

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Dissertation

**POLARIZATION TECHNIQUES FOR MITIGATION OF
LOW GRAZING ANGLE SEA CLUTTER**

by

MOLLY K. CRANE

B.A., Clark University, 2004
M.Eng., Boston University, 2012

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2017

© 2017 by
MOLLY K. CRANE
All rights reserved

Approved by

First Reader

S. Hamid Nawab, PhD
Professor of Electrical and Computer Engineering

Second Reader

W. Clem Karl, PhD
Professor of Electrical and Computer Engineering

Third Reader

Jeffrey B. Carruthers, PhD
Associate Professor of Electrical and Computer Engineering

Fourth Reader

Mabel D. Ramirez, PhD
Assistant Group Leader, Group 39, MIT Lincoln Laboratory

The author is an employee of MIT Lincoln Laboratory.

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

This material is based upon work supported under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Air Force.

*Give me six hours to chop down a tree,
and I will spend the first four sharpening the axe.*
Abraham Lincoln

*More complex does not mean better.
Something very complex may be better in theory,
but that does not mean it is better in practice.*
Dimitris Manolakis

Acknowledgments

One can undertake learning a vast, deep area of knowledge by entering its dark pool slowly from the shallow end, finding a solid footing on the stairway of fundamentals, then moving gradually deeper. This is how it's usually done in academia, of course. One can also stand on this stairway long enough to get acclimated to the water, then climb out, approach the deepest part of the pool, close her eyes, hold her breath, and jump. The latter approach is the one I've taken, and it has been incredibly challenging and immeasurably rewarding.

I am thankful to those who led me down the staircase but recognized the value of the deep dive (advisor Prof. S. Hamid Nawab in particular and Boston University's ECE Department in general); to the institution that offered me entrance to the deep end (MIT Lincoln Laboratory); to the great minds that dug out and filled the pool to begin with (see bibliography); to the group that recognized the value of swimming in this particular pool and funded my admission to it (MIT Lincoln Laboratory's Advanced Concepts Committee); to the lifeguards who kept watch over me but gave me much-needed leeway as I first learned to tread water, then to doggie paddle, then finally, to swim (the leaders of MIT LL's Group 39 in particular and Division 3 in general: Drs. Mabel Ramirez, Stephen Uftring, Sean Duffy, Nicholas Allan, and Katherine Rink); and of course, to my enduring and patient lifejackets, both professional (Mr. David Mooradd and Mr. Dennis Blejer) and personal (Ms. Gozde Guckaya - everything for you, for you are my everything). You all amaze and humble me every day. I aspire to one day swim as well as you, and when I do, it will be because of you. Infinite gratitude, always.

Even a reformed English lit major can only carry a metaphor so far. There are many people deserving of my deepest thanks who happen to swim mostly in other pools, but their involvement in this research and dissertation process has been pivotal. David Mooradd's name bears repeating here. David introduced me to the power of radar polarization, in which he has fervently believed for decades. It's in large part because of

his involvement that this project was funded. Dennis Blejer’s name also bears repeating. Dennis’s expertise in electromagnetics is truly rich. His many conversations with me were the “Polo” to my “Marco.” (Forgive me. There’s that pool metaphor again...)

There are plenty of new names, as well. In addition to David and me, the radar assembly and field test team included Mr. Timothy Navien, Dr. Jonathan Morse, Mr. Dean Mailhiot, and Mr. Daniel Callahan, all of MIT LL. Without them, no radar could have been assembled and no data could have been collected. Dr. Benjamin Woodacre, formerly of MIT LL, suggested exploring slow-time waveform encoding as a means of simultaneous scattering matrix measurement. The result of that suggestion appears in Section 3.2.1. Drs. Henry Thomas and Gerald Benitz, both of MIT LL, offered guidance, expertise, and a well-established starting point that were all critical to the development of the multi-channel transverse equalization algorithm described in Section 4.2. Ms. Sarah Lichtman of MIT LL implemented the vision for data labeling that is described in Section 4.5; her excellent work made expedient a critical task that would have otherwise proven incredibly tedious and time-consuming. My dissertation committee members—Drs. S. Hamid Nawab, W. Clem Karl, Jeffrey Carruthers, and Mabel Ramirez—have devoted valuable time to help guide this research. Dr. Samantha Weiss of MIT LL, a kindred literary spirit, provided invaluable editing of this dissertation. To all of you wonderful, wonderful people: again, infinite gratitude. I am truly blessed to call you my mentors, my colleagues, and my friends.

Finally, one more name bears repeating: Ms. Gozde Guckaya, you are my puzzle piece, my rock, my home. Your unwavering support and boundless patience through this process have been appreciated more than could ever be captured by mere words upon a page. I’ll never stop looking forward to spending life with you. This is for us.

Molly K. Crane
PhD Candidate, Boston University ECE Department
Associate Technical Staff, MIT Lincoln Laboratory

POLARIZATION TECHNIQUES FOR MITIGATION OF LOW GRAZING ANGLE SEA CLUTTER

(Order No.)

MOLLY K. CRANE

Boston University, College of Engineering, 2017

Major Professor: S. Hamid Nawab, PhD

Professor of Electrical and Computer Engineering

ABSTRACT

Maritime surveillance radars are critical in commerce, transportation, navigation, and defense. However, the sea environment is perhaps the most challenging of natural radar backdrops because maritime radars must contend with sea clutter. Sea clutter poses unique challenges in very low grazing angle geometries, in which typical statistical assumptions regarding sea clutter backscatter do not hold. As a result, traditional constant false alarm rate (CFAR) detection schemes may yield a large number of false alarms while objects of interest may be challenging to detect. Solutions posed in the literature to date have been either computationally impractical or lacked robustness.

This dissertation explores whether fully polarimetric radar offers a means of enhancing detection performance in low grazing angle sea clutter. To this end, MIT Lincoln Laboratory funded an experimental data collection using a fully polarimetric X-band radar assembled largely from COTS components. The Point de Chene Dataset, collected on the Atlantic coast of Massachusetts' Cape Ann in October 2015, comprises multiple sea states, bandwidths, and various objects of opportunity. The dataset also comprises three different polarimetric transmit schemes. In addition to discussing the radar, the dataset,

and associated post-processing, this dissertation presents a derivation showing that an established MIMO radar technique provides a novel means of simultaneous polarimetric scattering matrix measurement. A novel scheme for polarimetric radar calibration using a single active calibration target is also presented.

Subsequent research leveraged this dataset to develop Polarimetric Co-location Layering (PCL), a practical algorithm for mitigation of low grazing angle sea clutter, which is the most significant contribution of this dissertation. PCL routinely achieves significant reduction in the standard CFAR false alarm rate while maintaining detections on objects of interest. Moreover, PCL is elegant: It exploits fundamental characteristics of both sea clutter and object returns to determine which CFAR detections are due to sea clutter. We demonstrate that PCL is robust across a range of bandwidths, pulse repetition frequencies, and object types. Finally, we show that PCL integrates in parallel into the standard radar signal processing chain without incurring a computational time penalty.

Contents

1	Introduction	1
1.1	Radar signal processing chain	2
1.2	Radar signal information dimensions	2
1.3	Motivation for research	9
1.4	Scope of dissertation	12
2	Background	16
2.1	Radar polarimetry	16
2.1.1	Polarization fundamentals	16
2.1.2	Measuring polarization	22
2.1.3	The scattering matrix and its power representations	24
2.1.4	Measurement of the scattering matrix	28
2.1.5	The RCS-scattering information dimension in post-processed data	31
2.1.6	Relevant application areas	33
2.2	Sea clutter	33
2.2.1	Sea clutter phenomenology	34
2.2.2	Detection in sea clutter	36
3	Four Eyes and the Point de Chene Dataset	41
3.1	Four Eyes polarimetric radar system	41
3.1.1	Four Eyes' antennas	44
3.2	Waveforms and transmit schemes	45
3.2.1	Simultaneous SM measurement with Doppler division multiple ac- cess waveforms	48

3.3	Point de Chene field test, October 2015	52
3.3.1	Field test location	52
3.3.2	Persistent test objects	54
3.3.3	Objects of opportunity	55
3.3.4	Observed sea states	55
3.4	Point de Chene Dataset	60
3.4.1	Algorithm development data subset	61
4	Signal Post-Processing	63
4.1	Point de Chene calibration targets	64
4.2	Cross-channel transverse equalization	67
4.2.1	Signal model	68
4.2.2	Transverse equalization algorithm	69
4.3	Polarimetric calibration	75
4.3.1	Polarimetric calibration results	78
4.4	Radar data processing flow	79
4.4.1	Recorded data retrieval and signal conditioning	80
4.4.2	Basebanding	83
4.4.3	Downsampling	83
4.4.4	Pulse compression and Doppler processing	84
4.4.5	CFAR detection	85
4.5	Data labeling	88
5	Polarimetric Co-location Layering (PCL)	95
5.1	Sea clutter Doppler spectrum	96
5.1.1	On the polarization dependence of sea clutter Doppler spectrum .	96
5.1.2	On sea clutter mitigation using Doppler techniques in polarimetric radar	99
5.2	The fundamental principle behind PCL	102

5.3	PCL	106
5.3.1	Core algorithm	106
5.3.2	1-D performance metrics	114
5.4	PCL results	116
5.4.1	PCL performance across bandwidths	117
5.4.2	PCL performance across PRFs	125
5.4.3	PCL performance across object types	127
5.5	Integrating PCL into the standard radar signal processing chain	152
5.5.1	PCL detection filtering	153
5.5.2	Computational complexity	156
5.6	Summary	159
6	Summary and Future Work	161
6.1	Summary	161
6.2	Future Work	162
	References	164

List of Tables

2.1	The canonical polarization states	22
3.1	Four Eyes COTS system configuration	43
3.2	Point de Chene Dataset summary	61
3.3	Algorithm development data subset	62
4.1	Number of range bins to 4 km, by waveform	89
5.1	Empirical 1-D p_{fa} comparison across bandwidths	125
5.2	Empirical 1-D p_{fa} comparison across PRFs	127
5.3	Empirical 1-D p_{fa} comparison across object types and PRFs	152

List of Figures

1.1	Baseline noncoherent pulsed radar functionality	3
1.2	Baseline coherent radar functionality	5
1.3	A vertically polarized electromagnetic wave	7
1.4	RCS in fully polarimetric radar	8
1.5	The two-pronged sea clutter problem	11
1.6	The power of optical polarimetry	12
2.1	Time-independent electric field vector traces out wave polarization	17
2.2	Polarization ellipse	19
2.3	Poincarè's sphere	20
2.4	Simultaneous fully polarimetric radar	29
2.5	Alternating fully polarimetric radar	30
2.6	Single-polarization radar data products	31
2.7	Polarimetric radar data cube	32
2.8	σ_0 predictions across models and sea states	36
3.1	Four Eyes system block diagram	42
3.2	Four Eyes on location at the Atlantic Coast in Rockport, MA	43
3.3	Four Eyes' antennas in the anechoic chamber	44
3.4	Four Eyes' antenna beam patterns	45
3.5	Four Eyes' polarimetric transmit schemes	47
3.6	Nautical chart highlighting candidate locations	52
3.7	Progressive zoom maps to Cape Ann experiment location	53
3.8	Gong buoy	54

3.9	Persistent object viewing geometry	55
3.10	Objects of opportunity	56
3.11	Kayaker plus persistent objects	56
3.12	GIT and SIT sea state models	57
3.13	NOAA buoy 44098 location	58
3.14	Buoy 44098 SWH data mapped to Douglas sea state	59
3.15	Estimating Beaufort wind force scale from photographic evidence	60
4.1	Calibration delay line set-up	65
4.2	Calibration delay line	66
4.3	Determining $\mathbf{S}_{dl}$ by inspection	67
4.4	Four Eyes calibration set-up	68
4.5	Cross-channel transverse equalization block diagram	70
4.6	Coherently integrated, pulse-compressed delay line responses	71
4.7	Correction to Gaussian-edged frequency response over the passband . . .	73
4.8	Delay line point target responses following cross-channel equalization . .	75
4.9	Synthesis of circular trihedral response for polarimetric calibration quality assessment	79
4.10	Pre-polarimetric signal processing chain	80
4.11	ALT transmit scheme	81
4.12	IF data spectrum reversal	83
4.13	Basebanding desired spectrum	83
4.14	Oversampled range resolution cell	85
4.15	Range resolution cell sampled at twice critical Nyquist	85
4.16	Notional 1-D CFAR background sampling window	87
4.17	Notional 2-D CFAR background sampling window	88
4.18	Full file RTIs for HH and VV polarizations, File 325	91
4.19	Full file RTIs for File 325, zoomed in on object	92

4.20	File 325 data labels to 1.5 km range	92
4.21	File 325 data labels, zoomed in on object	93
5.1	Sea clutter Doppler differential as a function of look direction	98
5.2	The fundamental principle behind PCL	103
5.3	CA-CFAR detections across HH and VV polarizations	105
5.4	RTIs across HH and VV, buoy scene	107
5.5	Coherent averages across HH and VV, buoy scene	108
5.6	CA-CFAR detections on HH and VV coherent averages	110
5.7	PCL sequence on File 273, part A	112
5.8	PCL sequence on File 273, part B	113
5.9	Hits by detection operation, File 273, 1 st 100 CPIs	116
5.10	Detections and false alarm rate across detection operations, File 273, 1 st 100 CPIs	117
5.11	Performance comparison of standard single-polarization CFAR with PCL, $\delta_n = 1$ and $n_{CPI} = 3$, File 272 (40 MHz bandwidth)	119
5.12	Polarimetric signature variation results in dropped PCL detection	119
5.13	Performance comparison of standard single-polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 3$, File 272	121
5.14	Performance comparison of standard single-polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 272	121
5.15	Performance comparison of standard single-polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 271 (150 MHz bandwidth)	122
5.16	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 270 (400 MHz bandwidth)	123
5.17	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 8$, File 270	124

5.18	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 0$ and $n_{CPI} = 3$, File 273 (4 MHz bandwidth) at ~ 893 Hz PRF . . .	127
5.19	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 0$ and $n_{CPI} = 5$, File 273 at ~ 893 Hz PRF	128
5.20	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 3$, File 272 at ~ 893 Hz PRF	128
5.21	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 272 at ~ 893 Hz PRF	129
5.22	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 271 at ~ 893 Hz PRF	129
5.23	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 270 at ~ 893 Hz PRF	130
5.24	Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 8$, File 270 at ~ 893 Hz PRF	130
5.25	Max velocity at which clutter and object returns may be coincident . . .	131
5.26	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 267 (40 MHz bandwidth) at 6250 Hz PRF	134
5.27	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 267 at 6250 Hz PRF	135
5.28	Performance comparison with $\delta_n = 3$, $n_{CPI} = 6$, and $n_{prop} = 1$, File 267 at 6250 Hz PRF	135
5.29	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$, File 267 at 6250 Hz PRF	136
5.30	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 267 at ~ 893 Hz PRF	137
5.31	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 267 at ~ 893 Hz Hz PRF	137

5.32	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 1$, File 267 at ~ 893 Hz PRF	138
5.33	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 2$, File 267 at ~ 893 Hz PRF	138
5.34	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 2$, File 267 at ~ 893 Hz PRF	139
5.35	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 325 (150 MHz bandwidth) at 6250 Hz PRF	140
5.36	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 325 at 6250 Hz PRF	141
5.37	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 1$, File 325 at 6250 Hz PRF	141
5.38	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 2$, File 325 at 6250 Hz PRF	142
5.39	Performance comparison with $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 3$, File 325 at 6250 Hz PRF	142
5.40	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 325 at ~ 893 Hz PRF	143
5.41	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 325 at ~ 893 Hz PRF	144
5.42	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$, File 325 at ~ 893 Hz PRF	144
5.43	Performance comparison with $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$, File 325 at ~ 893 Hz PRF	145
5.44	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 266 (400 MHz bandwidth) at 6250 Hz PRF	146

5.45	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 266 at 6250 Hz PRF	147
5.46	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$, File 266 at 6250 Hz PRF	148
5.47	Performance comparison with $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$, File 266 at 6250 Hz PRF	148
5.48	Performance comparison with $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 266 at ~ 893 Hz PRF	149
5.49	Performance comparison with $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$, File 266 at ~ 893 Hz PRF	150
5.50	Performance comparison with $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$, File 266 at ~ 893 Hz PRF	150
5.51	Performance comparison with $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$, File 266 at ~ 893 Hz PRF	151
5.52	The standard coherent radar signal processing chain	153
5.53	Integration of PCL into the polarimetric radar signal processing chain . .	154
5.54	2-D CA-CFAR detections on RD data	155
5.55	2-D PCL detections on RD data	156

List of Abbreviations

1-D		one-dimensional
2-D		two-dimensional
ADC		Analog-to-Digital Converter
ALT		ALternating polarimetric transmit scheme
APF		Adaptive Polarization Filtering
AWG		Arbitrary Waveform Generator
BSA		Backscatter Alignment
CA-CFAR		Cell Averaging CFAR
CFAR		Constant False Alarm Rate
CPI		Coherent Processing Interval
COTS		Commercial Off-The-Shelf
CUT		Cell Under Test
DDMA		Doppler Division Multiple Access
EM		electromagnetic
FFT		Fast Fourier Transform
H		Horizontal polarization
HDD		Hard Drive Disk
iid		Independent and Identically Distributed
IF		Intermediate Frequency
L		Left-hand circular polarization
LFM		Linear Frequency Modulation
LNA		Low Noise Amplifier
MIT LL		MIT Lincoln Laboratory
MIMO		Multiple Input, Multiple Output
NOAA		National Oceanic and Atmospheric Administration
OS-CFAR		Order Statistic CFAR
PCL		Polarimetric Co-location Layering
PRF		Pulse Repetition Frequency
PRI		Pulse Repetition Interval
R		Right-hand circular polarization
RAID		Redundant Array of Independent Disks
RCS		Radar Cross Section
RD		Range-Doppler
RTI		Range-Time Intensity
SAR		Synthetic Aperture Radar
SCR		Signal-to-Clutter Ratio
SIM		SIMultaneous polarimetric transmit scheme
SIM-PHS		SIMultaneous PHaSe encoded polarimetric transmit scheme
SM		Scattering Matrix
SNR		Signal-to-Noise Ratio
STAP		Space-Time Adaptive Processing
SWH		Significant Wave Height
TWT		Traveling Wave Tube
V		Vertical polarization
VSG		Vector Signal Generator

X^+		$+45^\circ$ slant linear polarization
X^-		-45° slant linear polarization

Chapter 1

Introduction

Radar is a relatively old technology that continues to find new applications not only within the traditional fields of surveillance and object tracking, but also in areas like imaging, autonomous control, terrain mapping, atmospheric measurement, and weather pattern monitoring. Regardless of the application, at their core all radars operate according to the same fundamental process. Pulses of electromagnetic (EM) wave energy are emitted in a focused beam out of a transmit antenna. The waves interact with objects in the scene, many of which reradiate or scatter the incident wave energy in some direction. The radar then measures the portion of these scattered waves that both propagate in the direction of and couple into its receive antenna. Finally, the measured signals are analyzed using signal processing techniques, rendering an operator-interpretable or machine-interpretable representation of the radar scene.

The penultimate step of this process—coupling into the receive antenna—is hardly discussed in canonical radar texts (Skolnik, 2001; Stimson et al., 2014; Richards et al., 2010). Yet this step’s importance cannot be overstated. Whether a scattered wave couples into a receive antenna at all depends on the polarizations of both the antenna and the scattered wave. If the two are the same, all of the wave’s energy will couple into the antenna, and the radar will “see” all of the scatterer’s response. If the two are orthogonal, none of the wave’s energy will couple into the antenna, and the radar will be “blind” to the scatterer’s presence. Polarization and whether it can be leveraged in a practical manner to enhance a maritime surveillance radar’s ability to discriminate

between object and environment against the most challenging of radar backdrops, the dynamic ocean, is the focus of this dissertation.

1.1 Radar signal processing chain

It will serve the reader to be familiar with what this author will call the baseline radar signal processing chain. This chain comprises the components of radar functionality that lie between the digitization of signals coming into the radar receiver and the passage of detections to a tracker that attempts to establish tracks on objects over time (see Figure 1.1). In other words, this chain is a large part of the ultimate step in the fundamental radar process. At a minimum, the chain usually includes pulse compression (matched filtering), multiple pulse integration (noncoherent or coherent), and constant false alarm rate (CFAR) detection, the performance of which is dependent upon signal to noise ratio (SNR). The exact nature of these components depends in part on whether the radar is noncoherent, coherent, or polarimetric. That is, their nature depends upon the dimensions of information contained within the signals measured by the radar.

1.2 Radar signal information dimensions

It has long been understood that a monochromatic EM wave can be completely characterized by its amplitude, phase, and polarization. The earliest radars were noncoherent, measuring only the amplitudes of backscattered waves. Because radar transmit pulses encapsulate a known waveform, pulse compression—which amounts to correlating received amplitude signals against this waveform—enables localization in time of scatterers that reradiated energy towards the receive antenna. A scatterer’s partial radar cross section (RCS), which can be very roughly thought of as the electromagnetic “size” of the scatterer from the radar’s perspective, is captured in the SNR of the pulse compressed signal. Because EM waves propagate at the speed of light, the time delays of scatterers’ returns can be used to calculate the ranges of the scatterers from the radar in the direction of




Figure 1.1: Baseline noncoherent radar functionality for a pulsed radar is shown, with each step matched to its corresponding term(s) in the radar range equation given in Equation 1.1. Once the signal is digitized in the receiver, the radar signal processor performs pulse compression, which amounts to correlating the received signal against the transmitted waveform, improving single-pulse SNR by a factor of τB , where τ [sec] is the pulselength and B [Hz] is the waveform bandwidth. The range-time intensity (RTI) represents a series of pulse compressed returns, each of which is a function of range (fast time), stacked one atop the next according to pulse transmit time (slow time). Noncoherently integrating these returns along the slow time dimension yields a 1-D amplitude signal whose SNR is further improved by roughly $\sqrt{N}$, where N is the number of pulses integrated. 1-D CFAR detection is run on this integrated signal, and its range detections are passed to the tracker.

the antenna's beam. In sum, amplitude signals measured by noncoherent radar at any instant of time are real-valued scalars carrying partial RCS information in the range-time dimension.

A block diagram of noncoherent radar baseline functionality is shown in Figure 1.1. For such a radar, which uses both pulse compression and integration as shown, the radar range equation for a point target object can be written as

$$SNR_{coh} = \frac{PG^T}{4\pi R^2} \cdot \frac{\sigma}{4\pi R^2} \cdot \frac{G^R \lambda^2}{4\pi} \cdot \frac{\tau B \sqrt{N}}{L} \cdot \frac{1}{kT_s B} \quad (1.1)$$

where

- P = transmit power [W]
- G^T = transmit antenna gain, unitless
- R = range to scatterer [m]
- σ = scatterer RCS [m²]
- G^R = receive antenna gain, unitless
- λ = radar wavelength [m]
- τ = waveform pulsewidth [sec]
- B = bandwidth [sec⁻¹]
- N = number of pulses integrated, unitless
- L = losses, unitless
- k = Boltzmann's constant [W·sec/K] and
- T_s = system noise temperature [K].

A deeper examination of each term and its role in the radar range equation is beyond the scope of this document, but the interested reader is referred to Skolnik (2001) and Edde (1995). It is, however, instructive to consider what is gained in terms of SNR and how that gain is realized when radars can also measure other characteristics of impinging EM waves.

When a radar can lock transmit waveform phase to that of a known stable oscillator, it can measure the phases of waves coupling into the receive antenna by using the transmit waveform's phase as a reference. Such radars are appropriately dubbed coherent radars. They use phase information to measure the Doppler shift of scattered waves across multiple pulses at each range bin. From the Doppler shift, scatterers' radial velocities with respect to the radar platform are calculated. Thus, coherent radar signals at any instant of time are complex-valued scalars that carry partial RCS information in the range-time and Doppler-velocity dimensions. A block diagram of baseline coherent radar functionality is shown in Figure 1-2.




Figure 1.2: The baseline coherent pulsed radar utilizes Doppler processing, which in software amounts to computation of a Fourier transform for each range bin across slow time, as an effective form of coherent integration, which produces a gain of N , where N is the number of pulses integrated. 2-D CFAR detection is run on this integrated signal, and the velocities and ranges of detections are passed to the tracker.

Despite the additional cost and complexity of coherent radar systems as compared to noncoherent systems, virtually all modern operational radars are coherent because of the SNR gain that coherent integration offers. For a radar that utilizes pulse compression and coherent integration, the radar range equation for a point target object can be written as

$$SNR_{coh} = \frac{PG^T}{4\pi R^2} \cdot \frac{\sigma}{4\pi R^2} \cdot \frac{G^R \lambda^2}{4\pi} \cdot \frac{\tau BN}{L} \cdot \frac{1}{kT_s B} \quad (1.2)$$

and therefore

$$SNR_{coh} = \sqrt{N} \cdot SNR_{ncoh}.$$

Such a gain is clearly a boon to the detection process. Integrating just 10 pulses coherently yields an additional 5 dB of margin over noncoherent integration, with another 5 dB of margin gained for every order of magnitude by which N can be increased.

It is not immediately clear from Equation 1.2 how a radar system designer can further improve SNR without changing system gains, transmit power, or frequency. However, the keen reader will have noted use of the phrase *partial RCS information* in the preceding discussion. RCS is a scalar quantity, as shown in Equations 1.1 and 1.2, only from the point of view of systems that cannot measure polarization.

Polarization is defined by the path traced by the tip of an EM wave’s electric field vector over one period of propagation. The next section dives into the mathematics, but for the present it suffices to note that in the general case, the tip of the electric field traces out an ellipse over each wavelength. Thus, polarization of an EM wave at any time instant is a vector quantity comprising two complex scalar quantities that represent the instantaneous magnitude and orientation of the wave’s electric field in two orthogonal directions. Hence, in order to measure polarization, a radar must have two orthogonally polarized receive antenna elements that feed into a coherent radar system. These so-called dual-polarized (dual-pol) receive antennas measure all of the energy of EM waves scattered in their direction regardless of the waves’ polarizations.

However—and this is an important point that is often misunderstood within the general radar community—dual-pol on receive capability is insufficient for measuring the complete RCS information in a radar scene. This insufficiency is due to the scattered wave’s polarization being dependent upon the nature of the scatterer, the geometry of the scene, and the frequency and polarization of the EM wave that initially impinged upon that scatterer. In other words, the response a radar gets from interrogating a scatterer is a function not only of the object, but of how the radar interrogates that object.

As a simple example, consider a vertically polarized EM wave, as shown in Figure 1.3, impinging upon a horizontally oriented, perfectly conducting, infinitely thin wire. Despite the wire’s ideal electromagnetic properties, the wave will not induce a current in the wire because the wire is oriented orthogonally to the wave’s polarization. Thus, the wire will not reradiate EM energy toward the radar’s dual-pol receive antenna; the

radar will still be “blind” to the wire’s presence in the scene. It is in this sense that RCS as measured by non-polarimetric radars is dubbed “partial RCS” in the preceding discussion. The wire clearly does not have an RCS of zero. But the radar cannot perceive the wire’s true RCS because its method of interrogation—the polarization of its emitted wave—was mismatched to the geometry of the wire.




Figure 1.3: Assuming use of the Cartesian coordinate system and propagation in the $+z$ direction as shown, a vertically polarized electromagnetic wave has an electric field vector whose tip traces out a line along the vertical axis of the $x-y$ plane. The magnetic field vector (not pictured) traces out a line orthogonal to both the electric field vector and the direction of propagation.

Figure 1.4 illustrates this point. Shown are four synthetic aperture radar (SAR) images created with a fully polarimetric SAR system. The imaged scene contains a variety of objects, including a helix and several wires oriented horizontally across the radar’s field of view. The transmit and receive polarizations are indicated at the top left of each figure in sequence. So, for example, the figure labeled HH is imaged using horizontally polarized transmit and receive antennas; HV is imaged using horizontally polarized transmit and vertically polarized receive antennas; and so on.

The horizontal wires are absent from all but the HH image. The wires do not appear in the VH and VV images because impinging vertically polarized energy does not cause reradiation by horizontally oriented wires. The wires do not appear in the HV image because in response to a horizontally polarized impinging wave, the wires’ scattering properties are such that they reradiate horizontally polarized waves—to which the vertically polarized receive antenna is blind.

If the SAR whose imagery is shown in Figure 1.4 were merely single-polarized (single-pol), it would capture either HH or VV, depending on which of the fixed polarizations the antenna(s) possessed. If the radar were dual-pol on receive, then the SAR would measure




Figure 1-4: Fully polarimetric SAR imagery demonstrates the inability of a single-pol or dual-pol radar to capture complete RCS scattering information in a scene. Wires oriented horizontally across the radar’s field of view have no signature to a linearly polarized radar such as this except for in HH. Consequently, the objects are invisible to a vertically polarized radar and to a dual-pol radar that is vertical on transmit. *Imagery used courtesy of Dennis J. Blejer, MIT Lincoln Laboratory.*

either HH and HV or VH and VV, depending on the radar’s fixed transmit polarization. Only a system that is both dual-pol on receive and dual-pol on transmit can ensure that every object in a scene will interact with the radar’s transmitted energy. Radar systems that have this capability are dubbed *fully polarimetric*. Radar signals measured by fully polarimetric systems carry complete RCS-scattering information in the range-time and Doppler-velocity dimensions—that is, they measure the maximal amount of information that radar signals carry for a given scene geometry and radar frequency.

1.3 Motivation for research

Maritime surveillance radars have long been a critical component of global commerce, defense, and transportation. Shipboard, airborne, and ground-based radars keep watch over littoral waters, borders, and harbors. Maritime radars are also used for search and rescue, iceberg detection, and navigation, among myriad other applications.

Like most modern radars, maritime radars are usually coherent but not polarimetric, and thus cannot measure all the information contained in backscattered waves. The question begs asking: Given that the technology to build polarimetric radar systems has existed for decades—and it has, for at least three (Boerner, 2007)—why are not all modern radar systems built with polarimetric capability? As is often the case in engineering, the answer is that there is a cost-complexity trade-off.

A fully polarimetric radar system requires at least two orthogonally polarized antenna elements. So, even if the same dual-polarized antenna is used for both transmit and receive, the cost of the antenna may be as much as double that of standard single-polarization systems. If different antennas are used for transmit and receive, the cost may quadruple. Many modern radars utilize active phased array technology, and doubling the number of feeds and antenna elements in these systems quickly leads to rising costs. Beyond the antennas and receive chains, fully polarimetric systems also require additional hardware on the transmit side, ranging from a polarization switch for each antenna pair up to a full duplicate of the signal transmit chain for each antenna pair. Juxtaposed with the demands imposed by the constant push for greater bandwidth, larger and more capable arrays, and ever-present budgetary restrictions, fully polarimetric capability has a financial cost that rarely survives the need-to-have vs. nice-to-have chopping block in radar system design. Moreover, legacy issues with technological difficulties in polarimetric radar implementation (Root, 1982) have colored the view of much of the modern radar community even many years later (Boerner, 2007). Finally, because signals transmitted from and received by fully polarimetric radars must traverse different signal paths, careful

equalization in both the frequency response sense and the polarimetric calibration sense is necessary to capture the important between-channel phase relationships that characterize polarization state vectors. Polarimetric calibration can be challenging to implement in practice (Yueh et al., 1990; Zebker et al., 1991; Freeman, 1992). So, for many applications and in many environments, the fully polarimetric level of system cost and complexity is perceived as unnecessary for a radar to be able to perform its tasks reliably.

