Iceberg Detection Using Compact Polarimetric Synthetic Aperture Radar

Denbina, Michael

Denbina, M. (2014). Iceberg Detection Using Compact Polarimetric Synthetic Aperture Radar (Unpublished doctoral thesis). University of Calgary, Calgary, AB. doi:10.11575/PRISM/26410
http://hdl.handle.net/11023/1946
doctoral thesis

UNIVERSITY OF CALGARY

Iceberg Detection Using Compact Polarimetric Synthetic Aperture Radar

by

Michael Denbina

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE

DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF GEOMATICS ENGINEERING

CALGARY, ALBERTA

November, 2014

Abstract

The management and monitoring of icebergs and sea ice is an important consideration for any maritime operations that take place in the polar and near-polar regions. Spaceborne synthetic aperture radar (SAR) satellites can provide iceberg detection capabilities over vast areas of ocean with a quick revisit time. Canada's next Radarsat satellite, the Radarsat Constellation, will collect compact polarimetry, a polarimetric SAR configuration that transmits a circular polarization rather than the typical linear horizontal or vertical polarizations. In doing so, it can provide greater polarimetric information than typical linear dual-polarized SAR, without the swath width restrictions of fully polarimetric or quad-polarized SAR.

In this thesis, the use of compact polarimetry for iceberg detection is explored. The detection performance of Radarsat Constellation data, simulated using Radarsat-2 data, is assessed using validated iceberg locations collected by a survey vessel. Many new algorithms are tested for use with compact polarimetric SAR data, including detection methods that use pseudo quad-pol reconstruction, wherein quad-pol covariance matrix elements are approximated using the compact polarimetric covariance matrix. Detection using the Stokes parameters is also tested, both through the use of the skew-normal distribution to model the Stokes vector ocean clutter, and also by using the Stokes parameters to calculate the orientation and ellipticity of the polarization ellipse. A method is proposed where the polarization ellipse parameters are then used to vary the detection threshold for each pixel. Discrimination of ships and icebergs after target detection, using a support vector machine classifier, is also demonstrated.

Overall, the compact polarimetric data significantly outperformed the linear dual-polarized data, both in iceberg detection and ship/iceberg discrimination. Compact polarimetry shows much promise for maritime surveillance applications.

Acknowledgements

First I would like to thank my supervisor, Dr. Michael J. Collins, for providing support and guidance throughout my PhD studies at the University of Calgary. Without his mentorship this research would not have been possible.

I would also like to acknowledge the many people and organizations which provided funding, data, and support to this study: Paris Vachon of Defence Research and Development Canada (DRDC) provided funding and data; Desmond Power of C-CORE provided the field observations of icebergs; Chen Liu, also of DRDC, helped with implementation of the LRT; and François Charbonneau, of the Canada Centre for Remote Sensing, lead the RCM compact polarimetry project sponsored by the Canadian Space Agency, and created the software I used to simulate the Radarsat Constellation data. All Radarsat-2 data was provided through this Canadian Space Agency sponsored project. I would also like to acknowledge funding support from the Natural Science and Engineering Research Council of Canada, and scholarship support from the Alberta provincial government.

I would also like to acknowledge Dr. Bryan Mercer, my supervisor during my internship at Intermap Technologies, for first introducing me to polarimetric synthetic aperture radar, and showing me the potential of this technology.

I would also like to thank Ghada Atteia, my fellow PhD student, for her invaluable assistance and collaboration over the past four years.

Last but certainly not least, I would like to express my gratitude to my parents and my sister, for their love and support.

Table of Contents

List of Tables

List of Figures and Illustrations

List of Symbols, Abbreviations and Nomenclature

Symbol	Definition
χ	the ellipticity angle of the polarization ellipse
ψ	the orientation angle of the polarization ellipse
σ	radar cross section (RCS)
σ^0	normalized radar cross section (NRCS)
CFAR	constant false alarm rate
CTLR	circular transmit, linear receive
CL	circular-linear; same as CTLR
FAR	false alarm rate
HH	horizontal transmit, horizontal receive
HV	horizontal transmit, vertical receive
m	degree of polarization
PDF	probability distribution function
PFA	probablity of false alarm; same as FAR
PMD	probability of missed detection
RCM	Radarsat Constellation Mission
RH	right-circular transmit, horizontal receive
RL	right-circular transmit, left-circular receive
RR	right-circular transmit, right-circular receive
RS2	Radarsat-2
RV	right-circular transmit, vertical receive
SAR	synthetic aperture radar
S_0	total backscattered power Stokes parameter
S_1	linear horizontal/vertical Stokes parameter

S_2	$45°/135°$ orientation Stokes parameter
S_3	circular Stokes parameter
VH	vertical transmit, horizontal receive
VV	vertical transmit, vertical receive

Chapter 1

Introduction

1.1 Iceberg Detection And Synthetic Aperture Radar

The management and monitoring of icebergs and sea ice is an important consideration for any maritime operations that take place in the polar and near-polar regions. The National Snow & Ice Data Centre (NSIDC) reports that the Arctic sea ice extent has been steadily decreasing over the past few decades [2]. For example, the Arctic sea ice extent in January 2014 was around 14 million square kilometres, compared to levels of 15 million square kilometres commonly seen in January in the late 1970s and early 1980s [2]. Arctic sea ice extent has declined at an average rate of 3.2% per year from 1981 to 2012 [2], and as the sea ice has receded, shipping as well as oil and gas exploration has increased due to easier access to Arctic waters and rising fuel prices. Arctic shipping is attractive due to reduced travel distances and therefore fuel consumption. Arctic oil and gas exploration is attractive due to the large reserves which are thought to exist there. With estimates that 20% of the world's remaining fossil fuels lie in the Arctic [3], it seems likely that oil and gas exploration will continue to increase in the region despite the technical challenges involved in extracting these resources.

One of these challenges is the hazard posed to ships and sea structures by icebergs. While there is a great deal of variety in iceberg shapes and sizes, all icebergs, even relatively small ones, can present potential hazards to maritime operations and are therefore worth monitoring. The monitoring of icebergs, and their formation and drifting patterns, can also be of scientific interest, such as for measuring some of the effects of climate change. Rising temperatures have led to an increasing number of icebergs breaking off, or calving, from Greenland's continental ice shelf [4]. Greenland is the source of the majority of Arctic

icebergs, creating 10,000 to 15,000 new icebergs per year, most of which float South, through Baffin Bay and into the Labrador Sea [4].

Traditionally, icebergs were monitored using field observations from ships or aircraft, but the use of remote sensing for this purpose has become more common as the number of earth-observing satellites has increased. Satellite-borne synthetic aperture radar (SAR), with its wide area coverage and all-weather day-night operation, is a strong choice for iceberg detection applications. Remote sensing has serious benefits over field observations in terms of cost and scale, but is also less accurate. Detection accuracy is an important consideration, since false detections can result in ships being unnecessarily rerouted, causing delays and financial loss, while missed detections can be catastrophic.

The first SAR systems utilized a single polarization of radiation only, that is, the transmitted and received polarizations were the same (often this was the horizontal or vertical polarization). In recent years, the development of polarimetric SAR systems has provided a new dimension of information, allowing targets to be characterized based on their polarimetric scattering behaviour rather than their single-channel reflectivity. This additional information has also increased the potential accuracy with which icebergs can be detected and classified.

Microwave attenuation depths in freshwater ice generally range from 108-510 cm at X-band to 360-1801 cm at S-band [5], with scattering caused by both the iceberg's surface as well as by air bubbles and other impurities within the ice. Due to the relatively extensive microwave attenuation depth, most icebergs display significant volume scattering, at least in the fall and winter months. In summer, when there is likely to be wet snow or liquid water on the top of the iceberg, volume scattering is reduced, while surface scattering from the top of the iceberg is increased [6]. Imaging geometry also plays an important role in the types of microwave scattering observed—for example, tabular icebergs, whose sides are oriented perpendicularly to the sea, have strong double-bounce scattering, where radiation is reflected

from the iceberg wall and then from the sea (or vice versa) [6]. Due to the wide variety of scattering mechanisms presented by iceberg targets, and the fact these mechanisms are dependent on the size, dielectric properties, and imaging geometry of the iceberg, no generic model for iceberg backscattering has yet been developed.

Despite this lack of a polarimetric scattering model, polarimetric SAR has shown much promise in the detection of icebergs, with dual-polarization (dual-pol) systems, which have two polarimetric channels, being shown as superior to single-polarization (single-pol) systems, such as in the work of Howell et al. [7, 8, 9]. Quad-polarization (quad-pol) or fully polarimetric SAR data, with four polarimetric channels, can produce even higher detection accuracy than dual-pol SAR [10], but due to the fact that quad-pol SAR systems must alternate between the transmission of two orthogonal polarizations [11], the swath width of quad-pol SAR data is severely limited compared to dual-pol. For maritime surveillance applications, where a wide coverage area is a necessity due to the vast size of the ocean, this makes quad-pol SAR a less than ideal choice despite its high detection accuracy.

Shipborne radar has also been used extensively to detect icebergs, and has the advantage of being able to detect very small icebergs that are difficult to detect in remote sensing imagery. However, shipborne radar also has a limited detection range compared to the coverage of remote sensing data [12], measured on the order of nautical miles, compared to hundreds of kilometres for the swath width of satellite borne SAR.

There are a number of examples of previous work in the literature using SAR images to detect icebergs, using a number of different polarimetric configurations and methodologies. Lane et al. [13] calculated probability of detection curves for different icebergs of different size classes within sea ice areas, using a constant false alarm rate (CFAR) detector, using Radarsat-1 (single polarization) and Envisat ASAR (dual-polarization) data, using only the backscattered power, and not the phase. Wesche and Dierking [14] investigated the iceberg backscattering coefficients of icebergs in the Weddell Sea, in Antarctica, tracking how

the backscattering coefficients changed with different seasons using single-polarization data. Kim et al. [15] showed that icebergs could be detected in quad-pol Radarsat-2 data using the entropy/alpha decomposition parameters. Dierking and Wesche [16] investigated the polarimetric signatures of icebergs in sea ice, using quad-pol imagery to calculate parameters such as the phase difference between the HH and VV channels, as well as the entropy/alpha decomposition parameters. In general, any methods which rely on quad-pol data (such as the entropy/alpha) are hindered by the swath width reduction discussed in the previous paragraph. Much of the previous work done in this area has focused on either single-polarization data, with a lack of polarimetric information, or on quad-polarization data, with reduced swath width. Relatively little work has focused on coherent dual-pol SAR (such as the work done by Howell et al. [9]), which can present a compromise between the swath width of single-polarization data, and the maximum polarimetric information of quad-pol data.

1.2 Some Background on Compact Polarimetry

We therefore wish to produce the best detection performance possible using dual-pol SAR data. Traditionally, most dual-pol SAR systems transmitted a linear polarization (either the horizontal or vertical polarization), and received two orthogonal linear polarizations (both the horizontal and vertical polarizations). But this is not a requirement. In recent years, a new SAR architecture called compact polarimetry has been proposed. This architecture uses a dual-pol SAR system, but one that transmits a polarization other than the traditional horizontal or vertical, such as a wave oriented at $\pi/4$ radians (known as the $\pi/4$ mode), as proposed by Souyris et al. [17]. Stacy and Preiss proposed a compact polarimetric SAR configuration called the dual-circular mode, which transmitted a circular polarization (either right-circular or left-circular, where right and left describe the direction of rotation of the wave's electric field vector), and received both the right-circular and left-circular polarizations [18].

Raney proposed another approach to compact polarimetry, which he called hybrid-polarity [19], and which we will refer to here as the circular transmit, linear receive (CTLR) mode to clarify this mode's polarimetric structure, as in the work of Nord et al. [20] and others. In this mode, the transmitted polarization is circular and the received polarizations are the traditional horizontal and vertical [19]. The CTLR approach has a number of interesting advantages over a linear transmit polarization. First, since the transmitted polarization is circular, the resulting backscatter is rotationally invariant, meaning that the amount of power in the backscatted wave does not depend on the orientation of the target geometry around the radar's line of sight [19]. As well, since the received polarizations are the traditional horizontal and vertical, the transmitted and received polarization bases are different rather than the same, and therefore the signal-to-noise ratio (SNR) of the two channels is closer in magnitude than in a traditional dual-pol SAR system, where generally the cross-polarized channel (the channel with different transmit and received polarizations) will have a significantly smaller SNR than the co-polarized channel (the channel with the same transmit and receive polarization). But in a CTLR system neither channel is really co- or cross-polarized in the traditional sense, and therefore the two channels will have closer SNR, meaning that for a constant amount of cross-talk, the effects of this crosstalk will be less noticeable in the CTLR data than in traditional linear dual-pol data, where crosstalk from the co-polarized channel to the cross-polarized channel can have a significant effect on the cross-polarized channel unless the crosstalk is kept very low [19].

Compact polarimetry has been used in radar astronomy successfully for quite some time, and was in fact used to collect images of our Moon [21], but had not often been utilized in remote sensing of the Earth until recently. In the years since Raney's proposal of a CTLR SAR architecture for Earth-based remote sensing, much work has been done to investigate the viability of compact polarimetry for a wide variety of applications and projects. Yet the availability of actual compact polarimetric data is still very limited, with data of the

Earth currently only being collected by the Indian satellite Risat-1. However, Canada's next spaceborne SAR project, the Radarsat Constellation Mission (RCM), will collect CTLR SAR data at C-band in a wide variety of imaging modes [22]. As well, iceberg detection and Arctic maritime surveillance in general are stated objectives of the Radarsat Constellation [22]. Once launched (planned for 2018), the RCM will provide an unprecedented level of compact polarimetric SAR data, with daily revisit times for most of the Canadian Arctic [22]. In an effort to analyze the detection performance of CTLR SAR data as compared to traditional linear dual-pol data, we can use quad-pol SAR data collected by Radarsat-2 to simulate CTLR data at the spatial resolution and approximate noise floor of the RCM imaging modes. This can then give us an insight into the detection performance that can be expected from the RCM data once it is available. By working with simulated data, we can also experiment and validate algorithms that can be ready for use once the Radarsat Constellation has launched.

1.3 Objectives of this Research

While compact polarimetry shows much promise, the use of a circular transmit polarization poses some challenges. The majority of the ship and iceberg detection work done in the past using polarimetric SAR has focused on the linear polarizations. Therefore any algorithms or methods for iceberg detection taken from the literature will likely have to be adapted for use with CTLR SAR data. Similarly, there are many measurements and parameters which can be estimated using compact polarimetry that were previously unavailable to dual-pol SAR, or were less useful when derived from dual-pol SAR than from compact polarimetry (see, for example, chapters 4 and 5), and we therefore wish to assess their usefulness for the detection of icebergs, and how they can be used to improve the detection performance beyond what is possible using linear dual-pol data. It is therefore also important to assess the detection performance of compact polarimetry compared to linear dual-pol, to determine

what benefits compact polarimetry can provide.

Previous publications by the author have shown the strong potential of CTLR SAR data for iceberg detection [10, 23, 24, 25]. This thesis will both elucidate and build upon the author's previous work, as well as explore areas that have been previously unaddressed in the literature, such as the use of the Stokes and polarization ellipse parameters for use in iceberg detection.

There were two main objectives of this research: first, to determine the viability of compact polarimetric SAR data compared to traditional linear dual-pol; second, to improve upon the detection results of compact polarimetric SAR data by taking advantage of its strengths compared to linear transmit polarizations. Considering the increasing demand for iceberg monitoring, it is helpful to begin this work now, rather than attempt to catch up once the RCM has launched and compact polarimetric data of the Arctic becomes more widely available.

1.4 Structure of the Thesis

Chapter 2 gives a description of the data used in this study, including a summary of the Radarsat scenes and validated iceberg locations, as well as the simulated RCM imaging modes.

Chapter 3 gives some background information on the polarization of electromagnetic waves and polarimetric SAR that is necessary for understanding the latter sections of the thesis.

Chapter 4 gives some background information on compact polarimetric SAR specifically, focusing on the CTLR mode used in this thesis. This chapter also discusses the concept of pseudo quad-pol reconstruction, which involves the approximation of some quad-pol parameters using compact polarimetric SAR data and a backscatter model.

Chapter 5 discusses the Stokes parameters, a set of parameters that are well suited for

use with compact polarimetry. We also discuss the parameters of the polarization ellipse, specifically the orientation and ellipticity angles, and how these parameters are linked to the Stokes parameters through a parametrization called the Poincaré sphere. We then discuss the values for these parameters that we expect to see from a number of basic theoretical scattering types, then discuss the actual values observed in the data for ocean pixels and icebergs.

Chapter 6 discusses some methods used to detect targets in polarimetric SAR imagery, with emphasis on the Likelihood Ratio Test (LRT) method used in this study. The potential of using the Stokes parameters to detect targets will also be discussed, with a modified LRT being proposed for use with the Stokes parameters.

Chapter 7 discusses the relationship between false alarm rate (FAR) and detection threshold. Many detection methods used in the literature do this statistically, using a probability distribution function, but here we showcase a simple and effective method through which the $\log_{10}$ value of the FAR vs. the detection threshold is modelled empirically using an actual ocean subset of the SAR image.

Chapter 8 shows the detection results of the methods discussed so far using a variety of polarimetric modes: linear dual-pol, compact polarimetry, and reconstructed quad-pol elements from compact polarimetry. The detection performance of these various polarimetric modes is compared and discussed.

Chapter 9 describes a method through which the polarization ellipse orientation and ellipticity parameters are used to dynamically vary the false alarm rate (FAR) of the detection process from pixel-to-pixel. This method is well suited for use with compact polarimetry due to the rotational invariance of the circular transmit polarization. Results of this method are shown, and compared to the traditional constant FAR case.

Chapter 10 presents the results of discrimination between ship and iceberg targets after detection is performed. Discrimination is done using a support vector machine classifier,

and the results are compared for the CTLR and HH-HV polarization modes and the low resolution and medium resolution imaging modes.

Chapter 11 gives the conclusions of the study and some potential areas of future work.

1.5 Relevant Publications by the Author

M. Denbina and M.J. Collins, "Iceberg detection using pseudo quad-pol reconstruction of compact polarimetric SAR," *Atmosphere Ocean*, vol. 50, no. 4, pp. 437–446, 2012.

M. J. Collins, M. Denbina, and G. Atteia, "On the reconstruction of quad-pol SAR data from compact polarimetry for ocean target detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 1, pp. 591–600, 2013.

M. Denbina and M.J. Collins, "Iceberg detection using simulated dual-polarized Radarsat Constellation data," *Canadian Journal of Remote Sensing*, In Press, 2014.

M. Denbina and M.J. Collins, "Iceberg detection using analysis of the received polarization ellipse in compact polarimetry," *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2014.

Chapter 2

Data Summary

A summary of the Radarsat-2 scenes used in this study is shown in Table 2.1, sorted by beam position (where higher beam numbers correspond to higher incidence angles). The Radarsat-2 scenes were collected using the fine resolution quad-pol mode (fine-quad or FQ for short), which has range sample spacing of 4.73 m, and azimuth sample spacing varying from 4.8–5.0 m according to beam position. Nominal Radarsat-2 fine quad-pol resolution is approximately 5.2 m in range by 7.6 m in azimuth. These scenes were then used to simulate data in the proposed imaging modes of the Radarsat Constellation, as described in section 2.1, using simulator software developed by François Charbonneau of the Canada Centre for Remote Sensing.

The scenes are located around 70-71°N, 58-60°W. Scenes will be referred to in this paper based on the Scene ID (which is based on the date acquired, in the format *mmdd*, followed by a dash, then a number to differentiate multiple scenes collected on the same day). The "Beam" column in the table lists the fine-quad (FQ) beam position, with lower positions having lower incidence angles, and vice versa. "Inc. Angle" is the incidence angle range of the scene, in degrees. "W.S. (m/s)" lists the mean wind speed of each scene, in m/s, as calculated using the mean value of the HV σ_0, according to the linear regression developed by Vachon and Wolfe [26]. A value of less than 3 in the table denotes that the wind speed was too low to accurately estimate, and we can only state that it is less than 3 m/s. "$\overline{N}$" is the mean value of the constant of proportionality in each scene, as described in section 4.2 describing the pseudo quad-pol reconstruction process.

Table 2.1: An overview of the Radarsat-2 fine-quad scenes used for this project. Beam refers to the fine quad beam position, "Inc. Angle" is the incidence angle range of the scene in degrees, W.S. stands for the estimated wind speed of the scene, and $\overline{N}$ is an estimated reconstruction parameter (see section 4.2).

Scene ID	Date Acquired	Beam	Inc. Angle	W.S. (m/s)	$\overline{N}$
0905-1	2009-09-05	2	19.9° - 21.8°	5.0	51.0
0905-2	2009-09-05	2	19.9° - 21.8°	5.0	49.4
0805-1	2009-08-05	5	23.4° - 25.3°	1.3	4.5
0805-2	2009-08-05	5	23.4° - 25.3°	1.5	8.5
0829-1	2009-08-29	5	23.4° - 25.3°	<3	22.3
0815-2	2009-08-15	6	24.6° - 26.5°	5.9	18.8
0822-1	2009-08-22	7	25.8° - 27.6°	3.8	17.7
0822-2	2009-08-22	7	25.8° - 27.6°	4.4	17.5
0815-1	2009-08-15	9	28.0° - 29.9°	3.9	10.0
0825-1	2009-08-25	11	30.3° - 32.0°	5.9	11.1
0825-2	2009-08-25	11	30.3° - 32.0°	5.8	11.1
0818-1	2009-08-18	16	35.5° - 37.0°	6.8	7.6

Table 2.2: An overview of the three Radarsat Constellation Mission imaging modes simulated for this study. Simulated SLC (single-look complex) resolution is an average value taken from scene 0818-1.

RCM	Resolution (rng × az)		Looks	Swath	Incidence	Simulated
Imaging Mode	Near	Far	(rng × az)	Width	Angle	SLC Resolution
Ship Detection	10 × 38 m	17 × 77 m	5 × 1	350km	36° to 51°	1.5m × 38m
Medium Resolution	50 m	50 m	4 × 1	350km	19° to 54°	7.9m × 50m
Low Resolution	100 m	100 m	8 × 1	500km	19° to 58°	8.3m × 99m

2.1 Simulated RCM Imaging Modes

A summary of the wide swath RCM imaging modes simulated for this study are shown in Table 2.2. We used a simulator created by François Charbonneau, of the Canada Centre for Remote Sensing, to create simulated RCM single look complex (SLC) images using the Radarsat-2 fine-quad SLC images as input. The simulated data has a noise floor of approximately -22 dB (higher than the Radarsat-2 data, which has noise floors of approximately -32 dB and lower [26]).

As with the imaging modes of Radarsat-2, the RCM imaging modes can be collected using many different beam modes, each at a different incidence angle, and with slightly different spatial resolution. The "Resolution" column refers to the nominal multi-look spatial resolution, in the format of range $\times$ azimuth for the ship detection mode data. For the medium and low resolution modes, the pixels are square. Near range spatial resolution refers to the steepest beam mode available, and far range resolution refers to the shallowest beam mode available (note that the resolution is only dependent on incidence angle in the ship detection imaging mode case).

"Looks" refers to the number of looks that would be used for the multi-look products of each imaging mode, but since we are operating with simulated SLC imagery, the number of looks in our simulated data is equal to one. The "swath width" column contains the nominal swath width (the width of the image in the range direction), in kilometres.

The simulated SLC resolutions in the table are shown as an average for scene 0818-1, the highest incidence angle scene containing validated iceberg targets that we had available. Unfortunately, this was the only scene which corresponded to the planned ship detection incidence angle range. The planned ship detection mode only uses high incidence angle beam modes, so that the incidence angle of scene 0818-1 (36.5°-38.0°) roughly corresponds to the lowest incidence angle ship detection beam mode. Medium and low resolution modes were simulated for all Radarsat-2 scenes with validated iceberg targets. By simulating three

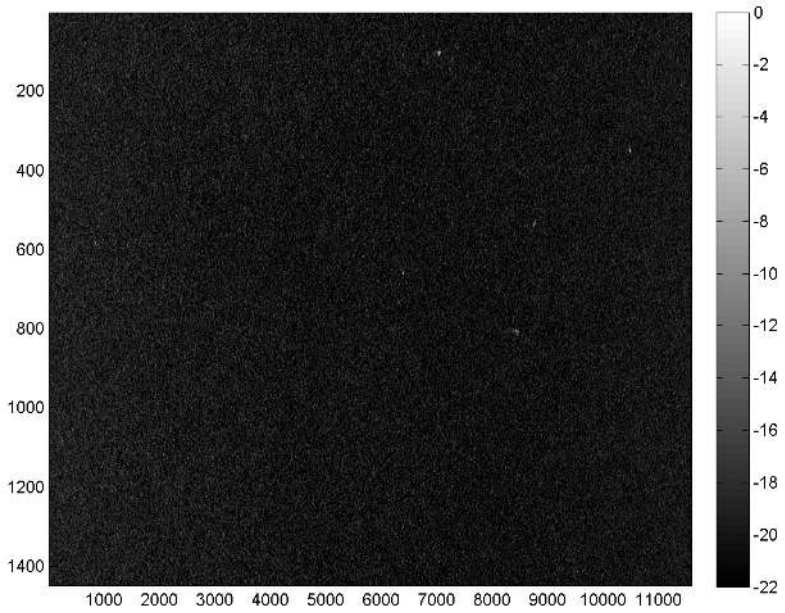

Figure 2.1: RH σ^0 image of scene 0818-1, for the RCM ship detection imaging mode.

different imaging modes, we are able to explore the effects of spatial resolution and pixel size on iceberg detection performance.

RH σ^0 images of scene 0818-1, for the three RCM imaging modes, are shown in Figures 2.1-2.3 by way of example. Note the different dimensions of each image, due to the varying pixel spacing of the different imaging modes. This scene contains one ship and six icebergs.

2.2 Validation Data

25 validated iceberg locations, based on field observations collected from a ship, were provided by Desmond Power of C-CORE. The validation data contained the iceberg's number, location (in latitude and longitude), time sighted, type (blocky, dome, tabular, etc.), and size (small, medium, and large), by which each iceberg was identified in the corresponding fine-quad Radarsat-2 scene (and therefore also in the various simulated RCM scenes). While a larger amount of validation data would be helpful, particularly for different seasons or years (since

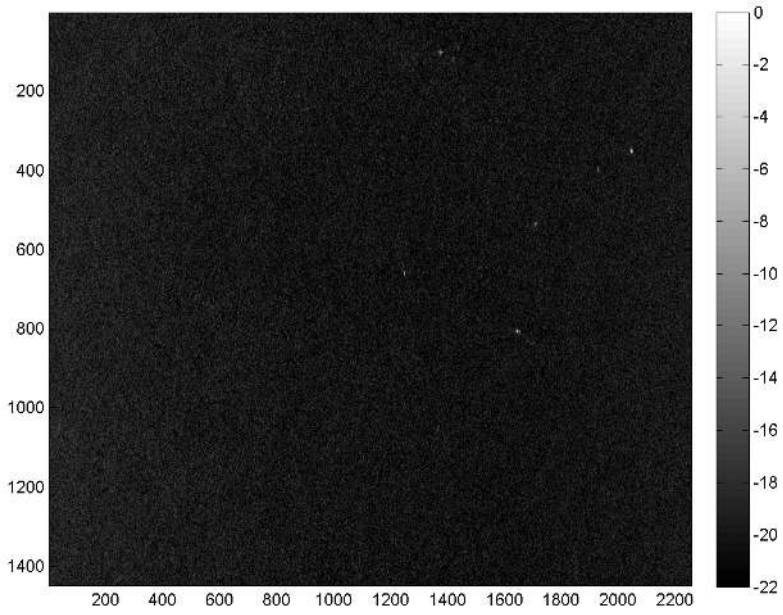

Figure 2.2: RH σ^0 image of scene 0818-1, for the RCM medium resolution imaging mode.

iceberg composition and geometry can vary greatly over time), it is difficult to obtain quad-pol SAR data that is coincident with surveyed iceberg locations. Using the classification system of the International Ice Patrol, small icebergs have heights of 5–15 m and lengths of 15–60 m, medium icebergs have heights of 16–45 m and lengths of 61–122 m, and large icebergs have heights of 46–75 m and lengths of 123–213 m [27]. In this study we are most concerned with the length of the icebergs—icebergs with lengths less than, or equal to, the spatial resolution of the SAR data can be difficult if not impossible to detect.

