

CHAPTER 4

MEASUREMENT RESULTS FOR RESPONSE DECOMPOSITION INTERFEROMETRIC DISTORTION MITIGATION

4.1 Measurement System

4.1.1 Single-Board 40 GHz Interferometric Radar

To validate the response decomposition distortion mitigation technique a custom, single-board, 40 GHz correlation interferometer was used. A single-board system was desired in order to be mounted on a vehicle for speed sensing testing as described later, and to be more compact than a typical connectorized system. The system was designed and built using commercially available RF MMIC components. The center frequency of 40 GHz was chosen as a trade-off between the cost and availability of IC's at the 40 GHz frequency and the desire to move to higher frequencies to support smaller physical baselines. A link budget was created in order to verify the received SNR for the given transmit power and components. The system is shown schematically in Fig. 4.1. This schematic represents the measurement setup which we describe in the following section, though the actual radar board contains three receive channels to allow for a pair of interferometric baselines, as opposed to the single baseline as shown in the image. The pair of baselines allows us to perform interferometric angular velocity in two dimensions with a pair of orthogonal baselines.

The printed circuit board (PCB) was designed using Rogers 4350 0.020" substrate with Ansys HFSS for transmission lines and Wilkinson power splitter (PS) [47]. The design consists of a 20 GHz HMC738 voltage-controlled oscillator (VCO), which acts as the continuous-wave (CW) transmitter, and the local-oscillator (LO) for the IQ downconverters. The signal is divided using a 1:4 Wilkinson PS which was designed using a microstrip transmission line and a full 3D wave simulation performed in HFSS as shown in Fig. 4.5. The Wilkinson PS has slightly higher loss than the lossless T-junction, but provides better




Figure 4.1: Schematic of millimeter-wave radar used for experimental measurements. System consists of two phase-quadrature, direct-downconversion receive channels with a continuous-wave transmitter operating at 40 GHz.

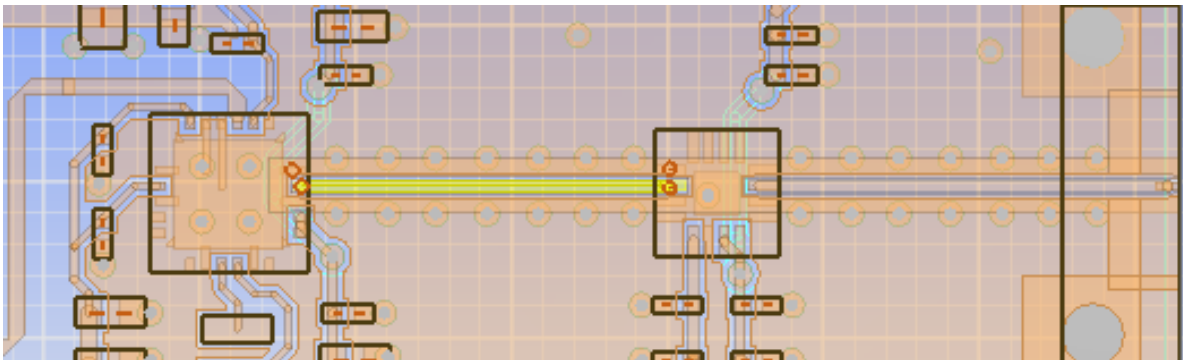


Figure 4.2: Testing the 40 GHz grounded coplanar waveguide transmission lines in Ansys SIwave.

isolation between the output ports [47]. The 1:4 PS consisted of three 1:2 PS. The individual 1:2 PS had a simulated insertion of 3.3 dB, while the complete 1:4 structure had a simulated insertion loss of 7 dB, slightly higher than the theoretical values of 3 and 6 dB respectively. The 20 GHz VCO signal is fed to an XX1000-QT frequency doubler to act as the 40 GHz continuous-wave (CW) transmitter. The 40 GHz transmission lines were designed as ground-coplanar waveguides (G-CPWG). The received signals are amplified by




Figure 4.3: The 40 GHz system on Rogers 4350 substrate has 3 IQ receive channels and a single transmit channel. The antennas are connectorized to allow for experimentation with various antenna types and arrays. Image [3] © 2020 IEEE.

HMC1040 GaAs pHEMT LNA's, which have a noise figure of 2.2 dB and a gain of 23 dB. The amplified signals are directly downconverted with HMC6147 IQ downconverters which have built-in LO x2 frequency multipliers. The baseband in-phase and quadrature signals are amplified with inverting op-amp amplifiers with a gain of 24.3 dB. The single TX and three RX channels use 2.4 mm connectors to allow for experimentation with various antenna types. After the PCB was designed it was imported in Ansys SIwave in order to ensure that the transmission line S_{11} 's were all below -10 dB. The S_{11} parameter represents how much power is reflected back at the input to the device or transmission line. If the system was perfectly matched there would be no reflections, though, in practice, a system will never be perfectly matched, and thus -10 dB represents an acceptable level, which is a common rule-of-thumb. Ansys SIwave is a best-in-class design platform for signal integrity, power




Figure 4.4: The measurement system used in this work consisted of a custom single board radar and L3-NARDA 15 dBi horn antennas with a 20λ interferometric baseline.

integrity, and EMI analysis of PCB's and electronic packages. An image of this analysis is shown in Fig. 4.2, where we can see the highlighted 40 GHz transmission line. After the transmission lines were verified, a resonant mode analysis was performed. This was done to detect areas on the PCB board with slight impedance mismatch which could result in unwanted resonances. In these areas, additional grounding vias were placed and the analysis was re-run.




Figure 4.5: Model of the 1:2 Wilkinson power splitter in Ansys HFSS (a), and the entire 1:4 structure (b) for full-wave simulation.

The assembled board is shown in Fig. 4.3. Since the board was only double-sided, jumper wires were used for certain low-frequency signals to allow for a solid ground plane under RF components and transmission lines. The IQ signals were sampled, off-board, with an NI USB-6002 DAQ. For the measurements shown in this section an additional HMC7229 power amplifier (PA) was used off board before connecting to the transmit antenna. The antennas used were L3-NARDA 15 dBi standard gain horn antennas, positioned to achieve a 20λ interferometric baseline. An image of the entire measurement system is shown in Fig. 4.4.

4.1.2 V-Band Patch Antenna Design

We designed 40 GHz microstrip patch antennas that could be used for certain applications, as discussed later, whereas other applications throughout this work used standard gain horn antennas. Rectangular microstrip patch antennas were designed using the expressions in [48].




Figure 4.6: The 40 GHz patch antenna with 2.4 mm connector. The design uses a quarter-wave transformer to match the transmission line impedance of 50Ω to that of the edge of the microstrip patch. The antenna has a dielectric superstrate included to shield the antenna from environmental conditions. Image [4] © 2020 IEEE.

The same Rogers 4350 substrate as the PCB was used, and the assembled patch antenna is shown in Fig. 4.6. The width of the 50 Ohm transmission line was close to the total width of the microstrip patch and so an inset feed was not possible. Instead, we used a quarter-wave transformer to match the 50 Ohm feed to the impedance at the edge of the antenna. This is calculated using

$$Z_0 = \sqrt{Z_L Z_{in}} \quad (4.1)$$

where Z_L is the impedance at the edge of the patch, and Z_{in} is the characteristic impedance of the transmission line. The characteristic impedance of the transmission line was designed to be 50 Ohms, and the impedance at the edge of the patch antenna was calculated using an open-source, web-based patch antenna impedance calculator. A 0.8 mil layer of solder resist with a dielectric constant of 3.1 and a loss tangent of 0.035 was applied for environmental shielding of the antenna. The S_{11} 's of the antennas were measured using a Keysight N9952A FieldFox, and the plots are shown in Fig. 4.6 for two individual antennas, a total of four




Figure 4.7: Measurements of the S_{11} from two antennas each with a resonance frequency of 42.4 GHz. The overall decrease in S_{11} at higher frequencies is likely an effect of the losses incurred by the dielectric superstrate. Image [4] © 2020 IEEE.

patch antennas were assembled in total. The measured resonance was calculated to be 42.4 GHz, and there is an overall decrease in the S_{11} that is likely caused by the layer of solder resist.

4.1.3 Ground-Truth System and Robotic Targets

Having accurate ground-truth position measurements is critical to validating our radar measurements. An OptiTrack motion capture system was used as the ground-truth system. This system comprises six Flex 13 infrared cameras with a 120 Hz frame rate and stated position accuracy of 0.2 mm. Each target had four asymmetric reflective markers placed on it to allow for accurate position and rotation measurements. A top-down diagram of the experimental setup is shown in Fig. 4.8. The six cameras, which are mounted from the ceiling, are shown which create the capture volume shown as the dashed circle. The origin is set as the center of the interferometric array, and the targets travel in various trajectories as shown. The




Figure 4.8: Top-down view of the experimental setup. The six motion capture cameras form the capture volume shown as the dashed circle. The radar is placed at the edge of the volume and we set the origin to the center of the interferometric array. The targets travel in various trajectories as shown.

ground-truth radial and angular velocity values were computed as the time-rate change of the range and angle measurements. The expected Doppler frequency shifts were calculated using

$$f_D = \frac{2v_r}{\lambda_c} \quad (4.2)$$

for each target, and the interferometric shifts were calculated using (1.7). The targets used were two Robotis Turtlebot3 Waffle platforms. Additional reflective targets made of polystyrene spheres coated in copper tape and placed on the robot platform, as shown in Fig. 4.9.

4.2 Measured Results

To validate the radar measurement setup and ground-truth position system we first performed single target measurements with the target moving at angles ranging from $0 - 360^\circ$




Figure 4.9: Robotis Turtlebot3 Waffle platforms used for measurements, the reflective targets consisted of polystyrene spheres coated in copper tape.

relative to the radar system. The I and Q radar signals were captured at 1920 Hz with the NI USB-6002 DAQ. The signals were high-pass filtered with a 3rd order Butterworth filter with a cutoff of 1 Hz. This was performed to remove any DC bias. The STFT was performed on each antenna response, which generated the Antenna One and Antenna Two Doppler responses. A window length and FFT size of 1024 samples was used with an overlap of 960 samples. The interferometric response was generated by conjugate multiplication in the time-domain between the two individual antenna responses. The time-frequency responses for various trajectories are shown in Fig. 4.10. In each column we see a different target angle trajectory, while the plots from top to bottom are the Doppler response from antenna one (a), the Doppler response from antenna two with the expected Doppler shift in a white dashed line (b), and the interferometric response with the expected interferometric shift in a white dashed line (c). Here we see good agreement between the measured and expected frequency shifts.

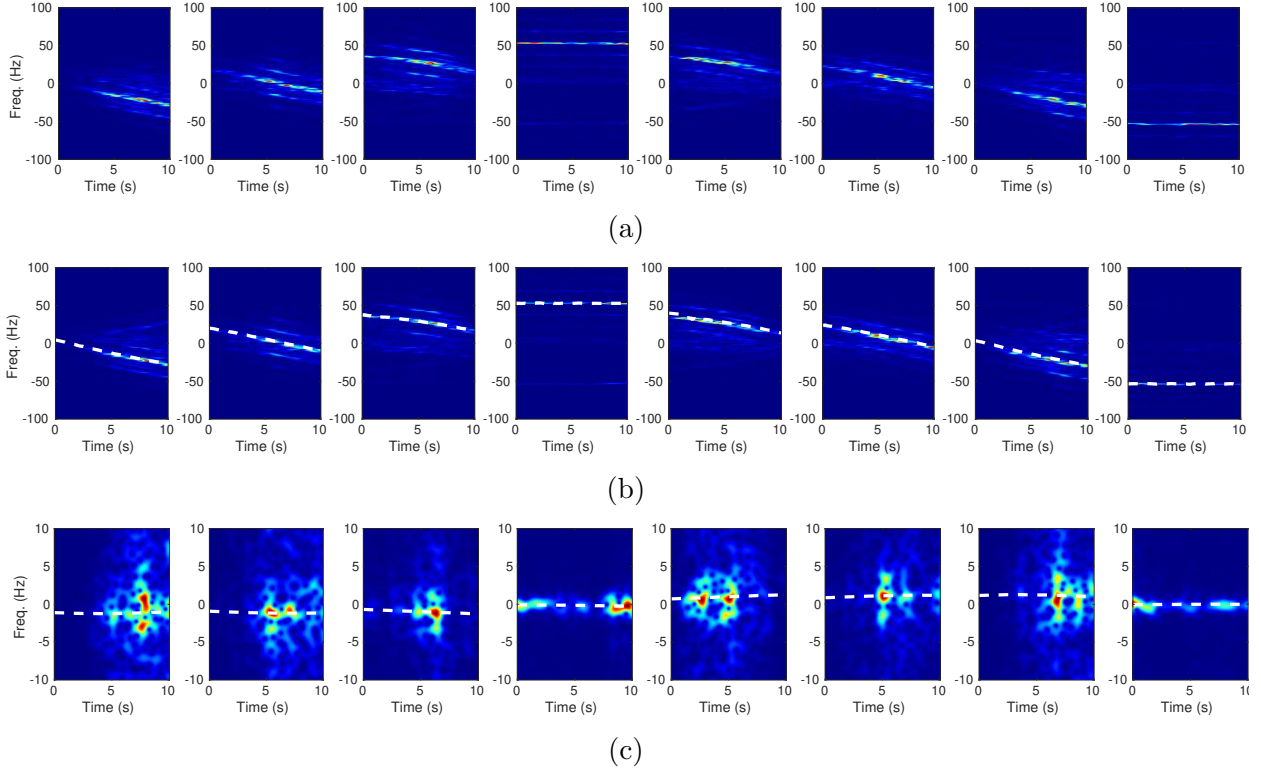


Figure 4.10: The time-frequency responses for a single target at various trajectory angles shown in each column. The plots are the Doppler response from antenna one (a), the Doppler response from antenna two with the expected Doppler shift overlaid in a white dashed line (b), and the interferometric response with the expected interferometric shift (c). The Doppler responses of columns two and six look identical, but we can notice a small negative interferometric shift in column two, and a small positive shift in column six, implying counter-clockwise and clockwise motion.