Yet, there are several applications and environments in which experimental polarimetric radar has proven invaluable. Most notably, the imaging and mapping applications of synthetic aperture radar (SAR) have leveraged polarimetry to produce the radar analog of optical imagery’s transition from black and white to color. With that transition has come a dramatic improvement in the ability to accurately classify radar image content (Lee and Pottier, 2009; Mather and Tso, 2009). It has also long been understood that polarization is the key to distinguishing different types of precipitation echoes in weather radar (Bringi and Chandrasekar, 2001) and to distinguishing objects from precipitation echoes in surveillance radar (Kennaugh, 1952).

The thrust of polarimetry in radar continues to be SAR-related despite the fact that there are numerous other applications and environments in which the amplitude and phase information of backscattered waves is simply not enough for a radar to perform at an acceptable level of reliability. It is the contention of this author and her advisors that, given the current trend toward multi-function radars expected to operate in a range of environments, the merits of full polarimetry deserve to be explored within each of these challenging applications and contexts. Only in so doing can a deeper understanding of what is lost when polarimetric capability does not survive the chopping block be gleaned, serving to either reinforce the present understanding of the trade space or to correct it.

One such application and environment is detection of objects of interest by a radar looking at near horizontal incidence out to sea, which is the focus of this dissertation. In all maritime surveillance settings, a radar must contend with *sea clutter*, or radar returns

from waves scattered by the undulating sea surface. This can be challenging regardless of the angle at which the radar’s beam intersects that surface. However, as this angle becomes small enough to fall into the so-called low grazing angle regime at less than 10° —or the very low grazing angle regime of less than a few degrees—the application of standard detection techniques in other than very low sea states often yields less than acceptable results.

In such a geometry, the sea surface creates many radar returns that look object-like, increasing the number of false alarms produced by the radar’s detector. Indeed, the false alarm rate due to sea clutter in only moderately rough sea conditions can become so high that a radar’s tracking system may be overwhelmed with the number of candidate





Figure 1.5: The range-Doppler map of a single small boat in low grazing angle sea clutter is shown. The color axis indicates relative power of returns in dB, normalized to the peak response in the map. CFAR detections are indicated by black circles, indicating the two-fold low grazing angle sea clutter problem. At such low grazing angles, the assumptions underlying a CFAR detector are not valid, so 1) the detector produces an overwhelming number of false alarms, while 2) objects of interest may be difficult to distinguish from the clutter-dense background.



(a) Without polarization filter

(b) With polarization filter

Figure 1-6: At left, the turtle is obscured by the glare of the sun’s reflections off of the water’s surface in the photograph taken without use of an optical polarizer. At right, a polarizer has been applied, filtering out the glare. The turtle is revealed. *Photographs taken by and used courtesy of David C. Mooradd, MIT LL.*

returns (Ward et al., 2006). Moreover, objects of interest may go undetected because of sea clutter’s glare. Figure 1-5 illustrates this two-pronged problem.

In optics, polarizers have long been leveraged as a means of filtering out unwanted glare, as shown in Figure 1-6. Because of this optical analogue, it has been thought for decades that polarimetry may hold the key to distinguishing between returns from sea clutter and returns from objects of interest (Haykin et al., 1994; Long, 2001). However, prior solutions have either been impractical in terms of incorporation into the standard radar signal processing chain or have lacked robustness. The lack of viable polarimetric solutions is likely due in large part to the dearth of available data for research.

1.4 Scope of dissertation

This dissertation supplements the current body of knowledge by exploring the efficacy of leveraging fully polarimetric radar to robustly enhance detection capability using practical approaches in a very low grazing angle maritime environment. The first part of the research comprises collection of appropriate data utilizing an X-band radar assembled from commercial off-the-shelf (COTS) components. The data collection features low grazing angle sea clutter, a variety of objects of interest, a variety of polarimetric

transmit schemes, a range of bandwidths, and a range of sea states. The second part of the research comprises development of an algorithm intended to enhance detection performance in low grazing angle sea clutter without restructuring the standard radar signal processing chain or imposing currently impractical computational complexity. Implicit in this work are development of a routine to equalize the channels in post-processing and application of metrics to quantify improved radar detection performance. A successful algorithm will be robust to variations in radar parameters, scene geometry, and the state of the sea surface. Ideally, the algorithm will also be practical in the sense that it will plug into the standard radar signal processing chain without requiring either restructuring of the chain or computational complexity that slows down the radar’s near real-time performance.

It should be noted at the outset that funding for this research was provided by MIT Lincoln Laboratory’s (MIT LL) Advanced Concepts Committee, and it is not intended that any collected data will be made available to the science and technology community at large. The data will, however, be available to the MIT LL community for further research efforts.

It should also be noted that, without loss of generality to other radar systems, the focus of this dissertation is on monostatic radar, for which the transmit and receive antennas are either the same antenna or are approximately co-located. Such systems utilize the backscatter alignment (BSA) coordinate system convention for received EM wave energy, and henceforth all reference to scattering by objects should be understood by the reader to be backscatter - i.e., energy that is reradiated or reflected from objects in the scene back in the radar’s direction.

Finally, the author has adopted the convention that where fundamental measurement units are applicable to a variable quantity, they are enclosed in rectangular brackets $[]$, as was done in Equation 1.1. Vectors and matrices are denoted using bold, capital letters,

though for ease of reading, vectors are additionally accented with the vector directional arrow, as in $\vec{\mathbf{E}}$, while matrices are not, as in $\mathbf{A}$.

The remainder of this document is organized as follows:

Chapter 2 gives the requisite background. The first part of the chapter features a discussion of radar polarimetry, including the mathematical foundations thereof and a brief discussion of the vast body of theoretical work completed by other researchers to date. The second section discusses the low grazing angle sea environment and its features and challenges from the perspective of object detection, including a brief discussion of past work on both polarimetric and non-polarimetric approaches to object detection in sea clutter.

Chapter 3 discusses Four Eyes, the X-band polarimetric radar assembled for a week-long field test conducted in October 2015 on the Atlantic Coast of Massachusetts' Cape Ann, as well as the dataset collected during that test. The radar's specifications and waveform design are discussed, as is the field test. In addition, a derivation is presented showing that Doppler division multiple access waveforms are an effective way of achieving simultaneous polarimetric scattering matrix measurement in radars with sufficiently high pulse repetition frequencies.

Chapter 4 describes the signal processing suite designed for Four Eyes' data. Included in this discussion are details of the objects used for calibration, the channel equalization algorithm, and the polarimetric calibration methodology. Details of the signal processing suite developed in MATLAB specifically for Four Eyes' data are discussed. The data labeling methodology used to label the radar data for performance quantification is also described here.

Chapter 5 motivates, describes, and quantifies the performance of Polarimetric Colocation Layering (PCL), the practical algorithm produced through this research that successfully mitigates the impact of low grazing angle sea clutter. PCL leverages a fundamental polarimetric characteristic of sea clutter to differentiate between detections

caused by objects of interest and detections caused by the sea surface. PCL performance is quantified using empirical probability of false alarm rates and continuity of detection metrics across a range of different bandwidths, PRFs, and object types. Implementation of PCL into the polarimetric radar signal processing chain is also discussed, as is the algorithm's computational complexity.

Chapter 6 summarizes the dissertation and suggests directions for future work.

Chapter 2

Background

2.1 Radar polarimetry

A brief discussion of the mathematical fundamentals of polarization is presented in this section, followed by a discussion of the polarization scattering matrix, which is synonymous with the RCS-scattering information dimension discussed in the previous chapter. Polarimetric scattering matrix measurement methodologies are also presented. The section concludes with a discussion of the application areas where work has been done that might translate well to the problem of object detection in low grazing angle sea clutter.

2.1.1 Polarization fundamentals

The mathematical development that follows has been well established in the literature, but is required in this dissertation to establish notation and necessary foundational concepts, including those that were stated without proof in the previous chapter. The interested reader is referred for a more exhaustive treatment to Huynen (1970), Boerner (2007), and Lee (2009), all of which were invaluable references for this author and from which the mathematics in Section 2.1.1 is derived.

Polarization descriptors: Jones vectors and the polarization ellipse

Consider a monochromatic constant amplitude EM wave emitted into free space by an antenna source. Such a wave has frequency $f = \frac{c}{\lambda}$ [Hz], where λ = wavelength and c = the speed of light. Without loss of generality, we can adopt the Cartesian coordinate system, along with the convention that the EM wave propagates in the $+z$ direction.




Figure 2.1: The time-independent electric field vector traces out wave polarization once per wavelength as the wave propagates through space. *Figure adapted from Wikipedia Commons.*

Because of the transverse nature of EM waves, this convention ensures that the electric and magnetic field components lie perpendicular to the direction of propagation—i.e., they comprise only x and y components. Further, because the proportional relationship of the electric to the magnetic fields in the far field of the antenna is well known, an EM wave can be completely described by a mathematical description of its electric field as a function of space and time:

$$\vec{\mathbf{E}}(z, t) = \begin{bmatrix} \alpha_x \cos(\omega t - kz + \delta_x) \\ \alpha_y \cos(\omega t - kz + \delta_y) \end{bmatrix} \quad (2.1)$$

where radian frequency $\omega = 2\pi f$; wavenumber $k = \frac{2\pi}{\lambda}$; α_x and α_y are the wave component amplitudes in the x and y directions, respectively; δ_x and δ_y are the phases of those wave components; and wave attenuation is disregarded. Equation 2.1 can be written in complex exponential notation as

$$\vec{\mathbf{E}}(z, t) = \text{Re} \left(\begin{bmatrix} \alpha_x e^{j\delta_x} \\ \alpha_y e^{j\delta_y} \end{bmatrix} e^{-jkz} e^{j\omega t} \right) = \text{Re} \left(\vec{\mathbf{E}}(z) e^{j\omega t} \right) \quad (2.2)$$

where $\vec{\mathbf{E}}(z)$ is the time-independent complex electric field vector phasor propagating in the $+z$ direction.

Polarization is defined by the path traced by the tip of this electric field vector phasor over one wavelength of propagation; three such wavelengths are shown in Figure 2.1. As shown by Equation 2.2, an EM wave is of course time-dependent, but its polarization is not. In pursuit of a polarization-only descriptor we can therefore drop the time-

dependent term. Moreover, the spatial dependence of the electric field vector phasor can be eliminated by projecting its path onto the $x - y$ plane. This projection yields two mathematically equivalent descriptions of a wave's polarization state.

Analytically, the projection amounts to evaluating $\vec{\mathbf{E}}(z)$ at $z = 0$, yielding

$$\vec{\mathbf{E}} = \begin{bmatrix} \alpha_x e^{j\delta_x} \\ \alpha_y e^{j\delta_y} \end{bmatrix}. \quad (2.3)$$

This complex vector is a succinct description of a wave's polarization state known as the *Jones vector*. The form of the Jones vector shows that a wave's polarization is dependent upon only the amplitudes of the electric field components in the x and y directions and the relative phase between them. To see the last part of this statement, consider that the phase of the first component can be eliminated by viewing it as an arbitrary phase factor common to both terms. Thus, Equation 2.3 can be rewritten as

$$\vec{\mathbf{E}} = e^{j\delta_x} \begin{bmatrix} \alpha_x \\ \alpha_y e^{j(\delta_y - \delta_x)} \end{bmatrix}, \quad (2.4)$$

where the vector still represents the wave's polarization state. When the electric field has components in both the x and y directions and either

$$\frac{\alpha_x}{\alpha_y} \neq 1 \quad \text{or} \quad \delta_y - \delta_x \notin \{0^\circ, 90^\circ\}$$


the polarization is neither circular, as shown in Figure 2-1, nor linear, as shown in Figure 1-3, but is elliptical.

Geometrically, in the general case the projection produces an elliptical pattern that is regularly traced out in the $x - y$ plane by the tip of the electric field vector. This projection, the *polarization ellipse*, is characterized by an ellipticity angle, an orientation angle, and a sense. Ellipticity angle χ is defined as

$$\chi = \arctan\left(\frac{a}{b}\right) \quad (2.5)$$

where a = the length of the ellipse’s semi-minor axis and b = the length of its semi-major axis as shown in Figure 2.2. Note that $0^\circ \leq \chi \leq 45^\circ$. When $\chi = 0^\circ$, polarization is linear ($a = 0$), and when $\chi = 45^\circ$, polarization is circular ($a = b$). Orientation angle Ψ is defined as the angle made by the positive x -axis and the ellipse’s semi-major axis, as shown in the figure; thus, $0^\circ \leq \Psi \leq 180^\circ$.¹ The sense of the ellipse, denoted by an arrow, indicates in what sense the vector rotates around the $+z$ axis when the wave is viewed from the source, propagating in the direction of the $+z$ axis. The sense is often called the “handedness” of the ellipse, because the well known “right hand rule” can be evoked to determine the sense. If, when pointing the thumb of one’s right hand in the direction of propagation, the fingers curl in the direction of rotation, the sense is called “right-hand”; if instead the fingers curl opposite the direction of rotation, the sense is called “left-hand.” Put another way, those ellipses whose patterns are traced out in a clockwise fashion (when looking in the direction of propagation) are *right-hand*. Those whose patterns are traced out in a counterclockwise fashion are *left-hand*.

Figure 2.2: Projection of the time-independent electric field vector phasor to the $x - y$ plane yields the *polarization ellipse*, a geometric representation of polarization. In this representation, polarization can be completely characterized by an ellipticity angle χ , an orientation angle Ψ , and a rotation sense, indicated by an arrow. *Figure adapted from Wikipedia Commons.*



¹Note that in polarimetry, many statements vary depending upon the definitions and conventions that are adopted. Orientation angle bounds are one such statement: Ψ can be defined as the angle the ellipse’s semi-major axis forms with the positive y -axis. Under this convention, $-90^\circ \leq \Psi \leq 90^\circ$.

These analytical and geometrical representations are equivalent. A polarization's ellipse maps to a Jones vector by the following relation:


$$\vec{\mathbf{E}} = Ae^{j\delta_x} \begin{bmatrix} \cos \Psi & \sin \Psi \\ \sin \Psi & \cos \Psi \end{bmatrix} \begin{bmatrix} \cos \chi \\ j \sin \chi \end{bmatrix} \quad (2.6)$$

where A is the ellipse amplitude as shown in Figure 2.2 and Ψ is understood to conform to the convention that $0^\circ \leq \Psi \leq 180^\circ$. The profound and detailed derivation of this relationship is available in Huynen (1970) for the keenly interested reader. For the purposes of this dissertation, it serves the reader to understand only that the two representations are, up to an absolute phase term, equivalent.

Poincarè's sphere and the canonical polarization states

Because all polarization ellipses have ellipticity $0^\circ \leq \chi \leq 45^\circ$, orientation $0^\circ \leq \Psi \leq 180^\circ$, and one of two senses, the infinite space of all polarization states with $A = 1$ can be mapped to the surface of a sphere of unit radius (Deschamps, 1949).² The Poincarè sphere representation of polarization has some consequences that make it a useful tool for considering polarization space, which can otherwise be difficult to conceptualize.

Figure 2.3: The infinite space of polarization states can be mapped to Poincarè's sphere. The northern hemisphere corresponds to right-hand sense polarizations and the southern to left-hand sense polarizations. Ellipticity and orientation map to latitude and longitude, respectively. The canonical polarizations are indicated on the sphere along with the abbreviations this dissertations will use for them. *Figure adapted from Boerner (2007).*



²Though Deschamps was the first to point out this relation for a radar-focused audience, the sphere had been used in optics since 1892, when it was introduced by Henri Poincarè—hence the name *Poincarè's sphere*.

The mapping is given by a few simple rules. Right-hand sense polarizations map to latitude $\phi_{lat} = 2\Psi^\circ$ on the sphere. Left-hand sense polarizations map to latitude $\phi_{lat} = -2\Psi^\circ$ on the sphere. Hence, left-hand polarizations map to the southern hemispherical surface, while right-hand polarizations map to the northern hemispherical surface.³ At the extremes are right-hand circular (R) and left-hand circular (L), which map to the north and south poles respectively as shown in Figure 2-3. Polarizations are mapped to a longitude using the relation $\phi_{lon} = 2\chi^\circ$, where the longitudinal coordinate indicates degrees east of horizontal polarization (H), which maps to $\phi_{lon} = 0^\circ$. Thus, vertical polarization (V) maps to $\phi_{lon} = 180^\circ$. Because their ellipticity angles are 0° , H and V map latitudinally to the equator, as do all linear polarizations. Of the infinitely many linear polarizations, two in particular other than H and V are worth noting: those whose orientations are 45° slant linear (X^+), and 135° slant linear (X^-), which map to $\phi_{lon} = 90^\circ$ and $\phi_{lon} = 270^\circ$, respectively.⁴

The endpoints of any diameter of Poincaré's sphere correspond to orthogonal polarizations. (H is orthogonal to V, R is orthogonal to L, and so on.) Hence, the endpoints of any diameter on Poincaré's sphere comprise a basis for polarization space and can thus be used to measure the polarization of any backscattered wave. The sphere also shows that six polarizations—H, V, X^+ , and X^- , R, L—comprise the simplest minimally dense uniform sampling of polarization space, as shown in Figure 2-3. These six polarizations are sometimes referred to as the canonical polarization states, and they are the ones most often utilized in practice.

³This is another of the cases in which adopted convention can change the mathematical statement (and confuse the newcomer to the field). If the convention is adopted that the y time-phase quadrature component lags rather than leads the in-phase x component for right hand circular polarization, then right-hand polarizations map to the southern hemisphere. The Jones vectors for R and L are also reversed under this convention.

⁴Note that 135° slant linear is equivalent to -45° slant linear; hence the ‘-’ superscript in the abbreviation for 135° slant linear.

Polarization State	Unit Jones Vector	Ellipticity χ°	Orientation Ψ°
Horizontal (H)	$\vec{\mathbf{E}}_{\mathbf{H}} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$	0	0
Vertical (V)	$\vec{\mathbf{E}}_{\mathbf{V}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$	0	90
Slant linear 45° ($\mathbf{X}^+$)	$\vec{\mathbf{E}}_{\mathbf{X}^+} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$	0	45
Slant linear 135° ($\mathbf{X}^-$)	$\vec{\mathbf{E}}_{\mathbf{X}^-} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}$	0	135
Right-hand circular (R)	$\vec{\mathbf{E}}_{\mathbf{R}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ j \end{bmatrix}$	45	[0,180]
Left-hand circular (L)	$\vec{\mathbf{E}}_{\mathbf{L}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -j \end{bmatrix}$	45	[0,180]

Table 2.1: The six canonical polarization states, their Jones vectors, and the associated polarization ellipse ellipticity and orientation angles are given.

Unit amplitude Jones vectors⁵ and associated polarization ellipse parameters for the canonical polarization states are given in Table 2.1. While polarization ellipse characteristics of backscattered waves are not directly measurable by radar, a radar that is dual-pol on receive can measure the Jones vectors of incoming waves.⁶

2.1.2 Measuring polarization

The amount of wave energy that will couple into a receive antenna is proportional to the degree to which the wave and antenna polarizations are aligned (Balanis, 2012). Poincaré’s sphere is a useful tool for conceptualizing polarization alignment and wave-antenna coupling. If the wave and antenna polarizations are on opposite sides of the sphere with 180° of angular separation between them, no coupling will occur; if the polarizations have an angular separation of 90° in any direction, half the energy will couple in; if the polarizations are perfectly aligned, all of the energy will couple in.

⁵Jones vectors normalized to unit amplitude are convenient mathematical descriptors of antenna polarizations, as will shortly be seen.

⁶The Jones vector characterizes polarization only for completely polarized waves. If the wave is only partially polarized (i.e., it has some constituent component that is randomly polarized), then the polarization must be characterized by a *Stokes vector*. The Stokes vector is a power representation of the wave’s polarization state, and hence does not preserve absolute phase.

Mathematically, the voltage induced in a receive antenna by an incoming wave is the dot product of the wave and antenna polarizations. Defining for use in the next sections the following Jones vectors:

$$\begin{aligned}\vec{\mathbf{H}}_t &= \begin{bmatrix} h_{t1} \\ h_{t2} \end{bmatrix} = \text{transmit antenna polarization} \\ \vec{\mathbf{E}}_{inc} &= \text{polarization of the EM wave incident on the target} \\ \vec{\mathbf{E}}_{scat} &= \text{polarization of the EM wave scattered by the target} \\ \vec{\mathbf{H}}_r &= \text{receive antenna polarization}\end{aligned}$$

we can write that

$$e = \vec{\mathbf{H}}_r^T \vec{\mathbf{E}}_{scat} \quad (2.7)$$

where the superscript T is the transpose operator and e is the complex-valued voltage that couples into the receive antenna.⁷

Because polarization is a vector-valued quantity, there must be an orthogonally polarized receive antenna $\vec{\mathbf{H}}_{r\perp}$ in order to measure any vector in polarization space such that

$$\vec{\mathbf{E}}_{meas} = \begin{bmatrix} e \\ e_{\perp} \end{bmatrix} = \begin{bmatrix} \vec{\mathbf{H}}_r^T \vec{\mathbf{E}}_{scat} \\ \vec{\mathbf{H}}_{r\perp}^T \vec{\mathbf{E}}_{scat} \end{bmatrix} \quad (2.8)$$

where each component of $\vec{\mathbf{E}}_{meas}$ captures exactly the backscattered wave energy that is not captured in the other component (Huynen, 1965). Each received voltage thus comprises one complex coefficient of a polarization vector measurement made using the basis $\vec{\mathbf{H}}_r, \vec{\mathbf{H}}_{r\perp}$.

An important consequence of having measured $\vec{\mathbf{E}}_{scat}$ in one polarization basis is that, because polarization space is spanned by these measurements, $\vec{\mathbf{E}}_{meas}$ can be transformed to any other polarization basis. In other words, any pair of orthogonal receive antenna polarization measurements can be synthesized in post-processing, allowing adaptive po-

⁷Note that Equation 2.7 shows why Jones vectors are normalized to unit length when they are antenna polarization descriptors: not doing so allows for mathematical amplification of the energy just by virtue of coupling into an antenna, which is clearly incorrect.

larimetric beamforming on receive (Poelman, 1977). While such a dual-pol on receive radar can measure and synthesize any receive polarization, it cannot measure completely the information in an object's RCS-scattering information dimension.

2.1.3 The scattering matrix and its power representations

To see why the preceding statement is true, consider a point target object located at a range r from the antenna. Factoring in the phase term related to propagation over this distance and attenuation of the EM wave energy in one direction we can write

$$\vec{\mathbf{E}}_{inc} = \frac{e^{jkr}}{r} \vec{\mathbf{H}}_t \quad (2.9)$$

where, as before, k is the wavenumber. Because the scattered wave polarization is also a two-element complex vector, the response of the target to the incident energy can be represented by a matrix $\mathbf{S}$ such that

$$\vec{\mathbf{E}}_{scat} = \mathbf{S} \vec{\mathbf{E}}_{inc} \quad (2.10)$$

from which it follows that $\mathbf{S}$ is 2 x 2 and complex-valued.

This matrix is the target *scattering matrix* (SM), and it completely characterizes the electromagnetic scattering properties of a target at a given aspect angle and frequency (Sinclair, 1950). Generically, the SM can be written elementwise as

$$\mathbf{S} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} = \begin{bmatrix} |\alpha_{11}|e^{j\phi_{11}} & |\alpha_{12}|e^{j\phi_{12}} \\ |\alpha_{21}|e^{j\phi_{21}} & |\alpha_{22}|e^{j\phi_{22}} \end{bmatrix} \quad (2.11)$$

where S_{11} and S_{22} are the co-polarized object responses and S_{12} and S_{21} are the cross-polarized object responses. That is, the former two elements quantify the proportion (in amplitude and phase) of each of two orthogonally polarized incident waves (with polarizations indicated by subscripts 1 and 2) that will retain polarization after being scattered by the object. The latter two elements quantify the proportion of the orthogonally po-

larized incident waves that will depolarize, or assume the orthogonal polarization, after being scattered by the object.

It is by now obvious to the keen reader that a single-pol on transmit, dual-pol on receive radar cannot fully measure the information in an object's RCS-scattering information dimension: one cannot determine four complex values by measuring only two. To show this explicitly, ignoring the phase and attenuation due to target distance and combining Equations 2.9 and 2.10, we have:

$$\vec{\mathbf{E}}_{scat} = \begin{bmatrix} S_{11}h_{t1} + S_{12}h_{t2} \\ S_{21}h_{t1} + S_{22}h_{t2} \end{bmatrix}. \quad (2.12)$$

Now, consider for example one of the most common cases: $\vec{\mathbf{H}}_t = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, corresponding to a horizontally polarized radar. The second terms in each summation in Equation 2.12 go to zero, so two of the SM elements are not measured at all. A similar argument can be made for the vertically polarized case. Even if all of the scattering matrix elements contribute to $\vec{\mathbf{E}}_{meas}$, as is the case for antenna polarizations that have non-zero amplitude in each component direction (like X^+ , X^- , R, or L), the system of equations is underdetermined. Thus, there is no single transmit polarization that will allow precise measurement of $\mathbf{S}$ by a radar that is dual-pol only on receive. However, if the radar can also interrogate the object using a pair of orthogonal transmit polarizations, the system of equations will be fully determined.⁸

Before proceeding, it should be noted that the elements of $\mathbf{S}$ depend upon the radar's transmit and receive polarization bases (Lee and Pottier, 2009). A given object will respond differently to differently polarized incident energy. Thus, $\mathbf{S}$ sensed in one basis will not be equal to $\mathbf{S}$ sensed in another basis, even for the same frequency and viewing

⁸There is ongoing research to explore computation of $\mathbf{S}$ to a tolerable degree of goodness using measurements from a radar that is single-pol (slant linear or circular) on transmit and dual-pol on receive (Nord et al., 2009). This area of study is called *compact polarimetry*. This dissertation focuses on precise measurement of $\mathbf{S}$ using a fully polarimetric radar, though the results presented in Chapter 5 indicate that substantial performance improvement may be had in maritime surveillance radar by using a form of compact polarimetric radar.

geometry. To make this polarization dependence explicit, it is customary to denote the elements of $\mathbf{S}$ by substituting for its subscripts the basis abbreviations. So, for example, $\mathbf{S}$ sensed by a radar that uses the linear basis is denoted

$$\mathbf{S}_{HV} = \begin{bmatrix} S_{HH} & s_{HV} \\ s_{VH} & s_{VV} \end{bmatrix}$$

while in the slant linear basis it is

$$\mathbf{S}_{X+X-} = \begin{bmatrix} S_{X+X+} & S_{X+X-} \\ S_{X-X+} & S_{X-X-} \end{bmatrix}$$

and in the circular basis it is

$$\mathbf{S}_{RL} = \begin{bmatrix} S_{RR} & S_{RL} \\ S_{LR} & S_{LL} \end{bmatrix}$$

and in general,

$$\mathbf{S}_{HV} \neq \mathbf{S}_{X+X-} \neq \mathbf{S}_{RL}.$$

As noted in Chapter 1, this dissertation focuses on the monostatic backscattering case. Thus, when the effects of Faraday rotation are absent, the SM in any basis is symmetric by reciprocity—that is, $S_{12} = S_{21}$ (Huynen, 1970).

Targets that correspond to spatially and/or temporally varying stochastic processes, such as the sea surface, can lead to partially polarized returns. Though averaging of SM measurements in time and space is often useful, the polarization of such scatterers can fluctuate rapidly, with absolute phase varying uniformly over $[0, 2\pi]$ (Huynen, 1970). The mean of the SM elements for such scatterers will be zero due to this randomness of phase. Thus, expression of the SM in an equivalent second-order power representation—which does not preserve absolute phase—can be useful under such conditions. The most commonly used power representations are covariance and coherency matrices. To form these

matrices, define the following vectorizations of constituent SM elements:⁹

$$\vec{\mathbf{L}} = \begin{bmatrix} S_{HH} \\ \sqrt{2}S_{HV} \\ S_{VV} \end{bmatrix} \quad (2.13)$$

$$\vec{\mathbf{P}} = \frac{1}{\sqrt{2}} \begin{bmatrix} S_{HH} + S_{VV} \\ S_{HH} - S_{VV} \\ 2S_{HV} \end{bmatrix} \quad (2.14)$$

where the scaling is such that the translation from matrix to vector leaves the total power in the elements unchanged. The vectorizations are chosen to be expansions of the SM in orthogonal matrix basis sets that are meaningful in terms of target scattering characteristics (Huynen, 1970; Cloude and Pottier, 1996). The covariance matrix is defined in terms of $\vec{\mathbf{L}}$ as

$$\begin{aligned} \mathbf{C} &= \langle \vec{\mathbf{L}} \cdot \vec{\mathbf{L}}^{*T} \rangle \\ &= \begin{bmatrix} \langle |S_{HH}|^2 \rangle & \sqrt{2}\langle S_{HH}S_{HV}^* \rangle & \langle S_{HH}S_{VV}^* \rangle \\ \sqrt{2}\langle S_{HV}S_{HH}^* \rangle & 2\langle |S_{HV}|^2 \rangle & \sqrt{2}\langle (S_{HV}S_{VV}^*) \rangle \\ \langle S_{VV}S_{HH}^* \rangle & \sqrt{2}\langle S_{VV}S_{HV}^* \rangle & \langle |S_{VV}|^2 \rangle \end{bmatrix} \end{aligned} \quad (2.15)$$

where $*T$ indicates the conjugate transpose and $\langle \rangle$ indicates averaging in either the spatial or temporal dimensions. The coherency matrix is then defined in terms of $\vec{\mathbf{P}}$ as

$$\begin{aligned} \mathbf{T} &= \langle \vec{\mathbf{P}} \cdot \vec{\mathbf{P}}^{*T} \rangle \\ &= \frac{1}{2} \begin{bmatrix} \langle |S_{HH} + S_{VV}|^2 \rangle & \langle (S_{HH} + S_{VV})(S_{HH} - S_{VV})^* \rangle & 2\langle (S_{HH} + S_{VV})S_{HV}^* \rangle \\ \langle (S_{HH} - S_{VV})(S_{HH} + S_{VV})^* \rangle & \langle |S_{HH} - S_{VV}|^2 \rangle & 2\langle (S_{HH} - S_{VV})S_{HV}^* \rangle \\ 2\langle S_{HV}(S_{HH} + S_{VV})^* \rangle & 2\langle S_{HV}(S_{HH} - S_{VV})^* \rangle & 4\langle |S_{HV}|^2 \rangle \end{bmatrix} \end{aligned} \quad (2.16)$$

where again $*T$ indicates the conjugate transpose and $\langle \rangle$ indicates averaging in either the spatial or temporal dimensions.

⁹Definition is under the monostatic BSA convention and in the H-V basis, though definitions can also be made under other scattering coordinate systems and in other bases.

2.1.4 Measurement of the scattering matrix

Receiving orthogonal polarizations, as indicated above, is as straightforward as having two orthogonal receive antennas (and as many receive channels). The antennas do the work of separating the incident energy into its constituent basis components. However, transmitting two orthogonally polarized waves to interrogate an object is a more difficult proposition. If two EM waves are emitted by a radar simultaneously, the waves are emitted in superposition and interact with the target simultaneously, so the scattered energy is not separable into individual SM contributions by the receive antennas alone.

To illustrate this with an example, consider without loss of generality to any other basis a linearly polarized radar from which H and V transmit waves are emitted simultaneously and in phase. Their energy combines in superposition, effectively forming $\sqrt{2}\vec{\mathbf{E}}_{X+}$. This incident wave interacts with and is scattered by the object, and results in a scattered wave identical to that in Equation 2.12 with $h_{t1} = 1$ and $h_{t2} = 1$. Again ignoring attenuation and round-trip phase delay, the scattered energy couples into the H and V receive antennas as

$$\vec{\mathbf{E}}_{meas} = \begin{bmatrix} e_H \\ e_V \end{bmatrix} = \sqrt{2} \begin{bmatrix} S_{HH} + S_{VH} \\ S_{HV} + S_{VV} \end{bmatrix}. \quad (2.17)$$

The antennas alone cannot fully distinguish between SM elements. Another degree of separability is required.