Unfortunately, the full iceberg location data cannot be included in this thesis due to intellectual property issues which prohibit disclosure of the precise iceberg locations (only the size and shape, not exact location, are given for each iceberg). For each iceberg a mask was manually created to identify the boundary between target and background. Masks were also created for a few ships (including the survey vessel), and some sea ice that was present in two of the scenes. The detection and discrimination performance of sea ice could not be investigated in this study due to a lack of sea ice validation data, but the masks ensured

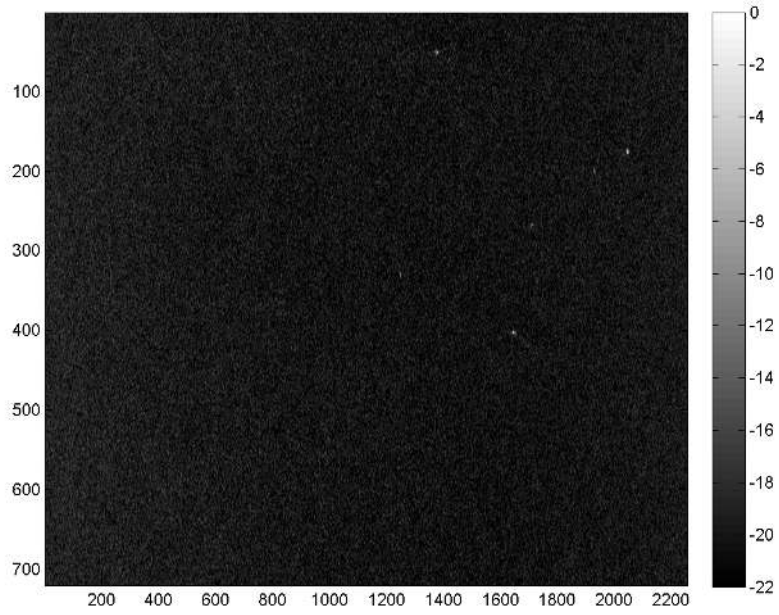

Figure 2.3: RH σ^0 image of scene 0818-1, for the RCM low resolution imaging mode.

that detections of sea ice would not count as false alarms.

In addition, a buffer region was created, covering a five pixel wide area around each of the iceberg masks. Detected clusters in this buffer region were not considered false alarms, since a number of icebergs were surrounded by debris fields which will generally be detected in the SAR image, but are difficult to precisely mask out.

The iceberg targets used in this thesis are summarized in Table 2.3, sorted alphabetically by shape category. Each iceberg was given a letter for identification. When the same iceberg was in multiple Radarsat scenes, it was given a number after the letter to differentiate the sightings.

Table 2.3: An overview of the iceberg validation data used in this study, sorted alphabetically by shape category. Iceberg F was part of a debris field around iceberg E, and did not have a classified size/shape in the validation data. Scene represents the corresponding Radarsat scene where the iceberg was sighted, using the scene IDs in Table 2.1. Number represents the iceberg's ID number within that scene (e.g., for scene 0818-1, which contains six icebergs, the icebergs are numbered from 1 to 6).

Iceberg ID	Scene	Number	Size	Shape
A	0805-2	1	M	Blocky
B1	0905-1	3	M	Blocky
B2	0905-2	2	M	Blocky
C1	0815-1	1	M	Dome
C2	0818-1	6	M	Dome
D1	0815-2	2	M	Drydock
D2	0815-2	3	M	Drydock
E	0818-1	3	M	Drydock
F	0818-1	4	-	-
G1	0829-1	1	S	Drydock
G2	0805-1	1	S	Drydock
H	0905-2	1	M	Drydock
I	0822-2	1	L	Pinnacle
J	0815-2	1	M	Pinnacle
K	0818-1	1	S	Pinnacle
L	0825-1	1	L	Tabular
M	0822-1	1	M	Wedge
N	0805-1	2	S	Wedge
O1	0815-1	2	M	Wedge
O2	0818-1	5	M	Wedge
P1	0822-2	2	L	Wedge
P2	0825-2	1	M	Wedge
Q	0818-1	2	S	Wedge
R1	0905-1	1	S	Wedge
R2	0905-1	2	S	Wedge

Chapter 3

Polarimetry Fundamentals

The polarization of an electromagnetic wave is a term used to describe the orientation of that wave's electric field vector, $\vec{E}$, and how this orientation changes over time (that is, as the wave propagates). Figure 3.1 shows an example of three electromagnetic waves with corresponding polarizations (from left to right): linear, circular, and elliptical. In the figure, the bold black line shows the electric field vector, $\vec{E}$, as it sweeps out a shape in the plane of polarization (the plane shown in blue at the bottom of each image). The plane of polarization is the plane orthogonal to the direction of the wave's propagation (in this case, the direction of propagation is towards the top or bottom of the page). The polarization state of the wave is essentially just the shape that $\vec{E}$ sweeps out in this plane. The red and light blue lines in each image show the x and y components of the electric field vector, E_x and E_y, which are orthogonal to each other. Note that both E_x and E_y are periodic functions which have an amplitude and a phase.

In subfigure (a) of Figure 3.1, note that E_x and E_y are in phase and have equal magnitude. This produces a wave that is linearly polarized at a 45° orientation. If E_y has an amplitude of zero and E_x is non-zero, this will produce a wave that is polarized in the x direction (that is, horizontally polarized). Similarly, if E_x has an amplitude of zero and E_y is non-zero, this will produce a wave that is vertically polarized. In general, as long as E_x and E_y are in phase, the produced wave will always have a linear polarization, with the orientation of that polarization depending on the relative magnitudes of E_x and E_y.

In subfigure (b), note that E_x and E_y are equal in magnitude but are now 90° out of phase. This produces a circularly polarized wave. The direction of the $\vec{E}$ vector's rotation (either right-circular or left-circular) depends on which component of $\vec{E}$ is leading in phase.

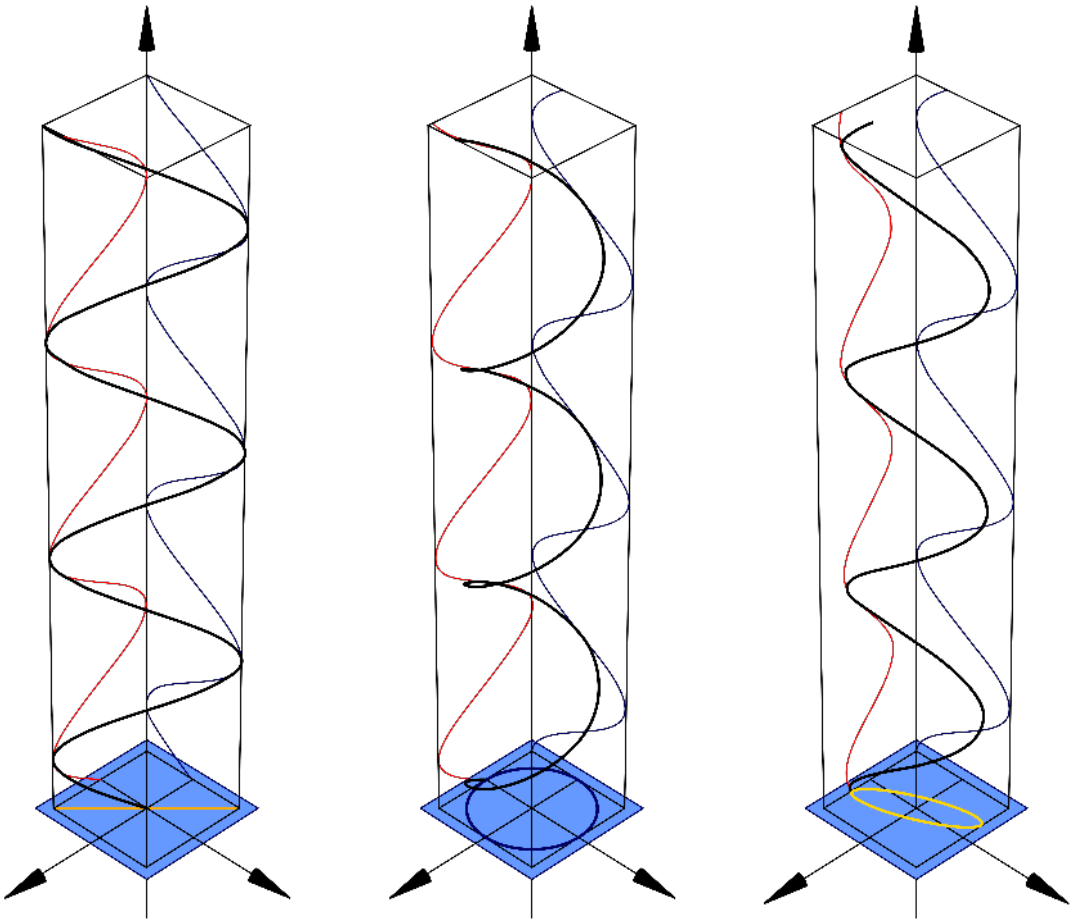

(a) Linear (b) Circular (c) Elliptical

Figure 3.1: Three different basic polarization states, from left to right: linear, circular, and elliptical. These images are in the public domain [1].

Finally, in subfigure (c), we see that E_x and E_y are not in phase, and are not equal in magnitude. This produces an elliptically polarized wave. The shape (or ellipticity) and orientation of this ellipse will depend on both the relative magnitudes of E_x and E_y and their phase difference.

Polarimetry, then, refers to the analysis and interpretation of a wave's polarization, and how changes in polarization can give useful information about the geometry or dielectric properties of targets and terrain. A polarimetric SAR sensor works by transmitting a wave with a known polarization state towards the ground. This wave then reflects, or backscatters,

off the ground and is then received by the SAR sensor. Since the wavelengths used by SAR systems have only minor interaction with the atmosphere, we know that a change in polarization state between the received and transmitted waves is due to the interaction between the wave and the ground. Different types of targets and terrain interact in different ways with a polarized wave, and react differently to different transmitted polarizations. By measuring these differences, we can extract useful information about the Earth's surface.

3.1 The Sinclair Matrix

There are many different ways of mathematically representing the polarization of radar backscatter, but a common method used in polarimetric SAR is that of the Sinclair matrix [11, p. 53]. Before discussing the Sinclair matrix, however, it is necessary to understand the Jones vector [11, p. 37], an easy way to describe the polarization of the waves transmitted and received by a SAR sensor. The Jones vector is essentially a vector containing the complex components E_x and E_y as described in the previous section. The electric field vector $\vec{E}$ will in general be a function of both time, t, and the spatial dimension in the direction of propagation, z. The Jones vector attempts to describe the polarization state of a wave by removing the dependence on time, by noting that $\vec{E}(z,t) = \Re\left(\vec{E}(z)\exp(j\omega t)\right)$, and then taking the value of $\vec{E}(z)$ at $z = 0$, so that the Jones vector $\vec{E}$ is equal to:

$$\vec{E} = \begin{bmatrix} E_x \exp(j\delta_x) \\ E_y \exp(j\delta_y) \end{bmatrix} \tag{3.1}$$

where E_x is the amplitude of the x component of $\vec{E}(z)$, E_y is the amplitude of the y component of $\vec{E}(z)$, and δ_x and δ_y are the phases of the x and y components respectively. Note that for a horizontal wave with unit power, $\vec{E}_H = [1,0]^T$, and for a vertical wave, $\vec{E}_V = [0,1]^T$, and for a right-circular wave, $\vec{E}_R = (1/\sqrt{2})[1,-j]^T$ [11]. The Jones vector is merely a more formal way of presenting the polarization states discussed in the previous section in terms of their x and y components.

In a polarimetric SAR, we are interested in how the polarization state of the incident (or transmitted) wave, represented by the Jones vector $\vec{E}_I$, is affected by the scattering process of the ground, resulting in a different Jones vector for the scattered (or received) wave, $\vec{E}_S$. $\vec{E}_S$ and $\vec{E}_I$ can be related as follows [11, p. 56]:

$$\vec{E}_S = \frac{\exp(-jkr)}{r}\mathbf{S}\vec{E}_I = \frac{\exp(-jkr)}{r}\begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix}\vec{E}_I \tag{3.2}$$

The fraction $\frac{\exp(-jkr)}{r}$ accounts for the propagation effects, both in amplitude and in phase, that result as the wave travels from the SAR sensor to the ground and back again. $\mathbf{S}$ is called the Sinclair matrix, or scattering matrix, and it represents the effects of the scattering process, or scattering mechanism, that takes place when the wave interacts with the ground cover.

Each element of the Sinclair matrix, S_{ij}, is a complex number that describes how the incident wave is reflected by the scatterers in a particular resolution cell of a SAR image (that is, a single Sinclair matrix represents the polarimetric backscatter for one pixel of the image). The subscript denotes the polarization state of the wave, with the first letter or number denoting the transmitted polarization and the second letter or number the received polarization. For example, "H" corresponds to the linear horizontal polarization, "V" corresponds to the linear vertical polarization, and "R" corresponds to the right-circular polarization. Here we use 1 and 2 to make it clear that while the use of the horizontal and vertical polarizations is common, it is not strictly required. Any orthogonal set of polarizations (for example, linear waves oriented at 45° and 135°, or right-circular and left-circular waves) can be used for this purpose. Note as well that the Jones vector implies that if the set of orthogonal components E_x and E_y are known, this is sufficient to describe the polarization state of the wave. In fact, any set of two orthogonal Jones vectors forms what is referred to as a polarization basis, which is sufficient to express the polarization state of any wave [11, p. 56]. That is, we can express any Jones vector $\vec{E}$ in the Cartesian form $\vec{E} = E_x\hat{x} + E_y\hat{y}$,

or indeed in the general form $\vec{E} = E_u \hat{u} + E_{u_\perp} \hat{u}_\perp$, where E_u and $E_{u_\perp}$ are complex scattering coefficients, and $\hat{u}$ and $\hat{u}_\perp$ are any set of orthogonal unit vectors (including complex ones). The use of the horizontal and vertical linear polarization basis is the most common one in the literature, due to simplicity of use and convention, but any set of orthogonal polarizations is valid.

For a typical quad-pol SAR, transmitting the horizontal and vertical polarizations [11]:

$$\mathbf{S} = \begin{bmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{bmatrix} \tag{3.3}$$

Note that in a quad-pol SAR, we transmit two orthogonal polarizations (e.g., H and V), and for each transmit polarization, we split the received wave into two orthogonal polarizations (e.g., the H and V polarization basis). Since we transmit two polarizations, and receive two polarizations for each transmitted wave, a quad-pol SAR has four channels in total.

For example, S_{HH} contains the information for both horizontal transmit and horizontal receive. S_{HV} contains the information for horizontal transmit and vertical receive. And so on. Note that when the transmitted and received polarizations are the same, we refer to these as the co-polarized (co-pol) channels. When the transmitted and received polarizations are different, we refer to these as the cross-polarized (cross-pol) channels.

Note that the Sinclair matrix is similar in some ways to the radar cross section σ or normalized radar cross section σ_0, parameters which have an extensive history in SAR and radar imaging that predates the use of polarimetry. The primary difference between the Sinclair matrix and these parameters is that the Sinclair matrix accounts for the phase of the waves in addition to their power. The radar cross section σ is calculated using the equation [11, p. 54]:

$$\sigma = 4\pi r^2 \frac{|\vec{E}_S|^2}{|\vec{E}_I|^2} \tag{3.4}$$

$\vec{E}_S$ and $\vec{E}_I$, as before, represent the Jones vectors of the scattered and incident waves, respectively, and r is the distance traversed by the wave. Note that we only use the squared

amplitude (that is, the intensity or power) of each wave's Jones vector, and therefore the phase information is discarded. Note the similarity to equation (3.2). While this equation is valid for a point target, in a polarimetric SAR we generally receive backscatter from many different targets within a resolution cell, and therefore we normalize the radar cross section over the illuminated area A_0 such that [11, p. 55]:

$$\sigma^0 = \frac{\langle \sigma \rangle}{A_0} = \frac{4\pi r^2}{A_0} \frac{\left\langle |\vec{E}_S|^2 \right\rangle}{|\vec{E}_I|^2} \tag{3.5}$$

Angle brackets represent that the quantity $|\vec{E}_S|^2$ is an average over the area. Note again that this equation depends on the power of the incident and scattered Jones vectors, but not their phases. Note that both σ and σ^0 (as well as the intensity of the complex scattering coefficients) represent ratios of received power divided by transmitted power. Therefore, they are often expressed in units of decibels.

3.2 The Scattering Vector and the Covariance Matrix

Often, rather than using the Sinclair matrix directly, we vectorize it, producing the scattering vector $\vec{k}$. For example, for a typical quad-pol SAR [11, p. 63]:

$$\vec{k}_{\text{quad-pol}} = V(\mathbf{S}) = [S_{HH}, S_{HV}, S_{VH}, S_{VV}]^T \tag{3.6}$$

By alternating the transmission of two orthogonal polarizations (in this case, H and V), quad-pol SAR provides the greatest amount of polarization information possible. When working with quad-pol data, in the monostatic case (where the transmitting and receiving antenna are the same), due to the reciprocity theorem (which states the transmitting and receiving antenna properties are the same), this implies that $S_{HV} = S_{VH}$ [11, p. 61], yielding the three element quad-pol scattering vector:

$$\vec{k}_{\text{quad-pol}} = \left[S_{HH}, \sqrt{2}S_{HV}, S_{VV}\right]^T \tag{3.7}$$

Note the use of $\sqrt{2}$, to keep the total backscattered power the same as in the four element case.

Unfortunately, for a given swath width, a quad-pol SAR system requires twice the pulse repetition frequency, data rate, and power usage of a dual-pol SAR system. This is due to the quad-pol system's need to alternate the transmission of two orthogonal polarizations, whereas a dual-pol SAR system transmits only one polarization. This is a particularly important consideration for spaceborne platforms, where the available power and data rate are extremely limited. To keep power usage constant, most spaceborne platforms use quad-pol modes with half the swath width of their dual-pol modes, greatly limiting the available coverage of quad-pol data. It is extremely important to image as large an area as possible when mapping the locations of icebergs. Because quad-pol modes have reduced swath width, the Canadian Ice Service does not use quad-pol SAR in their operations.

Since dual-pol SAR transmits only one polarization (e.g., *either* horizontal or vertical), and receives two orthogonal polarizations (e.g., *both* horizontal and vertical), this results in the following most common scattering vectors:

$$\vec{k}_{\text{dual-pol}} = [S_{HH}, S_{HV}]^T \tag{3.8}$$

$$\vec{k}_{\text{dual-pol}} = [S_{VV}, S_{VH}]^T \tag{3.9}$$

Note that when working with dual-pol SAR, we can no longer collect the full Sinclair matrix as in equation 3.3, and are restricted to a single row of the matrix representing a single transmit polarization.

The satellite Envisat also has an "alternating polarization" mode, where the two channels can be HH and HV; VV and VH; or even HH and VV. However, when operating in the HH/VV mode, the relative HH-VV phase coherence is not preserved [28]. The Radarsat Constellation will collect coherent HH-VV data in the medium resolution and low resolution modes, but with a reduced swath width.

Another important measurement in polarimetry is the covariance matrix of the scattering vector, calculated as the product of $\vec{k}$ with its complex conjugate transpose [11, p. 66]:

$$\mathbf{C} = \left\langle \vec{k}\vec{k}^{T*} \right\rangle \tag{3.10}$$

$\mathbf{C}$ is the covariance matrix, $\vec{k}$ is the scattering vector of the data, the superscripted $T*$ represents the complex conjugate transpose, and the angle brackets represent temporal or spatial averaging (e.g., multi-looking). Note that this equation for the covariance matrix of $\vec{k}$ implies that the real and imaginary components of the elements of $\vec{k}$ have mean values of zero, which is generally the case for the scattering coefficients but not necessarily any parameters derived from them, which will be discussed in more detail in section 6.3.

For a typical linear quad-pol system, $\mathbf{C}$ is then equal to [11, p. 67]:

$$\mathbf{C}_4 = \left\langle \begin{bmatrix} |S_{HH}|^2 & S_{HH}S_{HV}^* & S_{HH}S_{VH}^* & S_{HH}S_{VV}^* \\ S_{HV}S_{HH}^* & |S_{HV}|^2 & S_{HV}S_{VH}^* & S_{HV}S_{VV}^* \\ S_{VH}S_{HH}^* & S_{VH}S_{HV}^* & |S_{VH}|^2 & S_{VH}S_{VV}^* \\ S_{VV}S_{HH}^* & S_{VV}S_{HV}^* & S_{VV}S_{VH}^* & |S_{VV}|^2 \end{bmatrix} \right\rangle \tag{3.11}$$

Which in the monostatic, three element case, reduces to [11, p. 68]:

$$\mathbf{C}_3 = \left\langle \begin{bmatrix} |S_{HH}|^2 & \sqrt{2}S_{HH}S_{HV}^* & S_{HH}S_{VV}^* \\ \sqrt{2}S_{HV}S_{HH}^* & 2|S_{HV}|^2 & \sqrt{2}S_{HV}S_{VV}^* \\ S_{VV}S_{HH}^* & \sqrt{2}S_{VV}S_{HV}^* & |S_{VV}|^2 \end{bmatrix} \right\rangle \tag{3.12}$$

And for the two typical linear dual-pol systems:

$$\mathbf{C}_2 = \left\langle \begin{bmatrix} |S_{HH}|^2 & S_{HH}S_{HV}^* \\ S_{HV}S_{HH}^* & |S_{HV}|^2 \end{bmatrix} \right\rangle \tag{3.13}$$

$$\mathbf{C}_2 = \left\langle \begin{bmatrix} |S_{VV}|^2 & S_{VV}S_{VH}^* \\ S_{VH}S_{VV}^* & |S_{VH}|^2 \end{bmatrix} \right\rangle \tag{3.14}$$

The covariance matrix, particularly for the ocean pixels, is an important measurement that is used in the calculation of the target detection decision variable (described in section

6.1) and in the reconstruction of pseudo quad-pol data from compact polarimetry (described in section 4.2).

3.3 Speckle and the Nature of Ocean Backscatter

Speckle is a fundamental property of SAR images that must be considered when working with SAR data. Speckle is multiplicative noise that manifests due to the fact that the backscattered wave received by a SAR sensor does not represent a single reflection from the ground, but rather consists of many wavelets reflected from a large number of elementary scatterers which are randomly distributed throughout the resolution cell [11, p. 101]. This happens because the wavelengths used by a SAR sensor will generally be on a similar scale to that of the surface roughness of the ground cover [11], which is not the case for optical remote sensing. Due to the random locations of the elementary scatterers within a resolution cell, the distance between each scatterer and the SAR sensor will vary, which will manifest as phase differences in the received wavelets [11]. If the wavelets interfere with each other constructively, the signal received by the SAR sensor will be stronger than expected; if the wavelets interfere destructively, the received signal will be weaker [11]. This causes a variation in the amount of backscattered power received by the SAR sensor from pixel-to-pixel, even if the underlying radar cross section is constant (that is, the ground cover is homogeneous) [11]. What this means is that generally speaking, measurements or calculated parameters taken from a single pixel of a single-look SAR image are not reliable—some form of filtering or averaging must be performed in an attempt to reduce the effects of speckle.

The simplest model for SAR speckle is the Rayleigh model, where each complex scattering coefficient S will have real and imaginary parts that are independently and identically Gaussian (Normal) distributed with a mean of zero and variance $\sigma^2/2$ [11, p. 102]. If x is the real part of S, and y the imaginary part, their joint probability distribution function

(PDF) has the form [11, p. 102]:

$$p_{x,y}(x,y) = \frac{1}{\pi\sigma^2}\exp\left(-(x^2+y^2)/\sigma^2\right) \tag{3.15}$$

This results in an S with uniformly distributed phase between $-\pi$ and π, and a Rayleigh distributed amplitude. The PDF of the single-look amplitude of S, A, is then equal to [11, p. 103]:

$$p(A) = \frac{2A}{\sigma^2}\exp\left(-\frac{A^2}{\sigma^2}\right), A \geq 0 \tag{3.16}$$

And the PDF of the intensity of S, I, then has a negative exponential form that is equal to [11]:

$$p(I) = \frac{1}{\sigma^2}\exp\left(-\frac{I}{\sigma^2}\right), I \geq 0 \tag{3.17}$$

While simple, the Rayleigh model is not commonly used for real world applications in the literature, and is mostly used as the basis for other more complicated models. One of the most widely accepted models of SAR speckle in the literature, and one very commonly used for ocean imagery, is the product model [11, 29]. The product model multiplies the negative exponentially distributed intensity I of the Rayleigh model by a parameter that represents the texture variation throughout the image, g. The model is succinctly described by the equation [11, p. 108]:

$$\tilde{Y} = gI \tag{3.18}$$

where $\tilde{Y}$ is the product model distributed intensity. In the case where the texture parameter is gamma distributed, the intensity can be modelled using a probability distribution function known as the K distribution [11, 29, 30]. In this case, g has a PDF equal to [11]:

$$p_g(g) = \frac{1}{g\Gamma(\alpha)}(\alpha g)^{\alpha-1}\exp\left(-\alpha g\right), g \geq 0 \tag{3.19}$$

α is a texture parameter that determines the variance of g. g has variance equal to $1/\alpha$, such that for larger values of α the intensity $\tilde{Y}$ becomes more homogeneous, and vice versa [11]. In the extreme case, as α tends to infinity, g approaches a constant value of one, such that the K-distributed intensity reduces to that of the Rayleigh model ($\tilde{Y} = I$).

The K-distributed intensity can also be parametrized in terms of L, the number of statistically independent looks; ν, the shape parameter; and μ, the mean value of the K-distributed intensity [31, 29, 32]:

$$p(\tilde{Y}) = \frac{2}{\tilde{Y}} \left(\frac{L\nu\tilde{Y}}{\mu} \right)^{\frac{L+\nu}{2}} \frac{1}{\Gamma(L)\Gamma(\nu)} K_{\nu-L} \left(2\sqrt{\frac{L\nu\tilde{Y}}{\mu}} \right) \tag{3.20}$$

This is the model that will be used to model ocean clutter in this thesis. For a more detailed discussion of this, and how the clutter model is necessary to describe the functional relationship between detection threshold and false alarm rate, see chapter 7.

Worth noting is that the underlying RCS of the ocean varies on a spatial scale much larger than the pixel spacing of SAR imagery [29], so that essentially all of the pixel-to-pixel variation in ocean backscatter is due to speckle, and not to the underlying geophysical processes (e.g., changes in facet tilt angle, sea state, or salinity). Therefore, neighbouring ocean pixels can be averaged together in order to reduce the effects of the speckle without meaningfully degrading the ocean imagery. This is in contrast to backscatter from targets, such as icebergs or ships, which tends to be heterogeneous. If possible, it is best to avoid averaging targets together with the ocean pixels when speckle filtering, and many speckle filtering algorithms have been proposed that attempt to avoid degrading target backscatter while still performing smoothing over homogeneous areas, though generally these speckle filtering algorithms involve tradeoffs in terms of the amount of smoothing they perform, resulting in greater speckle noise in the ocean regions than a straight boxcar average, but greater preservation of targets. This issue will be discussed in more detail in chapter 6.