By looking at the plots we can see that columns four and eight comprise of purely radial motion where we see the highest positive and negative Doppler shifts. Here, the interferometric shifts are zero, implying no angular motion. The benefit of using the interferometric angular velocity measurement technique is demonstrated by looking at columns two and six. The Doppler responses look identical, and thus we would think these targets have similar trajectories. However, in their interferometric responses we notice a negative frequency shift in column two, compared to a positive shift in column six. This implies the target moves in a counterclockwise direction in one realization, compared to clockwise motion for other realization.

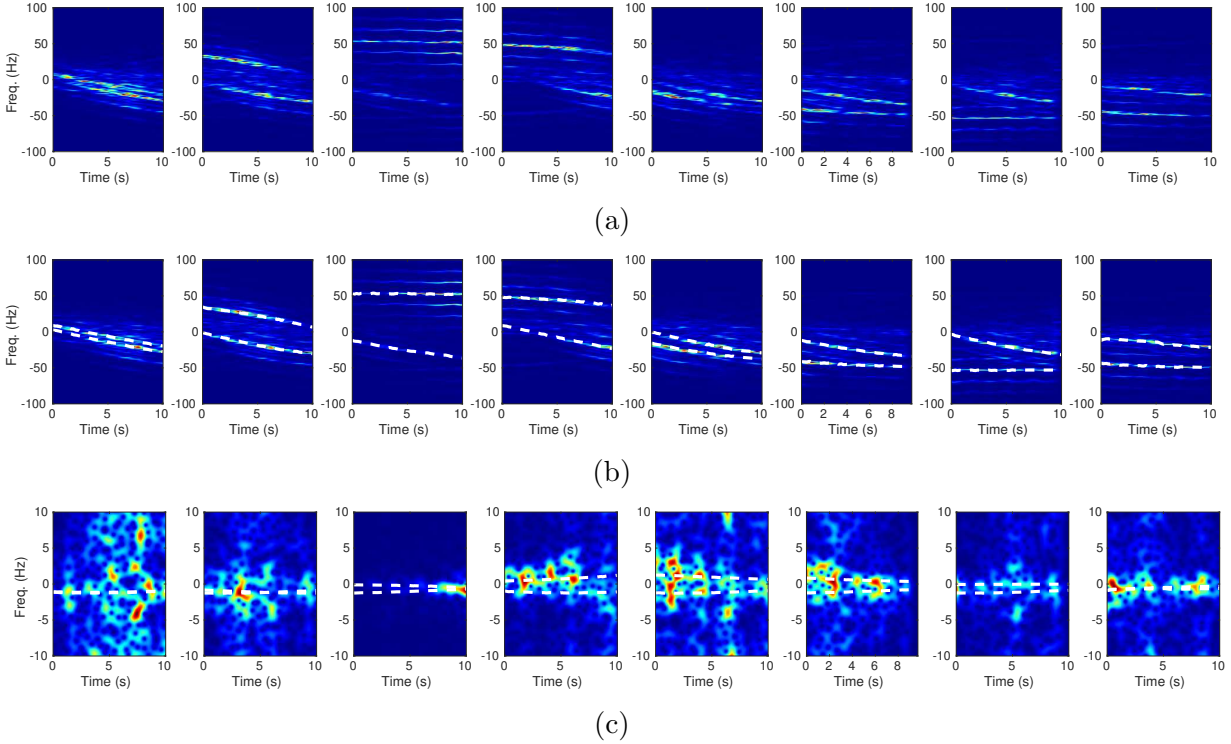


Figure 4.11: The time-frequency responses for two targets with the first target at a constant angle and the second target at various trajectory angles shown in each column. The plots from top to bottom are the Doppler response from antenna one (a), the Doppler response from antenna two with the expected Doppler shift overlaid in a white dashed line (b), and the interferometric response with the expected interferometric shift (c).

We also performed a similar series of measurements using two targets. Here we kept a single target moving in the same trajectory, and rotated the other target at various angles. The time-frequency plots are shown in Fig. 4.11. Here we notice that in some realizations the Doppler frequencies should be easily separable, while in other realizations they are very close together, and we should have little hope of disambiguating the two responses. We can also see that the interferometric response appear quite noisy, and there are not two distinct frequency responses as we would hope. Here the nonlinear distortion corrupts the underlying ideal signals which should motivate the use of our response decomposition distortion mitigation method.

To test the effectiveness of our method we focused on measurements with two targets in three distinct trajectories which were designed to demonstrate easy, medium, and




Figure 4.12: Robot trajectory directions for the easy (a), medium (b), and difficult case (c), with the antenna receiving pair shown at the bottom of each image. The large differential in target Doppler shifts will be easily separable in case (a), while the negligible Doppler shifts in case (c) will not be separable.

difficult cases for the distortion mitigation technique, as shown in Fig. 4.12. The motions shown in (a) result in positive and negative Doppler frequency shifts, which will be simple to isolate in the Doppler responses. There is no angular motion and thus the interferometric shifts will be zero for each target, thus motivating a non-zero interferometric shift case. Case (b) demonstrates trajectories which have both positive and negative Doppler shifts, along with positive and negative interferometric shifts. Case (c) has the largest positive and negative interferometric shifts, though the Doppler shifts are small and similar. In this case the distortion mitigation should not be able to isolate the individual target Doppler signals.

The time-frequency for the first (simple) case is shown in Fig. 4.13. The first antenna Doppler response is shown (a), where we see large positive and negative Doppler shifts due to the away and towards motion of the two targets. The second antenna Doppler response is shown (b) with the expected Doppler shift overlaid in the white dashed line. The interferometric response is shown (c) also with the expected interferometric shift. In this case the intermodulation distortion terms are much higher in frequency, and so we could estimate that each target has no angular motion from the interferometric response. The time-frequency responses for the second (medium) case are shown in (d), (e), and (f). We again see positive and negative frequency shifts due to the approaching and retreating targets, though this time they vary in time due to the angular motion. We can notice that

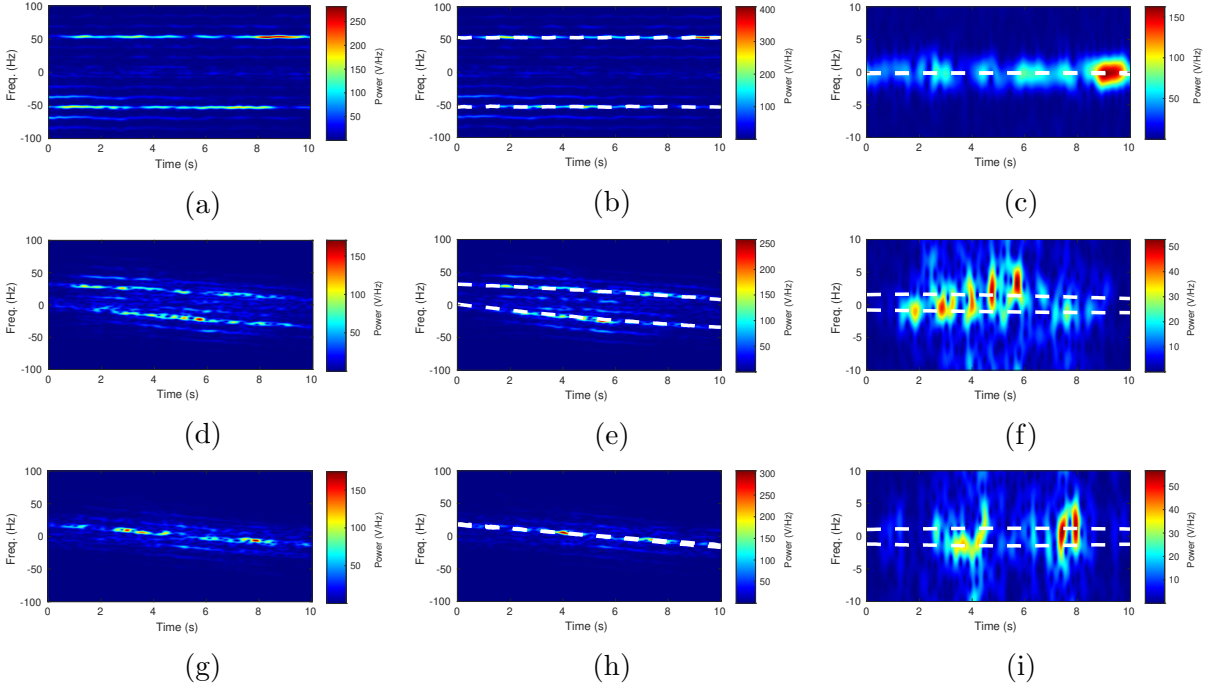


Figure 4.13: The measured time-frequency responses for the easy case include the antenna one Doppler response (a), antenna two Doppler response with ground-truth expected Doppler shift (b), and interferometric response with the expected interferometric shift (c). The same responses are shown for the medium case in (d), (e), (f), as well as the difficult case in (g), (h), and (i).

the interferometric response is corrupted by intermodulation distortion, and thus estimating the individual target angular velocities from the complete interferometric response in (f) would be very difficult. The time-frequency responses for the third (difficult) case are shown in (g), (h), and (i). Here we can notice that the Doppler response of each target is the same and overlaid on top of the other. There is an interferometric response in the region of the expected frequency shift, which is a result of the fact that each target has a similar Doppler shift and so the intermodulation distortion terms tend to zero.

4.2.1 Angular Velocity Estimation via Response Decomposition

In this section we apply the response decomposition distortion mitigation technique, described above, on the measured responses. We assume that the number of targets is known, which shown here is equal to two. The number of targets may be unknown and time-varying

in general, and a method may be used to estimate this parameter independently of the distortion mitigation method, though this is beyond the scope of this work. In the following section we assume that the location of the target in frequency is known and that we are able to perfectly associate like targets between the antenna elements. We can do this because we have the ground-truth radial velocities, and so can calculate the actual Doppler shift at any time instant without having to perform detection. After showing the ideal case we present results which are more practical, that is, the detection and target frequency location estimation is performed on the data.

4.2.1.1 Known Target Frequencies and Data Associations

As we see in Fig 4.13 (b), since we have measurements from the ground-truth system, we have access to the expected Doppler shift for each target. This is computed from the time-rate of change of the ground-truth position measurements. The next step is to decompose each antenna response into two separate responses, one for each target, by masking out the spectral response due to each target. We do this by masking out a portion of the spectrum using a passband of 10 Hz around the expected Doppler shift. This results in four responses, the target one and target two responses generated from each of the two antenna responses. Conjugate multiplication in the time-domain is cross-correlation in the frequency domain, thus we cross-correlate the antenna one, target one response with the antenna two, target one response. In the same manner, we also cross-correlate the antenna one, target two response, with the antenna two, target two response. It should be clear that by performing the cross-correlations individually in this manner we are mitigating the cross-terms which are generated when we perform the conjugate multiplication in the time domain.

One important aspect to mention here, is that, in this instance we are able to use the same mask for each antenna response. That is, that target one mask that we use to mask out the antenna one signal is the same mask we use to mask out the antenna two target one signal. This is possible in these measurements because the targets are moving at relatively

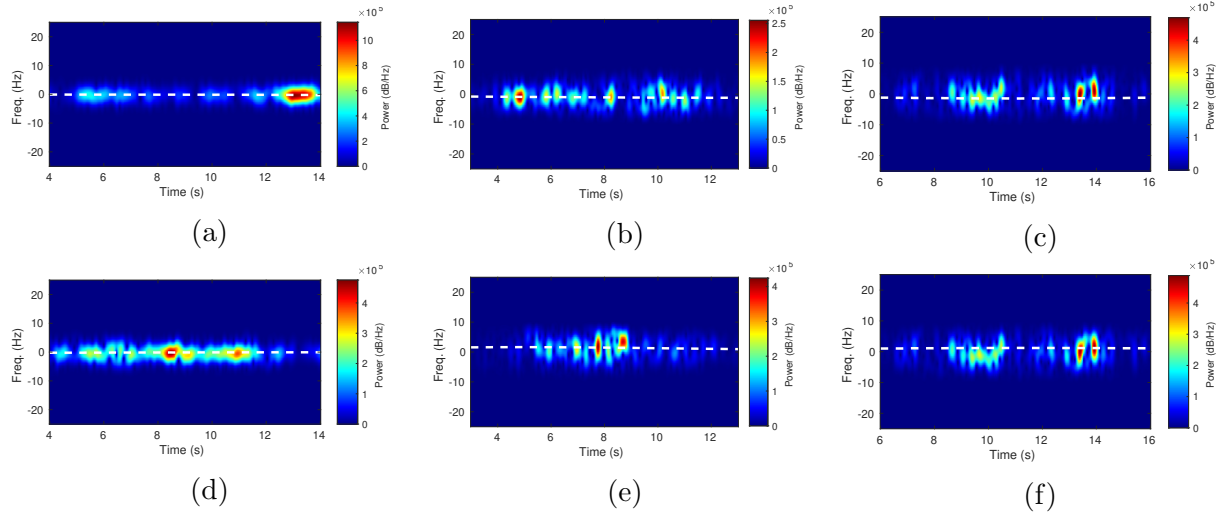


Figure 4.14: The reconstructed interferometric responses for the first case target one (a), and target two (d), the second case for target one (b), and target two (e), and the third case for target one (c), and target two (f). Here we can see good agreement for case one and two, while there is a visible bias in the responses for case three.

slow angular velocities and we are using a relatively modest interferometric baseline (20λ). This results in very small Doppler frequency differences between antenna one and antenna two for a single target response. Thus, as long as the mask width is large enough, we will capture the individual target response in both antenna responses. Though, as the baseline grows larger and the targets angular speed grows, there may be an appreciable difference in the Doppler frequencies between antennas for an individual target. Due to this, in general one should perform separate detection and masking for each antenna response and for each target. One benefit of using the same mask is that in this case there is no need for data association, because we know that a single mask is used for the same single target in each antenna response. It is only in the case where we are using different masks for each antenna element and each target, that we must consider the data association problem.