In general, there are two transmit schemes for fully polarimetric radar: *simultaneous* and *alternating*. Each scheme induces at least some signal separability in different ways. As shown in Figure 2-4, in the simultaneous case both polarizations are transmitted at the same time. However, the transmit waveforms w_{T1} and w_{T2} are also encoded with orthogonal codes; one such orthogonal encoding is linear frequency modulated (LFM) “chirp” signals with opposite chirp direction—i.e., an upchirp and a downchirp (Giuli et al., 1993). Each receive channel in the simultaneous case measures one co-polarized response and one cross-polarized response in superposition, as shown in the figure. But




Figure 2-4: A simultaneous transmit scheme for a linearly polarized fully polarimetric radar is shown. Orthogonally encoded transmit waveforms enable some additional separability of the superimposed co- and cross-pol signals in each receive channel, but that separability is often insufficient to accurately measure cross-polarized object responses that are close in range to large co-polarized object responses. *Figure adapted from work by David C. Mooradd, MIT LL.*

correlating the superimposed signal in each channel against both of the transmit waveform encodings w_{T1} and w_{T2} allows some degree of separation between the co- and cross-pol responses. However, encoding orthogonal waveforms in fast-time¹⁰ leads to limited signal separability (Krieger et al., 2012). That is, the cross-correlation between so-called orthogonal waveforms does not actually offer enough isolation to permit measurement of the cross-pol SM elements in the presence of strong co-pol returns. It can be shown that the cross-correlation isolation of fast-time orthogonal waveforms is proportional to the waveform’s time-bandwidth product (Mooradd, 2016), so this issue becomes increasingly egregious for waveforms that are narrowband or that use a short pulselength.

Because of this serious shortcoming, alternating fully polarimetric radar is the standard. As shown in Figure 2-5, on odd pulses only one of the polarizations is transmitted,

¹⁰The meaning here is that encoding is done in range-time dimension; that is, each transmitted pulse is orthogonal to the pulse in the other channel that is simultaneously transmitted.




Figure 2-5: An alternating transmit scheme for a linearly polarized fully polarimetric radar is shown. As shown in (a), only H is transmitted on odd pulses. On even pulses, V is transmitted. This scheme assumes the SM is stable over two pulse repetition intervals, but the antennas provide sufficient orthogonality to measure the full matrix. *Figure adapted from work by David C. Mooradd, MIT LL.*

so the co- and cross-pol responses do not come back in superposition. Note that the transmit waveforms w_{T1} and w_{T2} in this case need not be orthogonally encoded because the orthogonal antennas are able to separate the responses, resulting in measurement of

two of the SM elements. On even pulses, the other polarization is transmitted, and the process repeats. Because measurement of the SM in the alternating scheme requires two separate pulses, there is the potential that an object's scattering matrix does not remain stable from pulse to pulse. Pulse repetition intervals in maritime surveillance radar are typically less than a millisecond; the assumption that a target SM remains stable over that time is usually valid (Blejer, 2016a).

One of the contributions of this dissertation is theoretical application of a simultaneous transmit MIMO radar transmit scheme to achieve scattering matrix component separability in the slow time dimension; this derivation will be presented in Section 3.2.

2.1.5 The RCS-scattering information dimension in post-processed data

Now that we have established that the RCS-scattering information dimension is captured by the SM and have discussed measurement thereof, the reader will benefit from considering momentarily what this actually means in processed radar data. The baseline coherent radar signal processing chain typically produces two data products, as shown in




Figure 2-6: The data products generated by the baseline coherent radar signal processing chain for a vertically polarized radar (looking at sea clutter) are shown.

Figure 2·6: range-time intensity (RTI) data and range-Doppler (RD) data, upon which other signal processing operations are performed.

Because a fully polarimetric radar measures the SM at every resolution cell and across transmit pulses, such a radar’s RTI and RD also have a third dimension: the RCS-scattering information dimension. To form this dimension, the signal processing chain executes once on each channel—i.e., once on HH, HV, VH, and VV—rendering one RTI and RD per channel. For each coherent processing interval (CPI), or time period during which a given number of pulses is transmitted and their returns measured, the RTIs and the RDs can be stacked to form 3-D data matrices as shown in Figure 2·7. These matrices can be manipulated in various ways to achieve a number of different objectives.




Figure 2·7: The polarimetric radar data cube for one CPI that can be formed using single channel RTIs (or RDs) is shown graphically, as measured in the linear polarization basis. The cube is formed by stacking the processed RTIs (or RDs) from each channel along the polarimetric third dimension, which corresponds to the RCS-scattering information dimension.^a The elements of any vector along this dimension thus comprise the SM measurement at that range and for that pulse (or Doppler frequency).

^aNote that under the BSA, $S_{12} = S_{21}$, so $S_{HV} = S_{VH}$.

An important consequence of having measured the SM in a given polarization basis is that the SM can be transformed to any other basis. In other words, while a scattering matrix measurement depends on the antenna polarizations used, once the measurement is known in a known basis, it is possible to compute the SM that would have been measured by any other orthogonal radar polarizations (Poelman, 1981). This technique is known as *polarization synthesis*, and effectively, it renders the RCS-scattering matrix dimension infinite.

2.1.6 Relevant application areas

Radar polarimetry is a vast, deep, and complex field with a wide range of applications, many of which have yet to be tapped. While still not fully understood, the information contained in the RCS-scattering dimension has been leveraged effectively in some application areas, particularly SAR and Doppler weather radar. Part of the SAR success is due to the clear candidacy of SAR imagery for polarimetric analysis, as discussed in Chapter 1. Another part is due to the experimental SAR community’s collective recognition of polarimetry’s value, which has resulted in both polarimetric SAR system development and a fair amount of SAR data being made publicly available for additional research (Lee and Pottier, 2009). Yet another part may be due to the fact that it is usually not expected that SAR data will be optimally processed in real time.

Though polarimetric techniques can be computationally intensive, their potential to enhance radar performance in other applications—even those where real-time performance is expected—should not be overlooked. The computational power of hardware continues to increase as its cost decreases. In the meantime, suboptimal yet practical polarimetric approaches may offer the ability to improve radar performance in applications like object detection, tracking, and identification in a real-time capacity.

2.2 Sea clutter

Perhaps the most challenging of radar environments in which real-time performance is often required is the maritime setting. In any state other than a perfectly calm, mirror-like one, the complex sea surface interacts with radar energy in ways that are notoriously difficult to model or predict. The first part of this section presents some important aspects of sea clutter phenomenology. There is a vast body of literature on this topic; the interested reader is referred to Long (2001), which is a fairly exhaustive reference. The second part discusses the challenge of detection in low grazing angle sea clutter, includ-

ing a brief literature survey enumerating some of the polarimetric and non-polarimetric approaches to the sea clutter problem that have been put forth to date.

2.2.1 Sea clutter phenomenology

The visible surface of the sea consists of large waves, called *gravity waves*. Gravity waves can propagate as fast as the wind is blowing provided that the wind has blown for sufficient time and over sufficient distance. Riding atop these waves are smaller, ripple-like waves dubbed *capillary waves*, which result from surface tension forces as a wave propagates through water. Both gravity and capillary waves backscatter radar energy. Gravity waves backscatter energy from their faces—hence, look direction of a radar with respect to the wind is a key factor in this type of scattering—while the many small capillary waves are resonant Bragg scatterers. Gravity waves contribute a smooth, mean modulation to the surface called the *texture* component, while capillary waves contribute a *speckle* component. Sea surface roughness is quantified by the *sea state*, a value that depends on the average height of the highest one-third of gravity waves, with height measured wave crest to wave trough (Long, 2001; Ward et al., 2006).

Though sea surface scattering is not fully understood, there have been many dozens of experiments to date using horizontally and/or vertically polarized radars, yielding empirical data that reveals some clear trends (Nathanson et al., 1991). Chief among these trends in the low grazing angle regime is that the higher the sea state, the stronger the sea clutter. So, detection performance—which depends on the signal-to-clutter (SCR) ratio—is inversely proportional to sea state. Moreover, clutter is highly dependent on radar polarization. Nathanson observes that sea clutter returns measured with different polarizations are, to a certain extent, independent. Evidence for this view can be found in the Doppler shift of clutter, which is generally more significant for horizontal polarization

than for vertical, leading researchers to the conclusion that different polarizations yield returns from different types of scattering mechanisms (Ward et al., 2006).¹¹

It is also well known that the amplitude and spatial distributions of the returns are also polarization-dependent. Vertical polarization leads to higher mean backscatter levels than those measured with horizontal polarization; this phenomenon becomes more pronounced at lower sea states, grazing angles, and frequencies. Horizontal polarization, on the other hand, leads to a “spikier” texture in the clutter. *Sea spikes* are defined as unnaturally high, object-like amplitude returns that may persist in the worst case for several seconds. Spikes are particularly prevalent at low grazing angles, and a high number of spikes often leads to an inundation of false alarms (Ward et al., 2006). The spike problem becomes more prevalent for finer range resolution (higher bandwidth).

For the foregoing reasons, the maritime setting makes radar system design a difficult task. The situation gets bleaker: accurate models for mean radar backscatter from the sea, defined as the sea RCS per square meter of area illuminated by a radar’s beam and denoted σ_0 [dBsm/m²], are necessary in order to develop radar specifications based on a minimum detectable object RCS in a given setting. Existing models, while numerous, are disparate in the low and very low grazing angle regimes, as shown in Figure 2.8. Moreover, because available data comprises largely horizontal or vertical polarizations that are often collected independently during the same experiments, there are no existing mean backscatter models for polarizations other than horizontal or vertical. Hence, it is difficult to arrive at a clear choice of a fixed antenna polarization if optimum radar performance across a range of sea conditions and geometries is the objective (Nathanson et al., 1991).

¹¹The Doppler differential between HH and VV scattering is the cornerstone of the Polarimetric Co-location Layering algorithm. Past literature focused on this particular topic will be reviewed in greater detail in Chapter 5.




Figure 2-8: Shown are several clutter models' predictions over two sea states for mean radar backscatter per unit area of sea illuminated by an X-band radar beam. The variation across models is evident; however, they all reflect the fact that the sea's backscatter increases at higher sea states. The models included are the low grazing version of the Georgia Institute of Technology (GIT) model (Antipov, 1998); the Sittrop (SIT) model (Sittrop, 1977); the Royal Radar Establishment (RRE) model (Raynal and Doerry, 2010); the Technology Service Corporation (TSC) model (Antipov, 1998); the low grazing angle version of the Naval Research Laboratory (NRL) model (Gregers-Hansen and Mital, 2009); and the Hybrid model (Reilly and Dockery, 1990).

2.2.2 Detection in sea clutter

It is also difficult to predict detection performance in sea clutter for a given radar system. Application of statistical detection theory to radar performance prediction in any type of clutter requires good statistical models of the clutter and of the object being detected. The polarimetry-agnostic Swerling models are often used to capture the distribution of object amplitude fluctuations. There has been no effort to adapt these models to classes that encompass the polarization dependence of amplitude returns, nor has work been done toward developing a separate set of models that does account for this dependence (Watts, 2008).

As for sea clutter models, the aforementioned sets of empirical data have been used to identify families of probability distributions that accurately characterize amplitude statistics for both horizontal and vertical polarizations. The most promising of these to date is the compound K distribution (Ward et al., 2006), but it does not represent clutter spikes, so detection performance will be much degraded at lower grazing angles and/or for higher bandwidths (Watts, 2008). Moreover, at low grazing angles, the inhomogeneity of the clutter makes estimation of model parameters very difficult in general (Greco et al., 2008); the standard CFAR detection approach is unable to perform reliably (Watts, 2008). Hence, approaches to mitigate the impact of sea clutter in this low grazing geometries is an area ripe for research, a fair amount of which has been done to date.

Much of this research utilizes a single publicly available polarimetric data set: the IPIX radar Dartmouth data (Haykin et al., 2001). McMaster University’s generous publication of a subset of this data has allowed interested researchers access to a field not usually accessible other than by those with access to an experimental polarimetric radar. The IPIX data subset that is freely available comprises 17 files and a variety of sea states. However, the files feature a single type of canonical point target object (a foil covered beach ball), a single resolution that is low by current standards for the radar band in question (30 meters at X-band), and a very short range window (210 meters comprising 7 range resolutions and 14 range bins given the data’s sampling rate). While some researchers have reported modest success with enhancing object detectability in sea clutter using this data, whether these techniques would have any effect on distributed or extended objects at higher resolutions and across a range of viewing geometries is uncertain.

Polarimetric approaches

Early theoretical approaches to leveraging polarimetry for improved object detection assumed *a priori* knowledge of the statistical parameters of both the object and clutter

(Ioannidis and Hammers, 1979; Wanielik and Stock, 1990; Novak et al., 1993). Other approaches to polarimetric detection assume Gaussianity and/or homogeneity of the interfering signal. Still other approaches rely on secondary data that can be used as training data to estimate the parameters that will be used in computation of a detection statistic (De Maio and Alfano, 2003). All of these assumptions are unrealistic for sea clutter, whose inhomogeneity and non-stationarity is well known (Greco et al., 2008). Notable recent approaches leverage a polarimetric generalized likelihood ratio detector, which is robust to inhomogeneity and non-stationarity (Park and Wang, 2006; Hurtado and Nehorai, 2008). However, the assumption in this case is that the object is a deterministic point target object. This is not the case for the vast majority of real objects.

There are also several non-statistical approaches to improving detection in sea clutter. One such approach is the span algorithm, which simply sums the power received in each of the channels. However, Novak showed that because this approach does not leverage the phase information of polarimetric radar returns, it falls well short of the optimal polarimetric detector for which object and clutter parameters are known (Novak et al., 1989). There have also been attempts to use power matrix decompositions to classify returns as either object plus clutter or clutter only (Wu et al., 2011; Kim et al., 2014). However, the authors in these cases first reject all sea spike candidates before applying the algorithm, which clearly biases the results: The metric used to identify sea spikes for rejection does not preclude rejection of objects of interest.

Nathanson observes that it is not clear that polarization would be of any benefit in a general object detection in sea clutter application (Nathanson et al., 1991). Based on the research to date, this observation still largely rings true. Yet, van Zyl shows that the ocean’s radar returns are highly polarized, indicating that there is indeed a means to suppress these returns using polarimetric techniques (Van Zyl et al., 1987).

Non-polarimetric approaches

Non-polarimetric approaches to enhancing object detection in sea clutter abound. The range of techniques is impressive, though few have gained traction in the radar community at large (Watts, 2008); all have drawbacks that may outweigh any benefit. Chief among these is a class of algorithms known as “track before detect”; there are numerous pieces in the literature that offer up some variant of this class (Tonissen and Bar-Shalom, 1998; Boers and Driessen, 2001). The name is fairly self-explanatory in that detections are not declared as such until tracks corresponding to a spatially logical sequence of associated detections have been established on an object, placing the onus on the tracking algorithm in a radar rather than on its detector. The algorithms are promising, but have drawbacks: 1) they may be insufficient for fleeting objects; 2) they may fail for stationary objects that don’t “track”; 3) they may fail for high-speed objects that move faster than expected; 4) they may fail to distinguish between moving gravity wave components of the sea at low grazing angles; and 5) many of the tracking algorithms that show the most promise are much more computationally intensive than tracking algorithms currently used in operational radars (Watts, 2008; Davey et al., 2008).

Other approaches that show practical promise leverage the fact that sea clutter maintains coherence on only relatively short timescales (Antipov, 1998). One such approach is longer scale signal averaging (Panagopoulos and Soraghan, 2004). However, this approach risks averaging out transient object aspects that are desirable. Additionally, the texture component of the sea surface does not decorrelate as quickly as the speckle component (Antipov, 1998). Another approach involves CFAR thresholding by using prediction of sea clutter locations based on locations of clutter in previous radar scans rather than by model-based parameter estimation (Ward et al., 2007). However, this involves the introduction of memory into the radar’s range-Doppler image formation process and requires the ability to differentiate between sea and non-sea returns to enable prediction of sea clutter’s movement. The latter will be particularly challenging at low grazing angles.

There is also no shortage of unconventional approaches. Some call for leveraging time-frequency analysis to interpret different objects' frequency signatures against that of the ocean (Panagopoulos and Soraghan, 2004); however, the ocean's frequency signatures are not well understood themselves (Watts, 2008). Other approaches are quite clever, but move quickly outside the realm of applicability in the practical sense, at least for current and near-term systems. These include various transform methods (Carretero-Moya et al., 2009; Davidson and Griffiths, 2002; Guan et al., 2012) and multifractal analysis techniques (Hu et al., 2006). One unconventional area where there is great promise in the future is application of machine learning techniques to radar data, even in a real-time sense (Haykin, 2006; Vicen-Bueno et al., 2009).

The preceding literature review of both polarimetric and non-polarimetric approaches to enhancing object detection in sea clutter is not exhaustive, particularly in the non-polarimetric case. It is intended to emphasize to the reader that the object detection problem is one of ongoing study. The focus of this dissertation is on polarimetric techniques. Additional literature is reviewed in Chapter 5, where proximity to the topics discussed therein will be of greater service to the reader.

Chapter 3

Four Eyes and the Point de Chene Dataset

The first section of this chapter discusses Four Eyes, the X-band polarimetric radar assembled to collect a low grazing angle sea clutter dataset. The second section discusses the weeklong field test that took place in October 2015 in Rockport, Massachusetts. The third section discusses the waveforms and transmit schemes that were programmed into Four Eyes; this section includes a derivation showing that a slow-time phase encoding provides a means of precisely measuring the SM using a simultaneous polarimetric transmit scheme. The chapter concludes with a section summarizing the Point de Chene Dataset.

3.1 Four Eyes polarimetric radar system

Four Eyes is a fully polarimetric, transportable, X-band radar assembled almost entirely from COTS components. A system block diagram is shown in Figure 3-1. Four Eyes has two distinct but identical transmit chains and separate, dual-polarized, parabolic dish transmit and receive antennas. The channel 1 transmit chain consists of an Agilent N8241A arbitrary waveform generator (AWG), which has up to 400 MHz of bandwidth, feeding into an Agilent E8267D vector signal generator (VSG) configured to act as a local oscillator by taking in I/Q data and mixing with a continuous tone at 9.705 GHz. The upconverted signal feeds into a 1500 W traveling wave tube (TWT) amplifier, then into the horizontal transmit antenna port. The channel 2 transmit chain is identical, but feeds a vertically polarized transmit antenna port. The receive antenna is identical to the transmit antenna. The receive antenna signal outputs (one horizontally polarized, one




Figure 3.1: A block diagram for the Four Eyes X-band fully polarimetric radar is shown. The system features independent but identical transmit chains and separate but approximately co-located linear dual-polarized transmit and receive antennas. Timing hardware and signals are shown in red. Transmit hardware is shown in blue; receive hardware is green.

vertically polarized) feed into low-noise amplifiers (LNAs) and mix with a local oscillator tuned to 8.76 GHz for downconversion to intermediate frequency (IF) 945 MHz. The IF signals feed into a two-channel receiver that features a 400 MHz wide bandpass filter centered at IF. The receiver outputs are digitized by a two-channel analog-to-digital converter (ADC) at sampling frequency 1260 MHz. The data are quantized with 12-bit depth, though at the sampling frequency used and at the IF sampled, the effective bit depth reduces to just over 9. The data are recorded on 40 3TB hard drive disks (HDD) configured in a redundant array of independent disks (RAID) to allow for individual disk failure without incurring data loss.

System timing signals are generated on a Tektronix DTG5078, which triggers the AWGs, the VSGs, the TWTs, and the ADCs. The DTG, ADCs, and IF LO are also

Dual Transmitters:	
Carrier frequency	9.705 GHz (X-band)
Pulse repetition frequency	6250-12500 Hz
Pulse width	0.25-4 μ sec
Peak power	1500 W
Dual Receiver Chains:	
Intermediate frequency (IF)	945 MHz
IF bandwidth	400 MHz
Gain (inc. LNA)	29 dB
Noise figure	4 dB
Dual Antennas:	
Diameter	2 ft
Gain	33.75 dB
Beamwidth	3.7°
Cross-pol isolation, Ant. 1	20 dB
Cross-pol isolation, Ant. 2	26 dB
Polarization	Linear (H and V)
Data Acquisition:	
Sampling rate	1260 MHz
Vertical resolution	12 bits (9 effective)
Storage medium	40x 3TB HDDs in RAID-60

Table 3.1: Four Eyes COTS system configuration



(a) View of the full radar setup. System hardware is installed in a box truck. The antenna pedestal and transmitters are installed on a trailer behind the truck.



(b) Transmit and system timing hardware are in the far right rack; receiver, signal analysis, and power hardware are in the center; data acquisition hardware is at left.

Figure 3.2: Four Eyes on location in Rockport, MA. *Photographs taken by David C. Mooradd, MIT LL.*

slaved to a common 10 MHz reference oscillator provided by a Symmetricom GPS unit. Transmit-side operator control is via a laptop interface for loading transmit waveforms; the transmit operator also manually triggers system timing signals via the DTG. Operator control on the data acquisition side is via an interface that provides real-time feedback

as well as recording-control input. See Table 3.1 for a summary of system specifications as configured for the experiment. Photos of the system on location at the data collection site are shown in Figure 3.2.

3.1.1 Four Eyes’ antennas

Prior to field testing, Four Eyes’ twin two-foot parabolic dish antennas were measured in MIT LL’s compact antenna range, as shown in Figure 3.3, to ensure similar beamwidths and gains on both transmit and receive. The antenna range results in Figure 3.4 show that the antennas are nearly identical in beamwidth at 3.7° and gain at 33.75 dBi at the radar’s center frequency.

A common argument against using polarimetric radar involves what are perceived to be very difficult-to-achieve cross-polarization isolation requirements on the antennas (Blanchard and Newton, 1985). Because part of the




Figure 3.3: Four Eyes’ antennas mounted in the anechoic chamber. *Photograph by Paul Theophilakes, MIT LL.*

aim of this dissertation is to demonstrate that polarimetric radar can offer significant performance gains without imposing these kinds of impractical requirements, note that antenna range results show 20 dB of cross-polarization isolation on one of the antennas and 26 dB on the other; 20 dB, in particular, is a reasonable amount of cross-polarization isolation to expect from most practical antenna designs (Mooradd, 2016). Thus, any performance improvement resulting from polarimetric algorithms developed using Four Eyes data cannot be construed as due to unusually difficult-to-achieve cross-polarization isolation in the radar’s antennas.




Figure 3-4: Four Eyes' transmit and receive antenna beam patterns (left and right plots, respectively) at 9.8 [GHz] are shown. Agreement in beamwidth and gain is evident, as is 26 and 20 dB of cross-polarization isolation, respectively.

3.2 Waveforms and transmit schemes

All Four Eyes waveforms used in the October 2015 field test were *chirps*, or linear frequency modulated waveforms (LFMs). The LFMs were orthogonally coded across channels, so that the H channel transmitted downchirps, while the V channel transmitted upchirps. Specifically, after upconversion to the radar center frequency and using complex exponential notation,

$$\begin{aligned} w_{T1} &= e^{j(2\pi f_c t - \pi \frac{\beta}{\tau} t^2)}, |t| \leq \frac{\tau}{2} \\ w_{T2} &= e^{j(2\pi f_c t + \pi \frac{\beta}{\tau} t^2)}, |t| \leq \frac{\tau}{2} \end{aligned} \quad (3.1)$$

where t = time [sec]; β = bandwidth [Hz]; τ = pulselength [sec]; and f_c = center frequency [Hz].

The principal goal of Four Eyes waveform design was to create a suite of waveforms whose use on location would enable assessment of polarimetric algorithm robustness across a variety of radar range resolutions and pulse repetition frequencies (PRFs). Due to radar receiver and radio frequency authorization constraints, the bandwidth was capped

at 400 MHz. To get a large range of resolutions, each waveform was encoded with 4 MHz, 40 MHz, 150 MHz, and 400 MHz bandwidth versions, translating to range resolutions of 37.5 m, 3.75 m, 1 m, and 0.375 m, respectively.¹ Because the aim was to decimate pulses in slow-time in post-processing to simulate PRFs lower than the PRF employed during recording, the maximum possible PRF was used.

To determine this maximum PRF, the distance to the radar horizon is computed as

$$d_{hzn} = \sqrt{2Rh + h^2}$$

where h = height of the antennas above mean sea level and R = Earth's radius. The calculation yielded a line of sight horizon just beyond 12 km. However, initial link budget calculations using the most generous of sea clutter backscatter models showed that even at higher sea states, sea clutter return could not reasonably be expected beyond 6 km. Disregarding any potential for atmospheric ducting² and opting for the more conservative path, the 12 km maximum range was chosen, leading to the selected PRF of 12500 Hz using the familiar relation:

$$r_{max} = \frac{2}{c \cdot PRF}$$

where r_{max} = maximum range and c = speed of light.

Chosen pulselengths aimed to strike a balance between having visibility as near to the radar as possible while maximizing the sea clutter return by emitting more power. Initial calculations showed that pulselengths of 4 μ sec (resulting in a minimum range of 600 m) might be necessary to get strong sea clutter return as far out as the buoy at ~ 1.7 km, but seeing the strong clutter returns nearer the radar proved equally important. Hence, 2 μ sec pulselength LFM's were also encoded. Additionally, 0.25 μ sec pulselength

¹Note that this includes one resolution lower than and three resolutions higher than the 30 m resolution of the IPIX radar dataset files.

²Ducting is a regional phenomenon at low grazing angles whereby radar energy can get trapped in the atmosphere and thus propagate much further than the line of sight horizon.




Figure 3.5: Four Eyes’ polarimetric transmit schemes are depicted. Downchirps are in red; upchirps are in blue; upchirps with starting phase modulated according to pulse number are in gray. The schemes included the standard alternating scheme, a fast-time orthogonal simultaneous scheme limited by chirp waveform cross-correlation, and a variant of a slow-time phase encoding technique that has been used successfully in MIMO radar (Kantor and Davis, 2011). The implementation as shown for the latter is still limited by waveform cross-correlation in the fast-time dimension.

LFMs were encoded as calibration waveforms for recording files featuring the calibration trihedral, which was located only ~ 65 m from the radar.

The secondary goal of Four Eyes waveform design was that the dataset enable additional research—though not the focus of this dissertation—in understanding the limitations and benefits of simultaneous transmit schemes. As discussed in Section 2.1.4, simultaneous transmit schemes that employ code orthogonality in the fast-time dimension suffer from cross-correlation that is simply too high to yield meaningful cross-polarized returns in the presence of a strong co-polarized return (Krieger et al., 2012). While pulse compression is not sufficient for recovery of the SM, there may be other means that could

serve to reduce the cross-correlation.³ To meet the secondary goal, three different fully polarimetric transmit schemes were employed:

- transmission of orthogonal polarizations on alternate pulses (ALT);
- simultaneous transmission of orthogonal polarizations (SIM); and
- simultaneous transmission of orthogonal polarizations with a slow-time phase coding (SIM-PHS).

The waveform transmit schemes described above are depicted in Figure 3.5. In all cases, as indicated by w_{T1} and w_{T2} in Equation 3.1, orthogonally coded chirps were transmitted on each channel. Using upchirps (or downchirps) for both H and V transmit channels would have been sufficient in the ALT and SIM-PHS schemes and would have simplified post-processing. In fact, as the derivation in 3.2.1 will show, *not* using only upchirps (or downchirps) across both H and V transmit channels in the SIM-PHS scheme rendered the resulting scattering matrix data inseparable by Doppler processing alone.

3.2.1 Simultaneous SM measurement with Doppler division multiple access waveforms

The derivation presented in this section shows that a variant of the SIM-PHS scheme is indeed a valuable option for simultaneous polarimetric scattering matrix measurement under certain circumstances. The idea is that, if a system has sufficient Doppler bandwidth, then modulating pulses in slow time will allow separation of superimposed co- and cross-polarized scattering matrix components simply by Doppler processing (which amounts to taking the fast Fourier transform (FFT) of each range bin across the slow-time dimension). This technique has been applied in MIMO radar, where it has been dubbed Doppler division multiple access (DDMA), to obtain multiple measurements of

³One possibility is building a null of the orthogonal waveform's pulse compressed peak directly into the matched filter. This author made some headway toward implementing this approach, but further work is required.

a scene in order to execute space-time adaptive processing (STAP) (Mecca et al., 2006; Kantor and Davis, 2011).

Two channel DDMA signal model

Following the derivation in Mecca et. al. (2006), let

$$w_{T1}(t) = w_{T2}(t) \quad (3.2)$$

That is, let the horizontally polarized pulse waveform be encoded identically to the vertically polarized pulse waveform.⁴

Assume further that the pulses are transmitted at $PRF = \frac{1}{T_r} \geq Lv_{max}$; where $v_{max} =$ maximum desired unaliased Doppler velocity and $L =$ the number of transmit elements. In the case of a fully polarimetric radar with a single transmit antenna like Four Eyes, $L = 2$. Let M correspond to the number of pulses in a CPI. Choose M such that M/L is an integer, and let these M pulses be transmitted in a pulse train. Then, modulate the starting phase of each pulse in each channel such that the signals received from the pulse trains will be separated in the slow time Doppler dimension following Doppler processing. Mathematically, for $L = 2$, such pulse trains can be written:

$$P_{Ti}(t) = \sum_{m=0}^{M-1} w_{Ti}(t - mT_r) e^{j\pi(i-1)m}, \quad i \in \{1, 2\} \quad (3.3)$$

where i corresponds to transmit channel.