Chapter 4

Compact Polarimetry

This chapter gives an overview of some of the concepts surrounding compact polarimetry. As discussed briefly in the introduction, compact polarimetry is an umbrella term that describes a number of coherent dual-polarized SAR architectures, all of which involve the use of a transmit polarization other than the linear horizontal or vertical. By transmitting circularly polarized waves, or linear waves oriented at 45°, we hope to gain increased polarimetric information over horizontal/vertical transmit dual-pol, without the swath width reduction of full polarimetry.

This thesis presents work done using the CTLR (circular transmit, linear receive) mode of compact polarimetry, also sometimes referred to as CL (circular-linear), CL-pol, or hybrid-polarity [19]. This mode uses a circular transmit polarization, but splits the received waves into their horizontal and vertical components, yielding two channels: RH (right-circular transmit, horizontal receive) and RV (right-circular transmit, vertical receive). This is the mode that will be available on the Radarsat Constellation once it launches [22].

4.1 The CTLR Scattering Vector And Covariance Matrix

The Radarsat Constellation will transmit right-circular polarization when operating in the CTLR mode, and therefore a right-circular transmit polarization is used in this thesis. However, use of a left-circular transmit polarization is theoretically very similar, just with some changes in sign to accomodate the opposite rotation direction. Right-circular polarization has a Jones vector of $\vec{E_R} = (1/\sqrt{2})[1, -j]^T$, while left-circular polarization has a Jones vector of $\vec{E_L} = (1/\sqrt{2})[1, j]^T$ [11]. Note that right-circular and left-circular polarization waves are orthogonal, and therefore form a polarization basis [11].

For a CTLR SAR system, transmitting right-circular polarization, the scattering vector is equal to:

$$\vec{k}_{CTLR} = [S_{RH}, S_{RV}]^T \tag{4.1}$$

Comparing this vector to the typical linear dual-pol vectors (which contain either the HH and HV channels, or the VV and VH channels), we see that the transmitted polarization is different, but the received polarization basis is the same. Raney pointed out that the use of a different transmitted and received polarization basis meant that neither channel of a CTLR SAR system is co-polarized or cross-polarized in the traditional sense, and that therefore the signal-to-noise ratio of the two channels should be closer than in SAR systems where the transmitted and received polarization basis is the same [19]. In fact, the CTLR scattering vector can be rewritten in the linear basis as follows [20]:

$$\vec{k}_{CTLR} = \frac{1}{\sqrt{2}} [S_{HH} - iS_{HV}, -iS_{VV} + S_{VH}]^T \tag{4.2}$$

As this scattering vector shows, the RH and RV channels each contain a mixture of co-polarized and cross-polarized linear backscatter. Since the co-polarized channels generally have much higher signal levels than the cross-polarized channels, this means that often the RH channel is dominated by the HH component, and the RV channel is dominated by the VV component. This is particularly true for the Radarsat Constellation data, where the noise floor for most of the wide swath imaging modes is -22 dB, higher than the expected HV and VH σ^0 values for ocean backscatter [26]. Compared to common noise floor levels ranging from approximately -38 dB (at low incidence) to -32 dB (at high incidence) for the Radarsat-2 fine-quad mode data [26], this is a significant increase and makes the use of cross-pol channels for characterizing ocean backscatter problematic for RCM data (e.g., to calculate wind speed as in the work of Vachon and Wolfe [26]).

Equation 4.2 also demonstrates how the CTLR data can be simulated from full polarimetry (where the HH, HV, VH, and VV channels are all known).

The 2×2 covariance matrix of the CTLR mode is equal to:

$$\mathbf{C}_{CTLR} = \left\langle kk^{*T} \right\rangle = \left\langle \begin{bmatrix} |S_{RH}|^2 & S_{RH}S_{RV}^* \\ S_{RH}^*S_{RV} & |S_{RV}|^2 \end{bmatrix} \right\rangle \tag{4.3}$$

The angle brackets denote spatial or temporal averaging, and the superscript $*T$ is the complex conjugate transpose operation. Like the scattering vector, this covariance matrix can be rewritten in the linear basis as the sum of three smaller matrices [20], assuming reciprocity ($S_{HV} = S_{VH}$):

$$\begin{aligned} \mathbf{C}_{CTLR} = & \frac{1}{2}\Bigg\langle \begin{bmatrix} |S_{HH}|^2 & i(S_{HH} \cdot S_{VV}^*) \\ -i(S_{VV} \cdot S_{HH}^*) & |S_{VV}|^2 \end{bmatrix} + \begin{bmatrix} |S_{HV}|^2 & -i|S_{HV}|^2 \\ i|S_{HV}|^2 & |S_{HV}|^2 \end{bmatrix} + \\ & \begin{bmatrix} -2\Im(S_{HH} \cdot S_{HV}^*) & S_{HH} \cdot S_{HV}^* + S_{VV}^* \cdot S_{HV} \\ S_{HH}^* \cdot S_{HV} + S_{VV} \cdot S_{HV}^* & 2\Im(S_{VV} \cdot S_{HV}^*) \end{bmatrix} \Bigg\rangle \end{aligned} \tag{4.4}$$

Note that the first matrix consists of exclusively co-polarized terms, the second matrix consists of exclusively cross-polarized terms, and the third matrix consists of the products of co-polarized and cross-polarized terms.

4.2 Pseudo Quad-Pol Reconstruction

From equation 4.2, we see that the CTLR scattering vector contains a mixture of the four linear channels measured by a quad-polarized SAR system. Methods have been proposed in the literature that use a scattering model to calculate approximate values of some of these quad-pol channels from compact polarimetric data. Specifically, the 2×2 compact polarimetry covariance matrix is used to reconstruct some elements of the 3×3 quad-pol covariance matrix. This process is referred to as pseudo quad-pol reconstruction, and the resulting data as pseudo quad-pol or pseudo-quad.

The method was originally proposed by Souyris et al. [17], for use with the $\pi/4$ mode of compact polarimetry. Nord et al. [20] showed how the method could be also be used for

both the dual-circular polarization and CTLR modes of compact polarimetry, with similar reconstruction performance results.

Recall that the 3×3 quad-pol covariance matrix is as follows [11, 20]:

$$\mathbf{C}_3 = \left\langle \begin{bmatrix} |S_{HH}|^2 & \sqrt{2}S_{HH}S_{HV}^* & S_{HH}S_{VV}^* \\ \sqrt{2}S_{HV}S_{HH}^* & 2|S_{HV}|^2 & \sqrt{2}S_{HV}S_{VV}^* \\ S_{VV}S_{HH}^* & \sqrt{2}S_{VV}S_{HV}^* & |S_{VV}|^2 \end{bmatrix} \right\rangle \tag{4.5}$$

This covariance matrix contains nine unknowns: three real-valued diagonal elements (the HH, HV, and VV intensities), plus three complex off-diagonal elements, each of which has both a real and imaginary component ($\sqrt{2}S_{HH}S_{HV}^*$, $S_{HH}S_{VV}^*$, and $\sqrt{2}S_{HV}S_{VV}^*$). Unfortunately, the CTLR covariance matrix of equation 4.4 only contains four equations (one for each element of the matrix). To reduce the number of unknowns, Souyris et al. [17] made use of the reflection symmetry assumption, which implies that the cross-polarized channels have zero correlation with the co-polarized channels, resulting in the simplified covariance matrix [11, 20]:

$$\mathbf{C}_3 = \left\langle \begin{bmatrix} |S_{HH}|^2 & 0 & S_{HH}S_{VV}^* \\ 0 & 2|S_{HV}|^2 & 0 \\ S_{VV}S_{HH}^* & 0 & |S_{VV}|^2 \end{bmatrix} \right\rangle \tag{4.6}$$

The reflection symmetry assumption was shown to be valid at C-band in a paper by Collins et al. [24], co-authored by the author of this thesis. Note that the reflection symmetry assumption allows us to set the third matrix of equation 4.4 to zero, such that the CTLR covariance matrix consists of co-polarized and cross-polarized terms only, not their complex products. This reduces the number of unknowns to five: the three intensities, and the real and imaginary parts of $S_{HH}S_{VV}^*$. One more equation is still needed to constrain the solution space, and Souyris et al. [17] proposed the following equation to do so:

$$\frac{\langle |S_{HV}|^2 \rangle}{\langle |S_{HH}|^2 \rangle + \langle |S_{VV}|^2 \rangle} = \frac{(1 - |\rho|)}{N} \tag{4.7}$$

The angle brackets represent spatial averaging and ρ is the HH-VV coherence, defined as:

$$\rho \equiv \frac{\langle S_{HH} \cdot S_{VV}^{*} \rangle}{\sqrt{\langle |S_{HH}|^2 \rangle \cdot \langle |S_{VV}|^2 \rangle}} \tag{4.8}$$

Larger averaging windows will generally reduce the effects of reflection asymmetry, and therefore improve the accuracy of the reconstruction, up to a certain point. However, larger windows also reduce the resolution of the imagery. Because the application of this research is iceberg detection, where resolution affects the size of the smallest targets that can be detected, a 3×3 pixel averaging window was used to calculate the covariance matrix elements used in the reconstruction process.

Equation 4.7 provides a general relation between the cross-polarization ratio (the ratio of the HV intensity to the sum of the two co-polarized intensities) and the HH-VV coherence, tuned by the value of the parameter N. Originally, Souyris et al. [17] used a value of $N = 4$, which Nord et al. [20] showed was consistent with scattering from many types of natural terrain and vegetation. Nord et al. also proposed a method for refining the value of N during the reconstruction process, based on recalculating N as the ratio of double-bounce to cross-pol backscatter [20]. However, Collins et al. [24] demonstrated that these methods are not consistent with ocean backscatter. They proposed that for ocean imagery the backscattering is primarily dependent on incidence angle, and that the value of N used in the reconstruction process should take this into account. Using quad-pol data, they calculated the actual values of N, then plotted them vs. incidence angle. Figure 4.1 shows the mean value of N for each of the scenes in this study, $\bar{N}$, plotted vs. the mean incidence angle of the scene. We see that a power law regression can be used to relate $\bar{N}$ with incidence angle, with a relatively high r^2 value of 0.935.

Similarly, Collins et al. [24] developed the following regression for N vs. incidence angle (θ), which is the one used for pseudo quad-pol reconstruction in this thesis:

$$\bar{N} = 6.52 + 18305.73 \exp\left(-\theta^{0.60}\right) \tag{4.9}$$

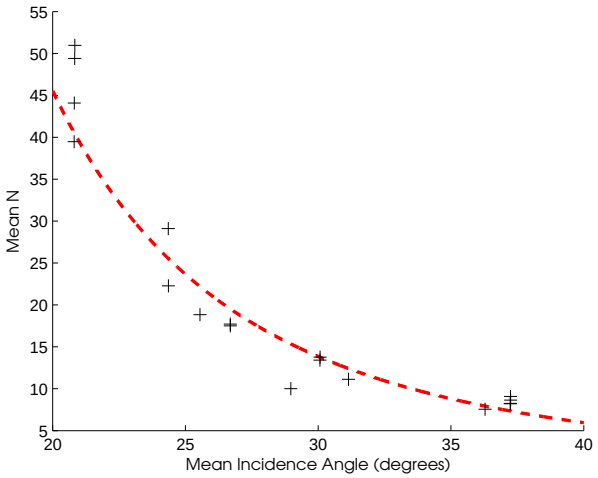

Figure 4.1: Mean N value versus mean incidence angle for each scene. The red line shows the regression $y = 300900x^{-2.936}$, with r^2 of 0.935.

Once a value for N is selected, equations 4.7 and 4.8 are solved iteratively using the CTLR covariance matrix elements to estimate $|S_{HV}|^2$ and ρ as follows [20]:

$$\rho_{(i+1)} = \frac{-i\mathbf{C}_{12} + (|S_{HV}|^2)_{(i)}}{\sqrt{\left(\mathbf{C}_{11} - (|S_{HV}|^2)_{(i)}\right)\left(\mathbf{C}_{22} - (|S_{HV}|^2)_{(i)}\right)}} \tag{4.10}$$

$$(|S_{HV}|^2)_{(i+1)} = (\mathbf{C}_{11} + \mathbf{C}_{22}) \cdot \left(\frac{\left(1 - |\rho_{(i+1)}|\right)}{N + 2\left(1 - |\rho_{(i+1)}|\right)}\right) \tag{4.11}$$

The subscripted (i) represents the iteration number and $\mathbf{C}_{ab}$ represents the element from the CTLR covariance matrix (equation 4.4) in row a and column b. A starting estimate is used of $(|S_{HV}|^2)_{(0)} = 0$. We iterate between updating the estimate of ρ and $(|S_{HV}|^2)$ until convergence (defined here as the HV intensity estimate changing by less than 1% of its previous value in one iteration), or until 20 iterations have been performed, whichever comes first. Note that in general, provided a value of N is chosen that is appropriate for the data, 99–100% of the ocean pixels will converge, often within three or four iterations.

Once estimates of ρ and $|S_{HV}|^2$ are calculated, they are used to calculate the values of

the reflection symmetric quad-pol covariance matrix, using the equations:

$$|S_{HH}|^2 = C_{11} - |S_{HV}|^2 \tag{4.12}$$

$$|S_{VV}|^2 = C_{22} - |S_{HV}|^2 \tag{4.13}$$

$$S_{HH} \cdot S_{VV}^* = -iC_{12} + |S_{HV}|^2 \tag{4.14}$$

C_{11}, C_{12} and C_{22} are the corresponding elements of the CTLR covariance matrix. The covariance matrix of the reconstructed pseudo quad-pol data can similarly be calculated:

$$\mathbf{C}_{PQ_{CTLR}} = \begin{bmatrix} C_{11} - |S_{HV}|^2 & 0 & -iC_{12} + |S_{HV}|^2 \\ 0 & 2|S_{HV}|^2 & 0 \\ (-iC_{12} + |S_{HV}|^2)^* & 0 & C_{22} - |S_{HV}|^2 \end{bmatrix} \tag{4.15}$$

Note that while we calculate approximate intensities for HH, VV, and HV, and the relative HH-VV phase difference, due to the reflection symmetry assumption we are not able to calculate any HV phase information, since we have assumed that the correlation between the HV channel and the co-polarized channels is zero. In practice, of course, this assumption is not 100% true—these elements of the covariance matrix are much smaller in comparison to the reconstructed elements, but for actual data they are not equal to zero. Collins et al. [24] showed that at C-band, these products could sometimes be as high as $5 - 10\%$ of the co-pol covariance matrix elements, with the reflection asymmetry increasing with incidence angle.

There are two main sources of error in the reconstruction process. First are errors arising from the reflection symmetry assumption, as described in the previous paragraph. Second, while we have modelled the behaviour of N by using the mean value of N across a range of incidence angles, it should be noted that real data has a distribution of N values that vary from pixel-to-pixel, something which cannot be accounted for by the model. Therefore, error will be introduced into the reconstruction process due to the discrepancy between the actual underlying N value and the value used as input to the reconstruction.

Typical values of the reconstruction error were investigated by Nord et al. [20] for terrain, and Collins et al. [24] for ocean backscatter. The values of the reconstruction errors in this study will be not be discussed here, other than to note that for this application they are, in fact, a secondary concern—what matters is whether the reconstruction process preserves the contrast between ocean pixels and target pixels, not the actual radiometric accuracy of its approximations. For iceberg detection, the real test of the reconstruction process is the detection performance of the pseudo-quad data, which will be assessed in chapter 8.

In this thesis, reconstructed pseudo quad-pol data will be used for iceberg detection in a number of ways. First, the full pseudo quad-pol dataset will be used (pseudo-HH, pseudo-VV, pseudo-HV, and the relative HH-VV phase). As well, the pseudo-HV intensity will be tested on its own, as it has been shown that the HV intensity is a particularly important parameter in iceberg detection [10]. Finally, the pseudo-HV intensity will be combined with the coherent RH and RV channels, to create a hybrid mode between the native CTLR and the pseudo quad-pol data.

There are other novel ways to use the compact polarimetry data which do not involve the reconstruction or approximation of quad-pol elements. One of the ways of characterizing the polarization of a wave is through the use of four real numbers called the Stokes parameters, a method which is well suited for use with compact polarimetry and which will be discussed in the following chapter.

Chapter 5

The Stokes Parameters and the Polarization Ellipse

There are a wide variety of ways in which we can represent the polarization of an electromagnetic wave, and therefore a wide variety of ways in which we can represent and analyze polarimetric SAR data. In the previous sections we have mostly used the Sinclair matrix $\mathbf{S}$, its vectorized form $\vec{k}$, and the covariance matrix $\mathbf{C}$. Here we present an alternative method of describing the polarization of a wave called the Stokes parameters, or Stokes vector. The Stokes parameters are closely related to the polarization ellipse, the shape swept out by the electric field vector $\vec{E}$ as it rotates in the plane orthogonal to the wave's direction of propagation. The relationship between the Stokes parameters and the polarization ellipse is artfully illustrated by the Poincaré sphere, a spherical parametrization of three of the Stokes parameters, which we will use to demonstrate how the polarization ellipse's orientation, ψ, and ellipticity, χ, can be easily calculated using Stokes data.

After presenting this theory, we derive the expected Stokes vector and polarization ellipse parameters for a number of "canonical scatterers"—basic scatterer types such as trihedrals, dihedrals, dipoles, etc., which have been widely discussed in the literature. We derive these values algebraically, by calculating the corresponding scattering vectors $\vec{k}$ for three different polarimetric modes: CTLR, HH-HV, and VV-VH, then calculating the Stokes vector elements and polarization ellipse parameters using basic trigonometric identities to simplify the expressions where necessary. In doing so, we demonstrate one of the strengths of a CTLR SAR system: that it is less affected by target orientation than a linear dual-pol SAR. For icebergs, where the target orientation is unknown and is essentially random, this is an important advantage.

Finally, we present plots of the Stokes vector and polarization ellipse parameters for the

actual data used in this study, giving us an idea of the types of values these parameters take on for both ocean and iceberg pixels.

5.1 The Stokes Parameters

The Stokes parameters (or Stokes vector) were proposed by George Gabriel Stokes in 1852 [33], and are a set of four real numbers that can be used to describe the polarization state of a wave. Stokes's work focused on the polarization of light waves, but the Stokes parameters can be used to describe the polarization state of any electromagnetic wave, including the microwaves transmitted and received by a SAR sensor. The Stokes parameters have been used with SAR data for a variety of applications, and are well discussed in the SAR literature (see, for example, [11, p. 43]). Raney, when proposing the use of CTLR as a compact polarimetric polarization mode, suggested that the Stokes parameters were a natural choice to represent the data collected by a CTLR SAR system, or indeed, by any coherent dual-polarized SAR [19].

The Stokes parameters for a dual-pol SAR can be calculated as follows, in the linear basis [19]:

$$S_0 = \left\langle |E_H|^2 + |E_V|^2 \right\rangle \tag{5.1}$$

$$S_1 = \left\langle |E_H|^2 - |E_V|^2 \right\rangle \tag{5.2}$$

$$S_2 = 2\Re \left\langle E_H E_V^* \right\rangle \tag{5.3}$$

$$S_3 = 2\Im \left\langle E_H E_V^* \right\rangle \tag{5.4}$$

E represents the complex voltage of the subscripted received polarization, and the angle brackets represent spatial or temporal averaging. Note that since calculation of S_2 and S_3 rely on the value of the complex product $E_H E_V^*$, the Stokes vector can only be calculated using a coherent SAR system, where the relative phase between E_H and E_V is known.

Note that the transmitted polarization is not included in the equation. The Stokes

parameter values for SAR data using a linear transmit polarization will be different from those calculated using a right-circular transmit polarization. Where applicable, we will mention the transmit polarization used to calculate a given Stokes vector, or denote a Stokes vector with its subscripted transmit polarization (e.g., $\vec{S}_H$ for horizontal transmission, $\vec{S}_V$ for vertical transmission, and $\vec{S}_R$ for right-circular transmission).

The naming of the Stokes vector elements varies in the literature. Sometimes the Stokes vector elements are numbered from S_1 to S_4, and other times from g_0 to g_3 or S_0 to S_3. Here we use the latter convention, so that S_0 represents the total power contained in the wave.

S_1, S_2, and S_3 then represent the amount of backscattered power within each of three degenerate polarization states [11, 19]: horizontal/vertical, 45°/135°-oriented, and right-circular/left-circular. A positive S_1 value represents horizontal polarization, and a negative value represents vertical polarization. Similarly, a positive value of S_2 represents 45°-oriented polarization, and a negative value represents 135°-oriented (or −45°-oriented) polarization. In this thesis, a positive value of S_3 represents left-circular polarization, and a negative value represents right-circular. Note that the sign convention of S_3 can vary in the literature, particularly in the field of optics, where a positive S_3 value represents right-circular polarization [34]. The convention we use here, however, is the generally accepted sign convention within the SAR literature, and is the one used by Lee and Pottier [11].

There are a wide variety of further parameters which can be easily computed using the Stokes parameters, but of these, arguably the most imporant (as emphasized by Raney [19]) is the degree of polarization, or DoP, here denoted with the letter m:

$$m = \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0} \tag{5.5}$$

m is the ratio of the polarized power of a wave to the total power of a wave. For a fully polarized wave, $m = 1$, while for a fully unpolarized wave, $m = 0$.

5.1.1 Conversion from CTLR to Dual-Circular Polarization

The Stokes parameters also provide a straight forward method of changing the received polarization basis from linear to circular (or vice versa), so that we can essentially convert from the CTLR mode of compact polarimetry to the dual-circular polarization (DCP) mode. In the previous section the equations for the Stokes parameters were presented using the linear receive basis, but we can also define the equations in the circular receive basis, as follows [19]:

$$S_0 = \left\langle |E_L|^2 + |E_R|^2 \right\rangle \tag{5.6}$$

$$S_1 = 2\Re \left\langle E_L E_R^* \right\rangle \tag{5.7}$$

$$S_2 = 2\Im \left\langle E_L E_R^* \right\rangle \tag{5.8}$$

$$S_3 = \left\langle |E_L|^2 - |E_R|^2 \right\rangle \tag{5.9}$$

As above, E represents the complex voltage of the subscripted received polarization, and the angle brackets represent spatial or temporal averaging. From these equations, we see that if we calculate the Stokes vector from right-circular CTLR data, the right-circular DCP channels can then be calculated using the following equations:

$$|E_{RL}|^2 = 0.5(S_0 + S_3) \tag{5.10}$$

$$|E_{RR}|^2 = 0.5(S_0 - S_3) \tag{5.11}$$

$$E_{RL}E_{RR}^* = 0.5(S_1 + iS_2) \tag{5.12}$$

In chapter 8, the detection performance of the native CTLR, and DCP calculated from the CTLR, will be assessed and compared.

5.2 The Polarization Ellipse

Another way of representing the polarization of a wave other than the Stokes parameters is to define the ellipse swept out by the rotation of that wave's electric field vector in the plane

of polarization [11, p. 34], which is the plane perpendicular to the direction of propagation of the wave. Figure 5.1 shows an example of a polarization ellipse, with two important parameters of the ellipse emphasized: ψ, the orientation of the ellipse (measured as the angle between the ellipse's major axis and the horizontal), and χ, the ellipticity of the ellipse (where $\chi = 0$ for a linear wave, and $\chi = \pm 45°$ for a circular wave, with the sign of χ denoting the direction of the wave's rotation).

Note that the electric field vector, $\vec{E}$, is generally defined by its two orthogonal components E_x and E_y [11], whose magnitudes and phases are directly related to the orientation and shape of the polarization ellipse. If E_x and E_y are in phase, the polarization of the wave is linear, and the polarization ellipse collapses into a line ($\chi = 0$). The magnitudes of E_x and E_y then determine if the wave is horizontally polarized ($\psi = 0$), vertically polarized ($\psi = 90°$), or polarized at some other orientation angle. If E_x and E_y are out of phase, the polarization then becomes elliptical ($\chi \neq 0$). A wave is circular if E_x and E_y are equal in magnitude and are exactly 90° out of phase ($\chi = \pm 45°$).

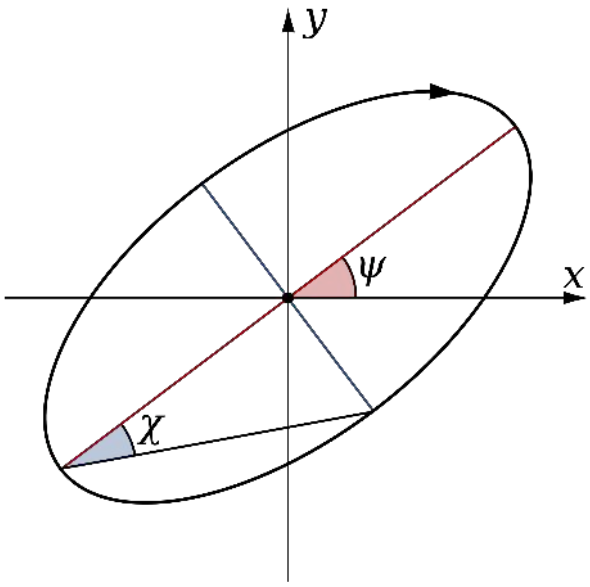

Figure 5.1: An example of the polarization ellipse, with orientation angle ψ and ellipticity angle χ. This image is in the public domain [1].

5.3 The Poincaré Sphere

While the Polarization ellipse is a convenient and natural way of illustrating the polarization state of a wave, for polarization states that are not one of the degenerate states (horizontal, vertical, circular, etc.), it can be unclear how to calculate ψ and χ from the polarization ellipse itself. Jules Henri Poincaré proposed the Poincaré sphere, essentially a spherical parametrization of the Stokes parameters [34]. The Poincaré sphere is shown in Figure 5.2, demonstrating that if the Stokes parameters S_1, S_2, and S_3 are used as the three Cartesian coordinates of a point in 3-D space, then 2ψ and 2χ can be calculated as two of the spherical coordinates of that same point.

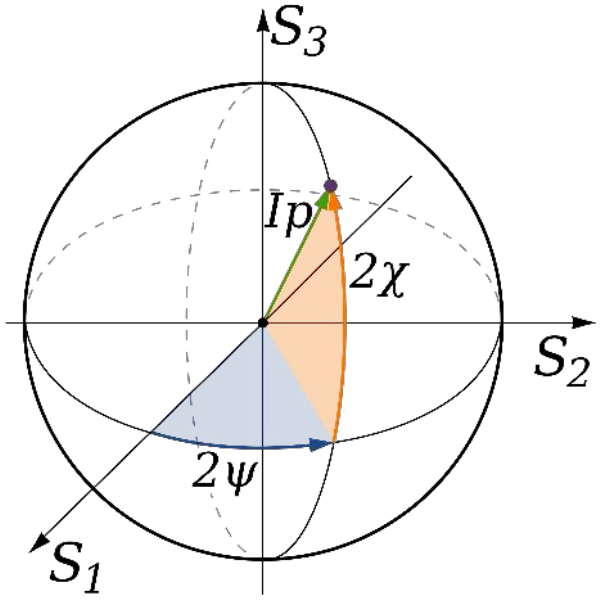

Figure 5.2: The Poincaré sphere, with the Stokes parameters S_1, S_2, and S_3 aligned along the Cartesian axes, with spherical coordinates 2ψ and 2χ. This image is in the public domain [1].