After computing the cross-correlations in the frequency domain for each time step we obtain two reconstructed interferometric responses, one for each target. These are shown for each case in Fig. 4.14. Here we can identify the slight positive and negative interferometric frequency shift differences between each target for the second case in (b) and (e). We

should be able to confirm visually that estimating angular velocity from these reconstructed responses will give us a more accurate estimate compared to the complete interferometric response shown in Fig. 4.13 (f). To actually estimate the interferometric frequency shift we simply take the maximum peak at each time instance of the reconstructed response with a cutoff threshold of -20 dB from the maximum peak voltage. The thresholding is performed to reduce noise in the estimates that is caused by areas in which peaks are detected but there is little signal response. We then perform smoothing with a window length of 60 samples or 0.5 s to reduce the variance of the estimates. The smoothing helps to reduce noise which is likely due to robot platform vibrations and areas where one target may be occluded. At each instance we use (1.7), though with the small angle approximation, to compute the angular velocities. The estimates for the first case are shown in Fig. 4.15 (a) and (b) for each target. As expected, the actual interferometric shift is close to zero because there is no angular motion for either target. The estimation statistics for all cases are shown in Table 4.1. The ground-truth mean angular velocity, μ , the estimated mean from the response decomposition method, $\hat{\mu}$, the percent error between the estimated and actual averages, and the standard deviation of the estimates are all shown. The percent error for target one is an outlier due to the very small actual and estimated values. We notice that the variance of the estimates is quite low for case one, as one would expect due to ease by which the individual Doppler responses are separable. The estimates for the second case are shown in Fig. 4.15 (c) and (d). Here we see an increase in the estimation error and the variance, though the estimates are still within 25% of the ground-truth values. Finally, the estimates for the third case are shown in Fig. 4.15 (e) and (f). As expected for the difficult case, where the Doppler responses are completely inseparable, the estimation error and variance increase substantially.




Figure 4.15: Angular velocity estimates for the first case with ground-truth angular velocity underlaid for target one (a) and target two (b), we see a low estimation error and variance for the easy case. The estimates for the second case are shown in (c) and (d) with a small increase in the variance as expected. The estimates for the third (difficult) case are shown in (e) and (f) where we see a larger increase in the estimation variance, as expected.

4.2.1.2 Unknown Target Frequencies and Data Associations

In the prior section we used the ground-truth system locate and mask out the responses due to each target in the individual antenna responses. In the general case, we will not have access to this information and so the target responses must be detected and decomposed in the response of each antenna element. The simplest method would be using peak detection, which works well if the target response resembles a sinusoid. In our measurements, we found

Table 4.1: Angular Velocity Estimation Statistics With Known Target Frequencies

Target:	μ (Rad/s):	$\hat{\mu}$ (Rad/s):	Std. (Rad/s)
Case 1, T1	-0.007	-0.014	0.009
Case 1, T2	-0.008	-0.011	0.005
Case 2, T1	0.072	0.067	0.019
Case 2, T2	-0.053	-0.066	0.024
Case 3, T1	0.059	-0.008	0.025
Case 3, T2	-0.071	-0.011	0.025

Table 4.2: Angular Velocity Estimation Statistics With Unknown Target Frequencies

Target:	μ (Rad/s):	$\hat{\mu}$ (Rad/s):	Std. (Rad/s)
Case 1, T1	-0.007	-0.010	0.008
Case 1, T2	-0.008	-0.008	0.004
Case 2, T1	0.072	0.023	0.020
Case 2, T2	-0.053	-0.036	0.015
Case 3, T1	0.059	-0.026	0.033
Case 3, T2	-0.071	-0.001	0.019

that the response of the robot targets did not resemble an ideal single frequency, but would have multiple peaks and harmonics which were caused by the vibration of the robotic base. Because of this often multiple peaks were detected a single target. To mitigate this, at each time instance we found the two non-overlapping windows, each with a passband of 10 Hz, which had the highest total integrated power. As we mentioned in the previous section, in general we should be performing this detection on each antenna response independently, then associate the passbands between the two antennas. Though in our measurements, because the angular velocities are relatively low and the interferometric baseline is modest (20λ), the differences in target frequency between the two antenna responses is quite small. Thus we were able to use the same mask for each antenna response and thus no data association was required. The results using this method were similar to the ideal scenario above, though with slightly higher estimation errors due to the uncertainty of the exact target locations. The estimation statistics for these cases are shown in Table 4.2.

CHAPTER 5

PRACTICAL APPLICATIONS OF ANGULAR VELOCITY ESTIMATION

In this chapter we discuss three practical applications of radar sensing using interferometric angular velocity estimation. In the first section we demonstrate a novel method of using interferometric millimeter-wave radar for ground-speed estimation of a moving vehicle. In doing so, we present two estimators, one Bayesian, and one neural network based, for point estimation of velocity from the measured radar spectrum. The estimators achieved a root-mean-square error (RMSE) of 0.403 m/s and 0.183 m/s respectively. This work was first reported in [5]. In the next section we seek to quantify the tracking improvement we are able to achieve by adding an additional angular velocity measurement to a state space tracker. We do this first by simulating a single point target in noise, and then by using a commercially available autonomous driving dataset, nuScenes, with measured-simulated data. In this work we found, on average, a 10-13% reduction in position error by adding a radial velocity measurement, and an additional 3-6% improvement by further adding an angular velocity measurement. Interestingly, we found that the additional angular velocity measurement had a greater improvement on accuracy in scenarios with high clutter, this work was first reported here [6]. In the final section we compared hand-gesture classification performance between systems containing micro-Doppler only and micro-Doppler plus interferometric responses using a CNN classifier. We found that including an interferometric measurement along with a Doppler measurement served to improve the classification in all cases, and this work was first reported here [49].

5.1 Vehicle Speed Estimation

For safety critical autonomous vehicles (AV), highly accurate vehicle position, velocity, and acceleration measurements are required. To improve the robustness and accuracy, velocity measurements are may be combined with position estimates via Linear Quadratic Estimation

[41], or more commonly, Kalman filtering [8]. This implies that accurate velocity measurements improve not only velocity controllers, but also the localization subsystem [50]. The AV literature describes both odometry and ego-motion estimation methods. Typically, odometry methods calculate velocity as an intermediate step, then integrate over time to obtain the total distance traveled. These methods typically state accuracy in terms of drifted distance from ground-truth in meters. The ego-motion estimation methods typically describe velocity estimation directly and tend to state accuracy in terms of root-mean-square error (RMSE) or average percentage of velocity error.

Most methods for measuring vehicle speed/odometry typically fall within one of four categories: using on-board wheel encoders, known steering angles, and tire pressure sensors with a dynamical vehicle model [51], using the already existing outward facing sensors, such as camera, lidar, or radar, to infer vehicle motion based on the relative motion of static objects in the scene [52, 53, 54, 55, 56, 57], using dedicated downward facing sensors to measure motion based on the ground-speed, which can also be done with camera, lidar, radar, or ultrasonic sensors [58, 59, 60, 61, 62], or using global navigation satellite system (GNSS) sensors combined with inertial measurement units (IMU's).

A major benefit of radar-based speed sensing is that it should provide accurate readings during times of other sensor outages, such as during wheel-slip for encoder based sensors, during low-visibility when camera and lidar based odometry have difficulty, and during times of GNSS outages, when INS systems will drift due to the integration of small errors. The work in this section focuses on radar based sensing which falls into the third category mentioned above, that is, methods which use downward facing sensors measuring the ground surface to infer vehicle speed. We present a new form of radar ground speed measurement that uses a correlation interferometer for the direct estimation of the ground angular velocity, and thereby the linear speed of the vehicle relative to the ground. Other microwave speed sensing approaches typically use forward-pointing, downward-angled radars to estimate the Doppler shift based on the radial velocity of the ground relative to the vehicle




Figure 5.1: A single point scatterer, n , with linear velocity $v_{l,n}$ which can be decomposed into radial and tangential components. Here $\theta_{1,n}$ is the angle from receiver 1 to point n , and $d_{1,n}$ is the distance from receiver 1 to n . Image [5] © 2020 IEEE.

[59, 60]. This approach is potentially impacted by road debris, and is also affected by the low signal-to-noise ratio (SNR) due to low look-angle which causes most of the energy to scatter in the forward direction.

Here we demonstrate high-accuracy ground-speed estimation from a moving vehicle using the 40 GHz interferometric radar hardware described above. This section is organized as follows. In Section 5.1.1 we provide a modification to complete interferometric signal response, which accounts for the backscattering effect of a downward-facing interferometric radar measuring ground-speed. In Section 5.1.2 we present Bayesian and neural net estimators we developed for performing point estimation to vehicle speed. In Section 5.1.3 we provide measurements performed on a test-vehicle, and evaluate the estimators on the captured measurements.

5.1.1 Signal Model

For an individual antenna element, the measurement geometry is shown in Fig. 5.1. The linear velocity of point target n is $v_{l,n}$, which is the parameter we would like to estimate. The radial and tangential velocities are calculated as $v_{r,n} = v_{l,n} \sin(\theta_{1,n})$ and $v_{t,n} = v_{l,n} \cos(\theta_{1,n})$ respectively, where $\theta_{1,n}$ is the angle from receiving element 1 to scatterer n . Although the road surface is continuous, it can be approximated as a summation of point scatterers. For a downward facing system, as shown here, an electromagnetic wave is transmitted and

reflected from the road surface which appears irregular due to the randomly rough texture of the pavement surface. The scattering of electromagnetic waves from randomly rough surfaces has been covered in the literature extensively, and is described in [63]. Fig. 5.2 demonstrates how the scattered response varies with surface roughness. Because the ground surface is random, the total electric field from a source at any given point cannot be solved for deterministically. Thus, a model of the backscattering effect is used based on the surface roughness and angle of reflection. This effect may be described by a diffuse component which scatters in all directions and a coherent component which scatters specularly in the direction opposite the angle of incidence. It may be approximated as [64]

$$B(\theta) = (1 - b)e^{-a\theta} + b \quad (5.1)$$

where a controls the width of the coherent portion of the beam and b controls the amount of backscattering from the diffuse component. Both of these parameters are determined by the ground surface roughness, which may be defined by the root-mean-square (rms) height of the surface. This expression can be useful for analytic calculations due to its simplicity, but it does not capture the continued attenuation in angle that manifest in measurements [65]. To account for this additional effect we multiply the entire response by an exponential term, resulting in the modified backscatter pattern

$$B(\theta) = ((1 - b)e^{-a\theta} + b)e^{-c\theta}. \quad (5.2)$$

In general, the parameters a , b , and c are unknown and will vary depending on the road surface roughness and also weather conditions. An estimator which includes these parameters as state variables may be able to correlate ranges of values with various road surface states.

Although a single downward facing antenna element can be used for velocity estimation based solely on the Doppler response, the interferometric aperture more closely approximates the ideal aperture for maximizing the accuracy in angular velocity measurements [66]. It is this fundamental result that motivated the use of a correlation interferometer for vehicle velocity estimation. The general interferometer response was given in (2.8), which




Figure 5.2: The backscattering pattern for a smooth surface (left) is primarily specular, a medium roughness surface (middle) has both a specular and a backscatter component, and the very rough surface (right) is primarily backscatter. Image [5] © 2020 IEEE.

should now be multiplied by the backscattering response. Thus the total response is

$$R(f) = \sum_{n=1}^N \sum_{k=1}^N A(\theta_{1,n})B(\theta_{1,n})A(\theta_{2,k})B(\theta_{2,k})\delta(f - f_{n,k}), \quad (5.3)$$

where the values of the frequency components $f_{n,k}$ may be either of the three models described in Chapter 2. In this case the exact model should be used because the far-field assumption does not strictly hold, but we have found that often even the approximate model provides sufficiently accurate results, as shown later.

5.1.2 Velocity Estimation

To estimate the vehicle’s velocity from a measured radar spectrum we must perform point estimation. Typically, if there is a model which describes the measurement, and a model for the dynamics which describe the time progression of the system, we can use Kalman filtering to obtain optimal state estimates under certain assumptions [67]. In instances where the measurement model is unknown, we can learn a mapping from the measurement to the point estimate with a neural network. In the following section we present two estimators, one using a traditional Unscented Kalman Filter (UKF) [68], which must be used due to the non-linear measurement model, and one using a feedforward neural network [40].

5.1.2.1 Bayesian Estimator

In state space tracking, the state variables are estimated from noisy measurements. In this work, the state variables are longitudinal vehicle velocity and acceleration. These methods

are considered Bayesian because prior knowledge about the distribution of the noise and the system kinematics is used to form posterior estimates which are more robust to sensor noise and environmental dynamics [69]. The system is modeled by a linear discrete-time process equation and a non-linear measurement equation given by

$$\begin{aligned}\mathbf{x}_k &= A\mathbf{x}_{k-1} + \mathbf{w}_{k-1} \\ \mathbf{y}_k &= h(\mathbf{x}_k, \mathbf{v}_k)\end{aligned}$$

where \mathbf{x}_k is the state-vector at time index k , A is the state-transition matrix, \mathbf{w}_k is the process noise, \mathbf{y}_k is the measurement vector, and \mathbf{v}_k is the measurement noise.