The choice of slow time phase modulation above is such that the return signal from transmit channel 1 will have its Doppler processed return centered at Doppler frequency $f_d = 0$ Hz, while the return signal from transmit channel 2 will have its Doppler processed return centered at $f_d = PRF/2$ Hz. Hence, the RD map of the transmit channel 1 image can be recovered by low pass filtering the Doppler spectrum of the combined image, while

⁴Note that this constraint was not imposed in the waveform scheme shown earlier in the SIM-PHS case. Also note that a larger number of channels requires that the waveforms on all channels be encoded identically.

the RD map of the transmit channel 2 image can be recovered by either high pass filtering the combined image spectrum or by first shifting that spectrum by π rad/sample, then low pass filtering.⁵

Simultaneous SM measurement with DDMA

We now show that this simultaneous transmit scheme can recover the SM in polarimetric radars. First, write the transmit pulse sequence across channels in vector form as:

$$\vec{\mathbf{P}}(t) = \begin{bmatrix} P_{T1}(t) \\ P_{T2}(t) \end{bmatrix} = \begin{bmatrix} \sum_{m=0}^{M-1} w_{Ti}(t - mT_r) \\ \sum_{m=0}^{M-1} w_{Ti}(t - mT_r)e^{j\pi m} \end{bmatrix}. \quad (3.5)$$

Without loss of generality, assume the radar is linearly polarized. Next, assume sufficient cross-pol isolation on the antennas followed by polarimetric calibration such that the transmit antenna vectors can be written after corrected distortion in matrix form as

$$\mathbf{T} = \begin{bmatrix} -\vec{\mathbf{H}}_{H_t}^T & - \\ -\vec{\mathbf{H}}_{V_t}^T & - \end{bmatrix} \approx \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \mathbf{I}_2 \quad (3.6)$$

and, similarly, the receive antenna vectors can be written in matrix form as

$$\mathbf{R} = \begin{bmatrix} -\vec{\mathbf{H}}_{H_r}^T & - \\ -\vec{\mathbf{H}}_{V_r}^T & - \end{bmatrix} \approx \mathbf{I}_2. \quad (3.7)$$

Consider a point scatterer at distance r from the radar, moving at constant radial velocity v with respect to the radar over the time of a CPI of M pulses. Ignoring attenuation and round-trip phase for simplicity, the response measured at the receive

⁵An equally valid phase coding for the two channel case—and indeed, the encoding suggested in (Mecca and Krolik, 2007) if $L = 2$ —takes the form

$$P_{Ti}(t) = \sum_{m=0}^{M-1} w_{Ti}(t - mT_r)e^{j\frac{\pi}{2}(1-2i)m}, \quad i \in \{1, 2\} \quad (3.4)$$

Demodulation in this case would involve either bandpass filtering the Doppler spectra centered at $f_d = \frac{\pi}{2}$ and at $f_d = -\frac{\pi}{2}$, or an appropriate spectral shift in each case followed by a lowpass filter.

antenna for such a scatterer can be written

$$\vec{\mathbf{U}}(t) = e^{j2\pi f_d t} \mathbf{R} \mathbf{S}^T \mathbf{T} \vec{\mathbf{P}}(t - t_0) \quad (3.8)$$

$$= e^{j2\pi f_d t} \mathbf{S} \vec{\mathbf{P}}(t - t_0) \quad (3.9)$$

$$= e^{j2\pi f_d t} \begin{bmatrix} S_{HH} P_{T1}(t - t_0) + S_{VH} P_{T2}(t - t_0) \\ S_{HV} P_{T1}(t - t_0) + S_{VV} P_{T2}(t - t_0) \end{bmatrix} \quad (3.10)$$

where the elements of $\vec{\mathbf{U}}$ are the signals received by the H- and V-pol antennas, respectively; T is the transpose operator, which is necessary here to preserve established subscript notation for $\mathbf{S}$; $t_0 = \frac{2r}{c}$ where r = the range to the scatterer and c = the speed of light in air; and $f_d = \frac{-2v}{\lambda}$ = the scatterer's Doppler frequency, where λ = radar wavelength.

We can write the signal component of the output of matched filtering a single pulse w_{T1} with a target return at time t_0 as

$$y_{T1}(t) = \int_{-\infty}^{\infty} w_{T1}(s - t_0) K w_{T1}^*(s + \tau - t) ds \quad (3.11)$$

where τ is the pulselength. By the constraint imposed in Equation 3.2, y_{T1} is equal to the output of matched filtering a single pulse w_{T2} with a target return at time t_0 . Thus, the output of matched filtering the pulse trains P_{Ti} with a target return at t_0 is

$$Y_{T1}(t) = \sum_{m=0}^{M-1} y_{T1}(t - mT_r) \text{ and} \quad (3.12)$$

$$Y_{T2}(t) = \sum_{m=0}^{M-1} y_{T1}(t - mT_r) e^{j\pi m} = e^{j\pi m} Y_{T1}(t), \quad m = \{0, 1, \dots, M-1\}. \quad (3.13)$$

Therefore, the output of appropriately matched filtering received signal $\vec{\mathbf{U}}(t)$ is

$$\vec{\mathbf{Z}}(t, m) = e^{j2\pi f_d t} \begin{bmatrix} (S_{HH} + S_{VH} e^{j\pi m}) Y_{T1}(t) \\ (S_{HV} + S_{VV} e^{j\pi m}) Y_{T1}(t) \end{bmatrix}, \quad m \in \{0, 1, \dots, M-1\}. \quad (3.14)$$

It follows that Doppler processing $Z_1(t)$ across the M received pulses and appropriately filtering in the Doppler dimension yields the RD images for S_{HH} and S_{VH} . Doppler

processing $Z_2(t)$ across the M received pulses and appropriately filtering in the RD dimension yields the RD images for S_{HV} and S_{VV} .

3.3 Point de Chene field test, October 2015

The field test location is discussed, followed by a discussion of both the persistent objects and objects of opportunity on which measurements were recorded during the week-long data collection campaign. The section closes with a discussion of the sea states under whose conditions data was recorded. Discussion of calibration targets used on location is presented in the next chapter.

3.3.1 Field test location




Figure 3-6: The nautical chart indicates all locations considered in Massachusetts' Cape Ann, along with their associated elevations. Lines of sight to persistent targets from each position are indicated by solid lines color-coded to each of the locations. The sea floor drops off fairly quickly off the tip of Cape Ann. *Adapted from National Oceanic and Atmospheric Administration's Office of Coast Survey Chart 13279 at <http://www.charts.noaa.gov/OnLineViewer/13279.shtml>; accessed 2015-05-20.*

During the design phase of the experiment, several locations on Massachusetts' Cape Ann were considered, as indicated on the nautical chart in Figure 3-6. At least one persistent target in the form of a navigational buoy was desirable, as were both means of ocean access by kayak and sufficient space to set up calibration targets in the far field of Four Eyes' antennas. The ocean-facing yard at 9 Point de Chene Avenue in Rockport best met all requirements. Four Eyes was transported during the last week of October 2015 to this east-facing location overlooking the Atlantic Ocean. See Figure 3-7 for a Google Earth progressive zoom aerial view of the location. The test site sits at 30' above mean sea level, with direct ocean access over a sloping rock face. The sea floor drops off steeply




Figure 3-7: Progressive zoom aerial view of Cape Ann experiment location. The bottom figure indicates the precise location of Four Eyes during the experiment. *Aerial images are Google Earth map data from Google and DigitalGlobe; inset photo taken by Ellen Ebacher, MIT LL.*

near the coastline in this region, as indicated in Figure 3-6, allowing for observation of the sea swell as it encroached on the shoreline.

3.3.2 Persistent test objects

As indicated in the previous section, the field test location offered access to two persistent and quasi-static objects. The first of these is a point target object in the form of a gong buoy, located at a slant range from the radar of approximately 1.7 km. As shown in Figure 3-8, such buoys consist of what are effectively a pair of trihedrals mounted atop a cylindrical base.⁶ However, in very low grazing angle geometries, the buoy can exhibit a fair amount of polarimetric variation. As the




Figure 3-8

sea undulates beneath the buoy, the buoy bobs slightly and the radar's energy can impinge upon the trihedrals horizontally. When this happens, the energy predominantly reflects off of only the two vertically oriented faces (which by themselves form a dihedral), and the co-polarized target returns can drop significantly.⁷ The second object, with an extent of ~ 75 m with respect to the radar's line of sight, is a breakwater wall. The polarimetric signature of the wall varies significantly along its length.

The viewing geometry of the persistent objects is shown in Figure 3-9. As shown, the antennas are situated at ~ 11 m above mean sea level and have a beamwidth of 3.7° . When the antenna's boresight is trained on the gong buoy located at ~ 1.7 km, the range of grazing angles θ_{grz} formed by the intersection of the mainlobe with the sea surface is $\leq 2.12^\circ$. The buoy itself is at $\theta_{grz} \approx 0.37^\circ$, while the breakwater wall's nearest point is at $\theta_{grz} \approx 0.24^\circ$. In this geometry, the main beam extends all the way out to the horizon at ~ 12 km.

⁶It is well known that a trihedral is a strong co-polarized (odd-bounce) radar target. Moreover, because commercial shipboard navigation radars are typically X-band, like Four Eyes, these trihedrals are calibrated in size to have the strongest signatures in response to incident energy at X-band wavelengths.

⁷A dihedral is a strong cross-polarized (even-bounce) target.




Figure 3-9: The side-looking geometry of the radar’s view to the persistent point target object, a gong buoy, and the persistent extended object, a breakwater wall with ~ 75 m of range extent from the radar’s perspective, is shown.

3.3.3 Objects of opportunity

Cape Ann is home to numerous commercial fishing outfits and lobstermen. Consequently, during the field test there were several small boats observable by the radar at various aspect angles and velocities. The boats frequently crossed through the radar’s mainlobe while their operators navigated to each of their lobster buoys. Upon arriving at each buoy, the boats floated while the lobstermen checked their traps. Several of the boats whose returns were recorded by the radar are shown in Figure 3-10.

The field test plan also included data collection on a kayaker who was part of the test team. The kayaker is shown in Figure 3-11 approaching the buoy and breakwater wall described in the previous section. The day on which kayak data was collected featured a low sea state and hence the data was not as compelling as originally hoped.

3.3.4 Observed sea states

Sea state was briefly mentioned in Section 2.2.1 because it is true that in low grazing angle geometries, the higher the sea state, the stronger the sea clutter. This is captured by all of the low grazing angle models for σ_0 that were shown in Figure 2-8. Predictions across sea states at low grazing angles from two of these models, the GIT and SIT models, are shown in Figure 3-12. Because sea backscatter increases at higher sea states, object



(a) Red fishing boat



(b) White lobster boat



(c) The “After Five”



(d) The “New Englander”

Figure 3·10: Several of the boats whose returns were recorded by the radar are shown. All photographs are zoomed in; the boats are at various distances from the radar. *Photographs taken by David Mooradd and Dean Mailhiot, MIT LL.*




Figure 3·11: The kayaker can be seen in the foreground paddling toward the persistent objects, the buoy and the breakwater wall, in the background. *Photograph taken by Dean Mailhiot, MIT LL.*

detection in sea clutter becomes increasingly challenging at higher sea states. Thus, a mixture of observed sea states is a valuable attribute for a sea clutter dataset.

The sea state models shown in Figures 2·8 and 3·12 refer to the hydrographic sea state, which is also known as the *Douglas Sea Scale* (Nathanson et al., 1991); it is this sea state

scale that is used in radar. Mariners, on the other hand, consider sea state in terms of the *Beaufort wind force scale*. It is the Beaufort scale to which the National Oceanic and Atmospheric Administration (NOAA) maps observations by its numerous weather buoys distributed throughout the world's oceans. The Douglas sea state is a function of significant wave height (SWH), which is defined as the average crest-to-trough height of the highest one-third of waves. The Beaufort scale, on the other hand, is a function of wind speed and *fetch*, defined as the distance over which a nearly constant wind has been blowing (Nathanson et al., 1991). Because of this difference in reference scales, determining observed sea state can be a bit of a tricky business.

While Four Eyes was on location, the remnants of Hurricane Patricia moved through the area, enabling observation of a range of higher sea states.⁸ Four Eyes was on location for a full week, but the days during which the bulk of data was recorded were October 27, 28, and 29. Without access to a portable wave buoy to determine SWH as a function

⁸These higher sea states are actually the reason a relaunch of the kayaker did not occur on a later day. The seas had quickly turned too treacherous to risk launching a small craft.




Figure 3-12: Low grazing angle models for mean RCS per square meter of area illuminated by the radar beam reflect the increase in sea clutter as sea state increases.

of time—and thus sea state—specifically for the field test area of interest, there are three methods by which sea state may be estimated.

The first is by computing σ_0 from processed data, then mapping to one of the many available empirical models. Because of the disparity in these models at low grazing angles, this approach is a better avenue for model validation than for determining true sea state. The second is by leveraging NOAA weather buoy data. This is an excellent option provided that the NOAA weather buoy is in close proximity to the field test area of interest.

In the case of the Point de Chene field test, the nearest NOAA weather buoy is the Jeffrey’s Ledge waverider buoy, #44098, located just over 35 km north by northeast of Four Eyes’ location, as shown in Figure 3-13. The buoy measures SWH every hour⁹; the data is made freely available (National Oceanic and Atmospheric Administration National Data Buoy Center, 2015). A plot of these measurements from October 27-29 is shown in 3-14.




Figure 3-13: NOAA buoy 44098’s location relative to the radar is shown. *Adapted from NOAA buoy location map at www.ndbc.noaa.gov; accessed 2016-09-13.*

The thresholds between Douglas sea states as a function of SWH (Nathanson et al., 1991), are indicated with red dashed lines. The plot shows that, over the course of the three most significant data collection days, the buoy recorded SWHs consistent with five different sea states. However, because of the distance between the NOAA buoy and the experiment location and because the

⁹Note that this buoy does not measure wind speed or direction.




Figure 3-14: NOAA buoy data for the waverider buoy nearest the field test site shows evidence for observation of five sea states as the remnants of Hurricane Patricia moved up the coast.

buoy sits over deeper ocean,¹⁰ the data provides only a rough estimate of the sea states observed and their times of occurrence.

The third alternative for estimating sea state involves mapping empirical observations and data to the Beaufort wind scale via photographs of the sea in the radar field of view. An example of this is shown in Figure 3-15. At left is a NOAA image that corresponds to Beaufort scale force 5; at right is a comparable photograph taken during the Point de Chene field test. This is largely a subjective mapping process, but the team was able to support the process with observations of an anemometer located on site that regis-

¹⁰It is well known that ocean wave velocity is related to the wavelength and the depth of the ocean. Wave velocity is actually not directly dependent on wind speed.



(a) Beaufort force 5: winds 17-21 knots *Image from Wikipedia Commons.* (b) Beaufort force 5 at Point de Chene *Photo by David C. Mooradd, MIT LL.*

Figure 3.15: An example of estimating Beaufort wind force based on photographic evidence is shown.

tered wind speeds in accordance with the Beaufort scale force estimations attained from a sequence of photographs. Once the Beaufort force is estimated, it is relatively straightforward to roughly map from Beaufort to Douglas, though the demarcation between the two is also inexact, as shown in Figure 3.14. The mapping from Beaufort wind force scale to Douglas sea state to SWH shown in Figure 3.14 is adapted from Nathanson et. al. (1991).

The result of the compendium of this analysis is that, while it can be difficult to say for certain which recordings correspond to which sea states, the Point de Chene dataset certainly comprises Douglas sea states 2, 3, and 4, and includes data that appears to be on the border of sea state 5.

3.4 Point de Chene Dataset

A summary of the full Point de Chene Dataset is given in Table 3.2.

Size:	
Number of files	488
Recording length	55 minutes
Total size	15.5 TB
<i>Calibration</i>	<i>1.3 TB</i>
<i>Clutter-only</i>	<i>4.0 TB</i>
<i>Clutter-plus-object (persistent)</i>	<i>5.8 TB</i>
<i>Clutter-plus-object (dynamic)</i>	<i>4.4 TB</i>
Waveforms:	
Polarimetric transmit schemes	ALT, SIM, SIM-PHS
Pulselengths	0.25, 2, and 4 μ sec
Bandwidths	4, 40, 150, 250, and 400 MHz
Pulse repetition frequencies	6250 and 12500 Hz
Radar scenes:	
Objects of interest	Nine total
<i>Persistent</i>	<i>Buoy and breakwater wall</i>
<i>Dynamic</i>	<i>Kayaker and fishing boats</i>
Sea states	2, 3, 4, 4+
Clutter-limited range swath	Waveform minimum range to 4 km

Table 3.2: Point de Chene Dataset summary

3.4.1 Algorithm development data subset

The reader will recall that the second major contribution of this dissertation is development of practical means to mitigate the impacts of low grazing angle sea clutter on the standard CFAR detection process. With this aim in mind, a handful of files were selected after post-processing as containing data that a) comprised various range resolutions, b) featured various objects of interest, and/or c) contained strong sea returns that occasionally obscured those objects. These files, summarized in 3.3, are the foci of the algorithm development and assessment research presented in Chapter 5.

File ID	τ μ sec	β MHz	Est. SS	Objects of Interest
266	2	400	4	White fishing boat
				The New Englander
				Buoy (in Sidelobe/Null)
267	2	40	4	White fishing boat
270	2	400	4	Buoy
271	2	150	4	Buoy
272	2	40	4	Buoy
273	2	4	4	Buoy
325	4	150	5	The New Englander

Table 3.3: The subset of files chosen from the larger Point de Chene dataset as foci for algorithm development is listed. Note that τ = pulsewidth and β = bandwidth in the table, while estimated sea state (SS) is from the Douglas sea state scale.

Chapter 4

Signal Post-Processing

It is critical in polarimetric radar, as with any form of multi-channel radar, that each channel be transverse equalized both respect to itself and to the other channels in the system. That is, the magnitude response of each channel must be flat so as to render distortionless output, while the phase response must be corrected such that the time delays incurred across signal paths through multiple channels are equivalent. Moreover, numerous tasks in polarimetry, including polarization synthesis and scattering matrix decompositions, also require that the channels be polarimetrically calibrated. That is, the system response must be corrected in order to render the correct polarimetric signature measurements for objects with known scattering matrices.

The first section of this chapter discusses the calibration targets used at Point de Chene that played key roles in post-processing for achieving channel equalization and polarimetric calibration. The second section discusses the algorithm developed for cross-channel transverse equalization of Four Eyes data and the associated equalization results. The third section discusses the polarimetric calibration methodology for Four Eyes and gives evidence of calibration quality. The fourth section describes the full signal processing chain that has been implemented for Four Eyes data, including details of signal pre-conditioning steps. The chapter closes with a section on the methodology used to create data labels for performance analysis purposes in the algorithm assessment portion of Chapter 5.

4.1 Point de Chene calibration targets

Each channel in a system consists of a chain of several cascaded components, each with its own frequency response. In Four Eyes' case, the frequency response of a given channel comprises the chain of frequency responses of an AWG, a VSG, a TWT, the transmit (Tx) and receive (Rx) antennas, an LNA, the receiver, and the ADC, as was shown in the system diagram in Figure 3-1. So, for the HH channel, we have frequency response:

$$H_{HH}(f) = H_{ADC_1}(f)H_{VSG_1}(f)H_{TWT_1}(f)H_{TxH}(f)H_{RxH}(f)H_{LNA_1}(f)H_{Rcv_1}(f)H_{ADC_1}(f) \quad (4.1)$$

where the subscripts on the right side of the equation indicate the components in the chain. Similar equations can be written for $H_{HV}(f)$, $H_{VH}(f)$, and $H_{VV}(f)$, the frequency responses of the other channels.

Performing transverse equalization requires measurement of the frequency response of each channel in the system. There are various ways this measurement may be obtained. Often, the measurement involves *direct inject*, or feeding the upconverted signal directly into the receiver. This approach ignores the contributions of several components to the channel's frequency response. In the case of Four Eyes, such a measurement would exclude the terms $H_{TWT_1}(f)$, $H_{TxH}(f)$, $H_{RxH}(f)$, and $H_{LNA_1}(f)$ in Equation 4.1. To avoid this, we used a delay line situated in the far field of the antenna. A delay line receives a signal on its input port, amplifies or attenuates it according to operator pre-sets, then re-transmits the signal through its output port. The delay line is given its name because it contains a fixed quantity of fiber optic line, traversing the length of which delays the re-transmitted signal. In the case of delay line used with Four Eyes, the delay is approximately 15 km.

The setup was as follows: The vertical port of an antenna identical to Four Eyes' antennas was hooked up to the delay line's input port; however, the antenna was rotated 45° about the z axis, yielding in Four Eyes' reference frame an X^+ polarized delay line




Figure 4.1: A depiction of the delay line setup with respect to Four Eyes is shown. The antennas on the delay line were rotated 45° so that half of Four Eyes’ transmitted energy would couple into the delay line regardless of whether Four Eyes was transmitting on the H or the V channel. Half of the delay line’s retransmitted energy would then couple into each of Four Eyes H and V receive channels.

input antenna. Another identical antenna rotated in the same fashion was connected to the delay line’s output port. This configuration is shown in Figure 4.1. The horizontal ports on the delay line antennas were unused. The consequence of aligning the antennas in this way is that, as discussed in Chapter 2, half of the energy in Four Eyes’ H-pol transmit signal couples into the delay line’s X^+ receive antenna; then, half of the X^+ polarized energy re-transmitted by the delay line couples into each of Four Eyes’ H and V receive antenna ports. The same is true for the case of Four Eyes’ V-pol transmit signal. Four Eyes’ antennas and one of the delay line antennas had rifle scopes mounted to their dish edges to enable careful alignment between the two sets of antennas. The delay line is shown in Figure 4.2.

Given this set-up, the frequency response of each of the four channels to a given waveform pair w_{T1} , w_{T2} can be measured by

1. transmitting one second of only H-pol w_{T1} pulses;
2. recording responses on both channels; and
3. repeating steps 1 and 2 for V-pol w_{T2} transmit.

Because nine different waveforms comprising different pulsewidths and bandwidths were programmed into Four Eyes, this process was repeated once for each waveform. Given that the frequency response of a system can vary somewhat over time, the full process was executed at the start and end of each data collection day.

As indicated above, measuring the channel frequency responses in this way enables inclusion of all components in the channel paths in the frequency response measurements. The keen reader will notice, however, that this also includes three components that will not be in the channel paths: the delay line’s antennas and the delay line itself. However, because the




Figure 4.2: A close up of the delay line set up in Cape Ann.

same input and output antenna ports were used on the delay line antennas, all signals through the delay line must follow the same path. Thus, any distortion induced by the components along that path will be imparted to all channels. It will therefore not impact the outcome of the cross-channel equalization.

The delay line served a second purpose: It was also a polarimetric calibration target. A polarimetric calibration target is one for which the scattering matrix $\mathbf{S}$ is known. The delay line’s SM in the radar’s linear polarization basis can be computed¹ as

$$\mathbf{S}_{dl} = \frac{1}{4} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad (4.2)$$

simply by considering the coupling sequence of the chain of antennas (Freeman et al., 1990; Allan, 1995), as shown in Figure 4.3.

¹That is, computed up to an absolute phase term that depends on the range of the delay line from the radar. Huynen showed that an object’s SM properties are independent of range, and hence the absolute phase can be disregarded (Huynen, 1970).

The second polarimetric target was a trihedral, which is well known in the linear basis to have SM

$$\mathbf{S}_{tri} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (4.3)$$

(Lee and Pottier, 2009), where an amplitude and arbitrary phase factor has been disregarded.

$$\begin{aligned} S_{HH} &= \text{[Diagram]} = \frac{1}{4} \\ S_{HV} &= \text{[Diagram]} = \frac{1}{4} \\ S_{VH} &= \text{[Diagram]} = \frac{1}{4} \\ S_{VV} &= \text{[Diagram]} = \frac{1}{4} \end{aligned}$$

Figure 4-3: The delay line scattering matrix can be determined (up to an arbitrary absolute phase factor) in the linear basis by inspecting the transmitted energy’s coupling sequence in Four Eyes’ reference frame.

The trihedral also served a second purpose: It was built to a known RCS at X-band such that the absolute gain of the system could be calibrated using its return. Trihedral recordings were taken with the trihedral in place and with the trihedral removed to enable background subtraction in post-processing, ensuring that land clutter and the target’s support tripod did not contaminate its measurement. As with the delay line, data was recorded on the trihedral at the beginning and end of each data collection day. The full calibration suite set

up, including Four Eyes’ perspective of the targets, is shown in Figure 4-4.

4.2 Cross-channel transverse equalization

It is well known that the matched filter that maximizes SNR in a received signal takes the form of the time-reversed complex conjugate of the original signal (Richards, 2014). That is, if we let $x(t)$ equal the original transmitted signal, then the filter that maximizes the SNR in the received signal is $m(t) = \alpha x^*(-t) \leftrightarrow \alpha X^*(f)$, where filter gain α has no effect on SNR and $\leftrightarrow$ indicates Fourier domain transformation. The scheme described below aims to modify these matched filters such that the time delay and frequency response of every channel through Four Eyes is equivalent across the bandwidth of interest, regardless of the channel, following pulse compression with the modified matched filters.




Figure 4.4: Four Eyes calibration set-up: Delay line is shown at left; the calibrated trihedral is shown at top; and the targets' positions with respect to Four Eyes on site at Point de Chene is shown at right.

4.2.1 Signal model

In the following sections, we adopt the usual convention that lowercase letters indicate time domain signals, while capital letters indicate their frequency domain counterparts.

Let the basebanded² signals received from the delay line be denoted

$$\begin{aligned}
 x_{11}(t) &= w_{T1}(t - t_d) * h_{11}(t) \leftrightarrow W_{T1}(f)H_{11}(f)e^{-j2\pi ft_d} \\
 x_{12}(t) &= w_{T1}(t - t_d) * h_{12}(t) \leftrightarrow W_{T1}(f)H_{12}(f)e^{-j2\pi ft_d} \\
 x_{21}(t) &= w_{T2}(t - t_d) * h_{21}(t) \leftrightarrow W_{T2}(f)H_{21}(f)e^{-j2\pi ft_d} \\
 x_{22}(t) &= w_{T2}(t - t_d) * h_{22}(t) \leftrightarrow W_{T2}(f)H_{22}(f)e^{-j2\pi ft_d}
 \end{aligned} \tag{4.4}$$

where $x_{ij}(t)$ for $i, j \in \{1, 2\}$ indicates the received signal from the delay line through the channel path indicated by the subscript sequence; $w_{T1}(t)$ indicates the horizontally polarized downchirp transmitted on channel 1 and $w_{T2}(t)$ indicates the vertically polarized upchirp transmitted on channel 2, as given in Equation 3.1; $*$ indicates the convolution

²The process of basebanding Four Eyes' IF data is discussed in Section 4.4. For the purposes of this section, without loss of mathematical accuracy, the data is assumed to be at baseband upon digitization by the ADC.

operator; $h_{ij}(t)$ for $i, j \in \{1, 2\}$ is the impulse response of the system through the channel path indicated by the subscript sequence; and t_d corresponds to the time delay induced by signal passage through the delay line. To simplify notation going forward, we will drop the statement that $i, j \in \{1, 2\}$, but the reader should understand that the statement is implied.

The signals in Equation 4.4 are, of course, digitized by the ADC at sampling frequency $f_s = 1260$ MHz so that the actual basebanded signal on which the equalization operations are performed is

$$x_{ij}(t)|_{t=nT_s} = x_{ij}[n] \leftrightarrow X_{ij}[k], \quad k = 0, 1, \dots, N-1 \quad (4.5)$$

where $T_s = \frac{1}{f_s}$ and $n = 0, 1, \dots, N-1$. The $x_{ij}(t)$ signals embody the channels' impulse responses in the frequency band of interest; thus, $X_{ij}[k]$ is the frequency response of the system, over the waveform's bandwidth, to a point target.

We seek a set of matched filters $m_{ij}[n] \leftrightarrow M_{ij}[k]$ such that matched filtering incoming pulses corrects the distortions and the time delays induced by different $h_{ij}[n] \leftrightarrow H_{ij}[k]$. Take as these initial matched filters $m_{0,ij}$ the time-reversed conjugates of the synthetic signals exactly as programmed into Four Eyes, where the additional subscript 0 indicates that these filters comprise the 0^{th} order matched filter, without any adjustments to its weights having been yet computed.

4.2.2 Transverse equalization algorithm

Starting with these 0^{th} order matched filters and using the difference between the actual and theoretical point target response as a correction factor for the matched filter weights, the equalization algorithm is a process of iteratively adjusting the matched filter weights for each of the four channels until the frequency response for each channel matches the ideal response over the passband of interest. Additional corrections are then made to the filters to time align the channels with respect to each other. Finally, a gain correction




Figure 4-5: A high-level block diagram of the cross-channel transverse equalization algorithm implemented for Four Eyes data is shown.

factor is built into the filters to align the peak amplitude responses. A high level block diagram of the algorithm is given in Figure 4-5.

Mathematically, this process is as follows. Matched filter P pulses of $X_{ij}[n]$ with the initial matched filters to yield


$$y_{0_{ij}}^{\{p\}}[n] \leftrightarrow Y_{0_{ij}}^{\{p\}}[k] = M_{0_{ij}}[k]X_{ij}^{\{p\}}[k], \quad p = 1, 2, \dots, P \quad (4.6)$$

where the superscript p indicates pulse number and the signals are sufficiently zero-padded to avoid aliasing due to circular convolution. To smooth out any timing jitter across pulses within each channel, coherently average the P pulse-compressed returns, yielding


$$z_{0_{ij}}[n] = \frac{1}{P} \sum_{p=1}^P y_{0_{ij}}^{\{p\}}[n] \quad (4.7)$$

as shown in Figure 4-6.

Because we seek to equalize only the received signal range swath that comprises the delay line's response, we can trim the coherently integrated signal $z_{0_{ij}}[n]$ in each channel to only those samples corresponding to the extent of the convolution of $m_{1_{ij}}[n]$ and $x_{ij}[n]$. This can be done using a rectangular window of length $l = 2L - 1$ to trim the signal in each channel, where L is the number of samples in $m_{0_{ij}}$ prior to zero padding and the window's midpoint is aligned with the peak of the delay line response in $z_{0_{ij}}[n]$. An



(a) Full-range view of integrated compressed delay line returns



(b) Zoomed-in view of integrated compressed delay line returns

Figure 4-6: The full-range swath of the P coherently integrated pulse-compressed delay line responses for the $2 \mu\text{sec}$ 150 MHz waveform is shown across each of the four channels in the top figure. In the legend, H corresponds to channel 1, while V corresponds to channel 2. The differences in the shape of the point target's response induced by frequency response variation across channels are not apparent on this range scale, but become apparent in the bottom figure, which is zoomed in to the peaks of the point target's response across channels.

example of these time-trimmed compressed responses for the 2 μ sec 150 MHz waveform is shown in Figure 4-6, which shows that the impulse response is not identical across channels. Moreover, the impulse response is distorted such that its response deviates from the ideal sinc response of a point target.

To correct these deviations, first transform the impulse response cutouts to the frequency domain. Compute the inverse matched filter as

$$M_{0ij}^{-1}[k] = \begin{cases} \frac{1}{M_{0ij}[k]} & \text{if } M_{0ij}[k] \neq 0 \\ 0 & \text{else} \end{cases} \quad (4.8)$$

and strip off the matched filter from $z_{0ij}[n]$, leaving just the frequency response averaged over P pulses:

$$\begin{aligned} S_{0ij}[k] &= Z_{0ij}[k] M_{0ij}^{-1}[k] = \frac{1}{P} \sum_{p=1}^P Y_{0ij}^{\{p\}}[k] M_{0ij}^{-1}[k] \\ &= \frac{1}{P} \sum_{p=1}^P X_{ij}^{\{p\}}[k] M_{0ij}[k] M_{0ij}^{-1}[k] \\ &= \frac{1}{P} \sum_{p=1}^P X_{ij}^{\{p\}}[k]. \end{aligned} \quad (4.9)$$

It is this estimate of the frequency response of each channel that we wish to correct to have a flat magnitude response across the passband and thereby produce the transverse equalized matched filter

$$M_{1ij}[k] = \frac{G[k] S_{0ij}^*[k]}{S_{0ij}[k] S_{0ij}^*[k] + .001} * e^{-j2\pi \frac{k}{N} L} \quad (4.10)$$

where $G[k]$ is a Gaussian-edged passband gate designed to mitigate the spectral ripple at the passband edge and the constant is added to avoid division by zero.³ The delay in

³An alternative approach involves using a window function designed to reduce the sidelobe level of the pulse compressed response (e.g., Taylor, Hamming, or Chebyshev windows). However, using such a function for $G[k]$ means the chosen window is “baked into” the matched filter implementation. In order to experiment with pulse compression results using different window functions, this must be avoided.




Figure 4-7: The transverse equalization process' third iteration results are shown for the HH and VV channels. The desired frequency response across the channels is a flat, Gaussian-edged passband, $G(f)$. The estimated channel frequency responses $S_{3_{ij}}(f)$ are shown in dark blue, while the frequency responses of the matched filters $M_{3_{ij}}(f)$ that correct these estimates to yield $G(f)$ are shown in red.