The Poincaré sphere artfully illustrates a number of connections and important ideas: first, that the Stokes parameters are directly connected to the polarization ellipse; second, that the three degenerate polarization states described by S_1, S_2, and S_3 are orthogonal and can describe any shape and orientation of polarization ellipse; third, that conversion

from the Stokes parameters to ψ and χ can be accomplished through a simple coordinate transformation. Therefore, ψ and χ can be derived from the Stokes parameters using the equations [35]:

$$\psi = \frac{1}{2}\arctan\left(\frac{S_2}{S_1}\right) \tag{5.13}$$

$$\chi = \frac{1}{2}\arctan\left(\frac{S_3}{\sqrt{S_1^2 + S_2^2}}\right) \tag{5.14}$$

Note that the distance between the origin and the point defined by S_1, S_2, and S_3 is equal to the total polarized power, $I_p = \sqrt{S_1^2 + S_2^2 + S_3^2}$, which is itself equal to the numerator of the equation for the degree of polarization, m.

For right-circular and vertical transmit polarizations, after calculating the ψ values, we add 180° to any negative values, such that the range of ψ values is from $0 - 180°$. We do this so that the majority of observed values will lie in the middle of the range. For horizontal transmit polarization we generally avoid this, since most of the ψ values will lie around zero in this case. Note that any ψ value x has the same meaning as $x + 180°$ (for x in degrees), since the polarization ellipse is symmetrical about both its major and minor axes.

5.4 The Stokes and Polarization Ellipse Response of Canonical Scatterers

Real world targets (such as icebergs) generally have a complicated geometric structure, which manifests in a similarly complicated scattering response. These scattering responses can be difficult to interpret, particularly in the case of icebergs where each target will have its own unique structure based on the way it was formed and its size. However, there are a number of canonical scattering mechanisms belonging to targets with very basic geometric properties that have been widely discussed in the literature. The backscattering behaviour of these simplified targets, such as flat plates, dipoles, or dihedrals, can be calculated theoretically rather easily. We can then use these canonical scattering mechanisms to aid in our analysis

of real world targets, for example, by noting the similarity of an ocean pixel's backscatter to a flate plate, or the similarity between the scattering of an iceberg wall and a dihedral.

Here we use Lee and Pottier [11] as a reference for some of these canonical scattering mechanisms. Using the Sinclair matrix of the elementary target, we can then calculate the corresponding Stokes and polarization ellipse parameters we would expect to receive from that target, for each of the three transmit polarizations under consideration (right-circular, horizontal, and vertical). To account for the fact that each scatterer can have an arbitrary orientation angle in the plane of polarization, ϕ, the Sinclair matrix for most of the targets will be a function of this angle (with the exception of trihedral scatterers, which produce the same backscatter regardless of orientation angle). The algebra in this section was done using Matlab, with simplification aided by basic trigonometric identities where necessary (the double-angle identity in particular).

5.4.1 Sphere, Flat Plate, Trihedral

A sphere, flat plate, and trihedral have the same scattering response. The quad-pol Sinclair matrix of these targets, in the Cartesian basis, is [11, p. 92]:

$$\mathbf{S} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{5.15}$$

The received Stokes vector for right-circular transmit polarization, $\vec{S}_R$, is then calculated to be:

$$\vec{S}_R = [1, 0, 0, 1]^T \tag{5.16}$$

This Stokes vector corresponds to an undefined ψ value, and a χ value of 45°. Note that the ψ value is undefined since both S_1 and S_2 are equal to zero. Since the received wave is completely circular, the orientation of the polarization ellipse is therefore arbitrary.

The received Stokes vectors for horizontal and vertical transmit polarization, $\vec{S}_H$ and $\vec{S}_V$,

are:

$$\vec{S}_H = [1, 1, 0, 0]^T \tag{5.17}$$

$$\vec{S}_V = [1, -1, 0, 0]^T \tag{5.18}$$

For the horizontal transmit case, $\psi_H = 0$ and $\chi_H = 0$. For the vertical transmit case, $\psi_V = 90°$ and $\chi_V = 0$. We note that in each case, the trihedral essentially reflects the transmitted polarization back at the sensor, leaving the polarization state of the wave unchanged.

5.4.2 Dipole

The quad-pol Sinclair matrix of a dipole oriented at ϕ degrees from the horizontal is equal to [11, p. 94]:

$$\mathbf{S} = \begin{bmatrix} \cos^2(\phi) & \frac{1}{2}\sin(2\phi) \\ \frac{1}{2}\sin(2\phi) & \sin^2(\phi) \end{bmatrix} \tag{5.19}$$

From which we calculate the received Stokes vectors:

$$\vec{S}_R = [0.5, 0.5\cos(2\phi), 0.5\sin(2\phi), 0]^T \tag{5.20}$$

$$\vec{S}_H = [\cos^2(\phi), \cos(2\phi)\cos^2(\phi), \sin(2\phi)\cos^2(\phi), 0]^T \tag{5.21}$$

$$\vec{S}_V = [\sin^2(\phi), \cos(2\phi)\sin^2(\phi), \sin(2\phi)\sin^2(\phi), 0]^T \tag{5.22}$$

This is an interesting case, where the orientation and ellipticity of the polarization ellipse for the three transmitted polarizations are the same ($\psi = \phi$, $\chi = 0$), but where the total power received by the sensor is quite different in each case. For right-circular transmission, the received power does not depend on the orientation of the dipole. For horizontal transmission the backscattered power depends on $\cos^2(\phi)$, so that we receive the maximum backscattered power for a horizontal dipole, and zero backscattered power for a vertical dipole. For the vertical transmission case the backscattered power depends on $\sin^2(\phi)$, so that this behaviour is reversed. This is an example of a case where an advantage of transmitting circular transmission is clear, in that it is agnostic to the orientation of the target in the plane of polarization (that is, the orientation about the radar's line of sight).

5.4.3 Dihedral

The quad-pol Sinclair matrix of a dihedral oriented at ϕ degrees from the horizontal is [11, p. 95]:

$$\mathbf{S} = \begin{bmatrix} \cos(2\phi) & \sin(2\phi) \\ \sin(2\phi) & -\cos(2\phi) \end{bmatrix} \tag{5.23}$$

From which we calculate the received Stokes vectors:

$$\vec{S}_R = [1, 0, 0, -1]^T \tag{5.24}$$

$$\vec{S}_H = [1, \cos(4\phi), \sin(4\phi), 0]^T \tag{5.25}$$

$$\vec{S}_V = [1, -\cos(4\phi), -\sin(4\phi), 0]^T \tag{5.26}$$

We see that again, the right-circular case is agnostic to the target orientation, with $\chi = -45°$, and undefined ψ (since the wave is perfectly circular).

For the horizontal and vertical polarizations, χ is zero in both cases, while $\psi = 2\phi$ in the horizontal case and $\psi = 90° + 2\phi$ in the vertical case. We note that for $\phi = 0$, this is the same polarization response as for a trihedral or flat plate. But note that the orientation angle of a trihedral or flat plate is not meaningful (since the geometry is symmetric in the plane of polarization), whereas here, the angle of the dihedral can have an affect on the backscattered polarization in the case of linear transmitted polarization.

5.4.4 Cylinder

Another basic target type is the cylinder, with the quad-pol Sinclair matrix [36]:

$$\mathbf{S} = \frac{1}{\sqrt{1.25}} \begin{bmatrix} 1 - 0.5\sin^2(\phi) & 0.25\sin(2\phi) \\ 0.25\sin(2\phi) & 1 - 0.5\cos^2(\phi) \end{bmatrix} \tag{5.27}$$

From which we calculate the received Stokes vectors:

$$\vec{S}_R = [0.5, 0.3\cos(2\phi), 0.3\sin(2\phi), 0.4]^T \tag{5.28}$$

$$\vec{S}_H = [0.6\cos^2(\phi) + 0.2, 0.05\cos(4\phi) + 0.3\cos(2\phi) + 0.45, ... \\ 0.05\sin(4\phi) + 0.3\sin(2\phi), 0]^T \tag{5.29}$$

$$\vec{S}_V = [0.6\sin^2(\phi) + 0.2, -0.05\cos(4\phi) + 0.3\cos(2\phi) - 0.45, ... \\ -0.05\sin(4\phi) + 0.3\sin(2\phi), 0]^T \tag{5.30}$$

For the right-circular case, $\psi = \phi$, and $\chi = 26.57°$. The χ values for the horizontal and vertical case are both zero, while the ψ value is a periodic function of ϕ in the horizontal and vertical cases.

5.4.5 Quarter Wave Device

Another basic target type is that of a quarter wave device. There are two types of quarter wave device, either a positive quarter wave device or a negative quarter wave device. First we look at the positive quarter wave device, represented by the quad-pol Sinclair matrix [36]:

$$\mathbf{S} = \begin{bmatrix} 1 + \sin^2(\phi)(-1+j) & 0.5\sin(2\phi)(1-j) \\ 0.5\sin(2\phi)(1-j) & 1 + \cos^2(\phi)(-1+j) \end{bmatrix} \tag{5.31}$$

From this we calculate the received Stokes vectors:

$$\vec{S}_R = [1, -\sin(2\phi), \cos(2\phi), 0]^T \tag{5.32}$$

$$\vec{S}_H = [1, 0.5 + 0.5\cos(4\phi), 0.5\sin(4\phi), \sin(2\phi)]^T \tag{5.33}$$

$$\vec{S}_V = [1, -0.5 - 0.5\cos(4\phi), -0.5\sin(4\phi), -\sin(2\phi)]^T \tag{5.34}$$

For the right-circular polarization, $\psi = \phi + 45°$, and $\chi = 0$. For the horizontal and vertical cases, both ψ and χ are again periodic functions of ϕ.

And for a negative quarter wave device, we see that the behaviour is very similar, only

with a $-j$ in place of j [36]:

$$\mathbf{S} = \begin{bmatrix} 1 + \sin^2(\phi)(-1-j) & 0.5\sin(2\phi)(1+j) \\ 0.5\sin(2\phi)(1+j) & 1 + \cos^2(\phi)(-1-j) \end{bmatrix} \tag{5.35}$$

From this we calculate the received Stokes vector for the right-circular case:

$$\vec{S}_R = [1, \sin(2\phi), -\cos(2\phi), 0]^T, \psi = \phi + 45, \chi = 0 \tag{5.36}$$

Note that S_1 and S_2 have changed signs, but that the vector is otherwise the same. Note also that $\psi = \phi - 45°$, and $\chi = 0$.

5.4.6 Discussion of Canonical Scatterers

There are a few things worth noting from the above theory. First, for all of the targets, when transmitting right-circular polarization, the amount of backscattered power (S_0) does not depend on the target's orientation about the radar line of sight. This is not the case for horizontal or vertical transmitted polarization, where both dipoles and cylinders produce different amounts of backscattered power depending on their orientation with respect to the angle of the transmitted polarization. This is an advantage of circular transmission, since we do not want a target to appear dimmer or brighter based on its geometry. This is an especially important consideration for icebergs, where the target orientation is unknown and can essentially be considered to be random. We therefore wish for the polarization response to depend as little as possible on target orientation. When using quad-pol data, one can extract the target orientation angle, removing its effects (see, for example, the work of Touzi et al. [37]), but this is not an option when using dual-pol SAR. We are therefore stuck with the effects of target orientation angle on our resulting data.

Second, for dihedral scatterers (such as the double-bounce backscatter from an iceberg wall and the ocean surface), when transmitting circular polarization, the target orientation has no effect on the received polarization ellipse, since the received wave is perfectly circular, and its orientation is therefore irrelevant. As well, dihedrals and trihedrals can be easily

distinguished when transmitting circular polarization due to the fact that while both these targets produce circular waves, for right-circular transmission a trihedral will produce a left-circular wave, while a dihedral will produce a right-circular wave. Contrast this with the transmitted horizontal and linear polarizations, where the χ angle is zero in both cases (representing a perfectly linear wave), and it is impossible to distinguish between a trihedral and a dihedral with $\phi = 0$.

Overall, then, we emphasize that these are compelling reasons to use compact polarimetry for this application rather than linear transmitted polarization. By transmitting circular polarization, we mitigate the effects of target orientation as much as possible with a dual-polarized SAR sensor.

5.5 Observed Stokes Parameter Values of Ocean Pixels

So far this chapter has focused on the theory behind the polarization ellipse and the Stokes parameters. Now we begin to present and discuss the actual Stokes parameter values as observed in the data used for this study. For each scene we used an image chip approximately 500 by 500 pixels in size, containing ocean pixels only, with which we plotted a density plot (2-D histogram) of the Stokes vector values S_1, S_2, and S_3. Figures for all scenes and imaging modes used in the study are given in Appendix A, for reference. Here we focus our discussion on a low incidence angle scene (0905-1), and a high incidence angle scene (0818-1), using the low resolution imaging mode only, for brevity. Since we are focusing on the Stokes vector behaviour across a sizable image chip, the use of different imaging modes should have only a minor affect on the resulting plots.

Figure 5.3 shows Stokes vector density plots of S_1, S_2, and S_3, for scene 0905-1, using right-circular transmit polarization. As might be expected, the circular Stokes parameter, S_3 has a higher value than the linear parameters S_1 and S_2, indicating a mostly circular wave. S_2 has a small positive bias, and S_1 has a small negative bias, so that the mean of the

ocean's ψ values will lie in the range of $45° - 90°$.

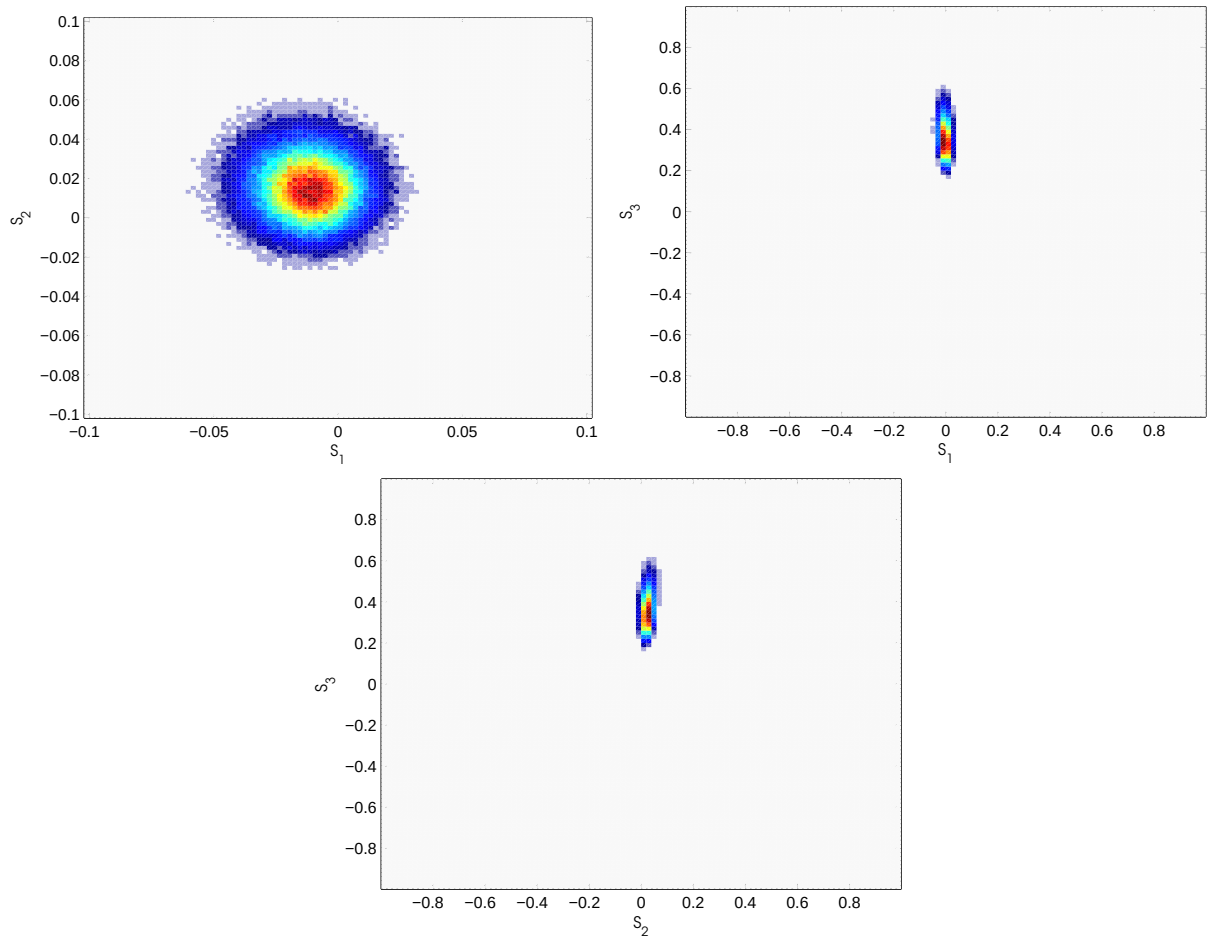

Figure 5.3: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-1, low resolution mode, for right-circular transmit polarization.

Figure 5.4 shows the Stokes vector density plots for scene 0818-1, using right-circular transmit polarization. We see that while S_3 is still larger than S_1 and S_2, it is not as large as for the steep incidence angle scene. This represents a reduction in the magnitude of χ, and therefore a less circular wave. We see that S_1 and S_2 have a similar bias as before, but this time the values are more widely distributed.

Now we look at the Stokes vector values for the linear transmit polarizations. Figure 5.5 shows the Stokes vector density plots for scene 0818-1, with the left column showing the

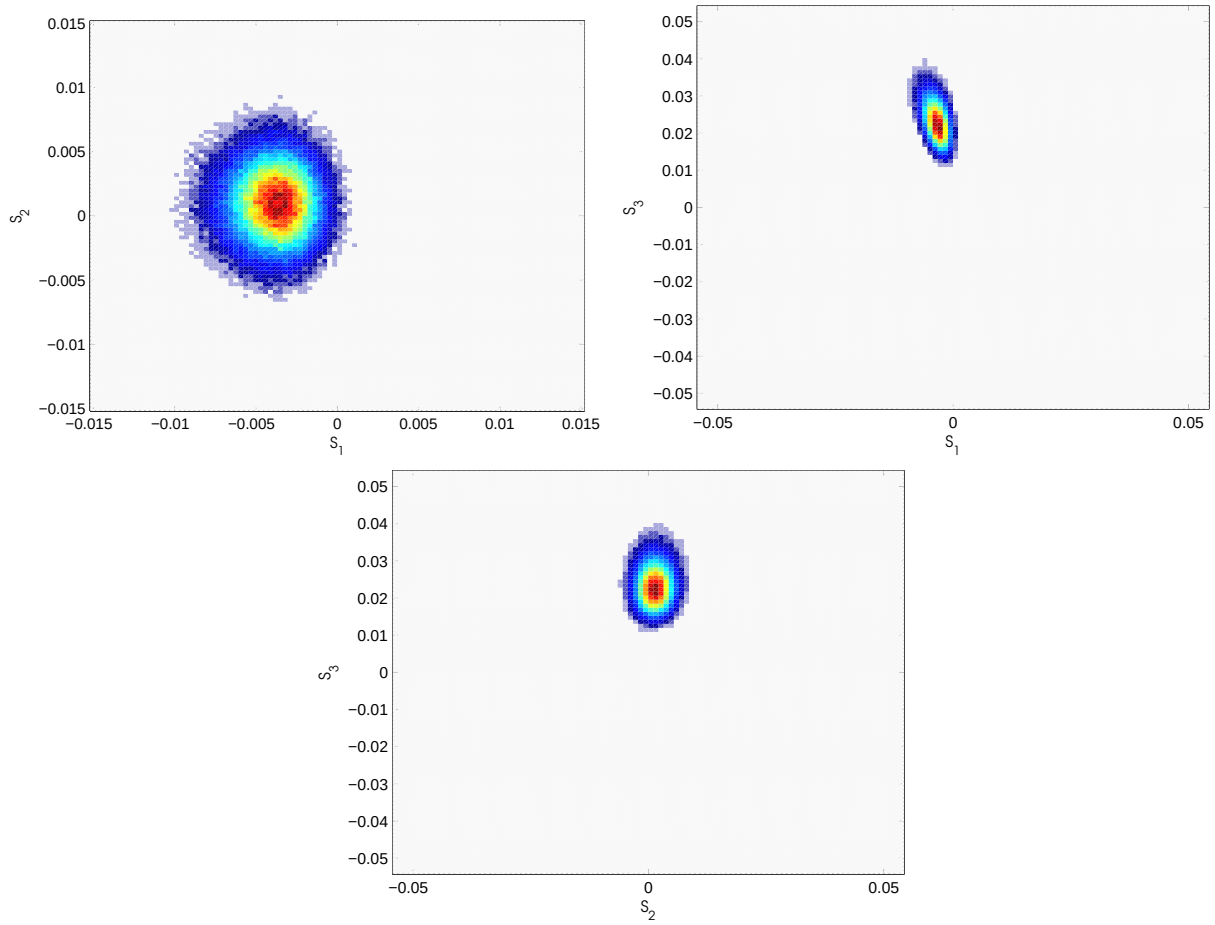

Figure 5.4: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, low resolution mode, for right-circular transmit polarization.

results for horizontal polarization, and the right column showing vertical polarization. We see that now S_1 has a greater magnitude than S_2 or S_3, and that S_1 has mostly positive values in the horizontal case, and the mostly negative values in the vertical case. This reflects the fact that when a linear transmit polarization is used, most of the resulting backscatter from the ocean will be aligned in the same direction as the transmitted polarization.

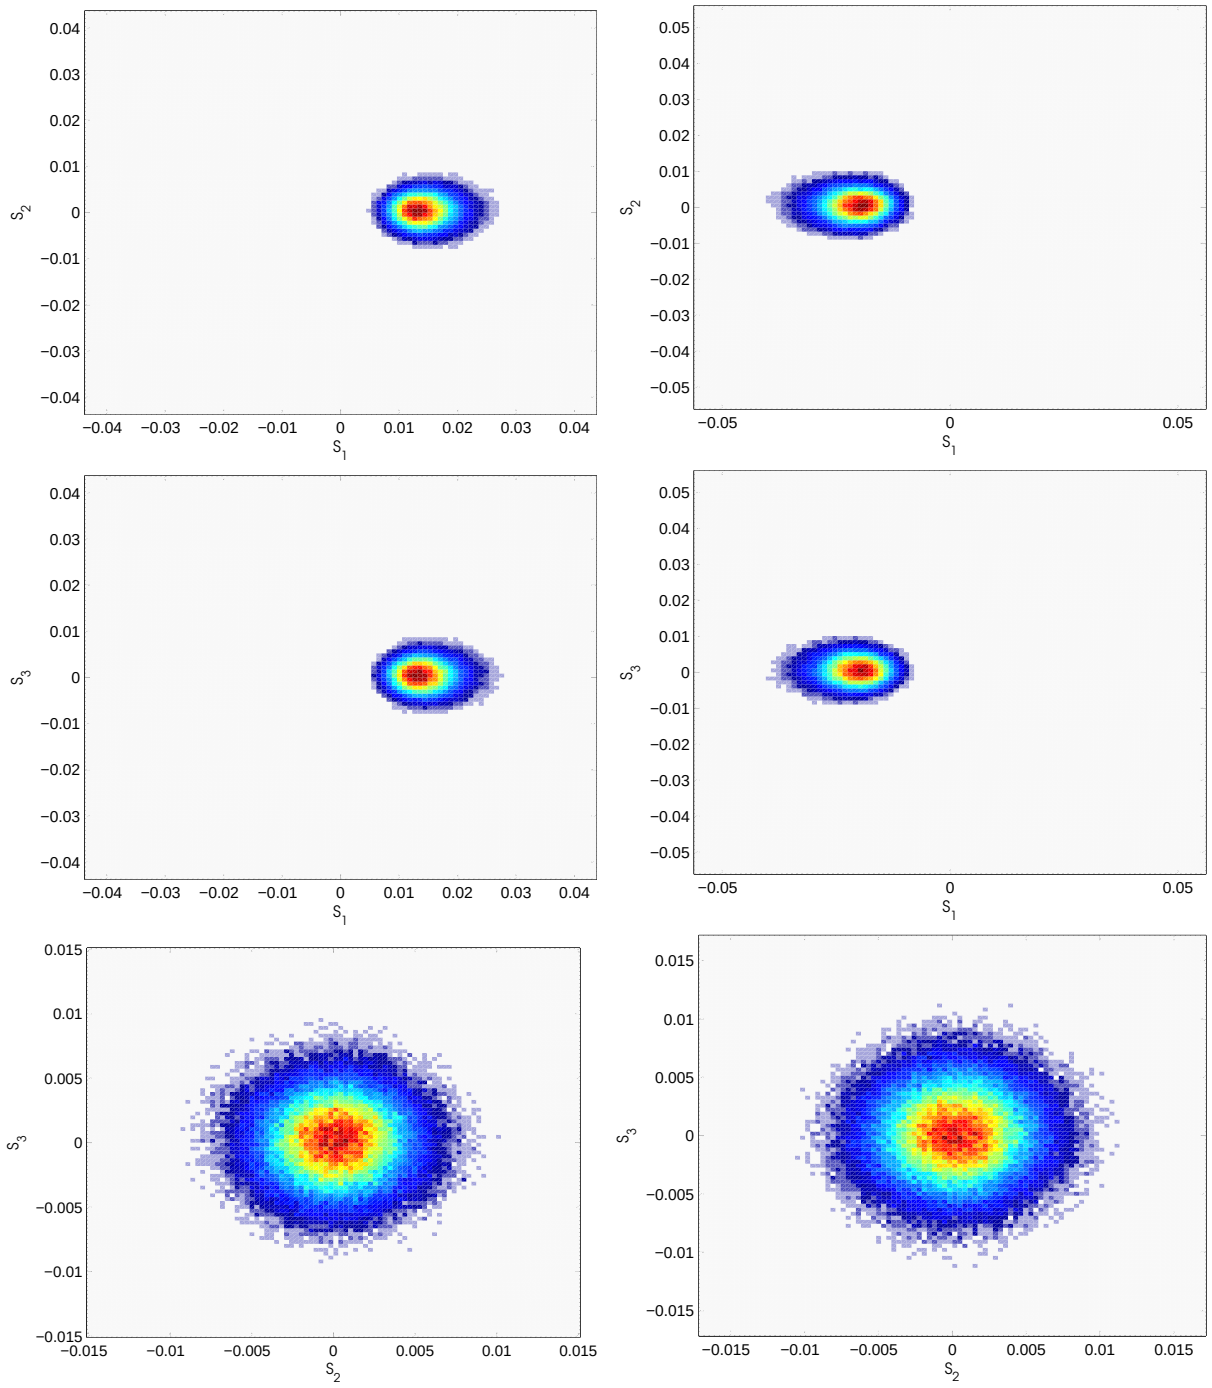

(a) HH-HV (b) VV-VH

Figure 5.5: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

5.6 Observed Polarization Ellipse Response for Ocean and Icebergs

Figure 5.6 shows the mean values of ψ, χ, and m vs. incidence angle for the ocean image chips in each scene, using right-circular transmit polarization. We see that ψ is relatively unaffected by incidence angle, and is generally in the range of $50° - 120°$. The ocean backscatter becomes more depolarized as the incidence angle increases, but the mean m stays above 0.8 even at the higher incidence angles, with a close to left-circular wave (mean $\chi > 35°$). These results suggest that at very low incidence angles, the ocean acts similarly to a trihedral or flat plate scatterer, while at higher incidence angles the ocean's scattering becomes more cylindrical, though falling short of what we would expect from a true cylinder scatterer (which would theoretically produce a χ angle of 26.57°, higher than the values observed in this data). In fact we would expect the ocean to exhibit Bragg scattering for most intermediate incidence angles [11], and our data is fairly consistent with that expectation.