The measurement noise models the uncertainty from the randomly rough road surface, environmental factors, and radar system noise. It was found experimentally that the spectral amplitude components were Rayleigh distributed, which implies that the underlying process is complex Gaussian, which is consistent with the central limit theorem. To increase the dynamic range we plot the log spectrum resulting in a Log-Rayleigh distribution [70] of the spectral amplitudes. When the system noise \mathbf{w}_k and the measurement noise \mathbf{v}_k are Gaussian, zero-mean, and uncorrelated, than the Kalman filter is optimal in the minimization of the mean squared error sense [67]. In this instance the errors are non-Gaussian; methods for handling non-Gaussian noise are beyond the scope of this work, so we simply use the best linear estimator, which in this case is the UKF. This method uses the unscented transform, which is a form of statistical linearization, to handle the non-linear measurement equation [71]. The measurement equation then becomes

$$\mathbf{y}_k = R(\mathbf{x}_k) + \mathbf{v}_k \tag{5.4}$$

where R is the model spectral response given in (5.3). To characterize the measurement noise \mathbf{v}_k , samples of the spectrum were gathered at a constant speed and the variance for each spectral component was computed. The overall variance was taken as the average over all spectral components.

An evaluation of empirical ego-motion tracking performance for different vehicle motion models is given in [72]. This work found that more complex models improve accuracy only when the additional parameters are observable in the measurement model. Our measurements are modeled as a function of velocity only thus a simple constant acceleration model was chosen. One method for selecting the process noise covariance matrix is through optimization, that is, finding the values which minimize the estimation error over some data. A simpler approach that is data independent is to select the maximum acceleration of the vehicle to be a number of standard deviations for the acceleration variance, the downside being that it is a conservative estimate [73]. From this value, the variance for the velocity is found by multiplying the acceleration variance with the size of the time-step.

5.1.2.2 Neural Net Estimator

When the measurement model is not known velocity estimation may also be performed using a fully-connected feedforward neural network (NN), for example, as shown in Fig. 5.3. In the figure, the 256 frequency amplitude values are shown as the inputs to the network as x_1 - x_{256} , and a single hidden layer of 16 neurons is shown. The benefit of this method is that no measurement model is required, only system measurements and ground-truth velocity values are needed for training the network. In the results shown in the next section, preprocessing of the spectrogram was performed to filter the noise on the spectral components. As mentioned, this was Log-Rayleigh distributed which is long-tailed, and so a peak-hold filter was used to reduce this effect.

Data was split into 75% training data and 12.5% for both the validation and test sets respectively. The validation data was used to determine when the network had been over-trained by examining the error at each learning iteration. The network was trained for 20k iterations with stochastic gradient descent (SGD) with momentum at a learning rate of 0.001 and a momentum value of 0.9 [74]. The training and validation errors vs. training iteration are shown in Fig. 5.4. The lowest validation error and corresponding training error




Figure 5.3: A sample network with 256 input features and 16 hidden neurons. The input features represent the amplitudes of the frequency components from the spectrogram for a given time slice. Various networks were tested and a network with 3 layers of 64 neurons each was selected. Image [5] © 2020 IEEE.




Figure 5.4: The training and validation error shown over 20k iterations. The network is selected based on the iteration with the lowest overall validation error, after this point the validation error increases and the network starts to overfit the training data. Image [5] © 2020 IEEE.

are shown on the plot, and this was the network which was selected to evaluate the test data. Notice that after this point the validation error begins to increase, while the training error continues to decrease, which implies that the network is overfitting the training data.

Various hold lengths were tested for the peak-hold filter, and a period of 15 time-

steps was selected based on the trade-off between filter throughput and reduction in RMSE. Also, various networks with varying numbers of neurons and layers were tested and a network with 3 layers of 64 neurons each was chosen based on the lowest validation error. It was found that the derivatives of the raw neural net estimates were outside the vehicle acceleration limits so smoothing was performed on the output with sliding-window regression [75].

5.1.3 Experimental Evaluation

5.1.3.1 Measurements

The radar system used for this work was the same single-board 40 GHz interferometric radar described previously. For data collection the radar system was mounted to the side of MSU's CANVAS research platform, which is a Lincoln MKZ outfitted with various sensors by AutonomousStuff, shown in Fig. 5.5. The ground-truth velocity was measured as a twist value in robot operating system (ROS) from the built-in controller from Dataspeed which uses the vehicle's wheel encoders to measure vehicle speed. Vehicle wheel encoders typically have an accuracy of around 40 cm/s, though an exact value for this system was not given. Prior tests were conducted with a Novatel PwrPak7-E1 GNSS/INS system which had a stated accuracy of 3 cm/s. The Dataspeed values appeared less noisy than the GNSS results, thus the actual accuracy is likely somewhere between 3-40 cm/s. The ground surface for all measurements was cement concrete of paved parking lots and public roads.

A 42.875 GHz CW signal was transmitted, and the received radar signals were digitized at 8 kHz using a National Instruments USB-6008 digital acquisition system (DAQ), having with an absolute voltage resolution of 2.56 mV. The antenna height from the ground surface was 0.6 m. For the results in this work, 20 minutes of data was captured while driving around a residential-like setting in MSU's Spartan Village. This was a mixture of medium to low-speed driving with brief stops in between. The sampled signals were converted to a time-frequency representation using the short-time Fourier transform (STFT) with a window



Figure 5.5: The 40 GHz radar hardware suction mounted to the MSU CANVAS autonomous vehicle test platform. Approximately 20 minutes of driving data was captured at MSU’s Spartan Village. Image [5] © 2020 IEEE.

length and FFT size of 1024 samples.

To characterize the system noise and verify the model response, the instantaneous frequency responses were examined from a single data capture, where the vehicle velocity was 10 ± 0.05 m/s, resulting in 472 samples. The samples were averaged in order to reduce the high variance associated with an individual measurement. This is shown in Fig. 5.6 along with the modeled response at 42.875 GHz. The modeled response was generated using a total of 500 simulated point scatterers distributed evenly in angle. The antenna pattern was modeled as a sinc function with a HPBW of 80° . The parameters for the backscattering pattern described in (5.2) were $a = 120$, $b = 0.1$, and $c = 5$. A fringing effect is apparent in both the modeled and measured responses, possibly caused by the non-linear mixing of the backscatter pattern. We note that the locations of the peaks are slightly misaligned,

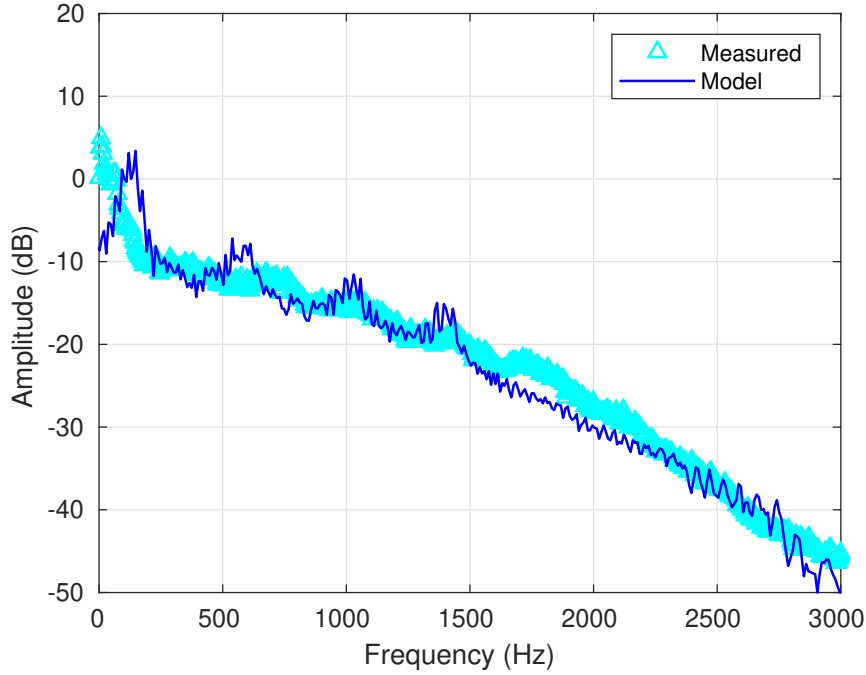


Figure 5.6: The average measured spectrum over 472 samples where the test vehicle was traveling at 10 ± 0.05 m/s, along with the modeled response using 500 point scatterers. Notice that the model seems to capture the fringing effect that is present in the measurements, though the locations of the peaks are slightly misaligned. Image [5] © 2020 IEEE.

and the measured peaks are noticeably less pronounced, however, the entire spectrum is representative of the angular velocity, and still provides a useful mechanism for estimating the angular velocity as demonstrated below.

In Fig. 5.6 the main peak of the model response is shifted slightly in frequency away from dc. The value of this shift is the theoretical interferometric frequency shift for a point source given in (1.7). It is unclear why this shift is not visible in the measured spectrum, though it is possible it is simply obscured by a larger dc component. Being able to recover this interferometric shift will enable not only the estimation of angular velocity, but the direction of travel as well, which would be useful for detecting forward and reverse motion. This is a novel measurement aspect of an interferometric system, which is not captured by a system using only a single antenna pointing normal to the road surface. This would have an equal number of positive and negative frequency shifts, and thus contain no information

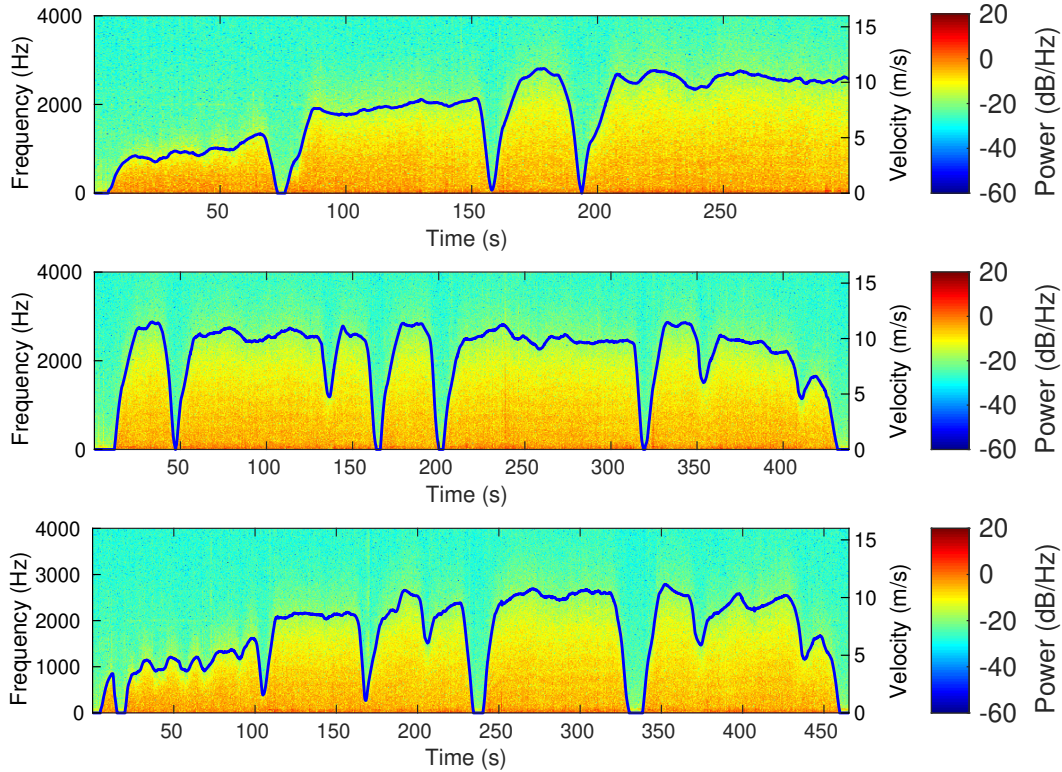


Figure 5.7: The time-frequency response for the three individual interferometric data captures. The ground-truth velocity from the vehicle CAN bus is shown overlaid in blue. Notice the fringing effect predicted by the model is faintly observable in an instantaneous time slice of the response. Image [5] © 2020 IEEE.

on directionality. Although the main peak contains important information, the shape of the entire spectrum is determined by the bulk vehicle velocity and so the velocity should be fully recoverable. The responses for the three data captures used in this work are shown in Fig. 5.17, with the ground-truth vehicle velocity overlaid in blue.

5.1.3.2 Estimator Results

The accuracy of each estimator is computed as the RMSE between the predicted velocity and the ground-truth velocity over the test dataset. Both estimators were tested end-to-end using simulated data based on the signal model described above. A set of measured velocity values from the test vehicle was used as the simulated trajectory. At each time instance

a spectrum was computed from (5.3) using 500 point scatterers. For the UKF estimator, the antenna pattern was taken from the simulated results in Ansys HFSS for the patch antennas, and the backscatter parameters were the same as described above. Rayleigh noise with a parameter $\sigma = 5$ was added to the spectral amplitudes to simulate noise from the randomly rough road surface and other system noise. The UKF estimator achieved an RMSE of 0.010 m/s on the simulated data, performing extremely well as expected. This is a direct result of perfect knowledge of the underlying system model in the simulated case. In reality, the system model will never perfectly describe the physical model so this will be a major source of estimator variance. The neural net estimator achieved an RMSE of 0.085 m/s on the same trajectory. As expected, the neural net performs well, but is not able to learn the complete measurement model that the UKF has prior knowledge of.

Next, we evaluated the estimators on the measurements as performed above. Although the Bayesian estimator could be evaluated over all data, we limit it to the test dataset seen by the NN estimator in order to provide an objective comparison. The results for the UKF estimator on the measured test data are shown in Fig. 5.8. The estimator achieved an accuracy of 0.403 m/s on approximately two minutes of test data. One reason the variance of the estimator is high is that the amplitude of the fringes is on the order of the variance of the noise on the spectral components, which is primarily caused by the randomly rough road surface. This noise is also not strictly independent as in the simulated case.

An image of the instantaneous spectral response and the model response for multiple UKF velocity estimates are shown in Fig. 5.9. Although the fringes are quite visible when integrating out the spectral noise as shown in Fig 5.6, for a given time instance the information containing the location of the fringes may be completely corrupted by noise. We also found that there was a slight mismatch between the modeled backscatter pattern and the actual response. One way to resolve this would be to physically measure the backscatter response at various speeds and use the measured response in place of the analytic model.