Equation 4.10 is chosen such that the peak of the pulse compressed response achieved using this matched filter will align with the center of the uncompressed pulse. In the average case, we have that

$$M_{1_{ij}}[k]X_{ij}[k] = G[k]. \quad (4.11)$$

That is, matched filtering an incoming received pulse with $M_{1_{ij}}[k]$ corrects the system frequency response to the desired Gaussian-edged flat passband response. Iterating over this sequence of steps three times, as shown in Figure 4-5, produces matched filters that are not further adjusted on subsequent iterations. The matched filters out of this process are thus $M_{3_{ij}}[k]$. Figure 4-7 shows frequency domain results of the third iteration of computation over the terms in Equation 4.10 for the HH and VV channels ($ij = 11$ and $ij = 22$, respectively).

Computing $G[k]$ such that the edges taper off quickly achieves the effect of edge ripple mitigation while retaining flexibility in pulse compression later on.

With the magnitude responses across channels thus equalized, we equalize the phase response of the system across channels. The algorithm for so doing is as follows:

1. pulse compressing P pulses in the HH channel (where again $ij = 11$) using the magnitude response transverse equalized matched filter $M_{311}[k]$;
2. finding the subsample location n_{pk} of the point target's peak response in the first pulse;
3. computing the number of samples $d_s^{\{p\}}$ required to shift each of the subsample peaks of the remaining P pulses to align at n_{pk} ;
4. coherently averaging the number of samples shifted across all P pulses to yield

$$d = \frac{1}{P} \sum_{p=1}^P d_s^{\{p\}};$$

5. shifting the matched filter by that corresponding quantity as

$$M_{411}[k] = M_{311}[k]e^{j2\pi\frac{k}{N}d};$$

6. repeating for each of the remaining channels, aligning to n_{pk} in every case.

At this stage, the matched filters are equalized. The final step is to render the gain and phase of the target's point response equivalent across channels. We accomplish this by again compressing P pulses in each channel using $M_{4ij}[k]$, then coherently averaging the result to determine $z_{4ij}[n_{pk}]$, the average complex scalar value at the location of each channel's delay line response peak. Each matched filter is then scaled by the complex value that compensates for any offset in gain and phase with respect to channel HH as

$$M_{5ij} = M_{4ij} \frac{z_{411}[n_{pk}]}{z_{4ij}[n_{pk}]} \quad (4.12)$$

The results of cross-channel transverse equalization in the time domain using this routine can be seen in Figure 4.8.⁴

⁴The figure also demonstrates that the Gaussian-edged frequency passband shown in Figure 4.7 almost ideally preserves the sinc-response expected of an ideal (brick-walled) matched filter. The -13 dB sidelobes typical of such a filter are evident, indicating that the preceding algorithm indeed mitigates spectral ripples at the edges of the passband whilst preserving the freedom to experiment with various sidelobe reduction windows in data post-processing.




Figure 4-8: Results of cross-channel transverse equalization on the delay line’s point response show that the channel’s frequency responses are indeed equalized, as are channel gains.

The preceding algorithm is executed on delay line calibration recordings for each of the nine waveforms, yielding a set of 36 transverse equalized matched filters for each day of data collection. These filters are stored in processor memory and retrieved for pulse compression of incoming data, as discussed further in Section 4.4.

4.3 Polarimetric calibration

To this point, the post-processing calibration work has ensured that the scattering matrix measurement in a given RTI range bin or RD cell is aligned across polarimetric channels. That is, we can be sure that the value in the range bin at, say, 2 km in the HH channel corresponds with the value in the corresponding range bin in each of the VH, HV, and VV channels. Thus, we can be sure that selecting a scattering matrix measurement along the third dimension of the data shown in Figure 2-7 is actually the data corresponding to the scattering matrix produced by the same scatterers across all channels. Moreover, we can be sure that the scattering matrix measurements have been filtered so as to render their

respective channel responses identical. Much can be accomplished using polarimetric radar that has not also been polarimetrically calibrated; however, polarimetric calibration ensures that important between-channel phase relationships are preserved, enabling numerous computational operations on measured scattering matrices (Freeman, 1992; El-Darymli et al., 2014). Polarimetric calibration is therefore highly desirable whenever possible.

A measured scattering matrix can be written as

$$\begin{aligned}\mathbf{M} &= Ae^{j\phi}\mathbf{R}\mathbf{S}^T\mathbf{T} + \mathbf{N} \\ &= Ae^{j\phi} \begin{bmatrix} R_{HH} & R_{HV} \\ R_{VH} & R_{VV} \end{bmatrix} \begin{bmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{bmatrix}^T \begin{bmatrix} T_{HH} & T_{HV} \\ T_{VH} & T_{VV} \end{bmatrix} + \begin{bmatrix} N_{HH} & N_{HV} \\ N_{VH} & N_{VV} \end{bmatrix}\end{aligned}\quad (4.13)$$

where the superscript T indicates the transpose; A encompasses radar system gains and losses and round trip attenuation; ϕ is the phase shift due to round trip delay; $\mathbf{T}$ and $\mathbf{R}$ are the transmitter and receiver polarimetric distortion matrices, respectively; $\mathbf{S}$ is the scatterer's true scattering matrix; and $\mathbf{N}$ is the additive noise induced on each channel path (Freeman et al., 1990). Ideally, $\mathbf{T}$ and $\mathbf{R}$ are identity matrices, so that no polarimetric distortion is introduced on transmit and receive. In reality, this is not the case. Thus, achieving polarimetric calibration amounts to determining and then compensating for distortion due to non-identity elements of $\mathbf{T}$ and $\mathbf{R}$.

For a monostatic radar, we can assume by reciprocity that $R_{HV} = R_{VH}$ and $T_{HV} = T_{VH}$. Thus, assuming measurement of objects with known scattering matrix $\mathbf{S}$, Equation 4.13 represents four equations in six unknowns that must be determined and corrected. Typically, this requires measurements with three different polarimetric calibration targets in order to determine the system of equations (Freeman, 1992). However, if the system has “good” cross-polarization isolation, then polarimetric calibration can be achieved using fewer calibration targets (Sarabandi et al., 1990). To see this, note that when a system has good polarimetric isolation, R_{HV} and T_{HV} are approximately 0. Thus, for a

monostatic system, we can rewrite Equation 4.13 as

$$\mathbf{M} = Ae^{j\phi} \begin{bmatrix} R_{HH} & 0 \\ 0 & R_{VV} \end{bmatrix} \begin{bmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{bmatrix}^T \begin{bmatrix} T_{HH} & 0 \\ 0 & T_{VV} \end{bmatrix} \quad (4.14)$$

where the distortion due to noise has been treated as negligible.

We now show that cross-channel transverse equalization calibration in the preceding section actually polarimetrically calibrates Four Eyes. Expanding Equation 4.14, we have that

$$\mathbf{M} = Ae^{j\phi} \begin{bmatrix} R_{HH}S_{HH}T_{HH} & R_{HH}S_{VH}T_{VV} \\ R_{VV}S_{HV}T_{HH} & R_{VV}S_{VV}T_{VV} \end{bmatrix}. \quad (4.15)$$

Vectorizing and reordering terms yields

$$\begin{aligned} \begin{bmatrix} M_{HH} \\ M_{VH} \\ M_{HV} \\ M_{VV} \end{bmatrix} &= Ae^{j\phi} \begin{bmatrix} T_{HH}R_{HH} \\ T_{VV}R_{HH} \\ T_{HH}R_{VV} \\ T_{VV}R_{VV} \end{bmatrix} \begin{bmatrix} S_{HH} \\ S_{VH} \\ S_{HV} \\ S_{VV} \end{bmatrix} \\ &= Ae^{j\phi} \begin{bmatrix} Q_{HH} \\ Q_{VH} \\ Q_{HV} \\ Q_{VV} \end{bmatrix} \begin{bmatrix} S_{HH} \\ S_{VH} \\ S_{HV} \\ S_{VV} \end{bmatrix} \\ &= Ae^{j\phi} \vec{\mathbf{Q}} \otimes \vec{\mathbf{S}} \end{aligned} \quad (4.16)$$

where $\otimes$ indicates the Hadamard product. Thus, we have that

$$\begin{bmatrix} S_{HH} \\ S_{VH} \\ S_{HV} \\ S_{VV} \end{bmatrix} = \frac{Q_{HH}}{Ae^{j\phi}} \begin{bmatrix} 1 \\ \frac{Q_{HH}}{Q_{VH}} \\ \frac{Q_{HH}}{Q_{HV}} \\ \frac{Q_{HH}}{Q_{VV}} \end{bmatrix} \otimes \begin{bmatrix} M_{HH} \\ M_{VH} \\ M_{HV} \\ M_{VV} \end{bmatrix}. \quad (4.17)$$

Up to an arbitrary scale factor and absolute phase, this is exactly the final step that was performed in the cross-channel transverse equalization algorithm detailed in Section 4.2.

Sarabandi et. al. (1990) achieve satisfactory calibration results for a monostatic system with 25 dB of cross-polarization isolation, but offer no further quantification of the criterion for good isolation. As was shown in Section 3.1.1, Four Eyes' antennas

were measured as having cross-pol isolations of 20 dB and 26 dB. Sarabandi et. al. also note that, ideally, calibration for systems with good cross-polarization isolation would be performed with a single calibration target with a known scattering matrix that has both co-polarized and cross-polarized signatures. Such a target is not easy to find or orient with respect to the radar's incident energy, but the delay line set up provides exactly such a calibration target. However, because the delay line antennas have their own distortion matrices and because it is unknown whether Four Eyes' cross-polarization isolation is good enough, it is also necessary to assess the quality of polarimetric calibration achieved by using the assumption of good cross-polarization isolation.


4.3.1 Polarimetric calibration results

A good metric for the quality of polarimetric calibration is the transformation of a trihedral's response in the linear H-V basis to the circular R-L polarization basis (Blejer, 2016b). The transformation can be computed via circular-from-linear polarimetric synthesis equations

$$\begin{aligned} S_{RL} &= \frac{1}{2}(S_{HH} + S_{VV}) \\ S_{RR} &= \frac{1}{2}(S_{HH} + 2jS_{HV} - S_{VV}) \\ S_{LL} &= \frac{1}{2}(S_{HH} - 2jS_{HV} - S_{VV}) \end{aligned} \tag{4.18}$$

where, because of reciprocity in monostatic systems, $S_{RL} = S_{LR}$.⁵ If the calibration is sufficient, the cross-polarization isolation shown in the H-V basis should be reasonably well preserved after the basis transformation (Blejer, 2016b). Figure 4-9, which shows the calibrated trihedral response in each basis, evidences the quality of polarimetric calibration for Four Eyes. Following synthesis of the linear basis response in the circular basis, <1 dB of cross-polarization isolation is lost. Thus, we conclude that the polarimetric calibration of Four Eyes sufficiently enables further computational operations on measured scattering matrices.

⁵Equations are straightforward to derive using S_{HV} and $\vec{E}_R$ and $\vec{E}_L$, the Jones vectors for right and left circularly polarized antennas.



(a) In the linear basis, ~ 28 dB of cross-polarization isolation is evident.

(b) Following transformation to the circular basis, ~ 27 dB of cross-polarization isolation is preserved.

Figure 4-9: Polarimetric calibration quality assessment using transformation of the trihedral response, calibrated via matched filters computed in cross-channel transverse equalization, from the linear basis to the circular basis via polarimetric synthesis. Note that in normalizing the amplitude, both responses were normalized to the peak in the linear basis.

4.4 Radar data processing flow

With matched filters in hand that achieve both cross-channel transverse equalization and polarimetric calibration for Four Eyes, attention turns now to the pre-polarimetric processing code suite that has been developed in MATLAB for Four Eyes data. The discussion below presumes the reader has familiarity with the signal processing operations implicit in the suite’s radar signal processing blocks. Thus, implementation details rather than their theoretical foundations are given in this section. The reader seeking exposition of the underpinnings for the operations below is referred to Richards (2014) and Edde (1995).

Four Eyes records real-valued data centered at IF. A copy of the Point de Chene Dataset comprising 15.5 TB of IF data is locally stored on a RAID connected for processing to a custom server that has 512 GB of RAM and dual 8-core CPUs. The server




Figure 4-10: A block diagram of the pre-polarimetric signal processing chain developed for Four Eyes data is shown. Computational blocks are in blue highlighted boxes, inputs and control parameters are in dark blue text, and outputs are in dark red text.

houses a signal processing suite that has been written specifically to accommodate Four Eyes’ data and its various transmission schemes, as well as to maintain flexibility in processing. A block diagram of the pre-polarimetric processing suite is shown in Figure 4-10. Each of these processing blocks will be addressed in turn.

4.4.1 Recorded data retrieval and signal conditioning

As discussed in Section 3.2, the maximum possible PRF for each waveform scheme was used to enable decimation of pulses in slow-time, thus allowing for exploration of any polarimetric algorithm’s efficacy over various PRF timescales. The focus in Chapter 5 is entirely on files recorded using the ALT transmit scheme. We therefore focus exclusively on this scheme in the following discussion.

In the ALT transmit scheme, the maximum PRF is 6250 Hz on each channel because interleaving of transmit pulses produces an effective PRF of 12500 Hz, which maps to the range of the Four Eyes' horizon. In other words, the ALT scheme assumes negligible decorrelation of the scattering matrix in a given cell over a staggered H and V transmit pulse pair. So, H and V returns from these staggered pulses are aligned in post-processing to produce full scattering matrix measurements. However, because the other channel's transmit occurs halfway through a given channel's PRI, the second half of the PRI produces unusable data in range. Thus, the usable data corresponds to an effective PRF of 12500 Hz, but the actual PRF of the retrieved data is 6250 Hz. To illustrate this concept, a diagram of this scheme is reproduced in Figure 4-11 with the pulse repetition interval $PRI = \frac{1}{PRF}$ [sec] denoted. We proceed with the remainder of the dissertation bearing in mind that the PRF of this data is 6250 Hz.

Several operator-input parameters are fed in at the recorded data retrieval stage to control later processing. Most notable among these is the slow-time decimation factor, m , which controls the quantity by which the actual PRF will be decimated. If $m = 1$, the PRF is not decimated; every pulse is retrieved and aligned to form the scattering matrix measurement across the CPI. If $m = 5$, for example, then every 5th pulse will be retrieved, reducing the ALT PRF to 1250 Hz. So, if H transmit pulses are odd-numbered




Figure 4-11: A diagram of the ALT transmit scheme is reproduced here, showing that while the actual PRF on each channel is 6250 Hz, the usable range data for a staggered pair of pulses corresponds to an effective PRF of 12500 Hz.

and V transmit pulses are even-numbered, as shown in Figure 4-11, then setting $m = 5$ retrieves pulse 1 from the HH and HV channels, pulse 6 from the VH and VV channels, pulse 11 from the HH and HV channels, pulse 16 from the VH and VV channels, and so on.⁶ Thus, decimating data recorded using the ALT scheme allows one to explore the effects of having used the ALT transmit scheme at a lower actual PRF.

The amount by which the data can be decimated is limited by the Doppler velocities in the radar scene. Clearly, choosing a decimation factor so high as to render the PRF too low to recover the full Doppler spectrum produces aliasing in the Doppler-velocity dimension. Because the maximum Doppler velocities of sea clutter are determined by the sea state, the maximum slow time decimation factor used depends upon the sea state featured in a given recording. We found that $m = 7$, corresponding to a PRF of ~ 893 Hz, was usually the slow time decimation limit for higher sea state recordings. CPI length is another user-input parameter. For our purposes, we typically use a CPI length of 100 pulses in later processing regardless of PRF. This decision enables evaluation of algorithmic performance across different scene-decorrelation timescales.

Once the data is retrieved according to slow time decimation and CPI length parameters, it is arranged in a fast-time (N range samples) by slow-time (P pulses) matrix for each channel, yielding four matrices that are stacked along the polarimetric dimension to form an $N \times P \times 4$ data matrix $\mathbf{D}_{in}$. The real-valued data at this point is in ADC counts. We scale it to voltage by multiplication with the constant

$$c = \frac{10^{\frac{a}{20}}}{2^{(b-1)} - 1}$$

where a is the maximum voltage input of the ADC [dBm] and b is the number of bits per sample used for quantization of the incoming analog signal. A final signal conditioning step removes DC bias by subtracting the mean voltage from each channel.

⁶Note that decimation by an even number of pulses is not allowed. Doing so would result in retrieving recorded data from the same two channels each time, prohibiting formation of the full scattering matrix measurements.




Figure 4.12: IF signals are at second Nyquist, so the spectrum centered at $\frac{f_s}{4}$ is reversed.

4.4.2 Basebanding

Because the data is digitized at IF 945 MHz and the ADC sampling frequency is $f_s = 1260$ MHz, the spectrum of the digitized data lies in the second Nyquist zone. That is, the spectrum centered at $\frac{f_s}{4} = 315$ MHz, which is folded in at folding frequency 630 MHz, is frequency-reversed, while the spectrum centered at $-\frac{f_s}{4}$ is not,




Figure 4.13: A representation of the discretized spectrum shows that basebanding the desired spectrum amounts to a counterclockwise rotation of $\frac{\pi}{2}$ [rad/sample].

as shown in Figure 4.12. In discrete normalized frequency space, one way of recovering the basebanded analytic signal amounts to rotating the spectrum by $\frac{\pi}{2}$ rad/sample, as shown in Figure 4.13. This amounts to taking the Hadamard product of each pulse in $\mathbf{D}_{in}$ with the vector $e^{j\frac{\pi}{2}n}$, $n = 0, 1, 2, \dots, N - 1$. Following conversion to baseband, the data is lowpass filtered to attenuate any frequencies outside the waveform's bandwidth.

4.4.3 Downsampling

As discussed in Section 3.2, the waveform bandwidths programmed into Four Eyes ranged from 4 to 400 MHz. Because the sampling frequency of the ADCs is $f_s = 1260$ MHz, the

incoming data is oversampled for all waveforms. In lower bandwidth cases, the data is very heavily oversampled. This results in the scattering matrix for a single radar range resolution mapping to many range bins, as shown in 4-14. In order to resolve this, the data are downsampled to twice the critical Nyquist rate.⁷ So, for example, data from the 4 MHz waveform is downsampled to $f_s = 8$ MHz; data from the 40 MHz waveform is downsampled to $f_s = 80$ MHz; and so on. Following the downsampling stage, we have data for which each scattering matrix measurement is repeated in exactly two contiguous range bins, as shown in 4-15, regardless of the waveform used. Without this step, application of any signal processing operations that rely upon sampling the background—e.g., CFAR processing or statistical estimation—will draw many samples from the same resolution cell. Such samples will clearly not be independent and identically distributed (iid), which is a key assumption underlying the success of these processing operations. Hence, downsampling is a critical enabler of subsequent processing techniques.

4.4.4 Pulse compression and Doppler processing

Following downsampling, the data is pulse compressed and Doppler processed. As was briefly mentioned in Section 4.2, not building a sidelobe reduction window directly into the matched filters enables experimentation in the pulse compression stage with various sidelobe reduction windows.⁸ Each such window significantly reduces the -13 dB range sidelobes a point response will have following matched filtering. Sidelobe reduction comes at the expense of SNR and range resolution degradations that depend upon the window used. The window functions implemented in the code suite include Hamming, Blackman, Chebyshev, and Taylor windows, where the latter two have -50 dB sidelobes. Neither the window functions nor mathematical details of pulse compression and Doppler processing are within the scope of this dissertation, as all are very well documented in the literature.

⁷Because the data are complex-valued following the basebanding operation, the critical Nyquist rate is equal to the bandwidth of the waveform.

⁸Sidelobe reduction in radar is necessary because without it, objects with small radar responses can be masked by the sidelobes of responses from nearby objects with larger signatures (Skolnik, 2001).

The interested reader is referred to Richards (2014) and Skolnik (2001). All results shown in Chapter 5 use a Taylor window with -50 dB sidelobes in the range dimension and a Blackman generalized cosine window in the Doppler dimension. The choice of windows was arrived at empirically, following exploration of results achieved with various window combinations.

For each CPI, the code suite’s pulse compression and Doppler processing steps produce two important data products: RTI and RD data stacked across polarimetric space in the linear basis, as was shown in Figures 2-6 and 2-7. Two additional products, the coherent and noncoherent averages of the RTI, are also computed. The averaged signals and the RTI are inputs into a 1-D CFAR detection algorithm, while the RD is input into a 2-D CFAR detection algorithm.

4.4.5 CFAR detection

A full exposition of CFAR detection is beyond the scope of this dissertation. The interested reader is referred for an excellent and thorough treatment of the topic to Richards (2014). At a high level, CFAR detection assumes that noise and clutter background signal power follows a particular probability distribution, then estimates a statistic of that distribution from samples of background data that are presumed to be




Figure 4-14: Prior to downsampling, lower bandwidth waveforms are heavily oversampled in range.




Figure 4-15: Following downsampling to twice critical Nyquist, exactly two range bins correspond to the waveform’s range resolution.

iid.⁹ Estimation of the power in the background is repeated for each of the cells undergoing detection, and hence the detector adapts to what may be changing backgrounds across a radar scene. The number of background samples N used for statistical estimation and a fixed probability of false alarm P_{fa} determine how many times larger than the estimated statistic the power in a given radar cell must be in order to declare that the data in the cell does not belong to the background. When the power in the cell exceeds this threshold, the detector declares a detection.

There are numerous types of CFAR detectors. Two were implemented in Four Eyes' processing suite: the cell-averaging CFAR (CA-CFAR) and the order statistic CFAR (OS-CFAR). Both variants make the usual assumption that the power of the background—whether the background is noise-only or clutter plus noise—is described by an exponential random variable (Rohling, 2006). That is, it assumes that the background's voltage amplitude is Gaussian-distributed. What varies is the estimated statistic and the factor α , by which that statistic is scaled in order to determine the overall threshold. The CA-CFAR statistic is the mean, computed from the set of amplitude samples as

$$z_{ca} = \left[\frac{1}{N} \sum_{i=1}^N x_i \right] \quad (4.19)$$

where i is the sample index. A CA-CFAR detection is declared in a given cell if its value is larger than $\alpha_{ca} z_{ca}$, where α_{ca} takes the form

$$\alpha_{ca} = N \left(P_{fa}^{\frac{1}{N}} - 1 \right) \quad (4.20)$$

(Richards, 2014). In the OS-CFAR, the statistic is the k^{th} order statistic. An order statistic k is simply the k^{th} value in a the set of N samples after the N samples have been sorted in order of increasing value. For example, if a set contains 21 numbers, then

⁹Note that these assumptions are the reasons that the standard CFAR detection approach fails in the presence of low grazing angle sea clutter. Not only does such clutter fail to obey the usual distribution assumptions, but it is non-stationary and inhomogeneous, so the iid assumption fails as well (Greco et al., 2008).

choosing $k = 11$ selects the median value in the set—i.e., x_{11} . Clearly, $1 \leq k \leq N$. In the OS-CFAR implementation for Four Eyes, k is chosen such that the order statistic is nearest the 75th percentile. An OS-CFAR detection is declared in a given cell if its value is larger than $\alpha_{os}x_k = \alpha_{os}z_{os}$ where

$$P_{fa} = \frac{N! (\alpha_{os} + N - k)!}{(N - k)! (\alpha_{os} + N)!} \quad (4.21)$$

(Richards, 2014). Rather than implementing this numerically, Four Eyes uses the look-up table of α_{os} values given in Rohling (1983).

In both CA-CFAR and OS-CFAR, the background is sampled using windows to determine which surrounding cells will be included in the background statistic estimate for a given cell under test (CUT). Figure 4-16 shows a notional window in a 1-D detection operation. Figure 4-17 shows a notional window in a 2-D detection operation. For CA-CFAR, a guardband is used as protection against contaminating the statistical estimate of the background with samples from an object in the CUT, which may be extended in range and/or Doppler. The OS-CFAR does not use a guardband. In the Four Eyes processing suite, the choice of N depends on the waveform used, and is made so as to keep the length of the windows in range roughly constant across waveforms even as the range resolution of the waveforms varies.




Figure 4-16: A notional 1-D CFAR window used for sampling the background of a CUT is shown. The window selects an equal number of samples to use in the leading and lagging halves of the window. In CA-CFAR, a guardband is employed to protect against including samples of objects with larger range extents in the background statistic estimate. In Four Eyes processing, the window size is determined based on the waveform.

The OS-CFAR is designed to be resilient to changes in the clutter background, and to allow detection of multiple objects that may be closely spaced without allowing the presence of another object in the window to skew the estimated statistic and artificially inflate the threshold value. We found that the OS-CFAR did not outperform the CA-CFAR in terms of fewer false alarms, because sea clutter returns do, in fact, look object-like in many cases. We also found that the OS-CFAR detector is so resilient to changes in the background that it did not detect extended objects without modification of the CFAR to include a guardband, which changes the essence of the detector. Given that the performance difference of the two CFARs in low grazing angle sea clutter was very small, we opted to focus exclusively on the more common CA-CFAR in the algorithm development research presented in the next chapter.

The output of the CA-CFAR processing stage are sets of detections. In the 1-D cases, the detections are functions of range only; in the 2-D cases, the detections are functions of range and Doppler. These CFAR detections are passed to the polarimetric processing stage along with the RTI data, RD data, and (non)coherently averaged data.

4.5 Data labeling

Now that discussion regarding the radar signal processing flow and products is complete, we turn our attention to the final pre-processing component. Data labeling is a critical step that enables quantification of an algorithm's performance. It is very challenging to label radar data, which is dynamic by nature, has different processing outputs depending upon user-input processing control




Figure 4-17: A notional 2-D CFAR window used for sampling the background of a CUT in an RD image is shown. The guardband, consisting of some number of guard cells, is used in CA-CFAR, and may include guard cells in the Doppler dimension as well as in the range dimension.

τ [μsec]	β [MHz]	Bins to 4 km	Range per Bin [m]
2	4	198	18.74
2	40	1975	1.874
2	150	7406	0.5
2	400	19749	0.187
4	4	182	18.74
4	40	1815	1.874
4	150	6806	0.5
4	400	18150	0.187

Table 4.1: The number of range bins that remain following downsampling, up to the 4 km range extent, depend upon the waveform and its minimum range. The table gives this number of range bins by waveform as well as the range extent of each bin.

parameters, and features objects whose range and Doppler signatures vary as functions of time and are themselves dependent upon processing control parameters. This section discusses the approach taken for data labeling.

The locations of objects in RTI data, which is a function of range and pulse time (slow-time), more readily offer labeling means than does RD data, which is a function of range and Doppler as well as of CPI length. Therefore, we focused on RTI data as a means of data labeling generation.¹⁰ The data recordings identified in Table 3.3 are all ~ 10 sec in length, comprising some 62500 pulses each. As noted in Table 3.2, the radar returns were really only clutter-limited (as opposed to noise-limited) to a range extent of 4 km, at which point sea clutter return power was reduced significantly enough that standard 2-D CFAR on RD maps easily yielded object detections without also producing numerous false alarms caused by sea clutter returns. Hence, we limited our data labeling focus to the more challenging range extent between the waveform minimum range and 4 km. The number of range bins this extent comprises for a given file following the downsampling operation described in Section 4.4.3 depends upon the waveform used. These figures are given in Table 4.1.

To create a meaningful set of labels for each file, every range bin must have an associated label—1 if the range bin contains an object, 0 if it does not—for every

¹⁰The consequence of this choice is that RD data is not a candidate for follow-on algorithm development work without first performing substantial additional data labeling work; however, RD data is still a key component of the PCL algorithm, as shall be seen in Chapter 5.


recorded pulse. In the best case, a file featuring the 4 μsec 4 MHz waveform, requires $(62500)(182) > 11 \times 10^6$ labels; in the worst case, a file featuring the 2 μsec 400 MHz waveform requires $(62500)(19749) > 1.2 \times 10^9$ labels. Clearly, this is a vast number of labels to determine, store, and quickly access.

Recognizing that the number of range bins in each pulse that actually contain an object is very small relative to the number of range bins contained in the pulse, sparse matrices were natural candidates for label storage. To populate these label matrices, rather than looking at 62500 pulses individually and selecting known object locations, we opted to process slow time decimated RTIs and interpolate between target start and end locations to yield regions corresponding to object labels for that RTI.


To see how this was done, first consider the RTIs of File 325 shown in Figure 4-18. The file features returns from the 4 μsec 150 MHz waveform and contains a single object of interest, The New Englander, whose returns over time are framed in black on each image and whose picture is shown on the HH image inset. All returns outside the black frame are due to sea clutter and/or noise. The RTIs show the entire 10 seconds of the recorded file and are formed with slow time decimation factor $m = 7$, corresponding to ~ 893 Hz PRF. In Figure 4-19, the images have been zoomed in to a 50 m range swath centered on the object.

Creating a label matrix populated with zeros equivalent in size to the full RTI provides a convenient starting point for label population. Imagine now dividing the full HH RTI into contiguous chunks across slow time, then using continuous lines to mark on each of these shorter RTIs the boundaries of the object's signature in range.¹¹ Mapping the continuous region between these lines to the corresponding discrete region in the label matrix, then changing labels therein to ones to indicate the presence of an object, allows data labeling across the full range extent in RTI space. The RTIs are divided into the

¹¹We utilize only the HH RTI because the RTIs are time-aligned. Any object responses will map in range to the object responses in other polarizations, even if the signatures in other polarizations may be stronger or weaker.



(a) Full file RTI, File 325, HH polarization



(b) Full file RTI, File 325, VV polarization

Figure 4.18: Full file RTIs show the returns from a single object of interest, The New Englander, in sea clutter. Variation in the range extent and relative power in the object's signature over time is apparent.




Figure 4-19: Full file RTIs show the path followed by the single object of interest, The New Englander with a close up view. Variations in signature and range extent across pulses are even more apparent.




Figure 4-20: The data labels for File 325 are shown. The gray area indicates the region labeled as containing an object's return; the black regions are labeled as not containing an object's return; and the white regions indicate areas that are unlabeled and will not be used in algorithm performance assessment.

minimal number of chunks necessary to capture variation in an object’s range extent and signature across time. This process is repeated for each object in each file of interest.

The results can be seen for File 325 in Figure 4-20. The zoomed-in range extent corresponding to that in Figures 4-19a and 4-19b is shown in Figure 4-21. The label images have in both cases been decimated in slow time identically to the RTIs shown in Figures 4-18 and 4-19 to enable easier comparison. The gray regions in these images correspond to the region in the label matrix that indicates an object is present. The black regions correspond to regions in the label matrix that indicate no object is present. The white regions correspond to “do not use” regions in the label matrix,

which are intended to prevent inaccurate labeling of returns near an object. The “do not use” regions together with the “object” regions form a rectangle across the full extent of the RTI for each target, which is partly an artifact of the overall label generation process.

Every file of interest for algorithm development, as delineated in Table 3.3, has an associated sparse data label matrix whose contents are determined according to the process described above. For each CPI processed in Four Eyes’ processing suite, a file’s sparse data label matrix is sampled according to the specifications of the CPI. That is, the labels for each pulse in the RTI data cube are extracted across the ranges the cube contains. These extracted labels are carried with the pre-processed data through subsequent processing stages to enable assessment of algorithm and classification performance.




Figure 4-21: The zoom-in on the data labels for File 325 demonstrates excellent agreement with the slow time decimated RTIs shown in Figures 4-19a and 4-19b.

Having addressed details of the achievement of the first major contribution of this dissertation, we turn attention now to the second major contribution: development of practical means to mitigate the impacts of low grazing angle sea clutter on the standard CFAR detection process.

Chapter 5

Polarimetric Co-location Layering (PCL)

This chapter presents a novel algorithm that is one of the major contributions of this dissertation. The aim of the algorithm development phase of this research was to leverage polarimetric dimensionality to find an approach to mitigating the impact of low grazing angle sea clutter that is robust across PRFs, bandwidths, and object types in the Four Eyes Point de Chene Dataset. As was shown in the brief review of polarimetric approaches to enhancing detection performance in low grazing angle sea clutter in Chapter 2, there are numerous creative approaches to sea clutter mitigation that are impractical in terms of the standard radar signal processing chain. Such approaches require introducing large computational complexity into the standard radar signal processing chain or require restructuring the processing chain entirely. Hence, the other aim of the algorithm development component of this research was design of an approach that would plug directly into the standard processing chain, imposing neither chain restructuring nor infeasible computation times.