Similar plots for horizontal and vertical transmit polarizations are shown in Figure 5.7. We note that for the horizontal case, the ψ values of the ocean are centred approximately around zero, while for the vertical case, they are clustered approximately around 90°. χ values are centred around zero for both cases, respenting a predominantly linear wave. ψ and χ values do not seem to be a strong function of incidence angle in either polarization case. As with the right-circular polarization, m decreases with increasing incidence angle, though this effect seems more pronounced than in the right-circular case.

Compared to the ocean, the iceberg targets in this study exhibit a much wider range of ψ, χ, and m values, signifying a wide variety of scattering mechanisms and target geometries. To illustrate this, we have created plots of χ vs. ψ and χ vs. m, with the ocean pixels shown as a density plot as in section 5.5, iceberg pixels plotted with green + symbols, and ship pixels (where present) plotted with black + symbols. Plots for all of the scenes and imaging modes in this study are shown in Appendix B. Here we show a plot for a low incidence angle scene (0905-1) and a high incidence angle scene (0818-1), using only the low resolution

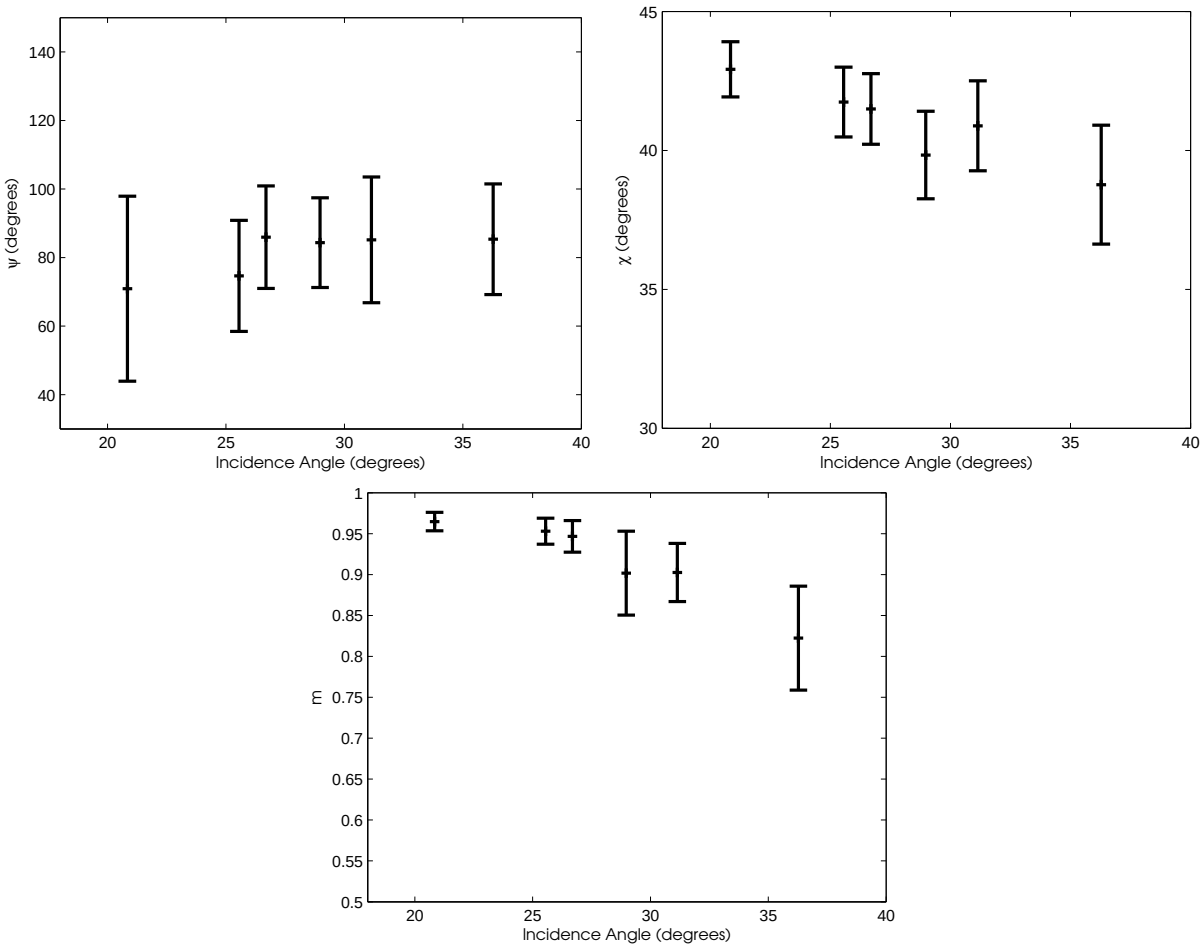

Figure 5.6: The mean values of ψ, χ, and m vs. incidence angle for the ocean chips in each scene, using right-circular transmit polarization. Multiple scenes with the same incidence angle have been combined. Error bars show standard deviation.

imaging mode, for brevity.

Figure 5.8 shows plots of χ vs. ψ and χ vs. m for scene 0905-1, low resolution imaging mode, for each transmit polarization. The values of ψ and χ appear largely as expected, with linear transmit polarizations resulting in mostly linear backscatter from ocean, ships, and icebergs. For right-circular transmit polarization, the ocean and icebergs result in mostly circular backscatter, while the ship target produces less circular backscatter (χ values closer to zero). Worth noting is that for right-circular polarization, there is a larger difference in m between the targets and the ocean than for linear polarization. The ship target in particular

exhibits low m values, denoting a mostly depolarized wave.

Figure 5.9 shows plots of χ vs. ψ and χ vs. m for scene 0818-1, low resolution imaging mode, for each transmit polarization. This is the highest incidence angle scene in the study, and contains the most iceberg targets (six) of any single scene. We see here that the iceberg pixels exhibit a wide array of ψ and χ values, and the ocean pixels similarly cover a larger area of the ψ-χ-m space than at lower incidence angles. Here the targets are less clustered around the ocean pixels, particularly for the linear polarization case, though it is difficult to say if this is due to the change in incidence angle or due to differences in the target geometry. The target pixels in the right-circular polarization case have a wider range of values than in the linear polarization case. This is a positive result, since it suggests that the targets would be easier to distinguish from the ocean when using right-circular polarization as opposed to a linear transmit polarization. While this analysis is mostly qualitative, we will see that the values of ψ, χ, and m can be used to generate an ocean model which can be utilized quantitatively to improve detection performance in Chapter 9.

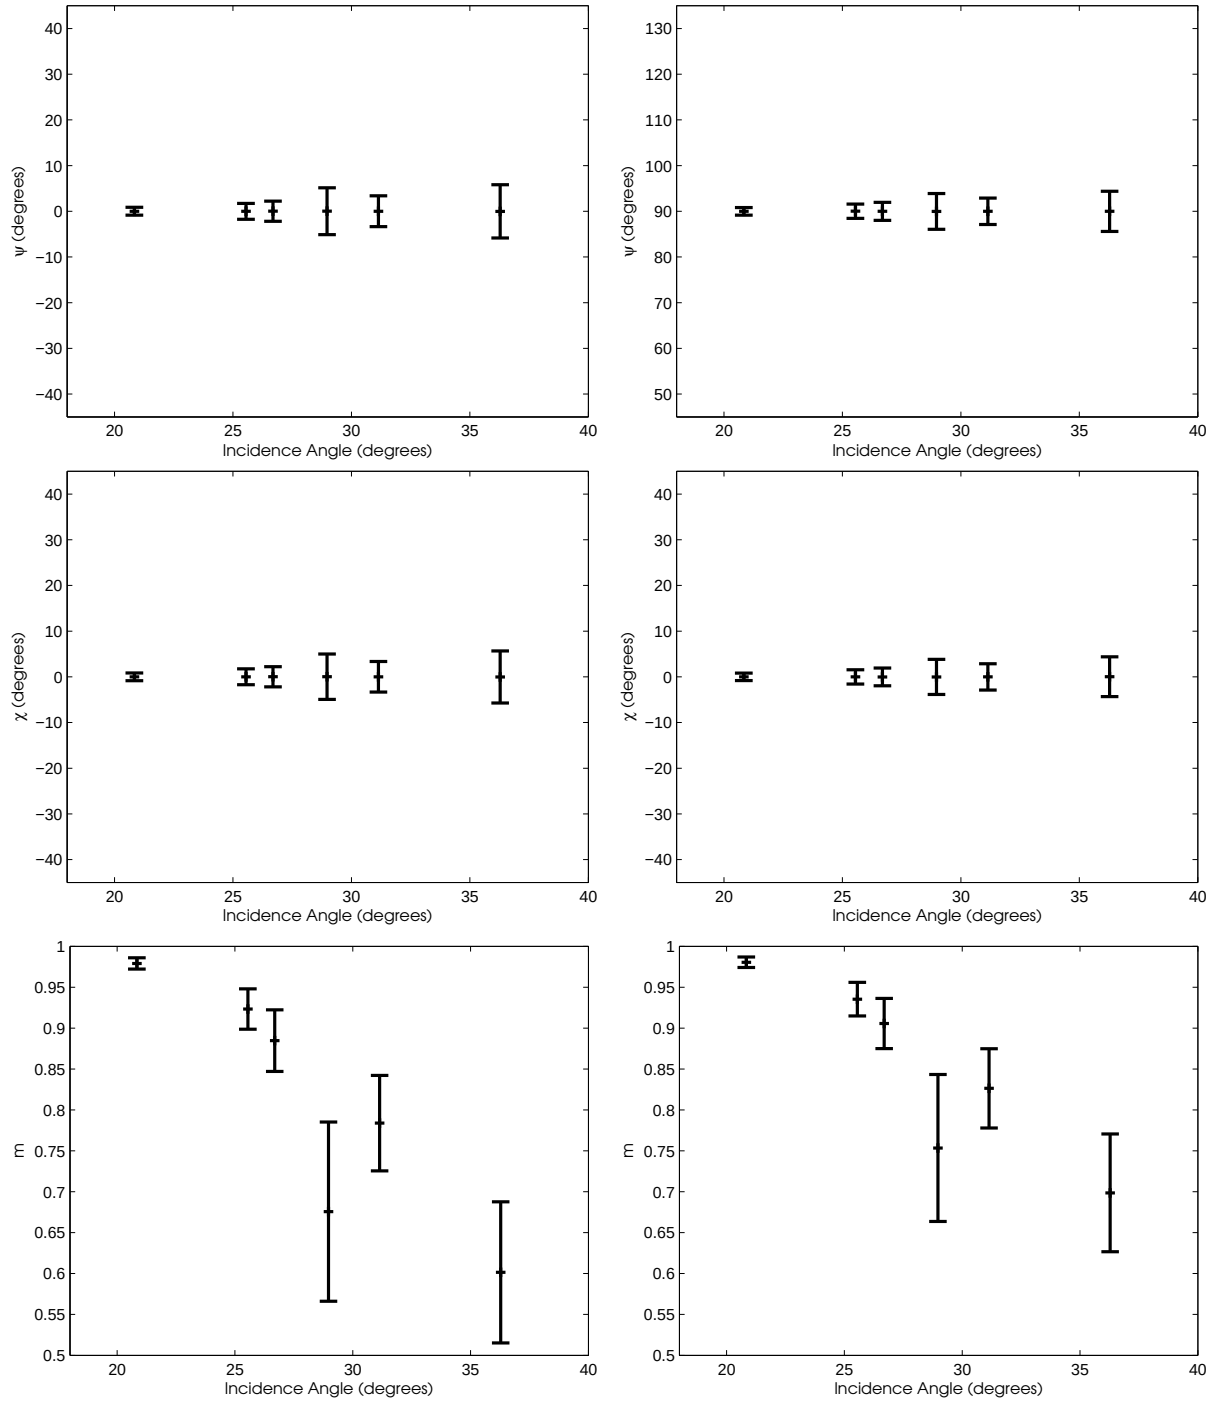

(a) HH-HV

(b) VV-VH

Figure 5.7: The mean values of ψ, χ, and m vs. incidence angle for the ocean chips in each scene, using horizontal polarization (left column) and vertical polarization (right column). Multiple scenes with the same incidence angle have been combined. Error bars show standard deviation.

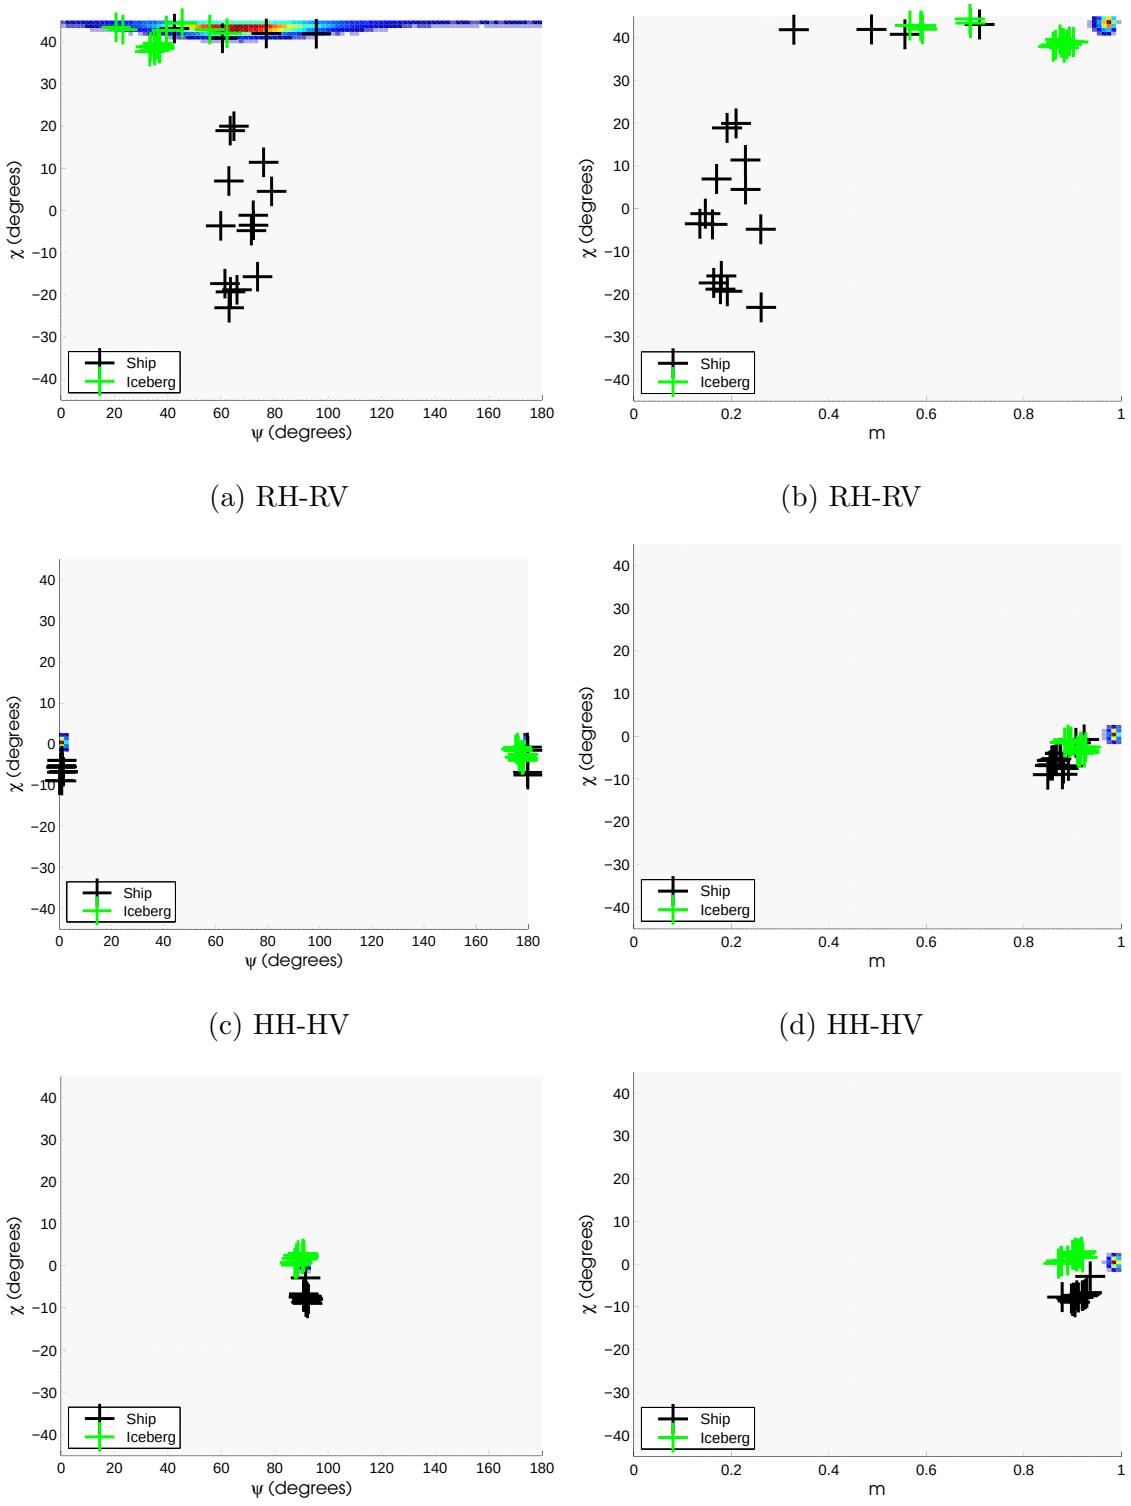

(a) RH-RV

(b) RH-RV

(c) HH-HV

(d) HH-HV

(e) VV-VH

(f) VV-VH

Figure 5.8: Plot of χ vs. ψ (left column) and χ vs. m (right column) for scene 0905-1, low resolution mode.

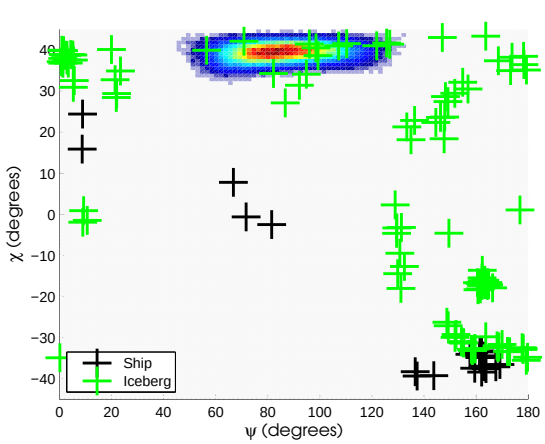

(a) RH-RV

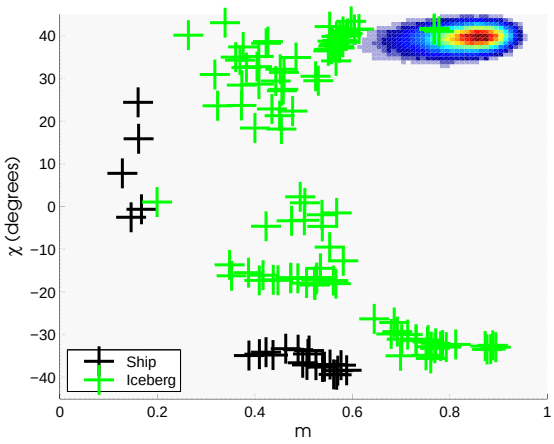

(b) RH-RV

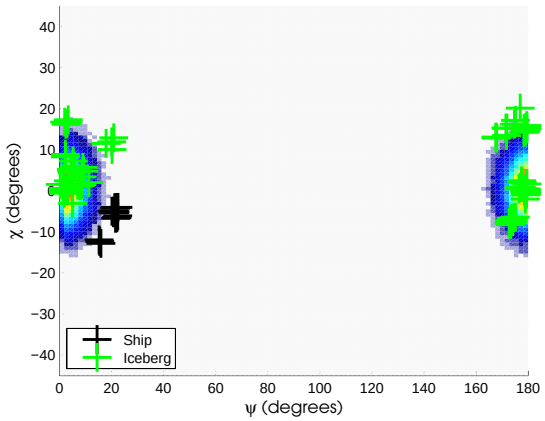

(c) HH-HV

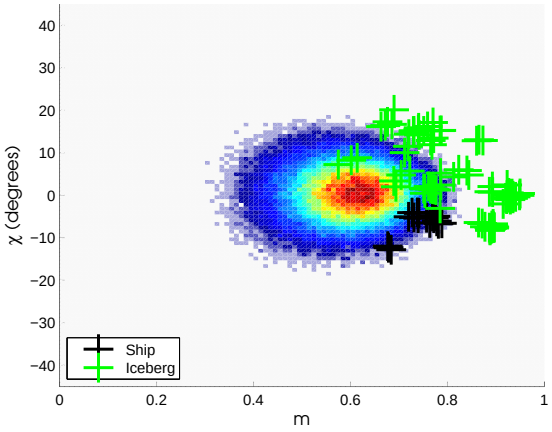

(d) HH-HV

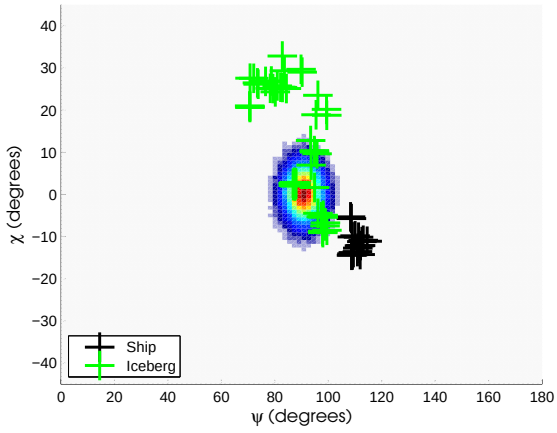

(e) VV-VH

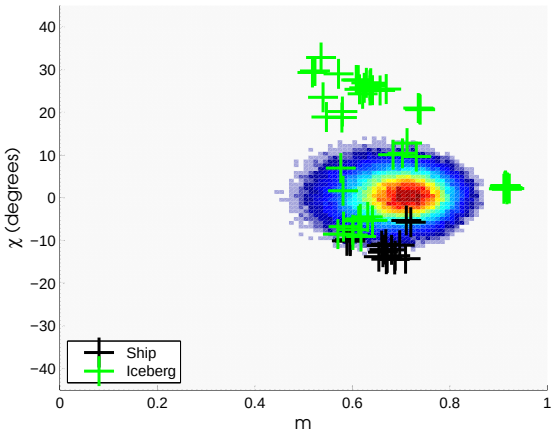

(f) VV-VH

Figure 5.9: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0818-1, low resolution mode.

Chapter 6

Ocean Target Detection In SAR Images

In very general terms, target detection in the SAR context involves discriminating between desired targets and background clutter in a SAR image. Detection problems can be approached similarly to classification problems, but whereas classification generally involves the use of a number of different possible classes, detection uses only two classes: target and non-target. When performing ocean target detection, classification between different types of targets (such as between icebergs and ships) is often performed as a later stage in the process. This is because the contrast between different types of target is often less than the contrast between the targets and the background. The backscattered power measured by a SAR is, in general, significantly less from the ocean than it is from any type of target, whether that target is a ship, a buoy, or ice.

For target detection to be successful, it is necessary for contrast to exist between the target and the background. Generally speaking, the greater the contrast, the greater the detection accuracy. When working with a single-channel SAR system, the idea of contrast between the targets and the background is fairly simple, conceptually — the target pixels should have a much greater backscattered power (σ^0, for example) than the ocean. For a single-channel image, we can set a certain intensity threshold, and classify all pixels above that threshold as targets, and all pixels below that threshold as background. This is the simplest target detection method possible, and is essentially the basis for the constant false alarm rate (CFAR) detection technique, which will be explained in further detail below. All a CFAR algorithm does is adjust the detection threshold for each pixel in the image, based on the statistical properties of the clutter surrounding that pixel, in order to obtain a desired false alarm rate (FAR).

However, when working with polarimetric SAR data, where we have more than one channel available, the detection problem becomes more complicated. When working with multiple channels of complex data, how do we define a single detection threshold? How do we define the contrast between the targets and the background clutter?

Since polarimetric SAR provides multiple channels (of often complex values), we either need to combine the different channels into a single image, or develop a detection algorithm that can use multiple channels as input. The advantage to combining the multiple channels into a single image (often referred to as a "decision variable" image) is that simpler detection algorithms, which only take a single channel as input (such as the constant false alarm rate, or CFAR, detector) can then be used. The disadvantage to this approach is that if the decision variable calculation does not accurately model the statistical distribution of the data, information can be lost in the process of converting the multiple polarimetric channels into a single value for each pixel.

There are a number of ways to calculate the decision variable, depending on how the statistics of the SAR ocean clutter are modelled, and how the polarimetric channels are used as input (e.g., is single-look data used as input, or a multi-look covariance matrix?). One decision variable calculation method that is well established in the literature is the work of Liu et al. [30, 31], who used the likelihood ratio test (LRT) to derive a decision variable from polarimetric SAR using the scattering vector $\vec{k}$ and the covariance matrix $\mathbf{C}$ of the ocean pixels, $\mathbf{C}_o$.

6.1 The Liu et al. Likelihood Ratio Test Decision Variable

The likelihood ratio test method was proposed by Liu et al. [30, 31], and is a process with which to combine different polarimetric channels into a single decision variable using the Neyman-Pearson criteria. We can quantify the detection performance at a particular threshold (i.e., pixels with a decision variable above the threshold are considered targets)

by calculating the probability of missed detection and probability of false alarm for that decision variable value.

The likelihood ratio is a general concept that has seen widespread use for many detection applications, even outside of the SAR context, and is defined as [30]:

$$\Lambda = \frac{P_T(\vec{U})}{P_O(\vec{U})} \tag{6.1}$$

L is the likelihood ratio, $\vec{U}$ is a vector of measured variables, P_T is the probability density function of the target, and P_O is the probability density function of the ocean. The greater the value of L for a given pixel, the more likely a given pixel represents a target. In the SAR context, $\vec{U}$ can be the scattering vector of a SAR system, $\vec{k}$. The likelihood ratio test, using a detection threshold equal to η, is defined [30] as:

$$\Lambda = \begin{cases} > \eta, & \text{for a target} \\ \\ \leq \eta, & \text{for ocean} \end{cases} \tag{6.2}$$

Liu et al. use a Gaussian distribution to approximate the probability density functions for both the target and the ocean. This is done because generally P_T is unknown, and P_O, while often modelled using a K-distribution in the literature (see the next chapter, for example), is more computationally complicated than desired. Use of the Gaussian distribution greatly simplifies the process, and while the Gaussian is not a perfect fit for the complex elements of the scattering vector $\vec{k}$, as discussed in section 3.3, the K distribution essentially begins with Gaussian scattering coefficients and merely multiplies them by the square root of a Gamma distributed texture variable. So the Gaussian distribution is a reasonable approximation for the data, provided that the variation in backscattering behaviour of the ocean from pixel-to-pixel is driven by speckle noise, which is generally the case [29].

A zero mean complex vector $\vec{k}$ that is Gaussian distributed has the following probability density function:

$$P(\vec{k}) = \frac{1}{\pi^p |\mathbf{C}|} \exp\left(-\vec{k}^{*T} \mathbf{C}^{-1} \vec{k}\right) \tag{6.3}$$

Where p is the number of elements in $\vec{k}$, $\mathbf{C}$ is the covariance matrix of $\vec{k}$, and the superscript $*T$ represents the complex conjugate transpose.