As mentioned previously, the derivatives of the raw NN estimates were outside the




Figure 5.8: The estimates from the UKF tracker on the interferometric response shown in blue with the ground-truth velocity shown in cyan. Here we show performance over the same test data used with the neural net estimator. Image [5] © 2020 IEEE.

range of possible vehicle acceleration values, so to reduce the variance, 3rd order sliding-window regression was performed with a window length of 300 samples. The smoothed output is shown in Fig. 5.10. Naturally, there is a tradeoff between estimator variance and sensor throughput, that is, if we require instantaneous measurements in the form of short time windows, the variance of the estimates will increase. The estimator results are shown in Table 5.1 for the UKF estimator, a baseline network with no peak-hold filter, a network with peak-hold filter of length 15, and the same network with a smoothed output. We see that the NN estimator with smoothing is able to outperform the Bayesian estimator. It should be noted that with a more accurate model of the antenna and backscattering pattern, and with methods for handling non-Gaussian noise, the performance of the UKF estimator could be improved. Although, in the case where the model is missing a key component of the physical response, it is possible that the NN is able to learn missing components and thus outperform an estimator which requires a system model.




Figure 5.9: The instantaneous interferometric spectral response shown in blue with the mean of the posterior state shown in magenta. Two sigma-point responses from the UKF are also shown in green and blue. Image [5] © 2020 IEEE.

Table 5.1: RMSE of Estimators

Network:	Interferometric:
Bayesian UKF	0.403 m/s
Baseline Network	0.267 m/s
Baseline w/ Peak Hold	0.245 m/s
Peak Hold & Regression	0.183 m/s

5.2 Multiobject Tracking with Orthogonal Velocity Measurements

As we have mentioned earlier, traditional automotive radars provide In this section we examine the effect of adding an additional, independent orthogonal measurement to an object tracker’s measurement model. We show this first in simulation using common single-object




Figure 5.10: The output of the neural net estimator after smoothing with 3rd order local regression, the ground-truth velocity values are shown in cyan. Image [5] © 2020 IEEE.

tracking (SOT) algorithms on a point-target in varying degrees of clutter. Next, we apply this method to measured-simulated data using a more recent random-finite-set (RFS) multiobject-tracking (MOT) approach. We use a publicly available autonomous driving dataset for range, angle, and Doppler measurements while synthesizing angular velocity measurements from the provided ground-truth labels. We also provide examples of interferometric system designs which should be capable of achieving the accuracy levels set forth in the simulated and measured-simulated results. In the following sections we describe the SOT and MOT state-space tracking frameworks used, and present our simulated and measured-simulated results.

5.2.1 Object Tracking

The difficulty in target tracking with radar systems is generally caused by random false alarms, missed-detections, and clutter. We define clutter as returns from objects in the scene which are not the targets we are care about tracking [76]. The problem becomes one of data association, where we must associate specific detections with the target of interest, and disregard others as clutter [46]. In this section we describe the single-object tracking (SOT) algorithms used in simulation and the random-finite-set (RFS) multiobject-tracking

(MOT) method use for the measured-simulated data.

5.2.1.1 Single Object Tracking

For single target tracking, the Bayes recursion is used, which consists of calculating the prior state density with the Chapman-Kolmogorov equation, the prediction step, and calculating the posterior using Bayes rule, the update step [77]

$$p(x_k|z_{1:k-1}) = \int \pi(x_k|x_{k-1})p(x_{k-1}|z_{1:k-1})dx_{k-1} \quad (5.5)$$

$$p(x_k|z_{1:k}) \propto p(x_k|z_{1:k-1})p(z_k|x_k) \quad (5.6)$$

where $\pi(x_k|x_{k-1})$ is the transition density for a single object, and z_k is the measurement vector containing known data associations. In the linear and Gaussian case it is possible to write the complete posterior analytically, though computing it becomes intractable after only a few time steps because of the large combination of data association hypotheses [77].

In the simulated results below we use three common SOT algorithms. Each use different methods for approximating the posterior. The nearest-neighbor (NN) algorithm prunes all data association hypotheses except the most likely one, probabilistic data association (PDA) merges the hypotheses at a single time instance into a single Gaussian distribution, and the Gaussian sum filter (GSF) uses both pruning and merging to represent the posterior as a mixture of weighted Gaussians.

5.2.1.2 Multiobject Tracking

Multiobject tracking deals with tracking N objects over time, in which the number of objects may be static or time-varying. When N is static and known we can use extensions to the single-object algorithms, such as global nearest neighbor (GNN), joint-probabilistic data association (JPDA), multi-hypothesis tracker (MHT). To deal with an unknown and time-varying number of targets, as is typically the case, extensions are needed to handle the birth and death of tracked objects [9].

A more recent approach, based on Mahler’s finite-set statistics (FISST), represents both targets and measurements as random-finite sets (RFS) [78]. An RFS is a set containing a finite number of elements, where both the elements and the cardinality are random variables. This formulation enables the entire multiobject tracking problem to be expressed as a single prediction and update step, almost identical to those given for the single object tracking problem. Here the states and measurements are represented as RFS’s and the integration becomes a reference measure over the set space [79]. This formulation is still intractable though, and so we must still use approximations. In this work we use the Gaussian mixture probability hypothesis density (GM-PHD) filter [80]. This approximation propagates the PHD of the posterior as a Gaussian mixture, where the PHD is the first-order statistical moment of an RFS, sometimes called its intensity function.

5.2.2 Experimental Results

To examine the effect of an additional angular velocity measurement we compared various tracking algorithms containing position-only measurements with those containing single and dual orthogonal velocity measurements. In the following we describe our simulated results using the SOT frameworks above, along with our measured-simulated results using the RFS MOT methods.

5.2.2.1 Simulated Results

The motion of a single point object was simulated in MATLAB using a constant-turn (CT) model as shown in Fig. 5.11. The state vector was

$$\mathbf{x} = [x, y, v, \phi, \Omega]^T \tag{5.7}$$

containing the target x and y positions, its velocity v , heading ϕ , and turn-rate Ω in the local-frame. The model noise was chosen to be $\sigma_v = 0.1$ m/s and $\sigma_\Omega = \frac{\pi}{360}$ rad/s. We assume uniform clutter and in this realization the clutter intensity is $\lambda = 10$, shown in cyan over the entire simulation time.

To test the tracking performance we examined the RMSE of trackers using three different measurement models. The first model contained measurements only of range and angle (RA), the second contained range, angle, and Doppler (RAD), and the third contained range, angle, Doppler, and angular velocity (RADW). Clearly, the addition of any independent measurement should improve the tracking error. Of interest though is how the clutter intensity affects this improvement and what level of angular velocity accuracy is required for a given improvement in tracking performance. The RADW measurement vector was

$$\mathbf{z}_{RADW} = [R, \theta, v_D, \omega]^T \quad (5.8)$$

containing the range, angle, Doppler velocity, and angular velocity in the global-frame. The RA and RAD measurement models were the same, except without their respective additional measurements. The measurement model noise components were the same for all models and chosen to be $\sigma_R = 5$ m/s, $\sigma_\theta = \frac{\pi}{180}$ rad., $\sigma_{v_D} = 1$ m/s, and $\sigma_\omega = \frac{\pi}{180}$ rad/s. Since the motion model and measurement model were both non-linear their Jacobians were calculated and used to perform the Extended Kalman Filter (EKF) prediction and update steps.

The sensor model contains parameters for the probability of detection, P_D , the clutter range, and the clutter rate. The clutter range values were chosen to be the hypothetical minimum and maximum values detectable by the sensor which for range, angle, Doppler, and angular velocity were chosen to be ± 1000 m, $\pm \pi$ rad, ± 100 m/s, and $\pm 2\pi$ rad/s. The clutter rate determines the expected value of the number of clutter detections per measurement.

We first simulated a point target with a perfect sensor ($P_D = 1.0$) and zero-clutter. The average RMSE of the state vector is shown for the three measurement models over 10,000 independent realizations in the top row of Table 5.2. As expected, the additional independent velocity measurements improve the tracking error. We see a reduction of around 10% and 4% for the Doppler and angular velocity measurements respectively. Next, we added clutter with $\lambda = 50$ and $P_D = 0.85$. The results are shown in Table 5.2 for the three measurement models using the nearest neighbor (NN), probabilistic data association (PDA), and Gaussian sum filter (GSF) algorithms. We see that at this particular P_D and clutter rate the improvements

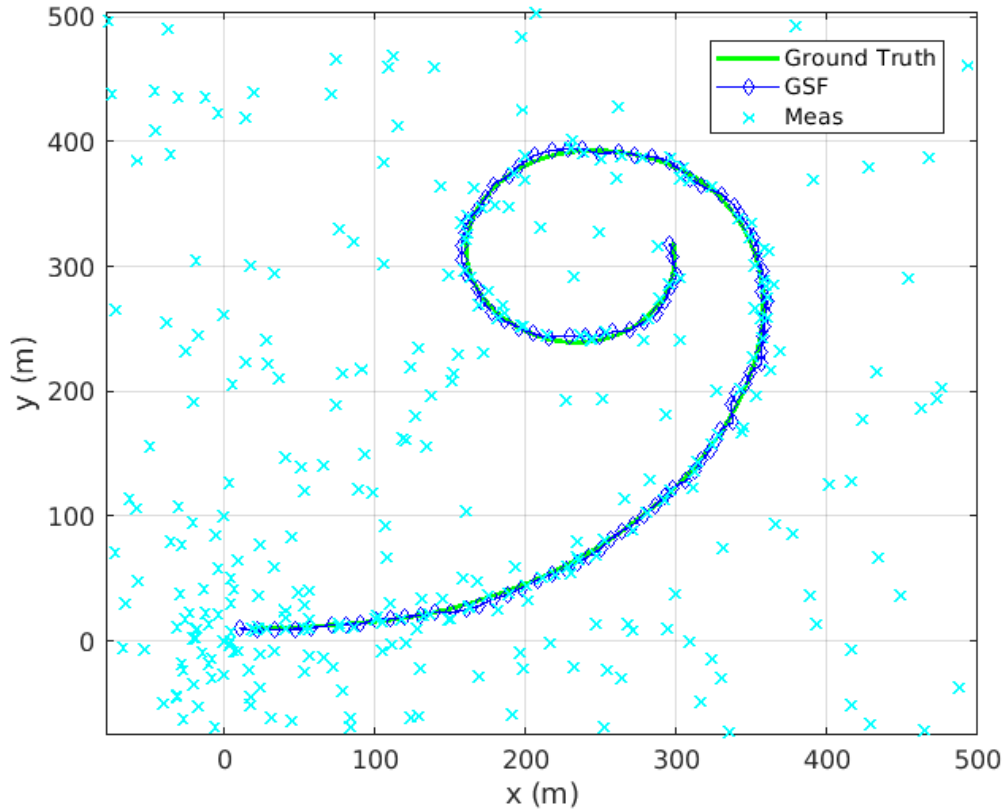


Figure 5.11: A single realization of a point-target with a constant-turn (CT) motion model with model noise of $\sigma_v = 0.1$ m/s and $\sigma_\Omega = \frac{\pi}{360}$ rad/s, and clutter intensity of $\lambda = 10$. The ground-truth position is shown in green along with the estimated position from the Gaussian sum filter (GSF). Image [6] © 2020 IEEE.

Table 5.2: Average RMSE over 10k realizations

Type:	RA (m):	RAD (m):	RADW (m):
No Clutter	2.19	1.97	1.90
NN	3.48	2.40	2.12
PDA	2.59	2.15	2.11
GSF	2.51	2.16	2.10

are similar. For the GSF we gain around 14% and 3% from the additional radial and angular measurements respectively.

To view the effect of the measurements at various clutter rates we swept the intensity from 0 to 1000, performing 1k realizations at each clutter rate and taking the average




Figure 5.12: The average RMSE of the state vector for each measurement model at various clutter rates. The value is taken as the average over 1k realizations, and a linear least-squares fit is shown to highlight the trend. Image [6] © 2020 IEEE.

RMSE of the state vector. The results are shown in Fig. 5.12 with a linear least-squares fit to highlight the overall trend. Interestingly, we see that as the clutter rate increases the effect of the additional angular velocity measurement also increases. For example, at a clutter rate of 100 detections the angular velocity measurement improves the RMSE by only 2%, while at 800 detections it improves by 7%.

5.2.2.2 nuScenes Dataset Results

To gain additional insight we performed a similar analysis using the nuScenes dataset [81] from Aptiv. There are many large-scale publicly available autonomous driving datasets but nuScenes is the only dataset that provides radar data, to the best of our knowledge. The dataset includes 1,000 driving scenes from Boston and Singapore of approximately 20 seconds

each. It provides sensor data from a single 32 line lidar, 6 cameras, 5 Continental ARS 408-21 radars, and controller-area-network (CAN) messages containing wheel-odometry and IMU data. Ground-truth bounding box labels are provided for various attributes, including vehicles and pedestrians.

The ARS 408-21 are 77 GHz frequency modulated continuous wave (FMCW) radars with far and near sensing regions of $\pm 9^\circ$ at 250m and $\pm 60^\circ$ at 70m respectively. They perform platform motion compensation by accepting vehicle longitudinal velocity and yaw-rate messages, and group detected clusters into dynamic categories such as moving, stationary, or oncoming. The raw target velocity values and motion compensated values are reported as x and y velocity components in the sensor frame, most likely to provide an orthogonal basis for adding compensation velocity components.

The ARS 408-21 radar modules are not designed to perform interferometric angular velocity estimation, and so we synthesized the measurement for detections falling within a ground-truth bounding-box using the box velocity, which is calculated as a central difference over time. This is shown in Fig. 5.13 where a tracked vehicle in the local-frame has 2 radar detections in blue near its rear bumper, shown with the x-components of its Doppler velocity. The ground-truth box velocity components are shown along with the total box velocity in green. To synthesize the angular velocity measurement we simply add the y-component from the box velocity plus Gaussian noise with $\sigma_y = 0.5$ m/s to the y component of each radar detection.