The result of this research is Polarimetric Co-location Layering (PCL), an algorithm that leverages a fundamental characteristic of the Doppler spectra of sea clutter and man-made objects to classify detections produced using the standard CFAR detection approach as either detections on objects or on clutter. The first section of the chapter discusses the sea clutter Doppler spectrum, including a review of work that has been done on the topic and a review of other detection techniques that leverage in tandem the Doppler and polarization characteristics of sea clutter. The second section motivates PCL, describing the fundamental polarimetric principles that are the basis of its efficacy.

The third section details the core PCL algorithm and introduces metrics by which the performance of PCL can be compared to that of standard CFAR detection processes. Included in this section is a discussion showing that the algorithm plugs directly into the standard radar signal processing chain, operating in parallel with the addition of efficient computational blocks. The third section shows results demonstrating that PCL is indeed robust across PRFs, bandwidths, and object types. The fourth section discusses PCL performance results across bandwidths and PRFs. The adaptation of PCL to detection of dynamic targets is also discussed in this section, as are the results of the adapted algorithm. The fifth section discusses integration of PCL into the standard radar signal processing chain. The section also includes an analysis of PCL’s computational complexity, showing that except in degenerate cases, use of the algorithm does not impose any processing delays as compared to standard radar processing times. The final section summarizes the chapter.

5.1 Sea clutter Doppler spectrum

This section first reviews the literature related to the polarization and look angle dependence of the sea clutter Doppler spectrum. Thereafter, we review other approaches that involve using Doppler techniques with polarimetric radar for mitigation of sea clutter; this review includes references to literature wherein such techniques have been suggested as a promising means of achieving this aim.

5.1.1 On the polarization dependence of sea clutter Doppler spectrum

The polarization dependence of the sea clutter Doppler spectrum has long been understood, with publications dating back nearly 50 years. In 1968, Pidgeon reported that at low grazing angles, the Doppler shift of HH sea clutter at C band was 2-4 times as large as the Doppler shift of VV sea clutter, but that the two orthogonally polarized Doppler spectra had approximately the same spectral width (Pidgeon, 1968). Pidgeon

posited that the difference was due to the responses of HH and VV to different sea surface scattering components. The following year, Valenzuela and Laing of the Naval Research Laboratory published data from a multi-band experiment aimed at more completely characterizing the Doppler spectra of radar sea echo (Valenzuela and Laing, 1970). Their data comprised P, L, C, and X bands; HH and VV polarizations; some azimuthal look variation with respect to the upwind direction; a few different sea states; and grazing angles ranging from a few degrees to 45° . Among other observations, the researchers reported that the average differential velocity of Doppler spectra in HH and VV first reported by Pidgeon appears to increase at lower grazing angles. Unsurprisingly, they found that sea clutter Doppler spectra are radar frequency as well as polarization dependent, and that the differential Doppler across HH and VV polarizations increases with increasing sea state.

There are numerous publications in later years that supplement the seminal work discussed above. The remainder of this brief discussion focuses on literature specific to very low grazing angle geometries. Chief among these are the reports of H.C. Chan (Chan, 1987; Chan, 1990), who analyzed an MIT LL multi-band sea clutter dataset collected on the Atlantic-facing coast of North Truro, MA. Chan also noted the polarization dependence of sea clutter Doppler spectra. He also defined metrics by which to quantify the spectral spread and offset from DC, then applied those metrics across HH and VV polarizations and look directions with respect to upwind. The results of his work supported the findings of the aforementioned publications and further added that the average Doppler differential is dependent on look direction. Specifically, he reported that the upwind and downwind directions have the largest polarization-dependent average differential Doppler, but that the differential approaches zero when looking in the crosswind direction.

The findings of these early large-scale experiments have continued to be borne out with smaller data collections reported in the literature. Werle reported an X-band data

collect verifying Pidgeon’s conjecture that different scattering processes are responsible for the difference in HH and VV Doppler spectra (Werle, 1995). Smith et. al. reported that wave groups and breaking waves had markedly different mean Doppler frequencies across HH and VV polarizations, again particularly in the upwind direction (Smith et al., 1996). Lamont-Smith used laboratory-generated wind waves in a wave tank to show that both the maximum and mean Doppler velocities show strong dependence on polarization at low grazing angles to frequencies as high as W band (Lamont-Smith, 2000).

Much of this early data was not collected with fully polarimetric radar in its modern sense. Instead, the data was collected for a given grazing angle, sea state, and look direction using HH polarization. Then, immediately afterward and before shifting to a new geometry, the antenna was manually rotated to provide the orthogonal VV co-polarization. It is not surprising, though, that the findings of researchers using modern polarimetric radars continue to support the large body of earlier work (Walker, 2001; McLaughlin et al., 1995). The polarization and look angle dependence of the sea’s Doppler spectrum is presently so well known that the preceding relationships feature prominently in the “Characteristics of Radar Sea Clutter” chapter of (Ward et al., 2006), which is currently among the most exhaustive available texts on the topic of sea clutter.




Figure 5-1: A corrected and generalized reproduction of a plot from (Ward et al., 2006) that was originally printed in (Ward et al., 1990) shows the trend followed by the mean Doppler differential of sea clutter in HH and VV polarizations as a function of look direction. HH typically exhibits more a more significant mean Doppler than does VV.

In that text, Ward et. al. present a plot showing that the Doppler differential of sea clutter echo has a co-sinusoidal dependence on look direction, with the average HH Doppler frequency being more significant than the average VV Doppler frequency. Neither a source for the data nor a frequency band is definitively given in the text or the paper from which it was reproduced (Ward et al., 1990). However, it is instructive to consider the general trend the data follows. Hence, a corrected version of that plot is reproduced in Figure 5.1 for the reader's consideration. Note the inversion of the Doppler frequency axis, which is the convention that has been adopted in RD images in this dissertation. The reader should not infer any relationship from the plot save for the general trend in Doppler shift as a function of look direction and the fact that the Doppler shift is more pronounced in the HH polarization.

Given the preceding discussion, it is not surprising that several researchers are indeed looking to Doppler techniques with polarimetric radar as a means of sea clutter mitigation, and have been doing so for decades.

5.1.2 On sea clutter mitigation using Doppler techniques in polarimetric radar

It bears reiterating at the outset of this section that, as discussed in Chapter 2, much of the open literature features approaches designed using data collected by McMaster University's IPIX X-band radar. A small subset of this data, recorded over two different outings in 1993 and 1998, was made publicly available on the web in 2001. Despite the system's original characterization by its designers as "dual-polarized," the radar in today's terminology is actually fully polarimetric. That is, it records the full polarimetric scattering matrix - i.e., HH, HV, VH, and VV - in each cell utilizing the ALT transmit scheme. The IPIX data features a single canonical point target floating in sea clutter across a range of sea states. It is the only publicly available polarimetric sea clutter dataset, though it is very limited in its look angles, range resolutions, and range extents.

The principal investigator for the McMaster research laboratory that stood up the IPIX radar was and still is Prof. Simon Haykin, though the lab is now called the Cognitive Systems Laboratory. Haykin was the originator of the now-popular notion of “cognitive radar,” and for over 25 years he has championed building remote sensing systems that can measure and leverage any combination of the dimensions of signal information—time, space, frequency, and polarization—to understand and discriminate between the objects it senses in its surroundings (Haykin, 1990; Haykin, 2006). Haykin has erected one corner of his vast body of work upon the study of radar sea clutter, a fair portion of which is focused on utilizing polarimetric Doppler techniques to detect targets in clutter.

As early as 1985, Haykin and his team posited that fusing polarimetric information together with Doppler signatures would prove useful for the application of sea ice and growler detection in shipboard radars (Haykin et al., 1985). Five years later, Haykin’s team published work using early IPIX data, postulating that shipboard radars have difficulty detecting small objects precisely because they ignore clutter’s two most distinguishing characteristics: Doppler and polarization (Currie et al., 1990). In this work, the authors present several images across polarization space in the H-V basis of a radar reflector in sea clutter, noting the Doppler spectrum differences between sea and object across looks in the polarimetric dimension. In 1991, the team published work indicating that for object detection, using the Doppler spectrum in only the co-polarized returns was at least as advantageous as using the Doppler spectrum in cross-polarized returns (Haykin et al., 1991).

There are more instances of such findings by Haykin and his team, but perhaps the most pointed statements Haykin makes are that “The [feature] most extensively studied has been dual [sic] polarization. Results were presented showing the differing properties of HH and VV for sea clutter, whereas for ice targets and rain clutter the HH and VV returns were highly correlated. This difference could be exploited to provide both the identification of the source of the clutter, and for improved target detection” (Haykin

et al., 1994). Haykin makes a similar note in a 1996 publication that the sea clutter Doppler spectrum will change over time and across polarizations more quickly than will that of a moving object. He goes on to develop a theoretical framework for detection across a higher number of dimensions (Jones and Haykin, 1996). It is Haykin’s early work that has laid the foundation for PCL.

While Haykin’s work is really the cornerstone of research in this area, numerous other researchers have been pursuing joint polarization-Doppler based techniques for sea clutter mitigation. Wanielik and Stock proposed execution of CFAR on a vector at each cell-under-test consisting of that cell’s polarimetric scattering matrix measurement (Wanielik and Stock, 1990). A research team out of the TNO Physics and Electronics Laboratory published work using range-Doppler data across the polarimetric dimension to discriminate between sea clutter returns and returns from a small sloop (Smith et al., 2002). There are numerous other publications that do not leverage mutual information across polarimetric channels, but do demonstrate the differences in Doppler power spectral density in sea clutter cells versus sea clutter cells containing objects (Greco et al., 2010; Li and Shui, 2016).

Additionally, there is at least one other application area in which the fusion of mutual information in the Doppler and polarization dimensions has led to greatly enhanced classification capability: polarimetric Doppler weather radar. The weather radar community has a near 30-year history of leveraging differential reflectivity—i.e., polarization ratio of HH returns to VV returns—and other polarization-based metrics to differentiate between precipitation types and sea clutter in Doppler radar weather maps (Husson et al., 1989; Bringi and Chandrasekar, 2001; Islam et al., 2012).

The preceding references are by no means exhaustive. Studies in this area are and will continue to move forward on a wide scale. Bearing this in mind, we now turn our attention to PCL.

5.2 The fundamental principle behind PCL

The fundamental principle behind PCL is that while sea clutter exhibits an average Doppler differential across polarizations, man-made objects do not. That is, the sea clutter scattering mechanisms that respond to horizontal polarization tend to move at different velocities *on the average* with respect to the radar than do those scattering mechanisms that respond to vertical polarization. For man-made objects, this is untrue. If an object is moving at Doppler velocity v m/s with respect to the radar, then that object will be moving at v m/s with respect to the radar in both horizontally and vertically polarized RD images, provided that the object has a signature in both images. The preceding statement is certainly valid for a rigid point target; consider a moving trihedral or the IPIX radar beach ball, for example, to see that this is true. As will shortly be shown, the statement is also true for extended objects, which comprise many rigid scatterers. However, accommodations must be made for variation in polarimetric signatures over the extent of the objects. Accommodating this variation is precisely what PCL does. Before explicating the means by which PCL accomplishes this task, we orient the reader using Four Eyes data to the fundamental principle we have just discussed and its manifestation in the outputs produced by the signal processing chain detailed in the preceding chapter.

For the sake of continuity, we constrain discussion for the present to the data file to which the reader has already been exposed: File 325—which contains a single object of interest, The New Englander—whose full file RTIs and data labels were shown in Section 4.5. The reader will recall that File 325 features the 4 μ sec 150 MHz waveform, and thus has a range resolution of 1 m. The RD images for a 50 pulse CPI of this file, taken from the center of the data recording, are shown in Figure 5-2. The images were formed using data slow-time decimated by $m = 7$ to ~ 893 Hz PRF; the images are zoomed in on the Doppler velocities that feature the strongest sea clutter return. The look direction for this data is within 45° of upwind, so the average differential maps to a location between




Figure 5-2: The fundamental principle behind PCL is captured by a pair of RD images, one in HH and one in VV. The images show that while the average Doppler velocity of sea clutter is stronger in HH than in VV, the return from the object is at the same Doppler in both polarizations.

90° and 135° region on the average Doppler differential cosinusoidal trend plot shown in Figure 5-1. In the RD images, the average Doppler differential is apparent, though clearly the Doppler spectrum across range shows significant overlap across images.¹ However, the radar return from The New Englander exhibits the same Doppler velocity in both polarizations.

Consider running a 2-D CFAR detector on these images independently. Note that an operation often layered onto a CFAR detector is peak-picking. That is, when the data are oversampled, as is Four Eyes’ data after downsampling to twice critical Nyquist, then contiguous range bins may surpass the detection threshold such that multiple detections are declared on the same scatterer. To mitigate this, one can select only the peak responses as detection candidates. Peak-picking is implemented simply by examining cells whose values have surpassed the CFAR threshold and checking whether those values are larger than their neighbors in both the range and Doppler dimensions. That is, if the

¹Some of the features of sea clutter that were discussed in 2 are visible here as well. Specifically, the HH returns look “spikier” than the VV returns, while the VV returns are generally stronger than the HH returns.


value in cell $y[m, n]$, where m = range index and n = Doppler index, surpasses the CFAR threshold, then a detection is only declared if also

$$y[m, n] > y[m \pm 1, n \pm 1] \quad (5.1)$$


The set of detections produced by this sequence of operations for the RD images in 5.2 is shown in Figure 5.3, where P_{fa} for CA-CFAR detection was set to 10^{-6} . The impact of sea clutter on CA-CFAR detection in both channels is clear: Despite there being only a single object of interest in the radar scene corresponding to a single detection in each channel, hundreds of detections are actually produced. Specifically, detection in the HH channel produces 226 false alarms, while detection in VV produces 213 false alarms; the fact that sea clutter background characteristics do not obey the exponential power distribution assumption imposed by CFAR detection results in a significantly higher P_{fa} than required.

The keen reader will observe that detections on the object of interest are approximately co-located in range and Doppler across channels. Therefore, a naive first step to false alarm mitigation might be requiring that detections be co-located in range-Doppler across polarizations to be declared object detections rather than sea clutter false alarms. In fact, this approach occasionally works well on very low resolution waveforms. We found that in the 4 MHz bandwidth (37.5 m range resolution) data, this approach was often sufficient to mitigate most false alarms due to sea clutter. However, at such low resolutions, the number of false alarms produced by sea clutter is negligible in comparison to the numbers of false alarms produced in 40 MHz or higher bandwidth data, as will shortly be seen. The more challenging cases are thus the higher resolution waveforms, with the 400 MHz waveform (0.375 m range resolution) being the most challenging.

The reasons higher resolution waveforms are more challenging are twofold. First, as discussed in Chapter 2, higher bandwidths result in spikier sea clutter textures. This phenomenon leads to more false alarms due to clutter because the background clutter-



(a) File 325, HH, CA-CFAR detections



(b) File 325, VV, CA-CFAR detections

Figure 5-3: CA-CFAR detections produced independently in HH (detections indicated by black markers) and VV (detections indicated by magenta markers) show the high number of false alarms produced in a 50-pulse CPI from File 325 across both polarizations. Only one of the detections in each case is on the signature of the object of interest, The New Englander.

plus-noise signal power sampled by a CFAR window is not uniformly increased. Second, higher bandwidth waveforms resolve finer object structure because of their better range resolution. Polarimetric signature variations across the extent of an object will thus produce detections on that object that are not necessarily co-located in range across polarizations. The consequences of requiring co-location across polarization in RD images for higher bandwidths, therefore, are that detections on objects may be thrown out because they are co-located in Doppler but are not co-located in range, while false alarms on sea clutter may be retained because false alarms in one channel happen to be co-located in both dimensions with some of the many false alarms in the other. Moreover, a detection in one channel may be approximately co-located with numerous detections in the other channel. In such cases, it is not clear how one should associate detections across polarizations to declare them “co-located,” for several different associations are possible.


The key to addressing these issues is to realize that the Doppler differential of sea clutter, while not necessarily evident from one detection to another, is evident *on the average*. With this in mind, we turn attention to the specifics of the PCL algorithm.

5.3 PCL

This section builds intuition for PCL by first examining the algorithm’s components as applied to the simplest case in the Four Eyes Point de Chene Dataset. This section culminates with the reader’s first exposure to the empirical false alarm and continuity of detection metrics, which are means of quantifying PCL performance in 1-D.

5.3.1 Core algorithm

To simplify illustration of the core PCL algorithm, we focus on File 273, which features the 2 μ sec 4 MHz waveform. Because this is Four Eyes’ lowest resolution waveform, the preliminary steps of PCL will be easiest to visualize. File 273 contains data collected




(a) File 273, HH, 100 pulse RTI (b) File 273, VV, 100 pulse RTI

Figure 5-4: RTIs show 100 pulses of compressed data from File 273 in each of the HH and VV channels.

on the buoy, which has a strong return relative to the background and is stationary in range at ~ 1.7 km. The buoy exhibits a Doppler signature over time that takes on values in the approximate range of $[-1, 1]$ m/sec as the buoy bobs atop the undulating ocean surface. Few false alarms are produced for a waveform of this bandwidth and the buoy is not obscured by a clutter background on an RD image because of its low Doppler and range from the radar. Hence, it is not a challenging case, but it is an instructive one.

Let us first consider RTIs from one CPI of this file across HH and VV polarizations. The CPI comprises 100 pulses, slow-time decimated by factor $m = 7$, yielding an ~ 893 Hz PRF. The RTIs for HH and VV are shown in Figure 5-4. The buoy's return is evident throughout the RTIs and is indicated on the images. All other returns are due to sea clutter.

Recalling from Chapter 1 that the SNR of a given object is dependent upon the number of pulses integrated (or, equivalently, averaged) in the radar processing chain, we compute the basic coherent average of each of these RTIs to yield 1-D range profiles



(a) File 273, HH, 100 pulse coherent average

(b) File 273, VV, 100 pulse coherent average

Figure 5-5: Coherent averages of 100 pulses of the buoy scene in File 273. Coherent averaging is a means of capturing the average Doppler response in each range bin.

of the radar scene. That is, we compute coherent averages

$$\begin{aligned}
 s_{hh}[n] &= \frac{1}{P} \left| \sum_{p=1}^P y_{hh}^{\{p\}}[n] \right| \\
 s_{vv}[n] &= \frac{1}{P} \left| \sum_{p=1}^P y_{vv}^{\{p\}}[n] \right|
 \end{aligned} \tag{5.2}$$

where p corresponds to the RTI pulse index, n corresponds to the range index, and subscripts hh and vv indicate the channel whose RTI is being coherently averaged. The signals s_{hh} and s_{vv} for the RTI data shown in Figure 5-4 are shown in Figure 5-5. The reader will recall that, for our purposes, we somewhat arbitrarily fix $P = 100$ for all CPIs, regardless of waveform. The rationale behind this choice is that it is well known that the transient characteristics of sea clutter decorrelate on timescales less than 10 msec at X-band (Chan, 1987; Antipov, 1998). Consequently, choosing $P = 100$ —which comprises ~ 16 msec at Four Eyes’ highest PRF—ensures averaging over timescales at least as long as the usual decorrelation time of transient sea clutter characteristics.

The effect of coherently averaging these P pulses is that we have averaged over the Doppler response of the scatterers in each range bin—that is, we have captured the average Doppler frequency of each channel in each bin. To see that this is true, consider a scatterer at range R_0 from the radar, moving at constant radial velocity v with respect to the radar. The location of this scatterer as a function of slow-time in either channel can be written as

$$\begin{aligned} y[p] &= Ae^{-j\frac{4\pi}{\lambda}\left(R_0 - v\left(pT + \frac{2R_p}{c}\right)\right)} \\ &= Ae^{-j\frac{4\pi}{\lambda}\left(R_0 - \frac{2vR_p}{c}\right)} e^{j2\pi\left(\frac{2v}{\lambda}\right)pT} \end{aligned} \quad (5.3)$$

where the “stop-and-hop” approximation, which assumes target motion stops while the radar pulse is in transit, has been used; A captures signal amplitude and attenuation; c = the speed of light; p = the pulse number in the CPI pulse sequence; R_p = the scatterer’s range to the radar on the p^{th} pulse; T [sec] is the pulselength. The first exponential term captures a phase shift relative to all terms of the slow-time sample sequence. The second exponential term captures the Doppler frequency of the scatterer (Richards, 2014).

The reader will recall that the operation typically performed by the coherent radar’s standard radar signal processing chain on a CPI of pulses is Doppler processing. Doppler processing amounts to computing the discrete Fourier transform of samples across slow-time in each range bin to determine the strength of each of the complex exponentials measured across the CPI time—at each of the velocities measured by the radar’s slow-time sampling frequency, the PRF—comprised by the second exponential term in Equation 5.3. By coherently averaging across slow-time instead of Doppler processing, we are still improving our signal SNR by a factor of P , but rather than doing so by integrating the strength of the individual Doppler velocities of all scatterers in each range bin, we are capturing the average Doppler velocity in each range bin in each channel.

We return our attention to the coherent averages s_{hh} and s_{vv} with the understanding that these signals capture the strength of the average Doppler response in each range bin.

Consider now passing s_{hh} and s_{vv} through a 1-D CFAR detection process, yielding two sets of detections that are functions of range-only in each channel. The detections produced by CA-CFAR for the signals in Figure 5.5 are shown in Figure 5.6. The probability of false alarm for the CA-CFAR results shown was again set to $P_{fa} = 10^{-6}$. (The reader is reminded that we are exploring the 4 MHz waveform as an instructive case; conclusions should not be drawn regarding the number of detections that 1-D CFAR operations will yield in the cases of finer resolution waveforms.) Two pairs of detections are co-located in range across HH and VV polarizations, as indicated. One of these pairs comprises detections on sea clutter while the other comprises detections on the buoy; the latter is expected per the discussion in Section 5.2 regarding object signatures across polarizations.

Imagine now that the radar is an observer, floating in the sea at the stationary range equal to the waveform's minimum range. The locations of detections on the moving sea surface will change from CPI to CPI. However, because HH and VV capture different scattering mechanisms of the sea surface with different mean Doppler frequencies and the mean Doppler of HH is more significant than the mean Doppler of VV, we expect that even those sea clutter detections that persist from CPI to CPI will move across range at different rates with respect to one another. In other words, sea clutter detections that are co-

located on one CPI will not remain co-located across a series of CPIs. Such detections in




Figure 5.6: The output of the 1-D CA-CFAR operation on coherently averaged signals yields a set of detections for each channel. In this CPI, two pairs of detections—one on the buoy and one on sea clutter return—are co-located in range.

HH will “wash over” the observer—or move away from the observer, depending on look direction—faster than will the corresponding detections in VV.

Thus, we can set a minimum number of CPIs across which a pair of detections must retain their co-location in range in order to be declared detections on objects rather than detections on the sea. Recognizing also that the strength of an extended object’s signature will vary somewhat in range as a function of polarization for higher bandwidth measurements, we can opt to relax the requirement that detections must be precisely co-located in range. This brings us to a definition of *polarimetric co-location*: a pair of detections on the coherently averaged signals s_{hh} and s_{vv} are considered polarimetrically co-located if they remain approximately co-located in range across a certain number of CPIs. “Approximate co-location” in range is captured by algorithm input parameter δ_n , which is the maximum permitted offset in range that will still allow a detection in HH to be considered co-located with an associated detection in VV. The number of CPIs across which a given detection pair must remain approximately co-located in range is captured by another algorithm input parameter, the PCL CPI criterion, which we denote n_{CPI} .

The easiest way to show the operation of PCL is with a video whose frames comprise images like Figure 5-6 across a contiguous sequence of CPIs. Since video does not lend itself well to a dissertation, we opt now to show a short sequence of these images using the instructive case of File 273. For this sequence, no slow-time decimation was used in formation of the 100-pulse CPIs; the PRF of the data underlying these images is thus 6250 Hz. At this PRF, the full sequence of CPIs shown in Figure 5-8 comprises just over one-third of a second of data. The sequence of images is in order row-wise; thus, the reader should read left-to-right and top-to-bottom, as is indicated by the CPI sequence number appearing in the caption for each image. The ordinate units are relative power in dB, normalized to the peak response over the entire sequence of CPIs. To generate this sequence, PCL parameters were $n_{CPI} = 3$ and $\delta_n = 0$. As the sequence shows, for the 4 MHz waveform, the resolution is low enough such that it suffices to allow no range




Figure 5.7: PCL sequence on File 273, part A: $\delta_n = 0$ and $n_{CPI} = 3$




Figure 5-8: PCL sequence on File 273, part B: $\delta_n = 0$ and $n_{CPI} = 3$

offset between detection pairs. Moreover, so few detections are generated on the 4 MHz waveform data that three CPIs sufficiently allows sea clutter detection pairs to separate over time due to sea clutter’s mean Doppler differential.

On this image sequence, HH and VV detections are indicated as usual. Detection pairs that have remained co-located in range for as long as n_{CPI} CPIs are indicated as PCL detections. We see from the sequence that on CPI 3, the buoy has been declared a PCL detection; it retains this status throughout the remainder of the sequence. Note also that a co-located detection pair on sea clutter appears in CPI 10, but because this detection pair does not remain co-located over three CPIs (or even over two, in this case), it never reaches PCL detection status. Moreover, approximately co-located pairs appear in CPIs 2, 4, 6, 15, and 18. Because $\delta_n = 0$, these pairs are never considered for PCL detection candidacy across CPIs.

We are, for the moment, constraining ourselves to the case of low bandwidth waveforms, high-SNR stationary objects, and a high PRF. The reader will bear with this example for a short time longer to allow introduction of metrics by which we can quantify PCL performance in one dimension. We will thereafter expand our consideration to more challenging scenarios before discussing integration of PCL into the standard radar signal processing chain.

5.3.2 1-D performance metrics

To quantify PCL performance relative to standard single-polarization CFAR processing in 1-D, we leverage the precomputed sparse data label matrices described in Section 4.5. The reader will recall that each such matrix contains one row for each pulse in the recorded file and one column for each range bin up to a distance of 4 km. Entries containing object returns are labeled 1 in the matrix; entries containing sea clutter and/or noise are labeled 0. There is also a small do-not-use region surrounding each object’s returns over time that guards against human labeling errors. This matrix is sampled to align with the

pulses and range bins included in a given CPI. From there, it is straightforward to map detections occurring in 1-D HH CFAR, 1-D VV CFAR, and PCL to detections either on objects (hits) or on sea clutter and/or noise (false alarms) for a given CPI. Repeating this process across a sequence of CPIs provides a means of comparing performance of PCL to single-polarization detection.

Examining Figure 5.9 will orient the reader to the means by which we visually capture object detections (hits) over a sequence of CPIs. The sole object of interest in the processed range extent to 2 km, the buoy, is shown along the ordinate. For each CPI during which at least one detection occurred on the object of interest in the 1-D CFAR operation in HH, a blue patch is plotted along the bottom for that CPI; likewise for 1-D CFAR in VV, shown along the middle in green; and PCL, shown at the top in red. If an object is detected continuously across the CPI extent by any of the three detection operations, then that operation's hit line is unbroken. If an object goes undetected by a given operation, then that operation's hit line will be broken at each CPI for which a detection was not made on the object. As shown, PCL has a transient period of at least n_{CPI} CPIs before a PCL detection can be declared on a co-located pair.

Empirical false alarm rate can also be quantified over a sequence of CPIs. As shown in Figure 5.10a, the total number of detections n_{det} , including both hits on objects and false alarms, is computed and plotted for each detection process. Again, the transient period equal to n_{CPI} for PCL is evident. The total number of detections for each process is mapped to the false alarm rate p_{fa} for each process shown in Figure 5.10b using the straightforward relation

$$p_{fa} = \frac{n_{fa}}{n_{dets}} \quad (5.4)$$

where n_{fa} is the number of detections whose data label is 0 out of the n_{dets} total detections made at that CPI. The thick dotted line shown atop the false alarm rate for each process is the mean false alarm rate produced by that process across all CPIs. The reader should note that the key takeaway of the false alarm rate plot is the relative reduction in empirical





Figure 5-9: Hits by detection operation, File 273, 1st 100 CPIs: The continuous detection bar metric is shown in the plot, with the initial latency to PCL detection indicated. An unbroken bar indicates that the detection process maintained continuous detections on the object indicated along the ordinate. A broken bar indicates a missed detection on the object, where the length of the break coincides with the number of CPIs in which the object's detection was missed.

p_{fa} from 1-D HH CFAR to PCL and from 1-D VV CFAR to PCL.² Also, Figure 5-10a is shown only to assist the reader in appreciation of the performance assessment process. Henceforth, only false alarm rate results like the one in Figure 5-10b will be presented.


5.4 PCL results

To this point, a single file has been examined in order to explicate the basics of PCL and to introduce the metrics by which performance is quantified in one dimension. This section extends the problem space by exploring PCL's performance on higher resolution waveforms; lower PRFs; and dynamic extended objects. As these areas are explored, necessary variations of PCL parameters are discussed, as are modifications to the core PCL algorithm that are needed to handle more difficult cases.

²There is no meaning in terms of absolute p_{fa} here; one can simply reduce the absolute probability of false alarm by extending the ranges considered to ranges too distant to produce any detections.



(a) Detections, File 273, 1st 100 CPIs: The total detections made by each process, including both false alarms and object detections, is indicated. The latency to PCL's original detection on the object is also indicated.



(b) False alarm rate, File 273, 1st 100 CPIs: The total number of detections shown in 5-10a above is used in conjunction with data labels to determine an overall empirical 1-D false alarm rate for each process in each CPI. Overlaid on each process with a thick dashed line is the mean empirical probability of false alarm for each process, where again HH CFAR is in dark blue, VV CFAR is in dark green, and PCL is in dark red.

Figure 5-10

5.4.1 PCL performance across bandwidths

For this section, focus remains on files containing persistent objects that were collected at approximately the same time as File 273, but that feature higher bandwidth (finer range

resolution) waveforms. As was the case with File 273, the radar beam was focused on the buoy for these files, so the buoy is a very high-SNR object. Because SNR is proportional to the time-bandwidth product of the waveform, the files now under consideration have even higher SNR on the buoy. Hence, these cases are not challenging in the SNR-based detection sense, but demonstrate that PCL filters out false alarms due to sea clutter across a wide range of waveform bandwidths.

Note that each of the plots in this section shows the first 100 CPIs of 100 pulses of data at 6250 Hz PRF. Consequently, each plot corresponds to the first 1.6 seconds of each recorded file. The results discussed in each of the subsections below are summarized in Table 5.1 at the end of this section.

File 272: 40 MHz bandwidth

For 40 MHz bandwidth data, the range resolution is 3.75 m. The data is oversampled in range at twice critical Nyquist as discussed in Section 4.4.3. At this finer oversampled resolution, pulse compression can result in a single range bin offset between the peak response of a point target object across HH and VV polarizations. Hence, the first adjustment made to PCL parameters is increasing δ_n to 1. The result is that any HH and VV detections offset by a single range bin will now be considered potential PCL detections in the next CPI. As shown in Figure 5-11a, the empirical mean p_{fa} of PCL is significantly lower than that of standard CFAR in either HH or VV.