If we take the log likelihood of L, and absorb all constant terms (e.g., any terms which do not depend on $\vec{k}$) into the threshold calculation, we end up with the following simpler calculation of the likelihood ratio:

$$\vec{k}^{*T}(\mathbf{C}_o^{-1} - \mathbf{C}_t^{-1})\vec{k} = \begin{cases} > \eta, & \text{for a target} \\ \leq \eta, & \text{for ocean} \end{cases} \tag{6.4}$$

$\mathbf{C}_o$ and $\mathbf{C}_t$ are the covariance matrices for ocean and target pixels, respectively. This equation is difficult to use directly, since it requires us to know C_t in advance, which requires us to have knowledge of the targets we wish to detect before we actually detect them. Liu et al. [30] simplified the equation by pointing out that generally, the covariance matrix for the target pixels is significantly greater than for the ocean pixels. Therefore, the inverse covariance matrix for the target pixels is significantly *less* than for the ocean pixels, and can be ignored. This allows us to approximate the decision variable for a multi-channel, coherent SAR system using the following equation[30, 31]:

$$DV = \vec{k}^{*T} \cdot \mathbf{C}_o^{-1} \cdot \vec{k} = \begin{cases} > \eta, & \text{for a target} \\ \leq \eta, & \text{for ocean} \end{cases} \tag{6.5}$$

$\vec{k}$ is the column scattering vector of the system, the supercript $*T$ is the complex conjugate transpose, and $\mathbf{C}_o^{-1}$ is the inverse covariance matrix of the background (ocean) pixels. Note that we do not account for the covariance matrix of the target (e.g., ship or iceberg) pixels. This is a simplification to increase the usability of the process. Since the target's covariance matrix will generally be much greater than that of the ocean (and therefore the inverse covariance matrix will generally be much smaller than that of the ocean), this has a negligible affect on the results [30, 31]. When calculating the decision variable, $\vec{k}$ will be different for each pixel, while $\mathbf{C}_o^{-1}$ is generally calculated once per scene or area (using a relatively

homogeneous image chip of ocean pixels). In this thesis, an approximately 500×500 pixel image chip from each scene was used to calculate $\mathbf{C}_o^{-1}$ for each scene, although the exact size of the image chip used is not particularly important. As long as the image chip's mean covariance matrix is similar to the mean covariance matrix of the scene's ocean pixels, the LRT should produce reasonable results.

For a non-coherent system (i.e., amplitude only, without meaningful phase information), the decision variable is calculated by dividing the intensity of each channel by the mean intensity of the ocean pixels for that channel, then summing the result. For example, for a system with non-coherent HH and HV channels, we can calculate the decision variable using the equation:

$$DV = \frac{|S_{HH}|^2}{E(|S_{HH}|^2)_o} + \frac{|S_{HV}|^2}{E(|S_{HV}|^2)_o} \tag{6.6}$$

$E(\cdots)_o$ is the expectation (mean) operator, applied to the ocean pixels only (e.g., applied to the same image chip that was used to estimate the ocean covariance matrix in the coherent case).

SAR systems are generally either coherent or non-coherent in their entirety — i.e., if one channel has phase information, so do the others. This is not the case with reconstructed pseudo quad-pol data. We have the HH-VV phase, but no HV phase information. We can calculate the decision variable for the pseudo quad-pol data as follows:

$$DV = (\vec{k}^H \cdot \mathbf{C}_o^{-1} \cdot \vec{k}) + \frac{|S_{HV}|^2}{E(|S_{HV}|^2)_o} \tag{6.7}$$

$\vec{k}$ is a two element scattering vector containing the pseudo HH-VV information, and $|S_{HV}|^2$ is the pseudo-HV intensity—simply add the decision variables for a coherent and non-coherent system together. Since pseudo quad-pol reconstruction gives us only the HH-VV phase, rather than the absolute phases of S_{HH} and S_{VV}, we set the HH phase to zero when forming our pseudo quad-pol scattering vector, then set the VV phase to satisfy the HH-VV phase calculated by the reconstruction.

Note that $\vec{k}$ in equation 6.7 does not necessarily have to contain the pseudo HH-VV information. A scattering vector of the RH and RV channels (the channels directly measured by a CTLR SAR) alongside the pseudo-HV can also be used, forming a hybrid CTLR/pseudo-HV system.

Since the pseudo quad-pol reconstruction process uses a covariance matrix as input (which has been spatially averaged) rather than the original quad-pol scattering vector, the pseudo quad-pol data is also spatially averaged by extension. If this is not accounted for, the pseudo quad-pol data will outperform all other modes [10]. Spatial averaging of a dataset's Hermitian (complex conjugate) products (such as when we average the CTLR covariance matrix at the beginning of the reconstruction process) improves detection performance, since this essentially smooths the speckle, with the tradeoff being that the resolution of the imagery is degraded.

In order to spatially average the non-pseudo data in a similar fashion as the pseudo quad-pol, see that equation 6.5 can be written as follows (for an n channel system):

$$DV = \sum_{i=1}^{n}\sum_{j=1}^{n}(S_j^* \cdot (C_o^{-1})_{ji} \cdot S_i) \tag{6.8}$$

S_j^* is the jth element of $\vec{k}^H$, S_i is the ith element of $\vec{k}$, and $(C_o^{-1})_{ji}$ is the element of C_o^{-1} in row j and column i.

By writing the equation in this form, it becomes clear that we can spatially average the Hermitian products of the data during the decision variable calculation as follows (where angle brackets denote spatial averaging, for an n channel system):

$$DV = \sum_{i=1}^{n}\sum_{j=1}^{n}(\langle S_j^* \cdot S_i\rangle \cdot (C_o^{-1})_{ji}) \tag{6.9}$$

For example, for a two channel system:

$$\begin{aligned} DV &= (\langle S_1^* \cdot S_1\rangle \cdot (C_o^{-1})_{11}) + (\langle S_2^* \cdot S_1\rangle \cdot (C_o^{-1})_{21}) \\ &+ (\langle S_1^* \cdot S_2\rangle \cdot (C_o^{-1})_{12}) \\ &+ (\langle S_2^* \cdot S_2\rangle \cdot (C_o^{-1})_{22}) \end{aligned} \tag{6.10}$$

Performing the necessary averaging for non-coherent data is easy by comparison—just spatially average the intensity.

In this thesis a 3×3 moving average is used due to the coarse resolution of the wide swath imaging modes, such that the spatial resolution is degraded as little as possible. There are many speckle filtering methods in the literature that are less aggressive than a simple boxcar spatial average, in that they attempt to compensate for the data so that regions of homogeneous speckle are smoothed while point targets and edges are preserved. Some preliminary testing was done using some of these filtering techniques, but these speckle filtering techniques (such as the refined Lee [38] and intensity-driven adaptive neighbourhood approaches [39]) reduced the detection performance (due to the fact that they performed less smoothing) for the data used in this study, and so they were not investigated further. It is possible they would produce superior results for very small targets which would be degraded by the 3×3 spatial averaging window used here—but most of the targets in this study spanned at least three pixels in length. Further investigation could be done in this area, but for the purposes of this thesis, a boxcar spatial average is used.

6.2 Stokes Vector Modelling Using the Skew-Normal Distribution

Previously we discussed how the Liu et al. LRT decision variable uses a Gaussian approximation for the probability distribution function of the ocean clutter. While this is a valid approximation for the polarimetric scattering vector $\vec{k}$, it is not necessarily valid for parameters derived from that vector. The Stokes vector, for example, does not appear to be normally distributed based on the investigations performed in section 5.5. If we want to perform detection using the Stokes vector, we can use the likelihood ratio test method, but we must use a probability distribution function to derive that likelihood ratio that is a better fit to the data.

Here a method is proposed that uses the skew-normal distribution to model the Stokes

vector, a distribution which was presented and developed by Azzalini [40] across a number of papers, sometimes with slightly different parametrizations or nomenclature. In general the skew-normal distribution takes the normal (Gaussian) distribution, and adds an additional skew parameter to account for skew in the data. This is helpful for characterizing data that appears somewhat close to Gaussian, but is also clearly asymmetric, like the Stokes vector elements. Here we use Azzalini's 1999 paper [40] as a reference for the necessary equations and nomenclature. First we investigate modelling the Stokes vector elements as a skew-normal distributed random variable. Then we propose a skew-normal LRT decision variable equation for use in target detection.

The skew-normal probability distribution function can be defined as [40]:

$$p(y) = 2\phi_k(y - \xi; \Omega)\Phi(\alpha^T \omega^{-1}(y - \xi)) \tag{6.11}$$

y is the skew-normal distributed random vector, k is the number of elements of y, ϕ_k is the k-dimensional normal distribution function, ξ is a vector of location parameters that serve a similar function as the mean of a normal distribution, but are not equal to the mean of the elements of y, Ω is a covariance matrix (though not the covariance matrix of y—see explanation below), Φ is the one-dimensional normal cumulative distribution function, α is a vector of skew parameters, and ω is a diagonal matrix of scale parameters (where each diagonal element is similar to the standard deviation, though not equal to it, and the off diagonal elements are equal to zero).

Ω is a covariance matrix calculated by subtracting ξ from y, rather than the actual mean of y:

$$\Omega = \frac{1}{n}\sum_{i}^{n}(y_i - \xi_i)(y_i - \xi_i)^T \tag{6.12}$$

where n is the number of points used in the covariance matrix calculation.

Note that equation 6.11 for the skew-normal PDF contains the function ϕ_k, the normal PDF. In the event of zero skewness ($\alpha = 0$), the normal CDF $\Phi(0)$ is equal to 0.5, and

therefore the skew-normal PDF reduces to the normal PDF. The skew-normal distribution therefore contains the normal distribution as a special case.

Before delving into multivariate skew-normal modelling, let us first examine each of the Stokes vector elements individually, through 1-D fitting of the skew-normal distribution.

For the one-dimensional case, the vectors and matrices reduce to scalars, and the mean of the distribution can be easily calculated using the formula [40]:

$$\mu = \xi + \omega\delta\sqrt{\frac{2}{\pi}} \tag{6.13}$$

where δ is related to the skew parameter α and is defined as [40]:

$$\delta = \frac{\alpha}{\sqrt{1+\alpha^2}} \tag{6.14}$$

Similarly, the variance of a single variable skew-normal distribution can be calculated using the formula [40]:

$$\sigma^2 = \omega^2\left(1 - \frac{2\delta^2}{\pi}\right) \tag{6.15}$$

Using these equations, for a given value of the skew parameter α, ξ and ω can be calculated such that the mean and variance of the fitted skew-normal distribution and the sample data are equal. As an experiment, before working with multivariable fitting, 1D fitting was performed by varying α between -1 and 1, in increments of 0.01. For each α appropriate ξ and ω values were calculated so that the fitted skew-normal distribution would match the mean and variance of the sample data. The results of this fitting are shown in Figures 6.1 and 6.2, for scenes 0818-1 and 0905-1, respectively, in the low resolution imaging mode. A logarithmic y-axis is used in order to better demonstrate the fit of the distributions at the tails. For scene 0818-1, the skew-normal visually appears to better fit the data for the S_0 and S_3 parameters. For the S_2 parameter, the skew-normal and normal distributions are almost identical—there is no skew estimated from the data. For the S_1 parameter, the skew-normal distribution has a better fit for the negative tail, but a worse fit for the positive tail. The results are fairly similar for scene 0905-1, with the exception that both S_1 and S_2

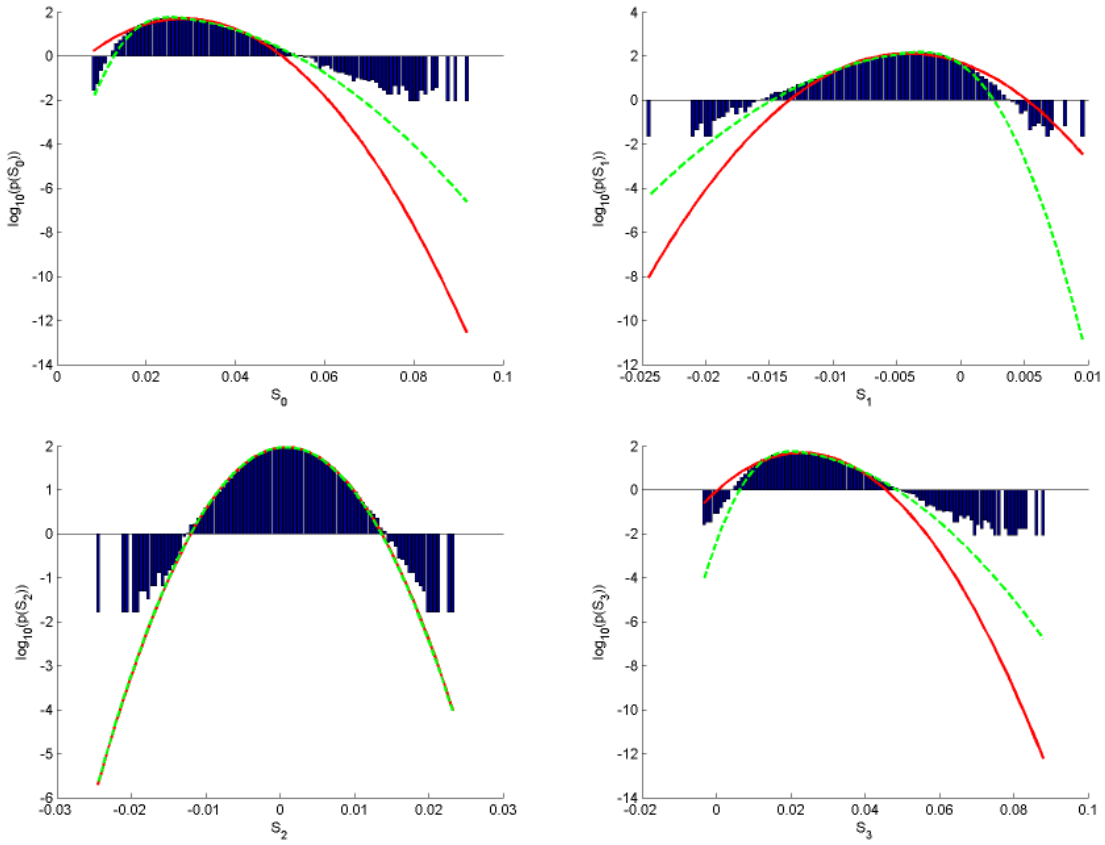

Figure 6.1: Log histograms of the four Stokes vector elements, for scene 0818-1, low resolution mode. Fitted normal distributions shown in red. Fitted skew-normal distributions shown in dashed green.

have minor but noticeable skew in this case, such that the normal distribution appears to be a slightly better fit for these parameters than the skew-normal distribution. However, the skew-normal distribution still corresponds more closely to the observed behaviour of the S_0 and S_3 parameters than the normal distribution. Worth noting is that the values of S_0 and S_3 are much higher than than that of S_1 and S_2—when transmitting a circular polarization, most of the received backscatter is also circularly polarized.

While 1-D fitting is a useful experiment, eventually we need to work with the full Stokes vector and perform multivariate fitting in 4-D. When working with multivariate data, the parameters can be estimated by maximizing the following log-likelihood equation, as discussed

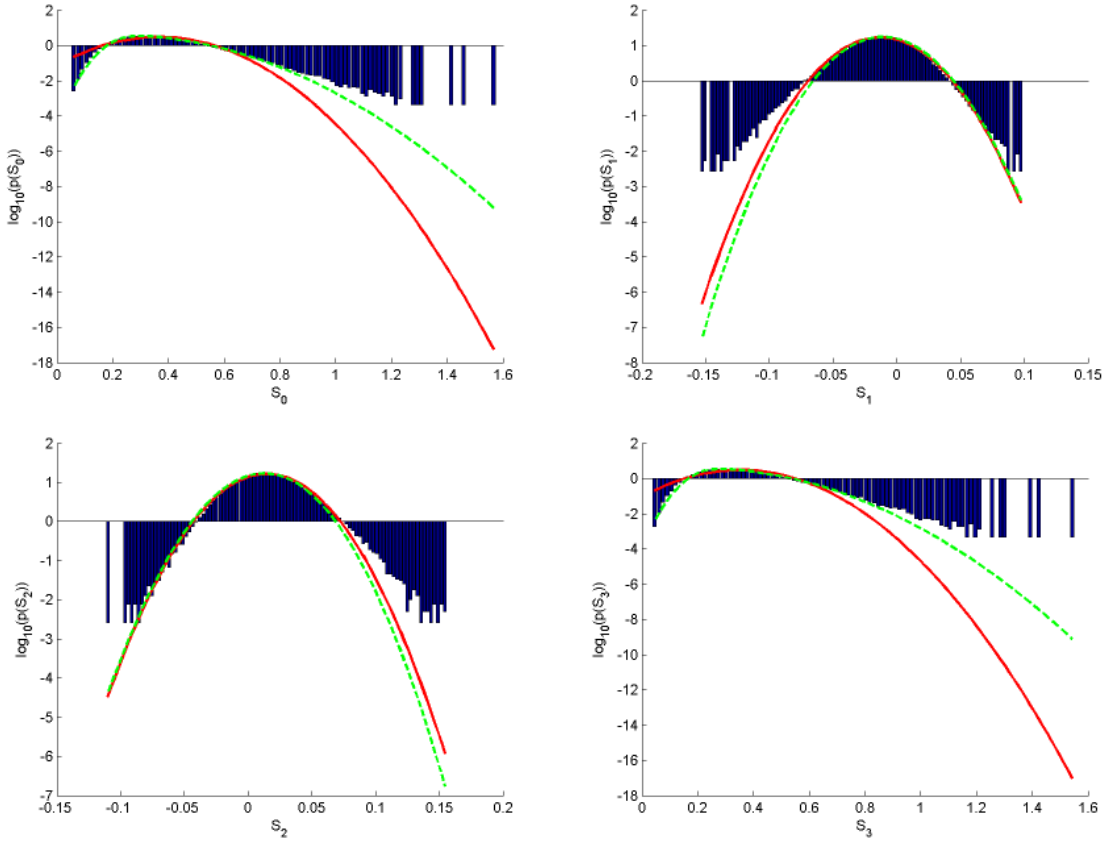

Figure 6.2: Log histograms of the four Stokes vector elements, for scene 0905-1, low resolution mode. Fitted normal distributions shown in red. Fitted skew-normal distributions shown in dashed green.

by Azzalini [40]:

$$l^*(\xi,\eta) = -\frac{1}{2}n\log|\Omega| - \frac{1}{2}nk + 1_n^T\zeta_0(u\eta) \tag{6.16}$$

y is an $n \times k$ matrix containing the data to be fitted (in this case, each row is a different pixel, for n pixels, and each column is a different Stokes vector element, for $k = 4$ elements), and ξ is similarly an $n \times k$ matrix containing the location parameters for each element of y, so that $u = y - \xi$.

Ω is then estimated as follows:

$$\Omega = \frac{1}{n}u^T u \tag{6.17}$$

where $|\Omega|$ is the determinant of Ω.

ζ_0 is defined by Azzalini as the function:

$$\zeta_0(x) = \log(2\Phi(x)) \tag{6.18}$$

applied separately to each element of x.

We are optimizing for the parameters ξ and η, where $\eta = \omega^{-1}\alpha$. Combining ω and α in this fashion reduces the number of variables involved in the optimization, resulting in a simpler optimization process. Since both ξ and η have a value for each of the k elements, this means we are optimizing the log-likelihood function for $2k = 8$ different variables in the full Stokes vector case.

The optimization was performed using the fminsearch() function within Matlab (instead of maximizing l^*, $-l^*$ was minimized). Since this function can potentially find local maximums rather than global ones, a wide variety of starting values for ξ and η were tested, which resulted in only minor differences in the final result. For the results shown here, the fitted 1D parameters were used as the starting estimates for the multidimensional ξ and η.

Since it is difficult to show the results of the fitting in four dimensions, the plots presented here show the fitting for each possible 2-D pair of Stokes vector elements, in the form of density plots (where the most likely values are shown in red, and less likely values in blue).

Detection results using the skew-normal LRT, however, will use 4-D fitting to produce the resulting decision variable.

Density plots showing the behaviour of the data and fitted skew-normal distributions of scenes 0818-1 and 0905-1, respectively, for the low resolution imaging mode, are shown as Figures 6.4 and 6.4.

For scene 0818-1, the S0-S3 plot has a shape that is difficult to represent using the skew-normal, due to the fact that the data cannot lie below the $y = -x$ line, due to the relationship between S_0 and S_3. Because the fitting process attempts to fit this line using the skew factor α, the other skew present in the data cannot be accounted for.

The distributions for the steep scene, 0905-1, are clustered together more closely, and therefore visually appear to be a better fit for the skew-normal distribution. For both of the scenes, though, the skew-normal appears visually to be a better fit than a normal distribution. Of course, regardless of the fit, for this application the real test of the skew-normal distribution is whether it can produce better detection performance than the normal LRT for the Stokes vector data.

6.3 A Stokes Vector Decision Variable Formulation

Similar to how Liu et al. [30] took the normal distribution and formulated a decision variable from it, here an equation for a skew-normal decision variable is derived, calculated from the parameters of a skew-normal distribution fitted to the full 4-D Stokes vector as described in the previous section. To begin we take 1 divided by the skew-normal PDF, so that the greater the probability of a pixel being from the ocean, the smaller the resulting decision variable, and vice versa. Like Liu et al., we perform no calculations regarding the expected PDF of the target, since in general the scattering behaviour, sizes, etc., of icebergs can be very unpredictable. Instead we focus on detecting pixels which are statistically very unlikely to be generated by normal ocean backscatter, as modelled using an image chip of ocean

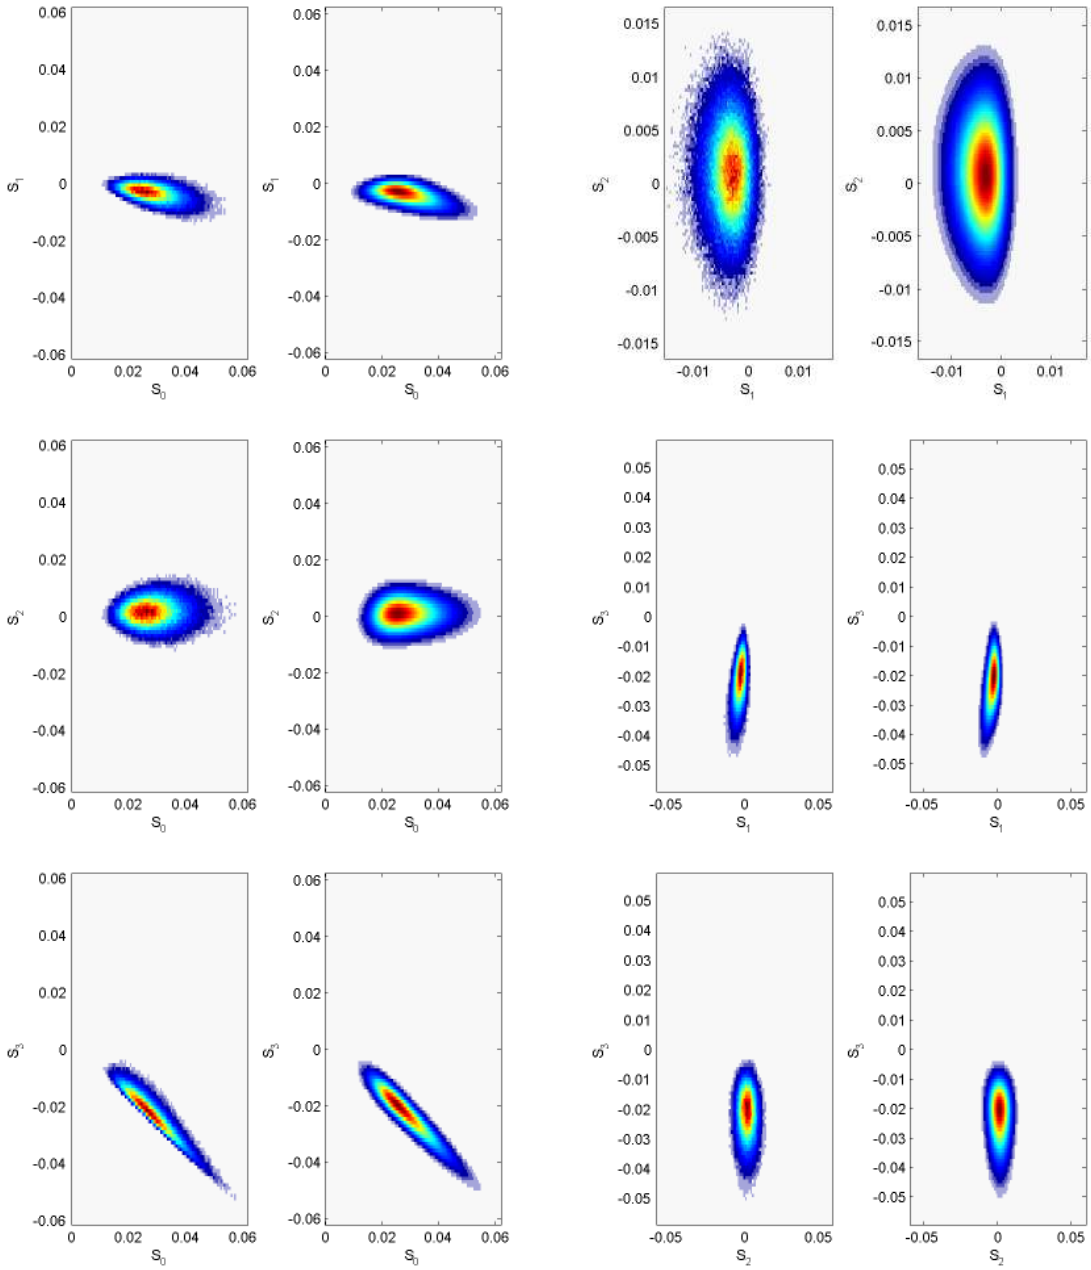

Figure 6.3: Density plots of the four Stokes vector components, for scene 0818-1, low resolution mode. For each pair of plots, the density plot of the ocean pixels is shown on the left, with the fitted skew-normal density function shown on the right.