We make a few major simplifying assumptions, the first is that each detection, no matter where it is spatially on the target, will have the same bulk angular velocity, which is generally not the case. Second, to compare the results with the simulated measurement models we must take the x-component to represent the bulk of the Doppler component and the y-component to represent the angular velocity component. This is somewhat less of a leap because in general the raw velocity components for *moving* targets from the ARS 408-21 modules report only an x-component, so it is likely the radial Doppler velocity is being




Figure 5.13: The ground-truth bounding box is shown in green along with its box velocity. The vehicle is slowing down and turning so in the local-frame the velocity is up and towards our vehicle which is shown with its base-link on the coordinate axes. The radar detections shown in blue are appended with the angular component of the ground-truth box velocity $v_{y,gt}$, plus Gaussian noise to emulate an angular velocity radar measurement. Image [6] © 2020 IEEE.

projected onto the x-axis directly in the radar firmware.

For evaluation we selected a scene from the dataset where a target vehicle was directly in front of the Aptiv vehicle for the whole duration of the scene. This is shown as the green bounding box in Fig. 5.14. The lidar scans are shown just for reference and the radar detections and velocity components are shown in blue. Notice that the radar detection that is falling within the box has an angular component equal to the box velocity. Most of the other detections are static objects and thus have angular components equal to the the projected relative yaw of the base vehicle. The state estimates and covariances from the GM-PHD filter are shown in magenta. The large light magenta lines are the covariances for the birth models, which are where new targets are most likely to appear. In this scenario we achieved an RMSE reduction of 13.1% and 5.8% as shown in Table 5.3, similar to the results shown in simulation.

In this scenario the results favorable, though we were able to find examples in which the results were less conclusive. This is, most notably, due to the point-object assumption

Table 5.3: State RMSE for target vehicle in scene-0061

Measurement Model:	RMSE (m):	Reduction (%):
Position only	0.99	-
Pos + v_x	0.86	13.1
Pos + v_x, v_y	0.81	5.8




Figure 5.14: The tracked vehicle shown with green annotation and the magenta state estimates with covariance. The ground-truth position was chosen to be the center of the rear edge of the box to represent the rear bumper where most of the radar detections hit. The large magenta lines represent the birth model covariances where new objects are most likely to appear. Image [6] © 2020 IEEE.

that the tracking algorithms make. This means that a target can generate at most one object detection at a time. For high-resolution sensors like lidar and automotive radar this is clearly not the case and so the state estimates move around the object spatially depending on where a particular detection strikes the object. This is random in some sense and thus

a model which takes into account the shape of the target is needed, for example by adding shape parameters to the underlying state vector and modifying the measurement likelihood and innovation calculations [82, 83].

5.3 Hand Gesture Recognition

Hand gesture recognition has been investigated thoroughly in the literature using many different sensor forms, including camera, lidar, sonar, and radar [84, 85, 86, 87]. Various trade-offs exist for each modality. For example, camera based systems require a source of lighting, lidar systems may be costly, and sonar systems may lack sufficient resolution. Millimeter-wave radar sensors provide a balance by being relatively low cost compared to lidar, while having sufficient Doppler resolution at millimeter-wave frequencies. These systems can operate in any lighting condition and do not have the privacy issues associated with cameras. In [87] a framework for radar based hand gesture sensing was described, which used range-Doppler maps from a custom 60 GHz antenna-in-package (AiP) chip. Around that time, the first radar hand gesture classification based solely on micro-Doppler was performed in [88]. Because Doppler radars measure only radial velocity, gestures performed at wide angles away from the broadside tend to reduce the classification accuracy [88]. Interferometric radar provides an interesting approach to this problem by allowing the direct measurement of angular velocity as well as radial velocity [22]. In [30] interferometric hand gesture classification performance was demonstrated using a support vector machine (SVM) classifier. It was found that adding the interferometric responses to the classifier input provided an average of 0.93% classification improvement over four measurement angles. The largest improvement was obtained at angles of 45° off-broadside where the improvement was 2.3%.

In this section we demonstrate classification results from a three-element, 40-GHz radar system containing two orthogonal interferometric baselines. Compared to the single interferometric baseline used in [30], our approach measures angular motion in both the azimuth and elevation dimensions, as well as radial motion, providing a three-dimensional

velocity measurement. Our classifier is based on a convolutional neural network (CNN), which has been shown to be effective for many types of classification problems [40]. The dual-axis millimeter-wave radar system provides five measurement outputs, which includes three micro-Doppler measurements and two interferometric measurements. In this section we explore various combinations of these radar modes to evaluate the relative impact of including the interferometric mode on the classification accuracy.

5.3.1 Experimental Measurements

The hardware used in this work was the same single-board 40 GHz radar described previously. The hand gestures used for classification are shown in Fig. 5.15, where each column is a different gesture and the temporal sequence is shown from top to bottom. Qualitatively, the gestures were: 1) swipe down then up, 2) swipe right then back, 3) turn dial down then up, 4) fist push out and back. We also recorded a control gesture in which no hand was present. The configuration of the antennas and the measurement locations are shown in Fig. 5.16. Each hand gesture was performed a total of 100 times each, at four different locations, broadside, $26 \pm 5^\circ$, $45 \pm 5^\circ$, and $56 \pm 5^\circ$ off broadside. The radar was then rotated $45 \pm 5^\circ$ and the measurements were repeated. The purpose of rotating the radar was to evaluate the efficacy of the dual interferometric baselines. A system with a single baseline can only measure angular velocity in a single dimension, thus the dual baseline system should perform better at arbitrary radar rotation angles. The measurements were performed over a period of two days for total of 4000 individual gestures, including the control gesture.

For processing, each 1.2s signal was first high-pass filtered with a 4th order Butterworth filter with a cutoff frequency of 3 Hz, then converted to the time-frequency domain by performing a short-time Fourier transform (STFT) with a Gaussian window and FFT size of 512 points. The overlap was 496 samples, forming an image of 512x569 for each gesture. The first 28 samples and last 29 samples were removed to form an even 512x512 image. The time-frequency response of a single realization for each of the gestures is shown in Fig. 5.17.




Figure 5.15: The four hand gestures used in this work. The temporal sequence is shown top to bottom for each of the gestures. Notice that each gesture ends in the same position as it begins; this was done to simplify processing the data. The fifth gesture is the control in which no hand was present.

5.3.2 CNN Classifier

Convolutional neural networks have found great success in many practical image processing applications due to their invariance to translations, illumination, and other types of noise [40]. Here we applied a CNN to the time-frequency responses of the radar output from multiple Doppler and interferometer responses and compared classification performance for four specific measurement cases. The first case used only the Doppler response from two receivers, the second case used these same responses plus the interferometric response from the baseline formed by the two receiver antennas. Case three used the Doppler only response from all three RX antennas, while case four used these responses plus the two interferometric

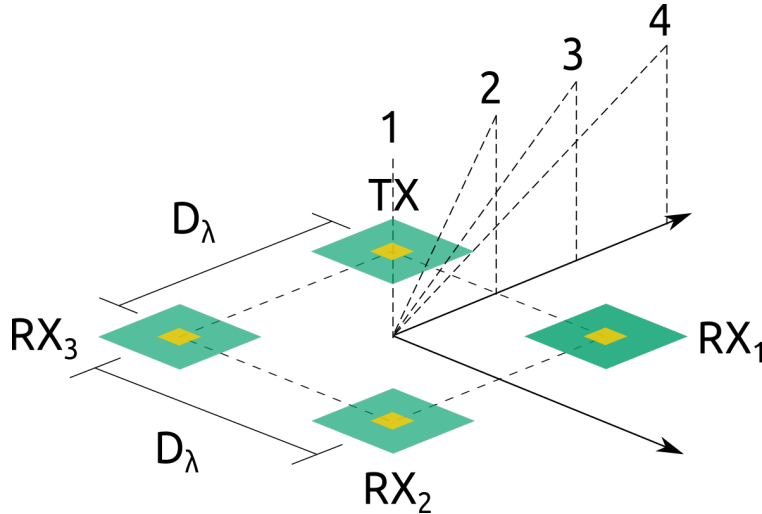


Figure 5.16: The radar antenna locations which form a grid separated by baseline $D_\lambda = 8$. The hand locations were measured at a height of 0.15 ± 0.05 m from the plane of the antennas, at angles $0 \pm 5^\circ$, $26 \pm 5^\circ$, $45 \pm 5^\circ$, and $56 \pm 5^\circ$.

responses from the two orthogonal baselines. The case of using only the interferometric responses without Doppler was examined, and the classification accuracy was only 85-90%, and so this case was disregarded. Data was split into subsets of 70% for training, and 15% for both validation and test. Stratified random sampling was used so that an equal proportion of all gestures and positions were present in each group.

The network topology is shown in Fig. 5.18, which shows five input images. The number of images in the first two layers changes based on the case being tested. The number of images in the second layer was chosen to be three times the number of input images. The section of the network in the dashed line remains the same between all four cases. The output classification is performed by 3 fully connected linear layers. For each case the network was trained with the cross entropy loss function using stochastic gradient descent, with a learning rate of 0.001 and a momentum parameter of 0.9 [74]. The network was implemented using PyTorch [89] on an Nvidia GeForce RTX 2080 GPU.

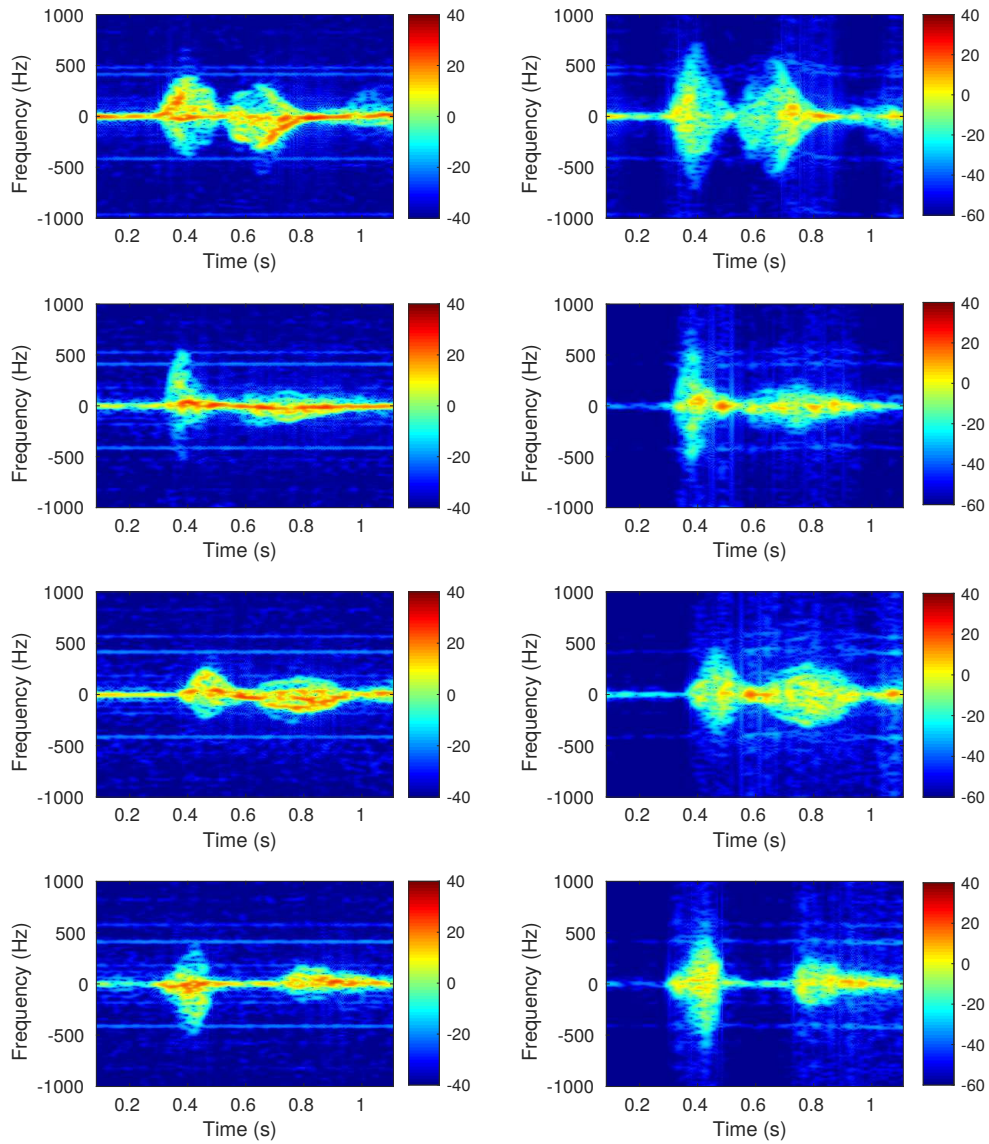


Figure 5.17: The time frequency responses for gestures one to four from top to bottom at the broadside location, the Doppler response is shown (left), and the interferometric response (right).