The reader will observe, however, that Figure 5-11b shows that PCL does not maintain continuous detection on the buoy. This is due to variation in the polarimetric signature of objects, which becomes increasingly evident at finer range resolutions. For File 272, as the buoy bobs atop the sea surface and its polarimetric signature varies, HH and VV CFAR detections on the buoy are offset by two range bins in CPI 72, as shown in the sequence of images in Figure 5-12. The images are zoomed in on the buoy's signature across polarizations. At left (CPI 71), the buoy's detections are polarimetrically co-




Figure 5-11: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 1$ and $n_{CPI} = 3$, on File 272 (40 MHz bandwidth) shows that PCL significantly outperforms single polarization detection in terms of false alarm rate. PCL drops the buoy at CPI 72, and it takes $n_{CPI} = 3$ CPIs to recover it.




Figure 5-12: Variation in the buoy's polarimetric signature results in a dropped PCL detection when the buoy detection in HH is offset by two range bins from its VV detection. The PCL CPI count restarts once co-location resumes in CPI 73.

located in HH and VV despite the variation in its polarimetric signature across HH and VV. In the center (CPI 72), the polarimetric signature has changed such that HH and VV detections on the buoy's signature are offset by two range bins; the buoy ceases to be

considered a PCL detection. At right (CPI 73), the variation again becomes less severe; HH and VV detections are offset by only one range bin, and the process of accumulating contiguous detections across CPIs begins anew. The PCL bar in Figure 5-11b shows the ramification of a dropped PCL detection: There is a latency of at least n_{CPI} CPIs before the object has had the requisite number of co-located detections occur and can again be classified as a PCL detection. This is because the core algorithm restarts its CPI criterion count any time a PCL detection is dropped. A modification that addresses this will be introduced in Section 5.4.3.

One means of resolving the polarimetric signature variation issue is to increase δ_n to 2, allowing the peaks of the object's response to be offset by a full range resolution cell. Because this means that any HH and VV detections offset by up to 2 range bins will be considered for polarimetric co-location in subsequent CPIs, more candidate pairs will pass PCL criteria. Consequently, the empirical false alarm rate will increase. Despite the increase, the mean p_{fa} for PCL is still significantly better than that of HH or VV CFAR alone, as shown in Figure 5-13a. The trade-off is that the buoy now passes the PCL criterion in all CPIs. As shown in Figure 5-13b, the PCL detection bar is unbroken.

However, PCL's increased false alarm rate incurred by increasing δ_n can be reduced by increasing n_{CPI} . That is, by requiring that candidate detection pairs be co-located across a greater number of CPIs, the additional spurious detection pairs considered because of a higher δ_n can be filtered out by requiring that all pairs maintain polarimetric co-location over a greater time. Figure 5-14a reflects the order of magnitude decrease in p_{fa} that results from changing n_{CPI} from 3 to 5. It is not generally expected that increasing n_{CPI} will impact object detections. As Figure 5-14b shows, increasing n_{CPI} does not impact detections on the buoy. Nevertheless, there is a trade-off for increasing n_{CPI} : greater latency in both initial declaration of PCL detections and in picking up any PCL object detections that are dropped due to polarimetric signature variation.




Figure 5-13: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 3$, on File 272 (40 MHz bandwidth) shows that PCL still significantly outperforms single polarization detection despite the increased false alarm rate due to incurred by increasing δ_n . PCL no longer drops the buoy.




Figure 5-14: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 272 (40 MHz bandwidth) shows that the increased false alarm rate incurred by increasing δ_n can be mitigated by increasing n_{CPI} . Increasing n_{CPI} comes at the price of longer latency in declaration of PCL detections, as shown.

The reader should now have an appreciation for the two basic PCL parameters and their impacts on empirical false alarm rate and object detections. Bearing these things in mind, attention turns now to even higher bandwidths.

File 271: 150 MHz bandwidth

On higher bandwidth files, the range resolution will of course be finer. The polarimetric signature of an object will therefore usually exhibit at least as much variation as it did with a lower bandwidth waveform. Hence, the last set of PCL parameters, $\delta_n = 2$ and $n_{CPI} = 5$, are preserved as we transition to analysis of File 271, which features the 150 MHz waveform. Figures 5-15a and 5-15b demonstrate that PCL is still significantly more effective than standard single-polarization CFAR detection, even at increasingly higher bandwidths.




Figure 5-15: Performance comparison of standard single-polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 271 (150 MHz bandwidth) demonstrates the superiority of PCL over single polarization CFAR at increasingly higher bandwidths.




Figure 5-16: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, File 270 (400 MHz bandwidth) demonstrates the superiority of PCL over single polarization CFAR at increasingly higher bandwidths.

File 270: 400 MHz bandwidth

The trend continues for the highest bandwidth data in the Point de Chene Dataset. The comparison of PCL with standard-single polarization CFAR for File 270, which features 400 MHz bandwidth buoy data, is shown in Figures 5-16a and 5-16a. The PCL parameters for the data in the figures have again been set to $\delta_n = 2$ and $n_{CPI} = 5$ in keeping with preservation of PCL parameterization from the next lowest bandwidth. PCL remains significantly more effective than standard CFAR in either HH or VV.

For higher bandwidth data, the empirical false alarm rate increases for PCL while it holds relatively stable for single-polarization CFAR. The increase is due to the increase in the raw number of detections produced by the HH and VV CFAR operations. That is, while the raw number of false alarms n_{fa} is increasing in HH and VV CFAR detection, so too is n_{dets} ; hence p_{fa} as computed in Equation 5.4 holds relatively steady. But the number of possible PCL detection pairs increases faster than n_{dets} increases in each channel, so PCL's n_{fa} increases at a higher rate than n_{dets} in HH and VV CFAR. Yet,




Figure 5-17: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 8$, on File 270 (400 MHz bandwidth) demonstrates the superiority of PCL over single polarization CFAR at increasingly higher bandwidths.

Figure 5-16a shows that the PCL’s false alarm rate clearly does not increase fast enough that PCL becomes inferior to single-polarization CFAR detection for bandwidths as high as 400 MHz.

Moreover, as Figure 5-17a shows, PCL p_{fa} can always be further reduced by increasing n_{CPI} as was done previously on File 272 data. The trade is again the longer latency required in order to declare a PCL detection. The longer latency is captured in Figure 5-17b.

Overall PCL performance across bandwidths

PCL performance across bandwidths is quantified for the results discussed above in Table 5.1. The superiority of PCL in terms of empirical false alarm rate is evident; PCL improves upon standard 1-D CFAR operations by at least 1 or 2 orders of magnitude. PCL parameters are denoted in the table’s rightmost columns for the reader’s convenience.

File	β [MHz]	Mean p_{fa}			PCL Params	
		HH CFAR	VV CFAR	PCL	δ_n	n_{CPI}
273	4	.02017	.01328	0	0	3
272	40	.01980	.01437	.00073	1	3
				.00114	2	3
				.00027	2	5
271	150	.01754	.01503	.00056	2	5
270	400	.01801	.01737	.00112	2	5
				.00055	2	8

Table 5.1: Empirical 1-D p_{fa} comparison across detection operations shows the performance improvement gained by using PCL rather than standard single-polarization CFAR. Moreover, the comparison shows that PCL retains its efficacy across a wide range of bandwidths.

5.4.2 PCL performance across PRFs

To this point, PCL results have demonstrated robustness across a range of bandwidths at the highest PRF in the Point de Chene Dataset, 6250 Hz. This section demonstrates that PCL is also robust across a range of PRFs. To show this, the same files considered in the preceding subsection are reconsidered here, but the CPIs are now slow-time decimated with factor $m = 7$, yielding a PRF of ~ 893 Hz. Files 270-273 feature a sea state estimated via the process described in Section 3.3.4 as Douglas sea state 4, so slow-time decimating by $m > 7$ results in aliasing in the Doppler dimension. Hence, PRF ~ 893 Hz represents the lowest bound on PRF.

Note that the CPI length is still fixed to 100 pulses. Decimating with $m = 7$ thus yields a CPI whose length in time is approximately 112 msec. The full 10 seconds recorded for each file comprise 89 such CPIs. The limit on the abscissa for all performance metric plots shown in this section is thus 89, but each plot now shows performance over the full length of the data.

In order to fairly assess PCL performance for this new lower PRF, we fix δ_n and n_{CPI} to the values set in the previous section, with two exceptions. First, we omit the results on File 272 (40 MHz bandwidth) that were shown previously with $\delta_n = 1$ and

$n_{CPI} = 3$ because it was already concluded in the previous section that increasing δ_n to 2 was necessary in order to avoid dropping the buoy detection. Second, we include additional results for File 273 (4 MHz bandwidth) to further demonstrate the reduction in false alarm rate that is achieved by increasing n_{CPI} .

The results for File 273 are shown in Figures 5-18 and 5-19. In both cases, $\delta_n = 0$ as before. In Figure 5-18, $n_{CPI} = 3$, again as before, while in Figure 5-19, $n_{CPI} = 5$. In both cases, PCL significantly outperforms single-polarization CFAR detection in HH and VV, even at this lowest possible PRF. Further, increasing n_{CPI} indeed reduces PCL's false alarm rate as shown.

Results for File 272 (40 MHz bandwidth) are shown in Figures 5-20 and 5-21. Results for File 271 (150 MHz bandwidth) are shown in Figure 5-22. Finally, results for File 273 (4 MHz bandwidth) are shown in Figures 5-23 and 5-24. In all cases, PCL maintains continuous detections on the buoy while achieving significantly lower empirical false alarm rates than those of CFAR in HH or VV alone. Because these results demonstrate that PCL retains its efficacy across all bandwidths at both PRF extremes, we conclude that PCL is indeed robust to variations in PRF.

Overall PCL performance across PRFs

The overall performance of PCL across bandwidths for the ~ 893 Hz PRF is shown in Table 5.2. The PCL parameters used to generate the preceding results are again shown in the rightmost columns for the reader's convenience. As the table evidences, PCL actually performs slightly better at lower PRF than at higher PRF for bandwidths above 4 MHz. This is likely due to the fact that using 100 pulses in the coherent averaging step of PCL amounts to averaging over 112 msec. While sea clutter's transient characteristics usually decorrelate in less than 10 msec at X-band (Chan, 1987; Antipov, 1998), aspects of sea clutter that cohere on longer timescales are likely averaged out by fixing the CPI length over lower PRFs.

File	β [MHz]	Mean p_{fa}			PCL Params	
		HH CFAR	VV CFAR	PCL	δ_n	n_{CPI}
273	4	.02654	.02189	.00387 .00116	0 0	3 5
272	40	.01541	.01143	.00028 .00002	2 2	3 5
271	150	.01775	.01418	.00046	2	5
270	400	.01354	.01476	.00106 .00057	2 2	5 8

Table 5.2: Empirical 1-D p_{fa} comparison at ~ 893 Hz PRF demonstrates that the improvement of PCL over standard single-polarization CFAR is evident even at the dataset’s PRF lower bound.

5.4.3 PCL performance across object types

Now that it has been established that PCL is robust across bandwidths and PRFs for persistent, high-SNR objects, attention turns to the extended dynamic objects of oppor-




Figure 5-18: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 0$ and $n_{CPI} = 3$, on File 273 (4 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates the superiority of PCL over single-polarization CFAR for low-bandwidth, low-PRF data.




Figure 5-19: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 0$ and $n_{CPI} = 5$, on File 273 (4 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates that increasing n_{CPI} further reduces PCL's false alarm rate at the expense of greater initial detection latency.




Figure 5-20: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 3$, on File 272 (40 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates the superiority of PCL over single-polarization CFAR for moderate bandwidth, low-PRF data.




Figure 5-21: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 272 (40 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates that increasing n_{CPI} further reduces PCL's false alarm rate at the expense of greater initial detection latency.




Figure 5-22: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 271 (150 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates the superiority of PCL over single-polarization CFAR for moderately high-bandwidth, low-PRF data.




Figure 5-23: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 5$, on File 270 (400 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates the superiority of PCL over single-polarization CFAR for high bandwidth, low-PRF data.




Figure 5-24: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$ and $n_{CPI} = 8$, on File 270 (400 MHz bandwidth) decimated to ~ 893 Hz PRF demonstrates that increasing n_{CPI} further reduces PCL's false alarm rate at the expense of greater initial detection latency.

tunity in the Point de Chene Dataset. Because these objects are not stationary, their ranges to the radar change as a function of time. Additionally, their view aspects and polarimetric signatures exhibit substantial variations. PCL must be modified to accommodate these variations.

The question we must first address is how to make the notion of polarimetric co-location meaningful for moving objects. To this point, new polarimetrically co-located pairs have had to be in the same location across slow-time as a co-located pair in the previous CPI in order for the two pairs to be associated from one CPI to the next. That is, in order to increment n_{CPI} for a co-located pair in a given range bin, a co-located pair must have been in the same range bin on the previous CPI.

To understand how PCL is adapted for moving objects, first consider the RD image shown in Figure 5-25, on which the approximate maximum velocity of the sea clutter response is indicated by the dashed white line. This maximum Doppler velocity of sea clutter is determined by the sea state. The sea state is in turn determined predominantly by wind speed, as discussed in Section 3.3.4. Thus, it is reasonable to conclude that

$$v_{max} \simeq w_{max}(s) \quad (5.5)$$

where v_{max} [m/sec] is the maximum sea clutter Doppler velocity, w_{max} is the maximum wind speed [m/sec], and s is a unitless variable representing the Douglas sea state.

Once sea state is estimated for the data by one of the methods discussed in Section 3.3.4 and mapped to v_{max} , we can derive a third PCL parameter: maximum target




Figure 5-25: The approximate maximum velocity at which sea clutter and object returns may be coincident is indicated by the dashed line.

displacement, denoted γ_{max} . This unitless parameter is computed as

$$\gamma_{max} = \frac{t_{cpi} v_{max}}{r_{bin}} \quad (5.6)$$

where t_{cpi} [sec] is the CPI time and r_{bin} [m] is the width of each range bin. The role of γ_{max} is to set the maximum number of range bins that a polarimetrically co-located detection pair may be displaced from one CPI to the next.³ That is, n_{CPI} for a given polarimetrically co-located pair will be incremented from one CPI to the next provided that another polarimetrically co-located pair is detected within γ_{max} range bins of the previous pair's location.

The second necessary adaptation of PCL accommodates the polarimetric signature variation of extended objects across HH and VV. In some sense, δ_n already captures this: Increasing δ_n , as discussed in Section 5.4.1, allows for a greater offset between the peaks of detected responses whilst still permitting them to be deemed polarimetrically co-located. However, in certain CPIs, peak detections for moving extended objects can be offset by a number of range bins that is large enough to render the requisite increase in δ_n too significant. In other words, because increasing δ_n is accompanied by an increase in empirical p_{fa} , there is a point at which increasing δ_n cancels out a fair amount of the reduction in empirical false alarm rate that PCL achieves. As previously discussed, this false alarm increase can be mitigated by increasing n_{CPI} in conjunction with the increase in δ_n , but there is also a point at which the latency induced by increasing n_{CPI} becomes undesirable.

Hence, the second adaptation of PCL instead takes the form of a fourth parameter: the propagation parameter, denoted n_{prop} . The propagation parameter is the number of CPIs over which a given PCL detection pair may retain PCL status without an associated PCL detection occurring. For example, consider a PCL detection in CPI 12 located at

³Some objects, like birds, may be moving faster than sea clutter, but we are not concerned with polarimetrically co-locating detections on such objects. Their returns will be in a noise-limited background as opposed to a clutter-limited one, and hence standard CFAR is sufficient for detection without producing a significant number of false alarms in the vicinity of the object.

range bin 200. If $n_{prop} = 3$, then even if that detection has no associated polarimetrically co-located pair in CPI 13, CPI 14, or CPI 15, if a PCL detection pair occurs within γ_{max} of range bin 200 on CPI 16, then that pair will automatically achieve PCL status. In other words, the new pair will not need to start the count toward n_{CPI} from scratch. It will instead inherit the n_{CPI} count from the co-located pair that was in CPI 12. Thus, the polarimetric signature of an object is allowed to vary significantly enough that a co-located pair may not be detected on that object for up to n_{prop} CPIs.

Notice that there is a trade-off in permitting PCL detection propagation in this fashion. Sea clutter false alarms in PCL can spawn other false alarms on sea clutter, provided that the new false alarms occur within $n_{prop} + 1$ CPIs and are within γ_{max} range bins of the earlier sea clutter false alarm. However, as will shortly be seen, PCL still significantly outperforms standard single-polarization CFAR detection.

Focus in the remainder of this section is on three files—Files 267, 325, and 266—comprising the three most challenging bandwidths previously examined. Each of the files features at least one dynamic extended target. The next subsections explore PCL results for each of these cases, with results shown for both high and low PRFs as well as for a variety of PCL parameters. In each case, γ_{max} was computed based on an estimated Douglas sea state of 4 and a range bin width determined by the downsampling process discussed in Section 4.4.3. The final section includes a tabular summary of PCL compared to standard single-polarization CFAR detection performance for all cases whose results are shown below.

File 267: 40 MHz bandwidth, high PRF

File 267 features the 40 MHz waveform and a single object of interest, the white lobster boat shown in Figure 3-10b, at ~ 850 m from Four Eyes and closing radially at ~ 6.6 m/sec. Considering first the higher PRF case where slow-time decimation factor $m = 1$ and thus the PRF is 6250 Hz, Figure 5-26 shows that the two PCL parameters used for PCL




Figure 5-26: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 267 (40 MHz bandwidth) at 6250 Hz PRF provides the baseline for performance without permitting propagation.

detections on persistent point objects are insufficient to produce continuous detections on dynamic extended objects. Note the breaks in both the HH and VV CFAR detection bars, which lead to missed PCL detections. Each time a PCL detection is missed, the count toward n_{CPI} starts over. Increasing δ_n from 2 to 3, as done in Figure 5-27, accounts for some of the polarimetric signature variation of such objects; for the increased cost of a higher p_{fa} , there are fewer breaks in PCL detection, while the remaining breaks in detection are shorter in duration.

Setting PCL parameter n_{prop} to 1, as shown in Figure 5-28, is sufficient to address the remaining detection breaks. That is, allowing for a missed detection in only one CPI mitigates the PCL missed detections incurred because of the inability of HH CFAR and VV CFAR to detect the object simultaneously in a given CPI. Consequently, PCL not only outperforms HH and VV CFARs in terms of p_{fa} , but it maintains continuous detection on the object when both HH CFAR and VV CFAR fail to do so. Note that in conjunction with increasing n_{prop} to 1, we increased n_{CPI} to 6 to mitigate the increase in




Figure 5-27: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 267 (40 MHz bandwidth) at 6250 Hz PRF demonstrates that allowing for a larger δ_n addresses part of the missed PCL detections by allowing for greater polarimetric variation across object signatures in HH and VV.




Figure 5-28: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 6$, and $n_{prop} = 1$ on File 267 (40 MHz bandwidth) at 6250 Hz PRF shows that allowing a single CPI's worth of propagation addresses the remaining PCL missed detections at the cost of a higher p_{fa} .




Figure 5-29: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$ on File 267 (40 MHz bandwidth) at 6250 Hz PRF shows that increasing n_{CPI} reduces p_{fa} at the cost of increased initial latency.

p_{fa} that results from increasing n_{prop} . Moreover, as was done previously on Files 270-273, increasing n_{CPI} even further to 8 reduces PCL's p_{fa} , as shown in Figure 5-29.

File 267: 40 MHz bandwidth, low PRF

Attention now turns to the lower bound PRF version of the same file, wherein slow-time decimation factor $m = 7$ and the resulting PRF is ~ 893 Hz. As shown in Figure 5-30, there is a lower false alarm rate in PCL at the reduced PRF—again, likely due to the longer integration time over sea clutter's transient characteristics. There are also fewer breaks in detection in HH and VV CFAR processes, which leads to fewer breaks in detection in PCL. The reader will note the second break in the PCL detection bar despite the lack of a break at that location in either the HH or VV CFAR detection bars. This is due to the object's polarimetric signature variation in HH and VV. As shown in Figure 5-31, increasing δ_n to 3—as was done in the high PRF case—mitigates this second break in PCL's detection bar.




Figure 5-30: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 267 (40 MHz bandwidth) at ~ 893 Hz PRF provides the baseline for performance at low-PRF without permitting propagation.




Figure 5-31: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 267 (40 MHz bandwidth) at ~ 893 Hz PRF demonstrates that allowing for a larger δ_n addresses part of the missed PCL detections by allowing for greater variation across polarimetric object signatures in HH and VV.




Figure 5-32: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 1$ on File 267 (40 MHz bandwidth) at ~ 893 Hz PRF shows that propagating by one CPI does not address remaining PCL missed detections.




Figure 5-33: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 2$ on File 267 (40 MHz bandwidth) at ~ 893 Hz PRF shows that further increasing n_{prop} addresses the remaining PCL missed detections at the cost of a higher p_{fa} .




Figure 5-34: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 2$ on File 267 (40 MHz bandwidth) at ~ 893 Hz PRF shows that increasing n_{CPI} reduces the false alarm rate incurred because of propagation. In this case, p_{fa} reduces to 0.

Shown in Figure 5-32 are the results achieved when n_{prop} is increased to 1. While the duration of PCL's detection break drops by one CPI, the propagation increase is insufficient to address the full break in PCL detection because the object is dropped by VV CFAR in two CPIs. However, as Figure 5-33 shows, increasing n_{prop} to 2 is sufficient to address the VV CFAR gap, and PCL retains its detections on the object for the duration of the file at the cost of a higher p_{fa} due to propagation of sea clutter false alarms. Also increasing n_{CPI} , as done in Figure 5-34, reduces the increased false alarm rate in exchange for longer initial detection latency. In this case, increasing n_{CPI} to 8 is sufficient to reduce PCL's false alarm rate to 0.

The reader should now have an appreciation for how PCL is adapted to accommodate moving objects. Moreover, the reader should appreciate the trades involved in varying the control parameters used to make this accommodation. Bearing these trades in mind,

we move to discussion of higher bandwidth files featuring dynamic extended objects of interest.

File 325: 150 MHz bandwidth, high PRF

File 325 features the 150 MHz waveform and the “New Englander,” shown in Figure 3-10d, at ~ 1030 m from Four Eyes and closing radially at ~ 2.75 m/sec. The New Englander is the most challenging dynamic extended object observed in that it exhibited the most significant polarimetric signature variation of any of the objects considered here when observed at the aspect angles in this file. As before, we consider first the higher PRF case where slow-time decimation factor $m = 1$ and thus the PRF is 6250 Hz.

Figure 5-35 shows the performance baseline achieved by setting $\delta_n = 2$, $n_{CPI} = 5$, and allowing no propagation, as was done for File 271, which featured the same waveform and the stationary buoy. Numerous detection breaks are evident due to both dropped VV CFAR detections and strong polarimetric signature variation. Setting $\delta_n = 3$ mitigates




Figure 5-35: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 325 (150 MHz bandwidth) provides the baseline for high-PRF performance without permitting propagation.




Figure 5-36: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 325 (150 MHz bandwidth) at 6250 Hz PRF shows that increasing δ_n to 3 addresses many of the missed detections.




Figure 5-37: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 1$ on File 325 (150 MHz bandwidth) at 6250 Hz PRF shows that increasing n_{prop} to 1 mitigates only a portion of the remaining detection breaks but incurs a substantial false alarm rate increase.




Figure 5-38: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 2$ on File 325 (150 MHz bandwidth) at 6250 Hz PRF shows that increasing n_{CPI} to 8 and n_{prop} to 2 addresses all but one dropped detection while mitigating much of the false alarm rate increase incurred by longer propagation.




Figure 5-39: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 3$ on File 325 (150 MHz bandwidth) at 6250 Hz PRF demonstrates that increasing n_{prop} to 3 addresses the remaining detection gap.

many of the detection breaks due to the latter, as shown in Figure 5-36. Figure 5-37 shows that increasing n_{prop} to 1 fails to address the remaining detection breaks while the accompanying increase in false alarm rate is significant. Increasing n_{prop} to 2 while also increasing n_{CPI} to 8, as in Figure 5-38, addresses all but one of the remaining detection breaks. However, it is necessary to increase n_{prop} to 3 in order to yield an unbroken PCL detection line, as shown in Figure 5-39, where n_{CPI} has also been increased to 10 to mitigate the false alarm rate increase caused by longer propagation times. Again, the trade-off in initial detection latency incurred as a result of a higher n_{CPI} is evident at the start of the PCL detection bar.

File 325: 150 MHz bandwidth, low PRF

It has now been shown that, even for dynamic extended objects with significant polarimetric signature variation, PCL outperforms standard single-polarization CFAR detection in the high-PRF case. Attention turns now to the low-PRF version of the same file,




Figure 5-40: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 325 (150 MHz bandwidth) at ~ 893 Hz PRF provides the baseline for low-PRF performance without permitting propagation.




Figure 5-41: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 325 (150 MHz bandwidth) at ~ 893 Hz PRF shows that increasing δ_n to 3 does not mitigate any PCL detection breaks.




Figure 5-42: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$ on File 325 (150 MHz bandwidth) at ~ 893 Hz PRF shows that one CPI of propagation addresses all PCL detection breaks. Concurrently increasing n_{CPI} helps to mitigate the propagation-induced p_{fa} increase.




Figure 5-43: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$ on File 325 (150 MHz bandwidth) at ~ 893 Hz PRF shows that increasing n_{CPI} to 10 further reduces the increase in p_{fa} incurred by propagation.

in which slow-time decimation factor $m = 7$ yields ~ 893 Hz PRF. We again begin by considering the baseline results achieved without allowing propagation.

Figure 5-40 shows this baseline performance. While there are no breaks in HH CFAR detections on the object, missed detections in VV CFAR lead to missed detections in PCL that are longer in duration than VV misses because $n_{CPI} = 5$ and no propagation is allowed. Increasing δ_n to 3 does not address the missed PCL detections, as shown in Figure 5-41. However, as shown in Figure 5-42, allowing one CPI of propagation and increasing n_{CPI} to 8 mitigates all PCL detection breaks while reducing the false alarm rate increase that propagation incurs. The increase in p_{fa} is further reduced in Figure 5-43, where n_{CPI} has been increased to 10.

So far, PCL retains its superiority over standard single-polarization CFAR detection in 1-D for all dynamic extended targets considered across both 40 MHz and 150 MHz

bandwidths at both ends of the dataset’s PRF spectrum. We turn attention now to the highest bandwidth case.

File 266: 400 MHz bandwidth, high PRF

File 266 features the 400 MHz waveform and three objects of interest: the white lobster boat at ~ 1025 m from Four Eyes, closing radially at ~ 6.6 m/sec; the “New Englander” at ~ 1440 m from Four Eyes, closing radially at ~ 1 m/sec; and the buoy at ~ 1711 m from Four Eyes, stationary in range but exhibiting Doppler variation in $[-1,1]$ m/sec. The buoy response in File 267 is in the radar’s sidelobes; thus, its response is much lower SNR than was true for previous buoy files. The reader will note that in all object detection plots in this section, ordinate axis labels for the objects have been abbreviated for the sake of legibility. Specifically, the white lobster boat is abbreviated as “W.L.B.”; the New Englander is abbreviated as “N.E.”; and the buoy is indicated as was done previously, though the reader should remember that the buoy has very low SNR.




Figure 5-44: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 266 (400 MHz bandwidth) at 6250 Hz PRF provides the baseline for high-PRF performance without permitting propagation.




Figure 5-45: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 266 (400 MHz bandwidth) at 6250 Hz PRF shows that increasing δ_n to 3 mitigates some of the PCL detection breaks on the W.L.B. and the N.E.

Figure 5-44 provides the PCL performance baseline using the PCL parameters $\delta_n = 2$ and $n_{CPI} = 5$ without allowing propagation. As shown, there are several dropped PCL detections on all objects. In the cases of the white lobster boat and the New Englander, because HH and VV CFAR both maintain continuous object detections, all missed PCL detections are due to polarimetric signature variation. Thus, many of the dropped PCL detections on these objects are mitigated by increasing δ_n to 3, as shown in Figure 5-45, where the accompanying increase in p_{fa} is also evident. The remaining missed PCL detections on the white lobster boat and the New Englander are mitigated by allowing propagation for one CPI, as shown in Figure 5-46. In Figure 5-46, n_{CPI} has also been increased to 8. However, the increase in p_{fa} incurred due to propagation is not reduced to its earlier level unless n_{CPI} is increased to 10, as shown in Figure 5-47.

These figures all show that while PCL significantly outperforms HH and VV CFAR for the white lobster boat and the New Englander in terms of both p_{fa} and continuity of




Figure 5-46: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$ on File 266 (400 MHz bandwidth) at 6250 Hz PRF shows that allowing one CPI of propagation mitigates the remaining detection breaks on the extended objects, but p_{fa} is elevated despite the accompanying increase in n_{CPI} .




Figure 5-47: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$ on File 266 (400 MHz bandwidth) at 6250 Hz PRF shows that increasing n_{CPI} to 10 cancels out most of the p_{fa} increase due to propagation.

object detection, the same is not true for the very low-SNR buoy. There are modifications that can be made to PCL to accommodate low-SNR objects, but modifications will not be addressed in this dissertation.

File 266: 400 MHz bandwidth, low PRF

Finally, attention turns to the low-PRF version of File 266, in which slow-time decimation factor $m = 7$ yields ~ 893 Hz PRF. As shown in Figure 5-48, PCL has a strong performance at low PRF on both the white lobster boat and the New Englander, achieving a substantial reduction in p_{fa} . The two breaks in detection on the white lobster boat are due to polarimetric signature variation. The variation is not accommodated by increasing δ_n to 3, as shown in Figure 5-49. However, as shown in Figure 5-50, increasing n_{prop} to 1 is sufficient to mitigate PCL's dropped detections; in this case, n_{CPI} was concurrently increased to 8 to mitigate the resultant increase in p_{fa} , as was done in the high PRF case. Figure 5-51 shows the reduction in p_{fa} achieved by further increasing n_{CPI}




Figure 5-48: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 2$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 266 (400 MHz bandwidth) at ~ 893 Hz PRF provides a performance baseline in the high-PRF case.




Figure 5-49: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 5$, and $n_{prop} = 0$ on File 266 (400 MHz bandwidth) at ~ 893 Hz PRF shows that increasing δ_n to 3 does not address the dropped PCL detections.




Figure 5-50: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 8$, and $n_{prop} = 1$ on File 266 (400 MHz bandwidth) at ~ 893 Hz PRF demonstrates that allowing one CPI of propagation is sufficient to address dropped PCL detections.




Figure 5-51: Performance comparison of standard single polarization CFAR with PCL, $\delta_n = 3$, $n_{CPI} = 10$, and $n_{prop} = 1$ on File 266 (400 MHz bandwidth) at ~ 893 Hz PRF shows the reduction in p_{fa} achieved by increasing n_{CPI} to 10.

to 10. Again, HH and VV CFAR detection operations both outperform PCL in terms of detection continuity on the very low-SNR buoy, as evidenced by all of the preceding figures.

Overall PCL performance across object types

The overall performance in terms of p_{fa} for PCL across bandwidths and PRFs for dynamic extended objects is shown in Table 5.3. PCL parameters used to generate the results are again shown in the rightmost columns for the reader's convenience. As the table shows, PCL again performs slightly better at low PRF than at high PRF. The improvement of PCL over that of standard single-polarization CFAR is evident: PCL's p_{fa} is consistently two orders of magnitude better than the p_{fa} of HH and VV CFAR operations.