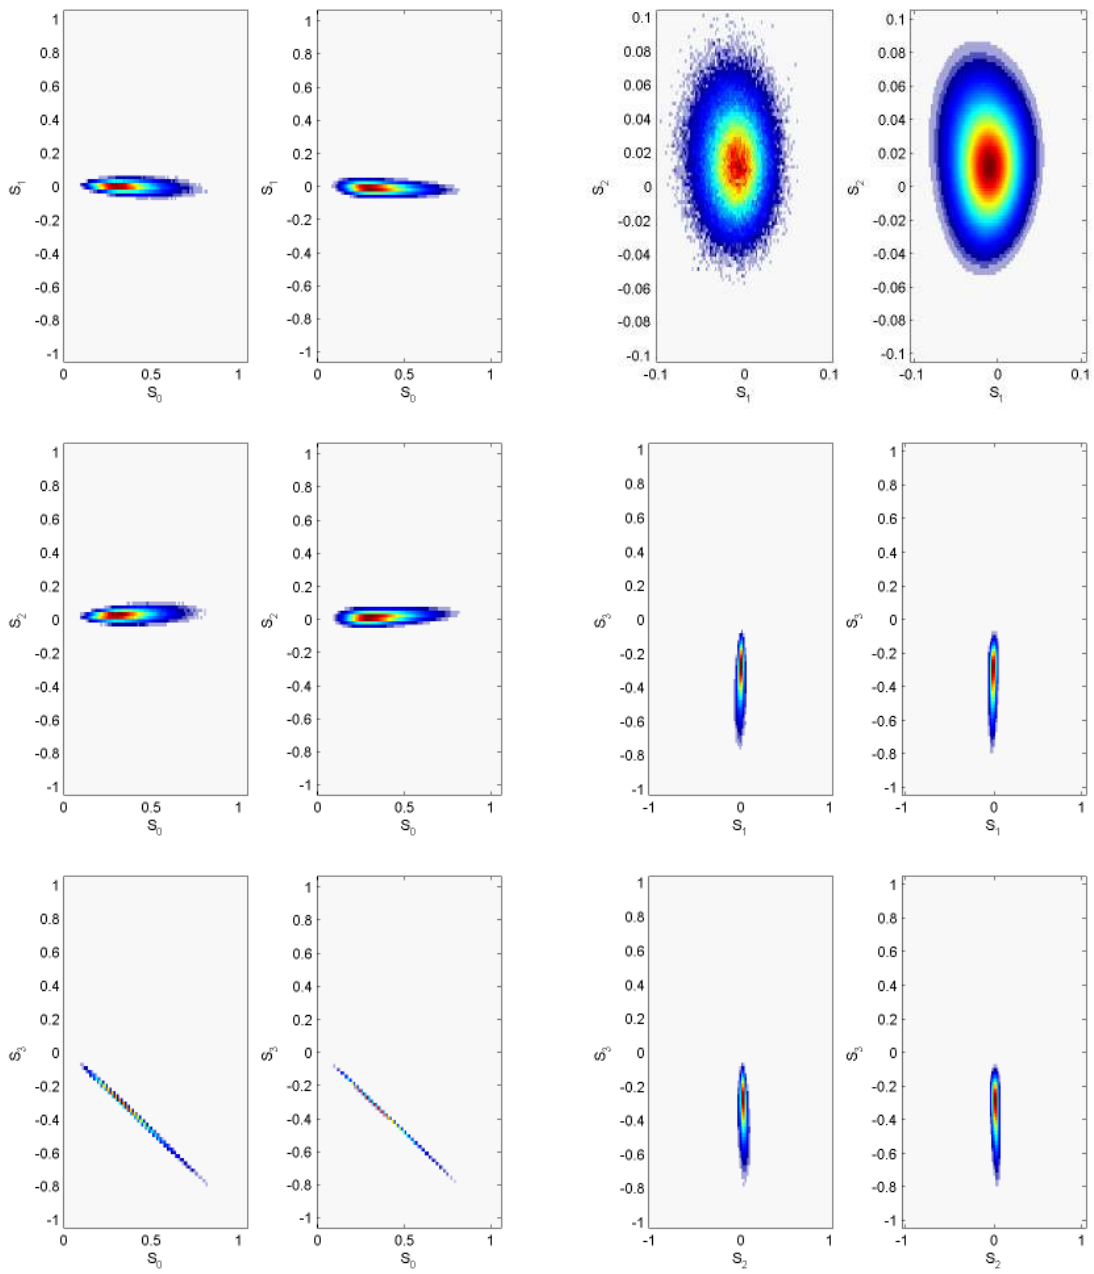

Figure 6.4: Density plots of the four Stokes vector components, for scene 0905-1, low resolution mode. For each pair of plots, the density plot of the ocean pixels is shown on the left, with the fitted skew-normal density function shown on the right.

pixels.

$$DV = \left[2(2\pi)^{-k/2}\left|\Omega\right|^{-1/2}\exp\left(-0.5(y-\xi)^T\Omega^{-1}(y-\xi)\right)\Phi(\alpha^T\omega^{-1}(y-\xi))\right]^{-1} \tag{6.19}$$

We absorb the constant terms into the decision variable value itself:

$$DV = \left[\exp\left(-0.5(y-\xi)^T\Omega^{-1}(y-\xi)\right)\Phi(\alpha^T\omega^{-1}(y-\xi))\right]^{-1} \tag{6.20}$$

Then we take the logarithm:

$$DV = -\log\left(\exp\left(-0.5(y-\xi)^T\Omega^{-1}(y-\xi)\right)\right) - \log\left(\Phi(\alpha^T\omega^{-1}(y-\xi))\right) \tag{6.21}$$

And simplify:

$$DV = 0.5\left((y-\xi)^T\Omega^{-1}(y-\xi)\right) - \log\left(\Phi(\alpha^T\omega^{-1}(y-\xi))\right) \tag{6.22}$$

Note that the first term is very similar to the normal LRT decision variable, but with the addition of the second skew term. As well, Ω is not necessarily equal to the covariance matrix of the ocean image chip, though it will generally be similar (for small values of ξ). Note that for data with no skew, $\alpha = 0$ and $\xi = 0$, and the skew-normal LRT decision variable reduces to the normal LRT decision variable. In chapter 8, the detection performance of the normal (Gaussian) LRT and skew-normal LRT for the Stokes vector data will be assessed and compared.

Chapter 7

Empirical Clutter Modelling

To actually perform detection on the calculated decision variable images, we must determine for each pixel in that image the desired detection threshold η. As mentioned previously, pixels with decision variable values greater than η are considered detected, while pixels with values below or equal to η are considered background. Generally the user specifies a desired false alarm rate, and the value of η corresponding to that false alarm rate must be determined and applied to the image. There is always a tradeoff between the false alarm rate or probability of false alarm (PFA) and the probability of missed detection (PMD). The higher the PFA the lower the PMD, and vice versa. Both the PFA and PMD for a particular detector depend on the detection threshold chosen, and by plotting PMD vs. PFA values for a range of detection thresholds, we can generate a plot referred to as the receiver operating characteristic (ROC) [41].

There are two potential ways to calculate the PFA value corresponding to a given detection threshold (or vice versa). We can do so using a statistical model for the ocean clutter (e.g., a speckle model in the case of SAR imagery, since as discussed near the end of section 3.3, the pixel-to-pixel variation in backscatter for ocean pixels is driven by speckle), or by calculating the PFA for a given detection threshold empirically, using the following equation:

$$\text{PFA} = \frac{N_{FA}}{N_o} \tag{7.1}$$

N_{FA} is the number of false alarms within an area of ocean containing N_o total ocean pixels. Empirical PFA values are simple to compute, but are limited by the large number of ocean samples necessary to calculate the very small PFA values needed for detection. For example, a 1000×1000 pixel image chip contains $N_o = 1,000,000$ pixels and can therefore only be used to calculate PFA values as low as 10^{-6}, when often PFA values of 10^{-9} or lower

would be desired. To calculate the detection threshold corresponding to a PFA value of 10^{-9} empirically would require one billion pixels, equivalent to a 31623×31623 pixel square image chip—clearly not practical.

However, while direct empirical calculation of PFA values is problematic, use of statistical clutter models come with their own challenges. Statistical calculation of the detection threshold for a given PFA requires the integration of a probability distribution function that has been fitted to the ocean clutter. While data from the ocean imagery is used in order to fit the distribution (i.e., the clutter model) to the data, once a particular distribution is fitted, detection thresholds can be calculated for any desired PFA value. As discussed in section 3.3, a common distribution used to model ocean clutter in SAR images is the K distribution. Liu et al. [31] used the K distribution to model LRT decision variables for both linear dual-pol and compact polarimetric SAR data, using the following form for $p(x)$, the probability of the K distributed random variable x:

$$p(x) = \frac{2}{x}\left(\frac{L\nu x}{\mu}\right)^{\frac{L+\nu}{2}} \frac{1}{\Gamma(L)\Gamma(\nu)} K_{\nu-L}\left(2\sqrt{\frac{L\nu x}{\mu}}\right) \tag{7.2}$$

L is the equivalent number of statistically independent looks, ν is the shape parameter, μ is the mean of x, Γ is the Gamma function, and $K_{\nu-L}()$ is the modified Bessel function of the second kind, of order $\nu - L$. Unfortunately there is no closed form integral of this equation (and no cumulative distribution function), and integration must therefore be done numerically. L is unknown, since the decision variable is calculated using a combination of multiple single-look polarimetric channels (which are not necessarily statistically independent), and therefore must be estimated along with ν and μ. μ can be estimated in a straight forward manner using the sample mean, but L and ν are more problematic. Liu et al. [31] estimated L and ν by calculating the K distribution for a wide variety of parameter values, then choosing the pair of L and ν that both minimized the mean square error between the K PDF and the data histogram, and passed a χ^2 goodness-of-fit test. The fitted K distribution was then integrated numerically to calculate the detection threshold corresponding to the

desired PFA value. Needless to say, this procedure is expensive computationally.

However, for target detection applications, it is not strictly necessary to fit the entire K distribution to the data histogram. What is of interest to us is that the detection threshold as a function of the PFA (which is related to the cumulative distribution function, not the PDF) can be calculated. As well, since we are only interested in the detection threshold for very small PFA values, we can ignore the peak of the distribution and focus directly on the tail probability.

Rather than attempting to fit the PDF of the K distribution to the data, an alternate method is proposed. First, let us investigate the behaviour of numerically integrated K distribution tail probabilities at high detection thresholds. Plots of log_{10}(PFA) versus the detection threshold are shown in Figure 7.1. The left plot shows the behaviour of the K distribution for a variety of ν values, $L = 4$, and $\mu = 1$. As the detection threshold increases, the behaviour of the curves grows increasingly linear. This behaviour becomes more pronounced as ν increases (or as L increases with ν held constant, though this is not shown for brevity). The right plot shows empirical PFA values as a solid black line, for simulated ship detection mode data of scene 0818-1. Note that these values only go down to a PFA of around 10^{-6}, too high for practical detection applications. For example, consider the fact that scene 0905-1, simulated in the medium resolution imaging mode, has dimensions of 1709 by 1357 pixels, for a total area of approximately 2.3 million pixels (covering a geographic area of roughly 331 square kilometres), resulting in an expected 2.3 false alarms for a PFA value of 10^{-6}. This image only contains three validated iceberg targets, such that the number of false alarms in this case would be almost equal to the number of targets. Consider also that the Arctic Ocean has a geographic area of approximately 14 million square kilometers, such that at this PFA value, and given a similar pixel spacing throughout (which is a simplification for the purposes of this example), we would expect approximately 98000 false alarms to be detected across the entirety of the Arctic Ocean, much greater than the annual number of

icebergs calved from Greenland. On top of the plotted empirical PFA values, a fitted K distribution is shown in blue, and a fitted line in red. We fit both the K distribution and the line using evenly spaced samples from the $\log_{10}$(PFA) versus threshold curve for $\log_{10}$(PFA) values between -2 and -6, in increments of -0.2. This evenly distributed sampling accounts for the fact that we have more samples for higher PFA values than for lower PFA values. By fitting the $\log_{10}$(PFA) data rather than the data histogram, we are fitting the function we are directly interested in calculating, rather than a PDF that will then need to be integrated later.

Note that the empirical PFA behaviour has a linear trend for PFA values of approximately 10^{-2} and below, consistent with a K distribution with a relatively high ν value. The behaviour for the other RCM imaging modes and polarimetric configurations had similarly high ν values. Notably, the fitted K distribution and the fitted line have the same r^2 values (0.998) over the fitted range. Therefore, rather than using a K distribution to calculate the detection thresholds used in this study, we instead fit a line to the empirical $\log_{10}$(PFA) vs. threshold data, and extrapolate when necessary to the desired PFA values (e.g., as low as 10^{-9} or 10^{-10}). This is a valid approach provided that the decision variable values are K distributed (as shown by Liu et al. [30]). For higher PFA values where such extrapolation is unnecessary (e.g., 10^{-5}), PFA values can be taken directly from the empirically calculated data, though this will generally only result in small differences from the linear regression.

The statistics of ocean speckle can theoretically vary greatly depending on the frequency and spatial resolution of the SAR sensor. It is possible that for higher resolution modes, or data from other satellites, the ocean distribution would be less consistent with K distributed speckle, and this approach would not be viable. It is also likely that if we were to use a different speckle filtering approach (that is, not a straight boxcar average during the decision variable calculation), the ocean clutter's distribution could be changed. However, the K distribution appears to produce good results for this data, and by using our knowledge of

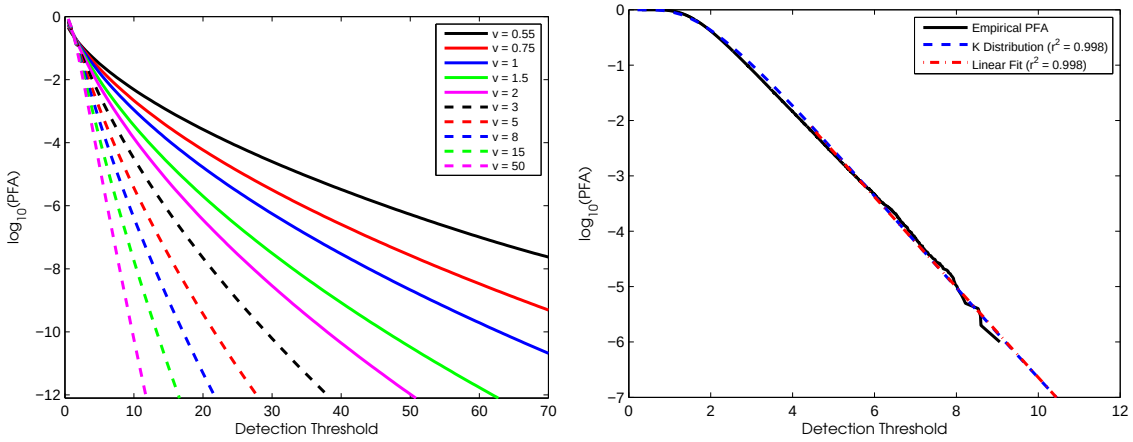

Figure 7.1: Left: $\log_{10}(\text{PFA})$ vs. threshold curves for a number of K distributions with varying ν, $L = 4$, and $\mu = 1$. Right: An empirical $\log_{10}(\text{PFA})$ vs. threshold curve from an ocean subset of scene 0818-1, for a decision variable calculated using simulated ship detection mode data of native CTLR (RH-RV). A linear fit for false alarm rates between 10^{-2} and 10^{-6} is shown in red, with equation $y = -0.815x + 1.51$. A similarly fitted K distribution is shown in blue, with $\nu = 23.75$, $L = 9.04$, and $\mu = 1.97$. The r^2 values shown in the legend are calculated over the fitted PFA range.

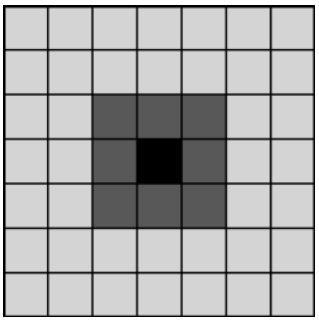

Figure 7.2: An example of the moving window used in the adaptive thresholding process. The light grey region is used to estimate the speckle statistics (and therefore the detection threshold) for the black pixel. The dark grey pixels are in the guard region, and are not used in the calculation. The window used in this study had a guard region of 31×31 pixels, with a total window size of 101×101 pixels.

the K distribution to surmise that the empirical PFA curve should have expected linear behaviour at low PFA values, we avoid having to actually fit the K distribution for every scene, which would have been a costly and potentially error-prone procedure.

7.1 Adaptive Thresholding and Clustering

Using the methods described in the previous section, we can calculate a detection threshold for any given PFA, using an ocean image chip selected from each scene (in this case, the same chip used to calculate the ocean covariance matrix for the LRT detection). However, this process yields a constant detection threshold across an area, and does not account for potential variations in ocean backscatter throughout the image. To account for some of this variation, we can slide a moving window across the image with a guard region at its centre (shown in Figure 7.2), as in constant false alarm rate (CFAR) detection [41], a method which is well established in the literature, and has been widely used to detect ships in SAR images [42, 43, 44].

In many cases, the moving window is used to recalculate the parameters of the clutter

model for every pixel in the image. However, recalculating our empirical model in this way would require a very large moving window. Consider that to calculate a PFA of 10^{-5} empirically requires a window area of at least 10^5 pixels, equivalent to a 316×316 square window. For this reason, rather than calculating a new empirical model for every pixel in the image, we instead take the sample mean of the values in the moving window (excluding the guard region), $\bar{x}_{i,j}$, and use it to scale the detection threshold for the pixel with coordinates (i, j), $t_{i,j}$, using the following equation:

$$t_{i,j} = t_0 \frac{\bar{x}_{i,j}}{\mu} \tag{7.3}$$

t_0 is the threshold calculated from the original empirical model, $\bar{x}_{i,j}$ is the sample mean of the clutter in the moving window (excluding the guard region), and μ is the mean value of the chip used to create the empirical model. This allows us to use a smaller window size for the adaptive thresholding process than if we were to recalculate the model for every pixel. Even if we were to use a K distribution to model the ocean clutter, re-estimating ν and L as the moving window travels across the image is both time consuming and inaccurate, due to the small number of samples available.

For this thesis a 31×31 pixel guard region is used, with a total window size of 101×101 pixels, to adjust the detection thresholds for each pixel in the image. Pixels with decision variable values above the threshold were then considered to be detected.

As a further step, the detected pixels in an image can be grouped into clusters, with each cluster consisting of 8-connected neighbouring detected pixels (i.e., pixels that are one pixel apart, either in range, azimuth, or both directions). A minimum cluster size can then be defined, such that clusters below this size (whether true detections or false alarms) are discarded. By increasing the minimum cluster size, the effective FAR can be reduced without actually changing the detection threshold (since both single-pixel false alarms, and single-pixel detections, for example, will be discarded). Of course, this also puts a lower limit on the size of targets that can be detected, though it should be noted that single-pixel detections

would be very difficult to verify or classify in an operational environment in any event. In this thesis, we use a minimum cluster size of four pixels for the low resolution and medium resolution imaging modes, and a minimum cluster size of eleven pixels for the ship detection imaging mode.

Chapter 8

Detection Performance Comparison of Various Polarimetric Modes

Detection was performed using a wide variety of polarimetric configurations in order to assess and compare their performance. The LRT was used to calculate a decision variable for the scattering vector of each mode as described in section 6.1, with the exception of detection using the Stokes vector, where detection was performed using both the normal LRT and the skew-normal LRT (as in section 6.3) in order to compare the results of these two methods. A 3×3 pixel spatial averaging window was used during the decision variable calculation process.

The detection thresholds were calculated for a false alarm rate of 10^{-8} as in chapter 7. This false alarm rate was low enough that for most of the modes, no false alarms were detected in any of the images. As described at the end of chapter 7, detected pixels were combined with their 8-connected neighbours into clusters, and clusters below a minimum size (four pixels for the low resolution and medium resolution imaging modes, and eleven pixels for the ship detection imaging mode) were discarded.

Due to the large number of polarimetric modes tested, in the following tables of results, shortened names for each mode are used, as follows:

HH-HV

Linear dual-pol mode using a transmitted horizontal polarization, resulting in coherent HH and HV channels. Probably the most common dual-pol mode used in the literature.

VV-VH

Coherent VV and VH channels. Another common dual-pol mode.

HH-VV

Incoherent HH and VV channels, such as in the "alternating polarization" available from Envisat. Here we test amplitude only HH-VV, since generally these HH-VV modes are either incoherent or have other limitations (e.g., in terms of swath width).

CL

Circular-Linear, or CTLR, compact polarimetry. Here the coherent RH-RV scattering vector is used directly, with no other parameters. These are the native channels collected by a CTLR SAR system.

PQ

The full pseudo quad-pol dataset reconstructed from the CTLR. The pseudo-HH and pseudo-VV are treated coherently, while the pseudo-HV has no phase information and is therefore only usable as an intensity.

PHV

Pseudo-HV intensity only.

CL-PHV

The coherent RH-RV combined with the pseudo-HV intensity.

DCP

Dual-Circular Polarization, that is, a scattering vector of RR and RL.

St.

The four Stokes vector elements used as input to the (normal) LRT.

St. SN

The four Stokes vector elements used as input to a skew-normal LRT as described in section 6.3.

Results are split into sections by the imaging mode used. For each imaging mode, we show tables of PMD values for each iceberg, with the icebergs grouped into four incidence angle categories, based on the beam mode of the original Radarsat-2 fine-quad scenes (denoted with FQ#). These four incidence angle categories are:

1. FQ2: $19.9° - 21.8°(1.9°)$,
2. FQ5/6/7: $23.4° - 27.6°(4.2°)$
3. FQ9/10/11: $28.0° - 32.0°(4.0°)$
4. FQ16: $36.5° - 38.0°(1.5°)$

The number in parentheses is the difference between the maximum and minimum incidence angle of that range.

As mentioned previously, the ship detection mode is planned to have high incidence angle beam modes only. Therefore, results are only shown here for simulated ship detection mode data of scene 0818-1, the highest incidence angle scene available, which roughly corresponds to the lowest planned incidence angle of the RCM ship detection beam modes.

For each iceberg the validation data contained the classified size category (small, medium, or large) and shape (tabular, wedge, blocky, drydock, pinnacle, or dome). The size categories are defined by the International Ice Patrol as follows [27]: between 15m and 60m in length for "small" icebergs, between 60m and 122m in length for "medium" icebergs, and greater than 122m in length for "large" icebergs.

The shape categories are defined [27]:

Tabular An iceberg with a flat top, and generally horizontal banding along the sides. These icebergs are formed from glaciers and ice shelves, and will often turn into one of the other types as they melt and break up in the water.

Wedge An iceberg with a wedge shape.

Blocky A flat-topped iceberg with steep vertical sides.

Drydock An iceberg that has a U-shaped area where the ice has melted, resulting in two columns or pinnacles on either side, giving the iceberg the shape of its namesake.

Pinnacle An iceberg with a spire or pyramid.

Dome An iceberg with a smooth, rounded top.

These iceberg shape classes are less quantitative than the size classes. Sometimes all the shape categories besides tabular icebergs will be lumped together into a single "non-tabular" class. While the shape categories are interesting, due to the wide variety of target geometries within each class, no analysis will be performed for them. However, we do report results categorized by size, in the form of tables of the median PMD values for each size class, as well as median PMD values for each incidence angle category for medium icebergs only (since medium icebergs represented the largest portion of our validation data). Note that while the data used here only contains three sizes, there are other size classes (e.g., growlers, very large, etc.) that were not contained in our dataset. In particular, icebergs less than 15 metres in length would be very difficult, if not impossible, to detect in our imagery, due to the fact that they are much smaller than the spatial resolution of the data.

While in the previous chapter we have defined the probability of false alarm and how it is related to the detection threshold, we have not defined the probability of missed detection. Here we calculate the PMD values using the equation:

$$\text{PMD} = 1 - \frac{N_D}{N_T} \tag{8.1}$$

N_D is the number of detected target pixels, and N_T is the total number of pixels in the target mask. The lower the PMD for a given target, the more pixels of that target were actually

detected. A PMD value of one corresponds to a completely missed target. As long as the PMD value is less than one, the target was detected in the image. Therefore, we can also think of the PMD as the ratio of undetected masked area (in pixels) to the total masked area. While the number of missed targets and number of false alarms for each polarimetric mode are the most conclusive ways of comparing the different modes, in a number of cases different modes produce the same number of missed targets and in these cases the PMD values can provide a more sensitive measure of the performance of each detector.

The results are in the form of tables of the PMD value of each iceberg, as well as median PMD values for each incidence angle or size category.

First, we present and discuss detection results for the low resolution imaging mode. Then, for the medium resolution imaging mode. And finally, for the ship detection imaging mode.

8.1 Low Resolution Imaging Mode Results

Table 8.1 shows the PMD values of each iceberg for the low resolution imaging mode. Three polarimetric modes are tied for the fewest number of missed targets (at five): the CTLR RH-RV, the DCP, and the Stokes skew-normal. Of these modes, we note that the CTLR and the DCP have exactly the same PMD values for every target, and indeed we find that this holds across the rest of the imaging modes. Therefore, it seems that changing the received polarization basis from linear to circular has no impact on the detection performance of a LRT detector—a result that makes sense intuitively, but had not previously been confirmed. Therefore, when we are discussing results for the CTLR mode in this section and in later sections, the analysis and conclusions also apply to the DCP mode, since the two modes are effectively the same.

The most interesting comparison here is perhaps between the CTLR and the Stokes skew-normal. Both modes miss the same number of targets (five), and have the same median PMD value (0.72). However, unlike the CTLR and DCP modes, there are some notable differences.

For the highest incidence angle category of $36.5° - 38.0°$, and the second lowest incidence angle category of $23.4° - 27.6°$, the Stokes skew-normal has a lower median PMD value than the CTLR. However, the Stokes skew-normal had a single false alarm detection, while the CTLR did not, making the CTLR overall more favourable according to this limited data set. However, we do see that use of the skew-normal LRT over the normal LRT improves the detection performance of the Stokes vector for this imaging mode.

Also of note is the fact that the compact polarimetric modes out perform the linear dual-pol modes fairly consistently, both in terms of the number of missed targets, as well as the median PMD values. The performance between the linear dual-pol and compact modes are fairly similar at the lowest incidence angle range of $19.9° - 21.8°$, but as the incidence angle increases the compact polarimetric modes, particularly the CTLR and Stokes skew-normal, appear to pull ahead in detection performance.

In previous work by the author using the Radarsat-2 fine quad mode data directly [10], the pseudo quad-pol reconstructed data seemed to perform better than the native CTLR. However, that is not the case here. One plausible explanation for this difference is due to the increased noise floor of the Radarsat Constellation compared to Radarsat-2. For the Radarsat-2 data, the noise effective sigma zero, or NESZ, was often below -30 dB, while the noise floor is -22 dB (referenced to the power of the transmitted radar pulse) for all three of the proposed imaging modes we are investigating in this thesis. While the co-pol channels will generally be above either of these noise floors, the HV (and therefore the pseudo-HV) σ^0 is generally below the noise floor of the Radarsat Constellation data for most ocean pixels. This affects the pseudo quad-pol data particularly severely because we had previously found that the pseudo-HV intensity was the most useful of the reconstructed parameters for iceberg detection [10], and with the higher RCM noise floor, the pseudo-HV parameter is dominated by noise for most of the ocean pixels in the image.

Table 8.2 shows the median PMD values for each size category. We find that for the

small icebergs, most of the compact polarimetric modes are tied for the lowest median PMD value of 0.79. In comparison the best of the linear dual-pol modes, the HH-HV mode, had a median PMD value of 0.81. For the medium icebergs both the CTLR and the combination of CTLR and pseudo-HV are tied for the lowest PMD value of 0.69. For the large icebergs, the lowest PMD value is actually held by the full pseudo quad-pol reconstructed dataset, a somewhat surprising result considering that overall it performed worse than the original CTLR. As we might expect, the small icebergs tend to have the highest median PMD values overall, implying that larger targets are easier to detect.

Table 8.3 shows the median PMD values in each incidence angle category, for the medium size icebergs only. This table attempts to eliminate the effect of iceberg size on the results, by looking only at medium icebergs. However, we are limited to a small number of iceberg targets in this case (five targets for the second incidence angle category, and three targets for each of the other categories). As before, the highest PMD values overall tend to occur for the lowest incidence angle category. It seems that very steep incidence angles have a negative effect on detection performance, as we might expect considering the fact that the lower the incidence angle, the greater the backscattered power from the ocean (and therefore, for a constant target backscatter, the smaller the contrast between the target and the ocean).