5.3.3 Classification Results

The classification accuracy for each of the four cases is shown in Table 5.4 for the 0° and 45° radar rotation angles, along with the total over both rotations. One interesting result is that

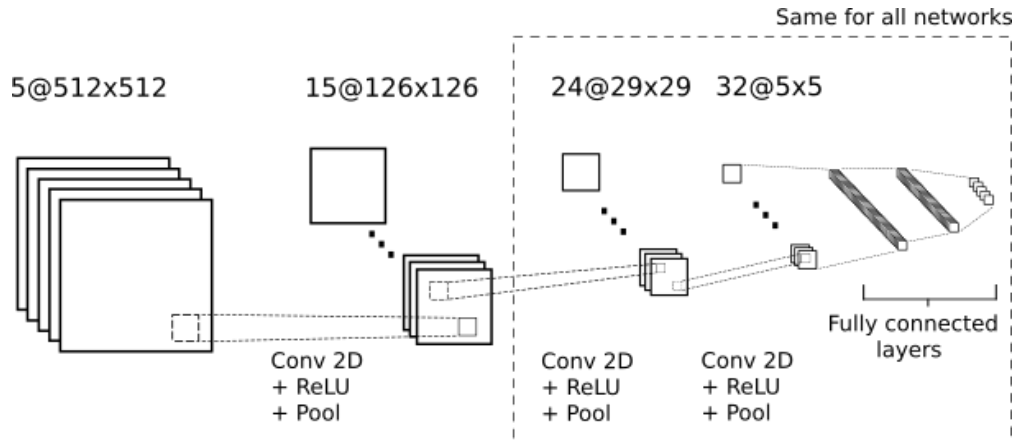


Figure 5.18: The topology of the CNN used in this work, each layer comprises a 2D convolution, rectified linear unit (ReLU), and max pooling. The outlined portion of the network remains the same for all of the cases being tested.

case two, with a single interferometric baseline, outperforms case three, with three micro-Doppler responses. This is a direct result of case two performing better on the measurements taken at the 45° rotation angle, where it achieved 99.33% accuracy vs. 97.33% for case three. In all situations case four, with three micro-Doppler responses and two interferometric responses achieves the highest classification accuracy. As expected, case four with dual baselines outperforms case two with a single baseline, though the improvement comes at the 0° radar rotation angle. It is interesting to note that, actually, all of the information in the interferometric response is contained in its two corresponding Doppler responses, since the interferometric response is simply the cross-correlation of the Doppler responses in the frequency domain. Ideally, the CNN should be able to learn and exploit the interferometric relationship from only the Doppler responses. In practice though, there is no guarantee that the optimizer will be able to learn such weights, thus the more signal, or explicit guidance, we give the network the better it will perform [90]. In the cases where the interferometric responses are given, the network gets extra signal, so learning is easier and accuracy is improved.

In [88] it was determined that the performance of Doppler only hand gesture classification varies with aspect angle. As the hand moves away from broadside, there is typically

Table 5.4: Classification accuracy over all aspect angles for 0° and 45° radar rotation

Network:	0° Rot:	45° Rot:	Total:
Case 1: 2uD	98.67%	98.67%	98.67%
Case 2: 2uD, 1int	99.00%	99.33%	99.17%
Case 3: 3uD	99.33%	98.33%	98.83%
Case 4: 3uD, 2int	100.00%	99.33%	99.67%

Table 5.5: Classification accuracy over 0° and 45° radar rotation for all aspect angles

Network:	0°:	26°:	45°:	56°:	Total:
Case 1: 2uD	98.67%	100.00%	98.67%	97.33%	98.67%
Case 2: 2uD, 1int	98.67%	100.00%	98.67%	99.47%	99.20%
Case 3: 3uD	99.47%	99.33%	99.33%	97.60%	98.93%
Case 4: 3uD, 2int	100.00%	100.00%	99.33%	99.47%	99.70%

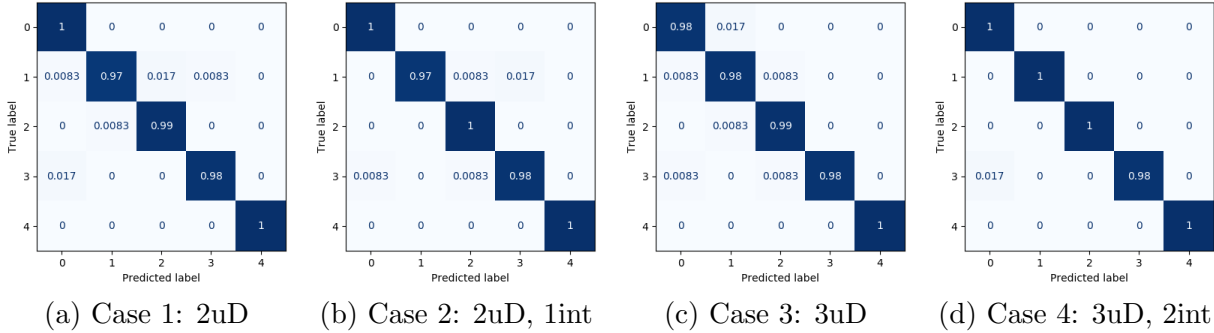


Figure 5.19: The confusion matrices for each of the four cases over all rotation and aspect angles, note that here the gestures are zero-indexed and thus gesture four is the control gesture.

less radial motion and thus less Doppler response, making classification more difficult. We addressed this issue in [91] where interferometric measurements which persisted in angle compared to Doppler only measurements were presented. We found later that for this phenomenon to hold, the radar must have some non-zero intermediate frequency (IF) [92]. Simply put, in a direct downconversion architecture, there must be some Doppler shift present in order to provide the carrier for which to perform phase-based interferometric processing. Nevertheless, the interferometric mode still yields better classification performance than

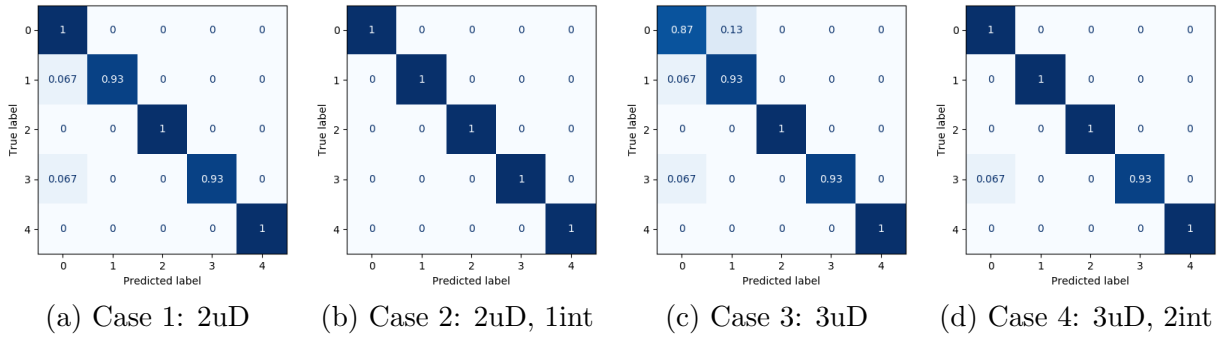


Figure 5.20: The confusion matrices for what we consider the hardest case, a radar rotation angle of 45° and an aspect angle of 56° . Here we see the strength of the interferometric systems where case two perfectly classifies all gestures, and case four confuses only a single gesture.

those cases where the interferometer is not used.

Classification results over both radar rotation angles for the various aspect angles are shown in Table 5.5. Both interferometric systems perform well even at angles of 56° off broadside. Our system uses direct downconversion, thus the performance comes even without the benefit of getting added interferometric signal persistence in the case of a non-zero IF. The confusion matrices for each case over all radar rotation and aspect angles are shown in Fig. 5.19. It is clear that system four, having three micro-Doppler responses and two interferometric responses is the least likely to misclassify a hand gesture. In Fig. 5.20 we provide the confusion matrices for what we consider to be the hardest case, which is a 45° radar rotation and a 56° aspect angle. We see that both interferometric systems outperform their Doppler only counterparts. In this instance case two perfectly classifies all gestures, while case four misclassifies only a single gesture.

CHAPTER 6

CONCLUSION

The fundamental questions this work sought to answer were: how practical are either of the proposed signal processing methods at jointly estimating the angular velocities of multiple targets, and are these methods scalable as the number of targets grows large? The results presented here confirm that in practice we are able to estimate the angular velocities of two targets with reasonable accuracy provided that the individual target responses do not overlap in the frequency domain. We also concluded that techniques which use the complete interferometric response (model-based methods) will not scale as the number of targets grows large. Doppler decomposition does become more challenging as the number of targets grows large, though more advanced methods of signal decomposition may be used. Similarly, the issue of data association becomes more challenging, though this is an area which is richly described in the literature.

6.1 Contributions

The presentation of novel, signal processing based interferometric distortion mitigation represents a significant contribution towards enabling multitarget angular velocity estimation with a distributed receiver. Specifically, we presented two classes of methods, one which operates on the complete interferometric signal with model-based methods, and a second which decomposes each antenna signal into individual target responses, which we refer to as response decomposition. Both methods provide valid techniques to detect multitarget angular velocity, though we find that the quadratic scaling of distortion terms makes model-based methods intractable as the number of targets grows large. Due to this we provide experimental results for a two target case using the response decomposition method.

Critical to the development of these methods was the derivation of the complete signal model of an interferometric radar to a general target consisting of N independent

scattering centers. The complete exact signal response had never been described in this manner, prior to this, only the response of two targets in the far-field to the array had been described. With this, we are able to compute the exact spectral response at any target location and compare this response to other more simplified models with varying degrees of approximations. Having a complete picture of the interferometric signal response is critical for future work, and the general N point target is extremely versatile, due to the fact that many targets may be represented as a summation of independent point scatterers.

We also presented a custom, compact, millimeter-wave interferometric radar which was developed using state-of-the-art microwave design tools. This system was developed using commercially available RF IC components. We were able to demonstrate multiple practical applications for interferometric angular velocity estimation, including vehicle velocity estimation, multiobject tracking, and hand gesture recognition. We hope that the presentation of our radar hardware design will motivate other researchers to build and design their own custom radars for their own individual applications.

6.2 Future Work

The most challenging issue associated with multitarget interferometric angular velocity estimation is the issue of intermodulation distortion. Although we presented promising results, this was performed in the case of only two targets present, which is the simplest case of interferometric distortion mitigation. Within the two target case, we looked at various trajectory cases, and in particular the difficult case in which the targets are not easily separable in the frequency domain. As the number of targets grows the likelihood of target overlap will increase. One area for future work will be to look at more advanced methods for separating the responses due to individual targets. The most common basis set we use is that of complex exponentials, which results in the Fourier transform. Though, there are many other methods which use other basis sets, such as the wavelet transform, empirical mode decomposition, orthogonal matching pursuit, and machine learning among others [93, 38, 94, 40].

As the number of targets increases the sophistication of the signal decomposition portion of the algorithm must grow as well, which represents an interesting and challenging area for future work.

We found that due to the relatively narrow system baseline and low angular velocity values we were able to use the same target filters on the both the antenna one and antenna two responses. This circumvents the need for data association methods, because we know exactly which filter we are applying to each antenna response. In the general case, we must perform separate detection and decomposition on each antenna response. Thus, another area for future work is to focus more on data association techniques. Luckily, there is much work in the MTT literature on various association methods and these should be able to be applied with relatively simple modifications. More sophisticated methods may also be applied if we are willing to assume that the targets we are measuring adhere to some common motion models. In this manner, we should be able to predict where in frequency, or other basis, the target should be located with respect to each antenna signal response.

The interferometric angular velocity measurements performed in this work were obtained using a CW transmitter. These methods and techniques are still applicable while using more complex waveforms. Future work might include performing joint measurement methods with FMCW waveforms. A system, in this respect, should be capable of jointly measuring range, Doppler, angle, and angular velocity, forming a complete basis of measurements in both position and velocity. There may be various accuracy and resolution trade-offs while using more complex waveforms, and so this is something that could be explored in the future from both a theoretical and practical perspective.

In summary, this work made a significant contribution towards enabling multitarget angular velocity measurements using a correlation interferometer. The methods described should provide a framework on top of which more complex algorithms may be developed. Interferometric angular velocity estimation holds extreme promise, and should be considered the fourth basic type of radar measurement. Indeed, forming a complete set of independent

radar measurements in position and velocity is of major importance for safety critical systems which require sub-centimeter position accuracy.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] C. Gu and J. Lien, *Short-Range Micro-Motion Sensing with Radar Technology*, ser. Control, Robotics and Sensors. Institution of Engineering & Technology, 2019.
- [2] E. Klinefelter, J. Merlo, R. Radha, and J. A. Nanzer, “A point target model for interferometric radar angular velocity estimation,” *IEEE Trans. Microw. Theory Techn.*, TMTT-2021-08-1038.R1, Accepted for Publication.
- [3] E. Klinefelter and J. A. Nanzer, “Automotive velocity sensing using millimeter-wave interferometric radar,” *IEEE Trans. Microw. Theory Techn.*, vol. 69, no. 1, pp. 1096–1104, Jan. 2020.
- [4] E. Klinefelter and J. A. Nanzer, “Millimeter-wave interferometric radar for speed-over-ground estimation,” in *2020 IEEE MTT-S IMS*, Jun. 2020, pp. 1–4.
- [5] —, “Automotive velocity sensing using millimeter-wave interferometric radar,” *IEEE Trans. Microw. Theory Techn.*, vol. 69, no. 1, pp. 1096–1104, Jan. 2021.
- [6] E. Klinefelter, J. A. Nanzer, and H. Radha, “Radar tracking with orthogonal velocity measurements for autonomous ground vehicles,” in *2020 IEEE Radar Conference (RadarConf20)*, 2020, pp. 1–5.
- [7] S. Shalev-Shwartz, S. Shammah, and A. Shashua, “On a formal model of safe and scalable self-driving cars,” *CoRR*, vol. abs/1708.06374, 2017.
- [8] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [9] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. III, R. Mahler, and B.-T. Vo, “Multitarget tracking,” *Wiley Encyclopedia*, pp. 1–25, 09 2015.
- [10] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011. [Online]. Available: <https://books.google.com/books?id=2aOiuAAACAAJ>
- [11] S. Blackman and R. Popoli, *Design and analysis of modern tracking systems*, ser. Artech House radar library. Norwood, MA: Artech House, 1999.
- [12] S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans, *Fundamentals of Object Tracking*. Cambridge University Press, 2011.
- [13] W. Koch, *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*, 10 2013, vol. 8.
- [14] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA: Artech House, Inc., 2007.