File	β [MHz]	PRF [Hz]	Mean p_{fa}			PCL Params		
			HH CFAR	VV CFAR	PCL	δ_n	nCPI	n _{prop}
267	40	6250	.02186	.01691	.00044	2	5	0
					.00054	3	5	0
					.00104	3	6	1
					.00056	3	8	1
		~ 893	.01263	.01040	.00013	2	5	0
					.00014	3	5	0
					.00038	3	5	1
					.00115	3	5	2
					.00000	3	8	2
325	150	6250	.01949	.01635	.00105	2	5	0
					.00164	3	5	0
					.00361	3	5	1
					.00419	3	8	2
					.00392	3	10	3
		~ 893	.01463	.01182	.00054	2	5	0
					.00099	3	5	0
					.00095	3	8	1
					.00067	3	10	1
266	400	6250	.01467	.01433	.00095	2	5	0
					.00147	3	5	0
					.00194	3	8	1
					.00156	3	10	1
		~ 893	.00956	.00963	.00065	2	5	0
					.00089	3	5	0
					.00069	3	8	1
					.00040	3	10	1

Table 5.3: Empirical 1-D p_{fa} comparison demonstrates the significant improvement of PCL over standard single-polarization CFAR across object types, bandwidths, and PRFs.

5.5 Integrating PCL into the standard radar signal processing chain

The reader will recall that in addition to robustness across bandwidths, PRFs, and object types, it was highly desirable that the algorithm also be feasible within the bounds of standard radar signal processing chain structure. Thus, neither excessive computational complexity nor restructuring of the processing chain is desirable. The question thus be-




Figure 5-52: A diagram of the radar signal processing chain shows the key computational steps that lie between data digitization and passing detections to the tracking/parameter estimation process: namely, pulse compression, Doppler processing, and 2-D CFAR detection.

comes how PCL may be made functional in terms of the standard radar signal processing chain, which is reproduced for the reader’s convenience in Figure 5-52.

This is the point at which the notion of *layering* comes into play; the ultimate word in PCL’s name derives from its method of incorporation into the standard processing chain. The reader will recall that in executing PCL, no unusual computational processes were required. Only coherent integration on already pulse-compressed data followed by 1-D CFAR detection operations in HH and VV channels are required for PCL to execute. Hence, we can “layer” PCL into the standard radar signal processing chain as a parallel process, as shown in Figure 5-53.

5.5.1 PCL detection filtering

The output of PCL in this structure and for each CPI is a set of range bins containing PCL detections. PCL detection filtering simply filters out 2-D CFAR detections (produced by the standard radar signal processing chain) whose ranges are not included in PCL range bin set. To see how this works, consider a 2-D CFAR that produces detections in all




Figure 5-53: A diagram of PCL as integrated as a parallel layer in the standard radar signal processing chain shows that PCL requires neither unavailable forms of data nor restructuring of the processing chain’s major computational blocks. Moreover, PCL is implemented as a parallel process that executes while Doppler processing and 2-D CFAR detection are performing computations. The only new block included serially in the chain is PCL filtering of 2-D CFAR detections.

positions indicated with a 1 in the notional binary matrix

$$\mathbf{D}_{cfar} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix} \quad (5.7)$$

where the row corresponds to the Doppler bin and the column corresponds to the range bin. A 1 in the matrix on the right-hand side of Equation 5.7 is analogous to 2-D CFAR producing a black circle indicating a detection on the CFAR detection RD images shown in Figure 5-53, though the size of the matrix here is significantly reduced to illustrate the concept of PCL filtering. Now consider that PCL, running in parallel, has determined

that only detections in range bins 5 and 9 are on objects, while the rest are on sea clutter. Following PCL filtering, the binary matrix in Equation 5.7 will become

$$\mathbf{D}_{pcl} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$




Figure 5-54: 2-D CA-CFAR detections on RD data for File 325, which contains The New Englander at the location and Doppler indicated on the figures, are shown for HH (detections indicated by black circles) and VV (detections indicated by red circles) for a single CPI. Hundreds of false alarms are evident in each case, along with a single detection on the object of interest.

To illustrate this with actual radar data, consider the 2-D CFAR detections produced for a single CPI of File 325 in the HH and VV channels shown in Figure 5-54. The figures show that there are hundreds of false alarms resulting from the 2-D CA-CFAR operations run on each channel's data. In each case, the object of interest—The New Englander, at the range and Doppler velocity indicated on the figures—is the source of a single detection. Following PCL filtering, detections in all but four range bins are




Figure 5-55: 2-D PCL detections on RD data for File 325, which contains The New Englander at the location and Doppler indicated on the figures, are shown for HH and VV for a single CPI. PCL detections are those that remain following PCL filtering of 2-D CA-CFAR detections by range bin. The decrease in false alarm rate is evident, as is retention of the detection on The New Englander.

removed by the filter, leaving the detections shown in Figure 5-55. Hundreds of false alarms are eliminated by the PCL filter, while the detections on the object of interest are preserved. As was the case with 1-D detection operations, PCL routinely achieves two orders of magnitude decrease in false alarm rate compared to standard 2-D CFAR on a single polarimetric channel.

5.5.2 Computational complexity

To integrate PCL in parallel without incurring additional computation time, the computational complexity of PCL should be at least as small as the combined complexity of the Doppler processing and 2-D CFAR operations with which PCL runs in parallel. We consider first the complexity of Doppler processing.

As previously discussed, Doppler processing amounts to taking the fast Fourier transform (FFT) at each range bin across slow-time. It is well-known that the computational complexity of the FFT is $\mathcal{O}(m \log m)$, where m is the length of the discrete signal be-

ing transformed. Therefore, the computational complexity of Doppler processing is each channel is $\mathcal{O}(np \log p)$, where n is the number of range bins being processed and p is the number of pulses in the CPI, because we must compute n FFTs each of length p .

In the time domain, CFAR is the convolution of the CFAR sampling window with the fast-time by slow-time data matrix on which detection is being performed. Convolution is usually most efficiently implemented as multiplication in the frequency domain. The 2-D FFT of the CFAR sampling window can be computed and stored offline; hence, CFAR window computation does not incur a time penalty. However, range-Doppler data must also be transformed to the frequency domain on each CPI. It is well-known that for an $n \times p$ data matrix, the 2-D FFT has computational complexity $\mathcal{O}(np \log p + pn \log n) = \mathcal{O}(np(\log p + \log n))$, because the FFT must be computed across each row and each column. In our case, n and p are again the numbers of range bins and pulses, respectively, in the CPI. The CFAR thresholding process then requires np comparisons. Moreover, if the CFAR uses the peak-picking process described in Section 5.2, an additional $4np$ comparisons are required. Hence the total 2-D CFAR detection process has computational complexity $\mathcal{O}(np(\log p + \log n + 5))$.

We thus have that the combined computational complexity of Doppler processing and 2-D CFAR detection is

$$\begin{aligned}
\mathcal{O}(np \log p) + \mathcal{O}(np(\log p + \log n + 5)) &= \mathcal{O}(np(2 \log p + \log n + 5)) \\
&= \mathcal{O}(np(2 \log p + \log n)) \quad (5.8) \\
&= \mathcal{O}(np(\log p + \log n))
\end{aligned}$$

If PCL has computational complexity less than or equal to the complexity shown in Equation 5.8, it is computationally efficient in the sense that it operates as fast or faster than the operations with which it runs in parallel. PCL requires two coherent averages over n range bins. Averaging requires p complex additions at each of n range bins, followed by one complex division. Across two channels, this process thus has complexity

$\mathcal{O}(2(2np + 2)) = \mathcal{O}(2np) = \mathcal{O}(np)$. PCL also requires 1-D CFAR detection in each channel. One-dimensional CFAR detection can be implemented as described for the 2-D CFAR detection process, but because it is 1-D, the overall computational complexity of both processes is $\mathcal{O}(2(n \log n + 3n)) = \mathcal{O}(n(2 \log n + 6)) = \mathcal{O}(n(\log n + 3))$, where the $3n$ is incurred because of comparison of each cell under test to the threshold as well as to its two neighboring values in the peak-picking operation.

The finer details of PCL's implementation are not given in this dissertation. However, PCL is presently implemented such that each CPI also requires: n real-valued additions and subtractions; $\mathcal{O}(2d_{hh}\delta_n)$ comparisons, where d_{hh} is the number of CFAR detections in HH; and $\mathcal{O}(2d_{pcl}\gamma_{max})$ comparisons, where d_{pcl} is the number of co-located detection pairs in the CPI. Consequently, PCL's overall computational complexity is

$$\begin{aligned} \mathcal{O}(np + n(\log n + 3) + n + 2d_{pcl}\gamma_{max} + 2d_{hh}\delta_n) = \\ \mathcal{O}(n(p + \log n + 4) + d_{pcl}\gamma_{max} + d_{hh}\delta_n). \end{aligned}$$

For any given PCL instantiation, γ_{max} and δ_n are both constants. Thus, the right-hand side of the preceding equation simplifies to

$$\mathcal{O}(n(p + \log n + 4) + d_{pcl} + d_{hh}).$$

Because neither d_{pcl} nor d_{hh} can exceed n , we can write that

$$\begin{aligned} \mathcal{O}(n(p + \log n + 4) + d_{pcl} + d_{hh}) \leq \\ \mathcal{O}(n(p + \log n + 4) + n + n) = \mathcal{O}(n(p + \log n + 6)). \quad (5.9) \end{aligned}$$

Comparing Equations 5.8 and 5.9, we see that when $n > 64$, the computational complexity of PCL is lower than the computational complexity of Doppler processing and 2-D CFAR detection unless $p = 1$. That is, PCL is more efficient than combined Doppler processing and 2-D CFAR detection operations unless Doppler processing is not

possible (because the CPI contains only one pulse) and thus 2-D CFAR detection in the radar signal processing chain is reduced to 1-D CFAR detection. In this degenerate case, the additional operations that PCL requires beyond its own 1-D CFAR processes result in a sequence of operations that will slow down the radar's overall processing time. Therefore, provided each CPI consists of more than one pulse, the only computational time penalty incurred by using PCL is the negligible amount of time required by the PCL detection filter.

5.6 Summary

We began discussion of PCL by first reviewing the extensive literature that has emphasized the mean Doppler differential of sea clutter in HH and VV. This was followed by a discussion of techniques designed to enhance detection performance in low grazing angle sea clutter by jointly leveraging information in both the Doppler and polarization dimensions. The fundamental principle behind PCL is that while sea clutter returns generally exhibit different mean Doppler signatures in HH and VV, objects will not. PCL thus executes a series of steps to determine which detection pairs produced by standard 1-D CFAR processes in HH and VV retain polarimetric co-location on the average over time. The algorithm deems such detection pairs potential object detections and filters out all others.

Extensive results shown in this chapter demonstrate that PCL is robust across bandwidths, PRFs, and object types, regularly achieving a reduction of two orders of magnitude in 1-D empirical false alarm rate over that of standard CFAR while maintaining continuous detections on objects. PCL is also practical. It can be layered in parallel into the standard radar signal processing chain without introducing undue computational complexity or requiring restructuring of the chain. For CPIs comprising more than one pulse and at least 64 range bins, PCL is more efficient than the standard radar processes with which it runs in parallel.

As a final point, the keen reader will have already noted that while PCL was designed using a fully polarimetric radar, only HH and VV measurements are actually required for the algorithm to function. Thus, a compact polarimetric radar—that is, one that can measure both HH and VV co-polarizations, regardless of its ability to accurately recover cross-polarized signatures—can implement PCL. It is also likely that PCL will be successful even without polarimetric calibration. Even though Four Eyes is polarimetrically calibrated (using the method described in Section 4.3), having antennas that are H- and V-polarized should provide measurements that are “good enough” for PCL, regardless of any reasonable distortions induced on transmit and receive.

Chapter 6

Summary and Future Work

The aim of this dissertation was to demonstrate that polarimetric dimensionality can be an effective means of mitigating low grazing angle sea clutter. Implicit in this demonstration were two major contributions: 1) collection of a low grazing angle sea clutter dataset using a fully polarimetric X-band radar assembled largely from COTS components and 2) development of a practical algorithm capable of leveraging polarimetric dimensionality to mitigate low grazing angle sea clutter. Both of these aims were achieved. While this dissertation made multiple contributions to research in the field, there are many paths forward for additional work.

6.1 Summary

The Point de Chene Dataset collected in October 2015 comprises a large dataset featuring multiple sea states and look directions with respect to the wind in a very low grazing angle geometry. The dataset comprises bandwidths ranging from 4 to 400 MHz, a high PRF that can be decimated in post-processing to form low-PRF data, and multiple objects including both stationary and dynamic extended objects of opportunity. Moreover, the dataset features three polarimetric transmit schemes that included the usual alternating polarimetric scattering matrix measurement scheme (ALT) as well as two simultaneous scattering matrix measurement schemes (SIM and SIM-PHS). While neither scheme proved a successful means of simultaneous scattering matrix measurement, their use led to a simultaneous measurement scheme that leverages Doppler division multiple access waveforms. A derivation showing that this scheme, which we will abbreviate

SIM-DDMA, is an effective means of simultaneous scattering matrix measurement was presented in Section 3.2; this derivation is one of the contributions of this dissertation.

A signal processing suite for Four Eyes' data was developed in MATLAB to accommodate the various polarimetric transmit schemes and to transverse equalize and polarimetrically calibrate Four Eyes' data, as discussed in Chapter 4. The polarimetric calibration scheme utilized is another contribution of this dissertation, as it leveraged a single active target with equal co- and cross-polarized scattering matrix signatures to achieve polarimetric calibration. The active target was a delay line with transmit and receive antennas both oriented in $\mathbf{X}^+$ slant linear configuration with respect to Four Eyes' dual-polarized ($\mathbf{H}$ and $\mathbf{V}$) transmit and receive antennas. To the author's knowledge, this work represents the first time that polarimetric calibration has been achieved using a single target in the field.

The third and most significant contribution of this dissertation is Polarimetric Collocation Layering, an algorithm that mitigates the impact of low grazing angle sea clutter on standard CFAR detection processes. Results demonstrated that PCL is robust across bandwidths, PRFs, and object types, regularly achieving a reduction of two orders of magnitude in 1-D empirical false alarm rate over that of standard CFAR while maintaining continuous detections on objects. PCL is also practical. It can be layered in parallel into the standard radar signal processing chain without introducing undue computational complexity or requiring restructuring of the chain. The computational complexity of PCL was derived, and proved to be lower than the computational complexity of the standard radar processes with which PCL runs in parallel.

6.2 Future Work

The Point de Chene Dataset collection, processing suite, and data labeling mechanism has opened the door for numerous avenues of study that may contribute to development of cognitive radar in the future. Chief among these is exploration of machine learning

techniques applied over polarimetric features to the problem of detection in low grazing angle sea clutter. This work is presently in process.

In addition, while we have shown that SIM-DDMA will be an effective means of simultaneously measuring scattering matrix components, an additional data collection will be required in order to demonstrate SIM-DDMA in practice. Such data will enable an understanding of what may be gained by instantaneous SM measurement (other than the obvious, which is effective radiated power) in terms of scene understanding.

Finally, while we have demonstrated that PCL is effective for a maritime radar in a littoral environment, it is not yet understood whether having a non-stationary radar will impact PCL performance. Intuitively, we anticipate that PCL will retain its efficacy, but cannot demonstrate this is so without additional data collections. Moreover, because the sea clutter mean Doppler differential is more pronounced at low grazing angles, further study using additional data is needed in order to understand the upper bound on grazing angle at which PCL performance degrades to the level of standard CFAR performance.

References

- Allan, N. (1995). *Remote sensing the ocean surface at low grazing angles using an X-band polarimetric radar*. PhD thesis, Northeastern University.
- Antipov, I. (1998). Simulation of sea clutter returns. Technical Report DSTO-TR-0679, Defense Science and Technology Organisation.
- Balanis, C. A. (2012). *Advanced Engineering Electromagnetics*. John Wiley & Sons, 2nd edition.
- Blanchard, A. J. and Newton, R. C. (1985). Demands on polarization purity in the measurement and imaging of distributed clutter. In *Inverse Methods in Electromagnetic Imaging*, pages 721–737. Springer.
- Blejer, D. J. (2016a). Scene correlation on alternating scattering matrix measurement timescales. Personal communications at MIT Lincoln Laboratory.
- Blejer, D. J. (2016b). Verification of polarimetric calibration quality. Personal communications at MIT Lincoln Laboratory.
- Boerner, W.-M. (2007). Basics of SAR polarimetry I. In *Radar Polarimetry and Interferometry*, pages 3–1–3–40. Neilly-sur-Seine, France: RTO.
- Boers, Y. and Driessen, J. (2001). Particle filter based detection for tracking. In *Proceedings of the 2001 American Control Conference*, pages 4393–4397. American Automatic Control Council.
- Bringi, V. N. and Chandrasekar, V. (2001). *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press.
- Carretero-Moya, J., Gismero-Menoyo, J., Asensio-Lopez, A., and Blanco-del Campo, A. (2009). Application of the radon transform to detect small-targets in sea clutter. *IET Radar, Sonar & Navigation*, 3(2):155–166.
- Chan, H. (1987). Multi-frequency measurement of radar sea clutter at low grazing angles. Technical report, Communications Research Centre, Dept. of Communications, Radar and Communications Technology Branch Canada.
- Chan, H. (1990). Radar sea-clutter at low grazing angles. In *IEE Proceedings*, volume 137:2, pages 102–112. IET.

- Cloude, S. R. and Pottier, E. (1996). A review of target decomposition theorems in radar polarimetry. *IEEE Transactions on Geoscience and Remote Sensing*, 34(2):498–518.
- Currie, B. W., Haykin, S., and Krasnor, C. (1990). Time-varying spectra for dual-polarized radar returns from targets in an ocean environment [X-band radar system]. In *Record of the IEEE 1990 International Radar Conference*, pages 365–369. IEEE.
- Davey, S. J., Rutten, M. G., and Cheung, B. (2008). A comparison of detection performance for several track-before-detect algorithms. *EURASIP Journal on Advances in Signal Processing*, 2008:41.
- Davidson, G. and Griffiths, H. (2002). Wavelet detection scheme for small targets in sea clutter. *Electronics letters*, 38(19):1.
- De Maio, A. and Alfano, G. (2003). Polarimetric adaptive detection in non-Gaussian noise. *Signal Processing*, 83(2):297–306.
- Deschamps, G. (1949). Geometrical representation of the polarization of a plane electromagnetic wave. *Proceedings of the Institute of Radio Engineers*, 37(2):172–172.
- El-Darymli, K., McGuire, P., Gill, E., Power, D., and Moloney, C. (2014). Understanding the significance of radiometric calibration for synthetic aperture radar imagery. In *2014 IEEE 27th Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–6. IEEE.
- Freeman, A. (1992). SAR calibration: An overview. *IEEE Transactions on Geoscience and Remote Sensing*, 30(6):1107–1121.
- Freeman, A., Shen, Y., and Werner, C. L. (1990). Polarimetric SAR calibration experiment using active radar calibrators. *IEEE Transactions on Geoscience and Remote Sensing*, 28(2):224–240.
- Giuli, D., Fossi, M., and Facheris, L. (1993). Radar target scattering matrix measurement through orthogonal signals. In *IET Proceedings of Radar and Signal Processing*, volume 140:4, pages 233–242. IET.
- Greco, M., Gini, F., Younsi, A., Rangaswamy, M., and Zoubir, A. (2008). Non-stationary sea clutter: impact on disturbance covariance matrix estimate and detector CFAR. In *2008 IEEE International Conference on Radar*, pages 558–562. IEEE.
- Greco, M., Stinco, P., and Gini, F. (2010). Identification and analysis of sea radar clutter spikes. *IET Radar, Sonar & Navigation*, 4(2):239–250.
- Gregers-Hansen, V. and Mital, R. (2009). An empirical sea clutter model for low grazing angles. In *2009 IEEE International Conference on Radar*, pages 1–5. IEEE.

- Guan, J., Chen, X.-L., Huang, Y.-P., and He, Y. (2012). Adaptive fractional Fourier transform-based detection algorithm for moving target in heavy sea clutter. *IET Radar, Sonar & Navigation*, 6(5):389–401.
- Haykin, S. (1990). Radar vision. In *Record of the IEEE 1990 International Radar Conference*, pages 585–588. IEEE.
- Haykin, S. (2006). Cognitive radar: a way of the future. *IEEE Signal Processing Magazine*, 23(1):30–40.
- Haykin, S., Bakker, R., and Currie, B. (2001). The McMaster IPIX radar sea clutter database - Dartmouth 1993. <http://soma.ece.mcmaster.ca/ipix/dartmouth/index.html>. Accessed: 2015-01-28.
- Haykin, S., Currie, B., Lewis, E., and Nickerson, K. (1985). Surface-based radar imaging of sea ice. *Proceedings of the IEEE*, 73(2):233–251.
- Haykin, S., Krasnor, C., Nohara, T. J., Currie, B. W., and Hamburger, D. (1991). A coherent dual-polarized radar for studying the ocean environment. *IEEE Transactions on Geoscience and Remote Sensing*, 29(1):189–191.
- Haykin, S., Lewis, E. O., Raney, R. K., and Rossiter, J. R. (1994). *Remote Sensing of Sea Ice and Icebergs*. John Wiley & Sons.
- Hu, J., Tung, W.-W., and Gao, J. (2006). Detection of low observable targets within sea clutter by structure function based multifractal analysis. *IEEE Transactions on Antennas and Propagation*, 54(1):136–143.
- Hurtado, M. and Nehorai, A. (2008). Polarimetric detection of targets in heavy inhomogeneous clutter. *IEEE Transactions on Signal Processing*, 56(4):1349–1361.
- Husson, D., Pointin, Y., and Ramond, D. (1989). Discrimination between hail and rain precipitation types from dual polarization radar, raingauge and hailpad data. *Theoretical and Applied Climatology*, 40(4):201–207.
- Huynen, J. R. (1965). Measurement of the target scattering matrix. *Proceedings of the IEEE*, 53(8):936–946.
- Huynen, J. R. (1970). *Phenomenological theory of radar targets*. PhD thesis, TU Delft, Delft University of Technology.
- Ioannidis, G. A. and Hammers, D. E. (1979). Optimum antenna polarizations for target discrimination in clutter. *IEEE Transactions on Antennas and Propagation*, 27:357–363.
- Islam, T., Rico-Ramirez, M. A., Han, D., and Srivastava, P. K. (2012). Artificial intelligence techniques for clutter identification with polarimetric radar signatures. *Atmospheric Research*, 109:95–113.

- Jones, G. and Haykin, S. (1996). Higher-dimension detection of targets from scanned radar. *IEE Proceedings-Radar, Sonar and Navigation*, 143(5):313–320.
- Kantor, J. and Davis, S. (2011). Airborne MIMO GMTI radar. Technical Report 1150, MIT Lincoln Laboratory.
- Kennaugh, E. (1952). Polarization properties of radar reflections. Master’s thesis, The Ohio State University.
- Kim, J. E., Lee, S. M., Lee, S.-P., Kim, S., Kim, Y.-S., and Kim, C. H. (2014). Target detection using radar in heavy sea clutter by polarimetric analysis and neural network. In *Advances in Neural Networks-ISNN 2014*, pages 479–488. Springer.
- Krieger, G., Younis, M., Huber, S., Bordoni, F., Patyuchenko, A., Kim, J., Laskowski, P., Villano, M., Romme, T., Lopez-Dekker, P., et al. (2012). Mimo-sar and the orthogonality confusion. In *2012 IEEE International Geoscience and Remote Sensing Symposium*, pages 1533–1536. IEEE.
- Lamont-Smith, T. (2000). Doppler spectra of laboratory wind waves at low grazing angle. *Waves Random Media*, 10(1):33–41.
- Lee, J.-S. and Pottier, E. (2009). *Polarimetric Radar Imaging: From Basics to Applications*. CRC Press.
- Li, D.-C. and Shui, P.-L. (2016). Floating small target detection in sea clutter via normalised doppler power spectrum. *IET Radar, Sonar & Navigation*, 10(4):699–706.
- Long, M. W. (2001). *Radar Reflectivity of Land and Sea*. Artech House, 3rd edition.
- Mather, P. and Tso, B. (2009). *Classification Methods for Remotely Sensed Data*. CRC Press.
- McLaughlin, D. J., Allan, N., Twarog, E. M., and Trizna, D. (1995). High resolution polarimetric radar scattering measurements of low grazing angle sea clutter. *IEEE Journal of Oceanic Engineering*, 20(3):166–178.
- Mecca, V. F. and Krolik, J. L. (2007). Slow-time MIMO STAP with improved power efficiency. In *2007 Conference Record of the Forty-First Asilomar Conference on Signals, Systems and Computers*, pages 202–206. IEEE.
- Mecca, V. F., Ramakrishnan, D., and Krolik, J. L. (2006). MIMO radar space-time adaptive processing for multipath clutter mitigation. In *Fourth IEEE Workshop on Sensor Array and Multichannel Processing*, pages 249–253. IEEE.
- Mooradd, D. C. (2016). Antenna cross-polarization isolation requirements. Personal communications at MIT Lincoln Laboratory.

- Nathanson, F. E., Reilly, J. P., and Cohen, M. N. (1991). *Radar Design Principles: Signal processing and the Environment*. McGraw-Hill, 2nd edition.
- National Oceanic and Atmospheric Administration National Data Buoy Center (2015). Buoy Data, Station 44098 - Jeffrey's Ledge, NH (160). http://www.ndbc.noaa.gov/station_page.php?station=44098. Data provided by the Scripps Institution of Oceanography. Accessed 2015-11-4.
- Nord, M. E., Ainsworth, T. L., Lee, J.-S., and Stacy, N. J. (2009). Comparison of compact polarimetric synthetic aperture radar modes. *IEEE Transactions on Geoscience and Remote Sensing*, 47(1):174–188.
- Novak, L. M., Burl, M. C., and Irving, W. (1993). Optimal polarimetric processing for enhanced target detection. *IEEE Transactions on Aerospace and Electronic Systems*, 29(1):234–244.
- Novak, L. M., Sechtin, M. B., and Cardullo, M. J. (1989). Studies of target detection algorithms that use polarimetric radar data. *IEEE Transactions on Aerospace and Electronic Systems*, 25(2):150–165.
- Panagopoulos, S. and Soraghan, J. J. (2004). Small-target detection in sea clutter. *IEEE Transactions on Geoscience and Remote Sensing*, 42(7):1355–1361.
- Park, H.-R. and Wang, H. (2006). Adaptive polarisation-space-time domain radar target detection in inhomogeneous clutter environments. *IEE Proceedings on Radar, Sonar and Navigation*, 153(1):35–43.
- Pidgeon, V. (1968). Doppler dependence of radar sea return. *Journal of Geophysical Research*, 73(4):1333–1341.
- Poelman, A. (1977). Reconsideration of the target detection criterion based on adaptive antenna polarizations. In *In AGARD New Devices, Tech. and Systems in Radar (SEE N77-22346 13-32)*, volume 1. AGARD.
- Poelman, A. (1981). Virtual polarisation adaptation: A method of increasing the detection capability of a radar system through polarisation-vector processing. *IEE Proceedings of Communications, Radar and Signal Processing*, 128(5):261–270.
- Raynal, A. M. and Doerry, A. W. (2010). Doppler characteristics of sea clutter. Technical Report SAND2010-3828, Sandia National Laboratories.
- Reilly, J. and Dockery, G. (1990). Influence of evaporation ducts on radar sea return. In *IEE Proceedings of Radar and Signal Processing*, volume 137:2, pages 80–88. IET.
- Richards, M. A. (2014). *Fundamentals of Radar Signal Processing*. IET, 2nd edition.
- Richards, M. A., Scheer, J. A., Holm, W. A., and Melvin, W. L., editors (2010). *Principles of Modern Radar*. SciTech Publishing.

- Rohling, H. (2006). Some radar topics: waveform design, range cfar and target recognition. In *Advances in Sensing with Security Applications*, pages 293–322. Springer.
- Root, L., editor (1982). *Proceedings, Second Workshop on Polarimetric Radar Technology*. GACIAC, IIT-RI.
- Sarabandi, K., Ulaby, F. T., and Tassoudji, M. A. (1990). Calibration of polarimetric radar systems with good polarization isolation. *IEEE Transactions on Geoscience and Remote Sensing*, 28(1):70–75.
- Sinclair, G. (1950). The transmission and reception of elliptically polarized waves. *Proceedings of the Institute of Radio Engineers*, 38(2):148–151.
- Sittrop, H. (1977). On the sea-clutter dependency on windspeed. In *IEE Conference Proceedings on Radar*, volume 1, pages 110–114. IET.
- Skolnik, M. I. (2001). *Introduction to RADAR Systems*. McGraw-Hill, NY, 3rd edition.
- Smith, A., Gelsema, S., Kester, L., Melief, H., Cabic, G., Theil, A., and Woudenberg, E. (2002). Maritime target and sea clutter measurements with a coherent doppler polarimetric surveillance radar. In *Proceedings of the IEEE Radar Conference, 2002*, pages 326–330. IEEE.
- Smith, M., Poulter, E., and McGregor, J. (1996). Doppler radar measurements of wave groups and breaking waves. *Journal of Geophysical Research: Oceans*, 101(C6):14269–14282.
- Stimson, G. W., Griffiths, H. D., Baker, C. J., and Adamy, D. (2014). *Introduction to Airborne Radar*. SciTech Publishing, 3rd edition.
- Tonissen, S. M. and Bar-Shalom, Y. (1998). Maximum likelihood track-before-detect with fluctuating target amplitude. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3):796–809.
- Valenzuela, G. and Laing, M. (1970). Study of Doppler spectra of radar sea echo. *Journal of Geophysical Research*, 75(3):551–563.
- Van Zyl, J. J., Zebker, H. A., and Elachi, C. (1987). Imaging radar polarization signatures: Theory and observation. *Radio Science*, 22(4):529–543.
- Vicen-Bueno, R., Carrasco-Alvarez, R., Rosa-Zurera, M., and Nieto-Borge, J. C. (2009). Sea clutter reduction and target enhancement by neural networks in a marine radar system. *Sensors*, 9(3):1913–1936.
- Walker, D. (2001). Doppler modeling of radar sea clutter. *IEE Proceedings-Radar, Sonar and Navigation*, 148(2):73–80.

- Wanielik, G. and Stock, D. (1990). Measured scattering-matrix-data and a polarimetric CFAR-detector which works on this data. In *Record of the IEEE 1990 International Radar Conference*, pages 514–519. IEEE.
- Ward, K., Baker, C., and Watts, S. (1990). Maritime surveillance radar. I. Radar scattering from the ocean surface. In *IEE Proceedings F-Radar and Signal Processing*, volume 137:2, pages 51–62. IET.
- Ward, K., Tough, R., and Shepherd, P. (2007). Sea clutter transient spatial coherence and scan-to-scan constant false alarm rate. *IET Radar, Sonar & Navigation*, 1(6):425–430.
- Ward, K. D., Watts, S., and Tough, R. J. (2006). *Sea Clutter: Scattering, the K Distribution and Radar Performance*. IET.
- Watts, S. (2008). Radar sea clutter: Recent progress and future challenges. In *2008 IEEE International Conference on Radar*, pages 10–16. IEEE.
- Werle, B. O. (1995). Sea backscatter, spikes and wave group observations at low grazing angles. In *Record of the IEEE 1995 International Radar Conference*, pages 187–195. IEEE.
- Wu, P., Wang, J., and Wang, W. (2011). A novel method of small target detection in sea clutter. *ISR/N Signal Processing*, 2011.
- Yueh, S., Kong, J. A., Barnes, R., and Shin, R. (1990). Calibration of polarimetric radars using in-scene reflectors. *Journal of Electromagnetic Waves and Applications*, 4(1):27–48.
- Zebker, H. A., Van Zyl, J. J., Durden, S. L., and Norikane, L. (1991). Calibrated imaging radar polarimetry: technique, examples, and applications. *IEEE Transactions on Geoscience and Remote Sensing*, 29(6):942–961.