- [15] J. A. Nanzer, “Millimeter-wave interferometric angular velocity detection,” *IEEE Trans. Microw. Theory Techn.*, vol. 58, no. 12, pp. 4128–4136, Dec. 2010.
- [16] H. Crew, *The Wave Theory of Light - Memoirs by Huygens, Young and Fresnel*. Read Books, 2009. [Online]. Available: <https://books.google.com/books?id=ftlipVx47cUC>
- [17] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Pergamon Press, 1984, vol. 6.
- [18] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and Synthesis in Radio Astronomy*. Hoboken, NJ: John Wiley and Sons, 2001.
- [19] H. A. Zebker and R. M. Goldstein, “Topographic mapping from interferometric synthetic aperture radar observations,” vol. 91, 05 1986.
- [20] R. Wohlleben, H. Mattes, and T. Krickbaum, *Interferometry in Radioastronomy and Radar Techniques*. Dordrecht, Netherlands: Kluwer Academic, 1991.
- [21] A. Doerry, B. Milesosky, and D. Bickel, “Tangential velocity measurement using interferometric mti radar,” *Sandia National Laboratories Technical Report*, Nov. 2002.
- [22] J. A. Nanzer and K. S. Zilevu, “Dual interferometric-doppler measurements of the radial and angular velocity of humans,” *IEEE Trans. Antennas Propag.*, vol. 62, no. 3, pp. 1513–1517, Mar. 2014.
- [23] X. Wang, P. Wang, and V. C. Chen, “Simultaneous measurement of radial and transversal velocities using a dual-frequency interferometric radar,” in *2019 IEEE RadarConf*, Apr. 22-26, 2019, pp. 1–6.
- [24] —, “Simultaneous measurement of radial and transversal velocities using interferometric radar,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 4, pp. 3080–3098, Aug. 2019.
- [25] V. C. Chen, *The micro-Doppler effect in radar*. Norwood, MA: Artech House, 2019.
- [26] X. Wang, P. Wang, X. Cao, and V. C. Chen, “Interferometric angular velocity measurement of rotating blades: theoretical analysis, modeling and simulation study,” *IET Radar, Sonar & Navig.*, vol. 13, no. 3, pp. 438–444, Nov. 2018.
- [27] M. Jian, Z. Lu, and V. C. Chen, “Experimental study on radar micro-doppler signatures of unmanned aerial vehicles,” in *2017 IEEE RadarConf*, May 8-12, 2017, pp. 0854–0857.
- [28] X. Wang, H. Liang, and P. Wang, “Detection and tracking of uavs using interferometric radar,” in *2019 Int. RADAR*, Sept. 23-27, 2019, pp. 1–6.
- [29] Q. Liu, V. Fusco, D. Linton, and D. Zelenchuk, “Interferometer for co-operating target motion detection,” in *2014 EuCAP*, Apr. 6-11, 2014, pp. 2935–2937.
- [30] H. Liang, X. Wang, M. S. Greco, and F. Gini, “Enhanced hand gesture recognition using continuous wave interferometric radar,” in *2020 IEEE Int. RADAR*, Apr. 28-30 2020, pp. 226–231.

- [31] J. Merlo, E. Klinefelter, and J. A. Nanzer, “A dual-axis interferometric radar for instantaneous 2d angular velocity measurement,” in *2020 IEEE International Symposium on Antennas and Propagation and North American Radio Science Meeting*, 2020, pp. 1661–1662.
- [32] J. A. Nanzer and K. S. Zilevu, “Distortion mitigation in interferometric angular velocity measurements,” *Electron Lett.*, vol. 50, no. 18, pp. 1316–1318, Aug. 2014.
- [33] S. Ishtiaq, X. Wang, and S. Hassan, “Detection and tracking of multiple targets using dual-frequency interferometric radar,” in *IET International Radar Conference (IET IRC 2020)*, vol. 2020, 2020, pp. 468–475.
- [34] P. Wang, H. Liang, X. Wang, and E. Aboutanios, “Transversal velocity measurement of multiple targets based on spatial interferometric averaging,” in *2020 IEEE Int. RADAR*, 2020, pp. 709–713.
- [35] J. Merlo, E. Klinefelter, S. Vakalis, and J. A. Nanzer, “A multiple baseline interferometric radar for multiple target angular velocity measurement,” *IEEE MWCL*, 2021.
- [36] J. A. Nanzer, “On the resolution of the interferometric measurement of the angular velocity of moving objects,” *IEEE Trans. Antennas Propag.*, vol. 60, no. 11, pp. 5356–5363, Nov. 2012.
- [37] V. C. Chen and H. Ling, *Time-Frequency Transforms for Radar Imaging and Signal Analysis*. Norwood, MA: Artech House, 2002.
- [38] L. Cohen, *Time-Frequency Analysis: Theory and Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1995.
- [39] E. Klinefelter and J. A. Nanzer, “Signal processing based distortion mitigation in interferometric angular velocity estimation,” *IEEE Trans. Microw. Theory Techn.*, Submitted for Publication.
- [40] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016, <http://www.deeplearningbook.org>.
- [41] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
- [42] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, Mar 1986.
- [43] R. Roy and T. Kailath, “Esprit-estimation of signal parameters via rotational invariance techniques,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, Jul 1989.
- [44] M. I. Skolnik, *Introduction to Radar Systems*. New York, NY, USA: McGraw-Hill, 2001.
- [45] N. Levanon and E. Mozeson, *Radar Signals*. New York, NY, USA: Wiley-IEEE, 2004.

- [46] Y. Bar-Shalom and X. Rong-Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, 1995.
- [47] D. Pozar, *Microwave Engineering, 4th Edition*. Hoboken, NJ: Wiley, 2011.
- [48] C. A. Balanis, *Antenna Theory*. New York, NY, USA: Wiley-Interscience, 2012.
- [49] E. Klinefelter and J. A. Nanzer, “Hand gesture recognition using a dual axis millimeter-wave interferometric-doppler radar and convolutional neural networks,” *IEEE 2021 EuMW*, Accepted for Publication.
- [50] S. Liu, L. Li, J. Tang, S. Wu, and J.-L. Gaudiot, *Creating Autonomous Vehicle Systems*. San Rafael, CA, USA: Morgan & Claypool, 2017.
- [51] Liang Chu, Yongsheng Zhang, Yanru Shi, Mingfa Xu, and Minghui Liu, “Vehicle lateral and longitudinal velocity estimation based on unscented kalman filter,” in *2010 2nd ICETC*, vol. 3, Jun. 2010, pp. V3–427–V3–432.
- [52] D. Nistér, O. Naroditsky, and J. R. Bergen, “Visual odometry,” *Proc. 2004 IEEE Comp. Soc. Conf. CVPR*, vol. 1, pp. 652–659, 2004.
- [53] M. Weydert, “Model-based ego-motion and vehicle parameter estimation using visual odometry,” in *2012 16th IEEE MEC*, Mar. 2012, pp. 914–919.
- [54] J. Zhang and S. Singh, “Loam: Lidar odometry and mapping in real-time,” in *Robotics: Sci. Sys. Conf.*, Jul. 2014.
- [55] K. Zindler, N. Geiß, K. Doll, and S. Heinlein, “Real-time ego-motion estimation using lidar and a vehicle model based extended kalman filter,” in *17th Int. IEEE Conf. ITSC*, Oct 2014, pp. 431–438.
- [56] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, “Instantaneous ego-motion estimation using multiple doppler radars,” in *2014 IEEE ICRA*, May 2014, pp. 1592–1597.
- [57] S. H. Cen and P. Newman, “Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions,” in *2018 IEEE ICRA*, May 2018, pp. 1–8.
- [58] Qifa Ke and T. Kanade, “Transforming camera geometry to a virtual downward-looking camera: robust ego-motion estimation and ground-layer detection,” in *Proc. 2003 IEEE Comp. Soc. Conf. CVPR*, vol. 1, Jun. 2003, pp. I–I.
- [59] W. Kleinhempel, “Automobile doppler speedometer,” in *Proc. 1993 VNIS Conf.*, Oct 1993, pp. 509–512.
- [60] P. D. L. Beasley, I. M. Simmons, and A. G. Stove, “High accuracy radar speedometer,” in *IEE Coll. on Auto. Sens.*, May 1992, pp. 8/1–8/5.
- [61] T. Kimura and H. Kobayashi, “An optimum calculation method for vehicle ground velocity measurement using the doppler effect of ultrasonic waves,” in *Proceedings 1991 IECON*, Oct 1991, pp. 2390–2393 vol.3.

- [62] C. Xu, L. Daniel, E. Hoare, V. Sizov, and M. Cherniakov, “Comparison of speed over ground estimation using acoustic and radar doppler sensors,” in *2014 11th Eur. Radar Conf.*, Oct 2014, pp. 189–192.
- [63] P. Beckmann and A. Spizzichino, *The Scattering of Electromagnetic Waves from Rough Surfaces*. New York, NY, USA: The Macmillan Company, 1963.
- [64] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing: Active and Passive*. Norwood, MA, USA: Artech House, 1986, vol. II.
- [65] W. Stiles and F. Ulaby, *Backscatter Response of Roads and Roadside Surfaces*. Sandia Laboratories, 1979.
- [66] M. D. Sharp and J. A. Nanzer, “Accuracy of angle rate measurements using a distributed radar with a correlation receiver,” *IEEE Antennas Wireless Propag. Lett.*, vol. 17, no. 2, pp. 209–212, Feb 2018.
- [67] D. Simon, *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches*. New York, NY, USA: Wiley-Interscience, 2006.
- [68] E. A. Wan and R. Van Der Merwe, “The unscented kalman filter for nonlinear estimation,” in *Proc. 2000 IEEE AS-SPCC*, 2000, pp. 153–158.
- [69] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. Cambridge, MA, USA: The MIT Press, 2005.
- [70] B. Rivet, L. Girin, and C. Jutten, “Log-rayleigh distribution: A simple and efficient statistical representation of log-spectral coefficients,” *IEEE Audio, Speech, Language Process.*, vol. 15, no. 3, pp. 796–802, Mar. 2007.
- [71] R. Zanetti and K. DeMars, “A comparison of linear attitude estimators,” in *2018 Space Flight Mech. Meet.*, Jan. 2018.
- [72] R. Schubert, C. Adam, M. Obst, N. Mattern, V. Leonhardt, and G. Wanielik, “Empirical evaluation of vehicular models for ego motion estimation,” in *2011 IEEE IV Symp.*, Jun. 2011, pp. 534–539.
- [73] J. E. Stellet, F. Straub, J. Schumacher, W. Branz, and J. M. Zöllner, “Estimating the process noise variance for vehicle motion models,” in *2015 IEEE 18th Int. Conf. ITS*, Sep. 2015, pp. 1512–1519.
- [74] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, “On the importance of initialization and momentum in deep learning,” in *Proc. 30th Int. Conf. ML*, vol. 28, no. 3, 2013, pp. 1139–1147.
- [75] W. S. Cleveland, “Robust locally weighted regression and smoothing scatterplots,” *Journ. American Stat. Assoc.*, vol. 74, no. 368, pp. 829–836, 1979.
- [76] M. A. Richards, J. A. Scheer, and W. A. Holm, *Principles of Modern Radar: Basic Principles*. Institution of Engineering and Technology, 2010.

- [77] L. Svensson and K. Ganstrom, “Multi-object tracking for automotive systems.” [Online]. Available: www.edX.org
- [78] I. R. Goodman, R. P. Mahler, and H. T. Nguyen, *Mathematics of Data Fusion*. USA: Kluwer Academic Publishers, 1997.
- [79] E. Brekke and M. Chitre, “Relationship between finite set statistics and the multiple hypothesis tracker,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1902–1917, Aug 2018.
- [80] B.-N. Vo and W.-K. Ma, “The Gaussian mixture probability hypothesis density filter,” *IEEE Transactions on Signal Processing*, vol. 54, pp. 4091 – 4104, 12 2006.
- [81] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusenes: A multimodal dataset for autonomous driving,” *arXiv preprint arXiv:1903.11027*, 2019.
- [82] K. Granström and M. Baum, “Extended object tracking: Introduction, overview and applications,” *CoRR*, vol. abs/1604.00970, 2016. [Online]. Available: <http://arxiv.org/abs/1604.00970>
- [83] K. Granström, C. Lundquist, and U. Orguner, “Tracking rectangular and elliptical extended targets using laser measurements,” *Fusion 2011 - 14th International Conference on Information Fusion*, pp. 1–8, 07 2011.
- [84] T. Starner, J. Weaver, and A. Pentland, “Real-time american sign language recognition using desk and wearable computer based video,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 12, pp. 1371–1375, 1998.
- [85] GM Cruise Holdings LLC, “Lidar based recognition of ride hailing gestures for autonomous vehicles,” Patent 16/365 094.
- [86] F. Zhou, X. Li, and Z. Wang, “Efficient high cross-user recognition rate ultrasonic hand gesture recognition system,” *IEEE Sensors J.*, vol. 20, no. 22, pp. 13 501–13 510, 2020.
- [87] J. Lien, N. Gillian, M. E. Karagozler, P. Amihoud, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, “Soli: Ubiquitous gesture sensing with millimeter wave radar,” *ACM Trans. Graph.*, vol. 35, no. 4, pp. 142:1–142:19, Jul. 2016.
- [88] Y. Kim and B. Toomajian, “Hand gesture recognition using micro-doppler signatures with convolutional neural network,” *IEEE Access*, vol. 4, pp. 7125–7130, 2016.
- [89] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in pytorch,” 2017.
- [90] V. Boddeti, private communication, Oct. 2020.
- [91] E. Klinefelter and J. A. Nanzer, “Interferometric radar for spatially-persistent gesture recognition in human-computer interaction,” in *2019 IEEE RadarConf*, 2019, pp. 1–5.

- [92] —, “A bound on radial velocity and observation time in interferometric angular velocity measurement,” in *2020 IEEE APS Int. Symp.*, 2020, pp. 1–2.
- [93] S. Mallat, *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*, 3rd ed. Orlando, FL, USA: Academic Press, Inc., 2008.
- [94] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: With Applications in R*. New York, NY: Springer Publishing Company, Incorporated, 2014.