Table 8.1: PMD values for the low resolution imaging mode, categorized by incidence angle. A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			Linear Dual-Pol			Compact Polarimetry						
ID (Scene)	Size	Shape	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
FQ2: $19.9° - 21.8°(1.9°)$												
R1 (0905-1-1)	S	Wedge	0.93	0.93	1.00	0.93	0.93	0.93	0.93	0.93	0.93	0.93
R2 (0905-1-2)	S	Wedge	0.81	0.86	1.00	0.79	0.79	0.83	0.79	0.79	0.79	0.79
B1 (0905-1-3)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
H (0905-2-1)	M	Drydock	0.78	0.78	1.00	0.79	0.81	0.80	0.79	0.79	0.79	0.77
B2 (0905-2-2)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Median			0.93	0.93	1.00	0.93	0.93	0.93	0.93	0.93	0.93	0.93
FQ5/6/7: $23.4° - 27.6°(4.2°)$												
G1 (0829-1-1)	S	Drydock	0.55	0.73	0.60	0.55	0.68	0.65	0.60	0.55	0.56	0.58
G2 (0805-1-1)	S	Drydock	0.47	0.50	0.46	0.51	0.56	0.66	0.56	0.51	0.40	0.41
N (0805-1-2)	S	Wedge	0.51	0.54	0.46	0.57	0.68	0.76	0.61	0.57	0.75	0.75
A (0805-2-1)	M	Blocky	0.66	0.63	0.60	0.72	0.77	0.82	0.70	0.72	0.83	0.30
J (0815-2-1)	M	Pinnacle	0.51	0.58	0.80	0.36	0.51	0.47	0.41	0.36	0.39	0.47
D1 (0815-2-2)	M	Drydock	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
D2 (0815-2-3)	M	Drydock	0.92	0.92	1.00	0.90	0.90	0.90	0.90	0.90	0.90	0.91
M (0822-1-1)	M	Wedge	0.74	0.74	1.00	0.40	0.49	0.51	0.46	0.40	0.40	0.37
I (0822-2-1)	L	Pinnacle	0.75	0.75	1.00	0.42	0.42	0.44	0.42	0.42	0.42	0.42
P1 (0822-2-2)	L	Wedge	0.73	0.76	1.00	0.61	0.67	0.67	0.67	0.61	0.69	0.71
Median			0.69	0.73	0.90	0.56	0.67	0.66	0.60	0.56	0.63	0.53
FQ9/10/11: $28.0° - 32.0°(4.0°)$												
C1 (0815-1-1)	M	Dome	0.66	0.68	0.76	0.59	0.63	0.71	0.63	0.59	0.54	0.64
O1 (0815-1-2)	M	Wedge	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
P2 (0825-2-1)	M	Wedge	0.77	0.83	1.00	0.71	0.64	0.76	0.76	0.71	0.69	0.68
L (0825-1-1)	L	Tabular	1.00	1.00	1.00	0.53	0.50	0.60	0.53	0.53	0.43	0.53
Median			0.89	0.91	1.00	0.65	0.63	0.73	0.69	0.65	0.61	0.66
FQ16: $36.5° - 38.0°(1.5°)$												
K (0818-1-1)	S	Pinnacle	0.87	0.91	0.88	0.76	0.88	0.82	0.79	0.76	0.73	0.72
Q (0818-1-2)	S	Wedge	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
E (0818-1-3)	M	Drydock	1.00	1.00	1.00	0.84	0.89	1.00	0.90	0.84	0.82	0.79
F (0818-1-4)	-	-	0.54	0.84	0.77	0.66	0.86	0.71	0.67	0.66	0.55	0.52
O2 (0818-1-5)	M	Wedge	0.69	0.71	0.80	0.65	0.70	0.77	0.69	0.65	0.65	0.64
C2 (0818-1-6)	M	Dome	0.92	1.00	1.00	0.90	1.00	1.00	1.00	0.90	1.00	0.84
Median			0.89	0.95	0.94	0.80	0.89	0.91	0.85	0.80	0.78	0.76
Overall Median			0.78	0.84	1.00	0.72	0.79	0.80	0.76	0.72	0.75	0.72
# of Missed Targets			7	8	16	5	6	7	6	5	6	5
# of False Alarms			0	0	2	0	0	0	0	0	1	1

Table 8.2: Median PMD values for the low resolution imaging mode, categorized by iceberg size (small, medium, or large). A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Size	Linear Dual-Pol			Compact Polarimetry						
	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
Small	0.81	0.86	1.00	0.79	0.79	0.83	0.79	0.79	0.79	0.79
Medium	0.75	0.81	0.94	0.69	0.79	0.78	0.69	0.69	0.74	0.71
Large	0.77	0.83	1.00	0.71	0.64	0.76	0.76	0.71	0.69	0.68
Overall Median	0.78	0.84	1.00	0.72	0.79	0.80	0.76	0.72	0.75	0.72

Table 8.3: Median PMD values for the low resolution imaging mode, for medium icebergs only, categorized by incidence angle. A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value in each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Incidence Angle Range	Linear Dual-Pol			Compact Polarimetry						
	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
$19.9° - 21.8°(1.9°)$	0.93	0.93	1.00	0.93	0.93	0.93	0.93	0.93	0.93	0.93
$23.4° - 27.6°(4.2°)$	0.55	0.63	0.60	0.57	0.68	0.67	0.61	0.57	0.69	0.58
$28.0° - 32.0°(4.0°)$	1.00	1.00	1.00	0.59	0.63	0.71	0.63	0.59	0.54	0.64
$36.5° - 38.0°(1.5°)$	0.87	0.91	0.88	0.76	0.88	0.82	0.79	0.76	0.73	0.72
Overall Median	0.78	0.84	1.00	0.72	0.79	0.80	0.76	0.72	0.75	0.72

8.2 Medium Resolution Imaging Mode Results

Table 8.4 shows the PMD values of each iceberg for the low resolution imaging mode. As with the low resolution imaging mode, the CTLR/DCP data yields the best results, with only one missed target and the lowest overall median PMD of 0.67. The full pseudo quad-pol and both Stokes vector modes also only missed one target, however, the two Stokes vector modes also had one false alarm.

As in the low resolution imaging mode, all of the linear dual-pol modes missed a greater number of targets and had higher median PMD values than the CTLR. Despite the small size of the data set available for this study, the strength of the compact polarimetry in this context seems quite clear. Even the weaker compact polarimetry modes (the pseudo-HV and the CTLR with pseudo-HV combination) missed a fewer number of targets than the HH-HV and the VV-VH.

Comparing the Stokes vector data using a normal LRT with the skew-normal LRT, the improvement from using the skew-normal LRT is less clear for the medium resolution imaging mode than in the low resolution mode. Here, the skew-normal LRT actually produces a slightly higher overall median PMD value than the normal LRT. For the two lowest incidence angle categories, the normal and skew-normal LRTs produce the same median PMD value. For the third incidence angle category the skew-normal LRT produces a slightly higher median PMD value, while for the highest incidence angle category, the skew-normal produces a slightly lower PMD value. It appears that the skew-normal LRT provides a better fit to the data in the low resolution imaging mode than in the medium resolution case. It is possible that finer resolution data is more difficult to model. Work could possibly be done improving the fitting process of the skew-normal distribution—it is likely that the skew-normal distribution was attempting to account for the fact that the three Stokes vector elements S_1, S_2, and S_3 must be less than S_0 (i.e., the four Stokes vector elements are not completely independent) through use of the skew factor, in the process not accounting for

some of other the skew present in the data.

Table 8.5 lists the median PMD values for each size category. As expected the values are highest for the small category. Somewhat surprising is that for small icebergs the lowest median PMD is obtained using the Stokes vector with a normal LRT. For the medium and large icebergs, the best performance was from the CTLR/DCP data as usual.

Table 8.6 lists the median PMD values of each incidence angle category, for the medium icebergs only. As before, we see that the PMD values are highest for the lowest incidence angle category, as expected. For the first three incidence angle categories, the normal LRT Stokes vector actually has the lowest median PMD, while for the highest incidence angle category, the CTLR and skew-normal Stokes have the lowest median PMD. It's possible that the skew-normal LRT is a better fit for the Stokes vector data at higher incidence angles, though in any event, for the incidence angle categories where the skew-normal LRT is an improvement for the Stokes vector data, the CTLR also performs just as well, making the skew-normal LRT somewhat superfluous for the medium resolution imaging mode.

Table 8.4: imaging mode, categorized by incidence angle. A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			Linear Dual-Pol			Compact Polarimetry						
ID (Scene)	Size	Shape	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
FQ2: $19.9° - 21.8°(1.9°)$												
R1 (0905-1-1)	S	Wedge	0.89	0.89	1.00	0.85	0.84	0.85	0.85	0.85	0.83	0.83
R2 (0905-1-2)	S	Wedge	0.90	0.90	1.00	0.87	0.85	0.90	0.87	0.87	0.85	0.85
B1 (0905-1-3)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
H (0905-2-1)	M	Drydock	0.83	0.84	1.00	0.83	0.88	0.82	0.81	0.83	0.83	0.80
B2 (0905-2-2)	M	Blocky	1.00	1.00	1.00	0.85	0.89	0.89	0.85	0.85	0.85	0.85
Median			0.90	0.90	1.00	0.85	0.88	0.89	0.85	0.85	0.85	0.85
FQ5/6/7: $23.4° - 27.6°(4.2°)$												
G1 (0829-1-1)	S	Drydock	0.62	0.73	0.65	0.50	0.71	0.61	0.53	0.50	0.52	0.48
G2 (0805-1-1)	S	Drydock	0.50	0.54	0.44	0.54	0.58	0.71	0.58	0.54	0.37	0.37
N (0805-1-2)	S	Wedge	0.50	0.53	0.44	0.58	0.64	0.71	0.59	0.58	0.64	0.66
A (0805-2-1)	M	Blocky	0.64	0.62	0.60	0.66	0.72	0.81	0.68	0.66	0.73	0.77
J (0815-2-1)	M	Pinnacle	0.48	0.54	0.80	0.32	0.39	0.39	0.35	0.32	0.34	0.39
D1 (0815-2-2)	M	Drydock	0.74	0.74	0.74	0.77	0.77	1.00	1.00	0.77	0.66	0.80
D2 (0815-2-3)	M	Drydock	0.89	0.91	0.98	0.74	0.78	0.82	0.79	0.74	0.69	0.72
M (0822-1-1)	M	Wedge	0.65	0.83	0.86	0.55	0.55	0.62	0.59	0.55	0.50	0.50
I (0822-2-1)	L	Pinnacle	0.61	0.80	1.00	0.41	0.44	0.42	0.41	0.41	0.41	0.41
P1 (0822-2-2)	L	Wedge	0.72	0.73	0.83	0.54	0.58	0.55	0.55	0.54	0.52	0.53
Median			0.63	0.73	0.77	0.54	0.61	0.66	0.58	0.54	0.52	0.52
FQ9/10/11: $28.0° - 32.0°(4.0°)$												
C1 (0815-1-1)	M	Dome	0.55	0.60	0.66	0.59	0.61	0.62	0.59	0.59	0.58	0.59
O1 (0815-1-2)	M	Wedge	0.82	1.00	0.86	0.81	0.81	0.82	0.81	0.81	0.83	0.93
P2 (0825-2-1)	M	Wedge	0.78	0.83	0.96	0.67	0.74	0.75	0.70	0.67	0.75	0.77
L (0825-1-1)	L	Tabular	0.71	0.67	1.00	0.48	0.54	0.50	0.50	0.48	0.42	0.40
Median			0.74	0.75	0.91	0.63	0.67	0.68	0.64	0.63	0.66	0.68
FQ16: $36.5° - 38.0°(1.5°)$												
K (0818-1-1)	S	Pinnacle	0.83	0.88	0.87	0.72	0.90	0.79	0.76	0.72	0.84	0.81
Q (0818-1-2)	S	Wedge	1.00	0.88	0.91	0.88	0.91	0.90	0.88	0.88	0.90	0.91
E (0818-1-3)	M	Drydock	0.85	0.94	0.88	0.75	0.85	0.84	0.77	0.75	0.73	0.72
F (0818-1-4)	-	-	0.66	0.86	0.80	0.62	0.84	0.67	0.63	0.62	0.58	0.55
O2 (0818-1-5)	M	Wedge	0.70	0.74	0.76	0.64	0.73	0.72	0.65	0.64	0.63	0.62
C2 (0818-1-6)	M	Dome	0.92	0.97	0.95	0.86	0.96	0.89	0.86	0.86	0.82	0.80
Median			0.84	0.88	0.87	0.74	0.88	0.82	0.77	0.74	0.78	0.76
Overall Median			0.74	0.83	0.87	0.67	0.77	0.79	0.70	0.67	0.69	0.72
# of Missed Targets			3	3	7	1	1	2	2	1	1	1
# of False Alarms			0	0	0	0	0	0	0	0	1	1

Table 8.5: Median PMD values for the medium resolution imaging mode, categorized by iceberg size (small, medium, or large). A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Size	Linear Dual-Pol			Compact Polarimetry						
	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
Small	0.74	0.83	0.95	0.77	0.77	0.89	0.85	0.77	0.66	0.80
Medium	0.72	0.79	0.84	0.64	0.76	0.75	0.66	0.64	0.68	0.69
Large	0.78	0.83	0.96	0.67	0.74	0.75	0.70	0.67	0.69	0.72
Overall Median	0.74	0.83	0.87	0.67	0.77	0.79	0.70	0.67	0.69	0.72

Table 8.6: Median PMD values for the medium resolution imaging mode, for medium icebergs only, categorized by incidence angle. A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Incidence Angle Range	Linear Dual-Pol			Compact Polarimetry						
	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
$19.9° - 21.8°(1.9°)$	0.89	0.89	1.00	0.85	0.88	0.85	0.85	0.85	0.83	0.83
$23.4° - 27.6°(4.2°)$	0.62	0.62	0.60	0.54	0.64	0.71	0.58	0.54	0.52	0.53
$28.0° - 32.0°(4.0°)$	0.71	0.67	0.86	0.59	0.61	0.62	0.59	0.59	0.58	0.59
$36.5° - 38.0°(1.5°)$	0.83	0.88	0.87	0.72	0.85	0.79	0.76	0.72	0.73	0.72
Overall Median	0.74	0.83	0.87	0.67	0.77	0.79	0.70	0.67	0.69	0.72

Table 8.7: PMD values for the ship detection imaging mode, categorized by incidence angle. A FAR of 10^{-8} was used, with a minimum allowed target cluster size of four pixels. The lowest value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			Linear Dual-Pol			Compact Polarimetry						
ID (Scene)	Size	Shape	HH-HV	VV-VH	HH-VV	CL	PQ	PHV	CL-PHV	DCP	St.	St. SN
FQ16: $36.5° - 38.0°(1.5°)$												
K (0818-1-1)	S	Pinnacle	0.82	0.91	0.90	0.75	0.93	0.85	0.76	0.75	0.92	0.91
Q (0818-1-2)	S	Wedge	0.90	0.92	0.95	0.88	0.93	0.93	0.89	0.88	0.89	0.88
E (0818-1-3)	M	Drydock	0.88	1.00	0.93	0.76	0.92	0.90	0.79	0.76	0.89	0.90
F (0818-1-4)	-	-	0.81	0.88	0.88	0.76	0.84	0.83	0.77	0.76	0.76	0.77
O2 (0818-1-5)	M	Wedge	0.73	0.78	0.81	0.63	0.71	0.80	0.64	0.63	0.70	0.70
C2 (0818-1-6)	M	Dome	0.90	0.94	0.94	0.82	0.89	0.87	0.82	0.82	0.82	0.81
Overall Median			0.85	0.91	0.92	0.76	0.91	0.86	0.78	0.76	0.86	0.84
# of Missed Targets			0	1	0	0	0	0	0	0	0	0
# of False Alarms			1	0	0	0	0	0	0	0	0	0

8.3 Ship Detection Imaging Mode Results

Table 8.7 shows the PMD values of each iceberg in scene 0818-1 for the ship detection imaging mode. Again we see that the CTLR mode has the lowest median PMD value (0.76), well below that of the best linear dual-pol mode (0.85 for the HH-HV). None of the polarimetric modes miss targets, except for the VV-VH mode, where one target was missed. The HH-HV mode had one false alarm, while the other modes had none.

Use of the skew-normal LRT slightly lowers the median PMD value for the Stokes vector detection (from 0.86 to 0.84), but still produces worse performance overall than that of the CTLR mode. However, it does have the lowest PMD value for two of the six iceberg targets, while the CTLR mode has the lowest PMD value for the other four iceberg targets.

Due to the limited amount of data, no analysis of iceberg size will be done for the ship detection imaging mode.

8.4 Imaging Mode Comparison

Figure 8.1 provides a visual comparison of RH σ^0 values across the three different imaging modes, for iceberg E from scene 0818-1. As the spatial resolution becomes coarser, starting from the ship detection mode (top left), to the medium resolution mode (top right), and finally to the low resolution mode (bottom), it's clear that the detail with which the iceberg can be characterized visually in the image is reduced. Worth noting is that both the ship detection and medium resolution modes have a 350 km swath, while the low resolution mode has a 500 km swath. To achieve this increased swath width, the low resolution mode has a significantly coarser azimuth resolution. As shown in the previous sections, this coarser resolution results in a greater number of missed iceberg targets (five for the CTLR) than the other imaging modes.

In comparison, the ship detection and medium resolution modes are closer in spatial resolution, with the main difference lying in the fact that the ship detection mode has a significantly finer range resolution than the medium resolution mode. Also worth noting is that this resolution is designed to vary with incidence angle for the ship detection mode. As well, the ship detection mode is limited to high incidence angle beam modes, whereas the medium resolution mode has a wider range of incidence angles for which data can be collected. This flexibility could potentially be of use operationally, if it was necessary to have coverage of a particular area at a certain time.

While the ship detection mode is certainly more interesting visually, this study has provided no evidence that the ship detection mode actually produces better iceberg detection performance than the medium resolution mode—here the two modes produce very similar results. We note that for scene 0818-1, the CTLR data had almost the same median PMD value for the ship detection mode (0.76) as for the medium resolution mode (0.74), and both imaging modes had no missed targets for this scene. Due to the limited incidence angles of the ship detection imaging mode, results for the other scenes besides 0818-1 have been

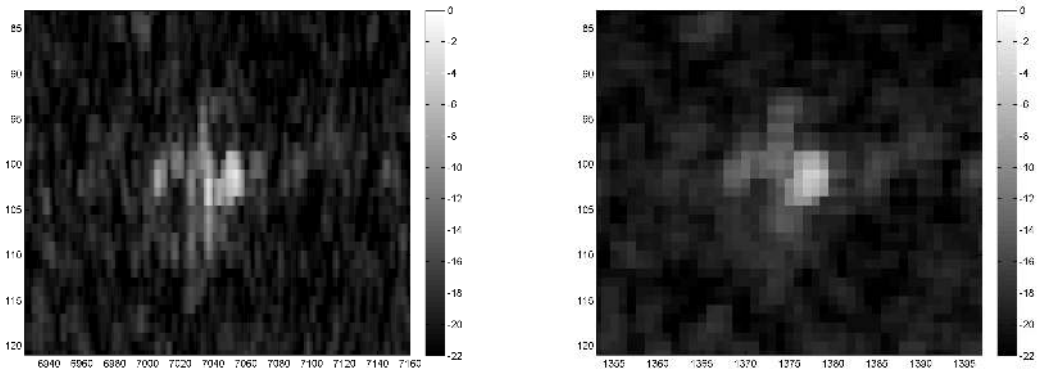

(a) Ship Detection Mode

(b) Medium Resolution Mode

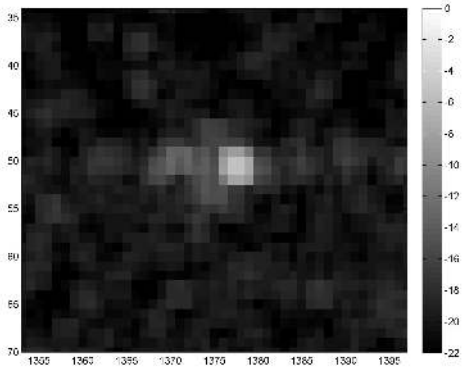

(c) Low Resolution Mode

Figure 8.1: RH σ^0 images of iceberg E (medium drydock) from scene 0818-1, for the three different imaging modes.

omitted. However, when a test was performed using the full set of scenes with the ship detection imaging mode, it was found that the same iceberg target missed by the CTLR data in the medium resolution mode was also missed in the ship detection mode. To actually determine which of these modes is superior in terms of detection performance would likely take a significantly greater amount of data. What is clear is that use of the low resolution imaging mode greaterly reduces the detection performance for all of the polarimetric modes, with the tradeoff being a much wider coverage area.

Chapter 9

Adaptive False Alarm Rate Adjustment

In this chapter a method is presented that utilizes the shape and orientation of the received polarization ellipse in an attempt to improve the detection results of a coherent dual-polarized SAR system. Though it can be used for any system which is able to calculate the received Stokes parameters, it is particularly well suited to a circular transmit polarization, since as shown in chapter 5, the orientation angle ψ is a better discriminator of ocean and targets when the transmitted polarization is rotationally invariant, as it is in the circular transmit case [45]. For a linear transmit polarization, the orientation of the received polarization ellipse tends to be aligned to the transmit polarization's orientation. But for circular transmit polarization, the transmitted polarization has no orientation, and therefore the received polarization ellipse has no preferred orientation, with the value of ψ depending more strongly on the target geometry.

The motivation for this method is that as shown previously, use of the skew-normal LRT to detect icebergs using the Stokes vector yielded mixed results. While the skew-normal LRT was an improvement over the normal LRT for many of the incidence angles and imaging modes tested, it produced inferior results in some cases, and in general fell short of the detection performance of the native CTLR RH-RV used as input to the normal LRT.

Yet the Stokes parameters, and by extension the polarization ellipse parameters, contain useful information that is not utilized by the LRT decision variable calculation. This method attempts to use the Stokes parameters to augment the detection performance of the LRT decision variable, not through the decision variable calculation itself but by refining the choice of the detection threshold for each pixel.

One reason the ψ and χ parameters are worth investigating is that they are both phases,

whereas the LRT process uses mostly power-based parameters (e.g., the covariance matrix elements) as input. This is also the reason that the ψ and χ angles are used here rather than the raw Stokes vector directly, since the Stokes vector elements are also measures of power, and are therefore somewhat redundant to the information already used as input to the decision variable calculation.

9.1 False Alarm Rate as a Function of ψ and χ

Recall from chapter 5 that the polarization ellipse orientation, ψ, and ellipticity, χ, are two of the spherical coordinates that define the same point in space as the Stokes vector parameters (if the Stokes are used as Cartesian coordinates). The third spherical coordinate is the radius of the sphere, or the distance between the point and the origin, which is equal to the total polarized power, $I_p = \sqrt{S_1^2 + S_2^2 + S_3^2}$. Note that the degree of polarization, m, is equal to I_p/S_0. For pixels with low degree of polarization values, I_p will be quite low, making the angles ψ and χ more sensitive to small fluctuations in the Stokes vector parameters S_1, S_2, S_3 (e.g., fluctuations from speckle), and therefore less reliable. In the extreme case, as m (and therefore I_p) approaches zero, the values of ψ and χ are almost certainly dominated by the effects of speckle and instrument noise, making their values meaningless. However, for high m values (such as from the majority of ocean pixels), the polarization ellipse orientation and ellipticity are key indicators of the dominant scattering mechanism of the area under study, and provide valuable information that can be utilized in the target detection process.

In chapter 5 plots of χ vs. ψ and χ vs. m were shown that confirmed that the majority of ocean pixels are well confined within a relatively small area of the ψ–χ space, consistent with the fact that ocean pixels should theoretically exhibit mostly surface and Bragg scattering, with their exact scattering behaviour mostly depending on incidence angle (for relatively calm sea). In contrast, most of the iceberg and ship targets in this study appeared to exhibit a wide array of scattering mechanisms, including volume (e.g., depolarized) and double-bounce

backscatter (e.g., dihedral scattering as in section 5.4.3, indicated in the CTLR imagery by a negative χ angle).

We therefore wish to vary the detection threshold from pixel-to-pixel, based on the values of ψ, χ, and to a lesser extent, m. We do this by first building an empirical model of maximum values of the ocean backscatter for each combination of ψ and χ values. We split the ψ-χ plane into one degree wide bins. For each bin we find the maximum decision variable value within the same ocean chip as was used to estimate the ocean covariance matrix for the decision variable calculation. We use the maximum values rather than the mean because these max values represent the pixels which are most likely to be falsely detected. This ψ–χ grid is then smoothed using a 5×5 moving average. We convert the smoothed decision variable values for each bin into a corresponding $\log_{10}(\text{FAR})$ value using our previously calculated linear $\log_{10}(\text{FAR})$ vs. detection threshold function as in chapter 7.

For this method, rather than the user specifying a single desired false alarm rate value (e.g., 10^{-8}), they instead specify the range of values they would like to be used for highly polarized pixels (e.g., 10^{-5} down to 10^{-9}), as well as a single value they would like to use for depolarized, low m pixels (e.g., 10^{-8}, same as in the constant FAR case). We then scale and shift our grid of $\log_{10}(\text{FAR})$ to these desired FAR values using the simple equations:

$$\text{FAR}_{scaled} = \text{FAR}_{grid} \times \frac{\log_{10}(\text{FAR}_{high}) - \log_{10}(\text{FAR}_{low})}{\max(\text{FAR}_{grid}) - \min(\text{FAR}_{grid})} \tag{9.1}$$

$$\text{FAR}_{final} = \text{FAR}_{scaled} + \log_{10}(\text{FAR}_{high}) - \max(\text{FAR}_{scaled}) \tag{9.2}$$

$\max(\cdot)$ and $\min(\cdot)$ are functions which take the maximum and minimum values of their argument, respectively. FAR_{high} is the maximum FAR value for high m pixels as specified by the user (e.g., 10^{-5}). FAR_{low} is the minimum FAR value for high m pixels as specified by the user (e.g., 10^{-9}). FAR_{grid} is the original $\log_{10}(\text{FAR})$ grid, smoothed but not yet scaled to the appropriate FAR values. FAR_{scaled} is a grid that has been scaled but not yet shifted. FAR_{final} is the final $\log_{10}(\text{FAR})$ grid used in the detection process.

We still need a way to deal with pixels that have a low m value (and where ψ and χ

are less meaningful regarding that pixel's dominant type of backscatter). For pixels with $m < 0.3$ we do not use the ψ-χ grid, and instead use a constant FAR, generally somewhere between the values of FAR_{high} and FAR_{low}, close to what we expect the mean FAR value to be (e.g., 10^{-8}).

An example of a final FAR grid, scaled and shifted and ready for use in detection, is shown in Figure 9.1 for the CTLR, HH-HV, and VV-VH modes. FAR grids for all of the scenes in this study are included in Appendix C for reference.

Note that here we derive a 2-D function of FAR vs. ψ and χ, but in theory there is nothing preventing the creation of a 3-D function (including m as a direct parameter to the FAR grid, rather than simply using it as a lower bound for the process). However, the more the data in the ocean image chip is split up in this fashion, the less samples we have available to calculate the desired detection threshold for each bin. Use of a 3-D FAR function produced worse results for this dataset than the 2-D function. It is possible that a 3-D function could be used, but this would require a greater amount of data, or possibly the use of a statistical, rather than an empirical, model.

9.2 Summary of the Method

A step-by-step breakdown of the method is as follows:

1. Calculate the decision variable using the normal LRT, and derive an empirical function for $\log_{10}$(FAR) vs. detection threshold, as before.

2. Calculate the ψ, χ, and m values for each pixel in the ocean image chip. The ψ and χ angles tended to be particularly noisy for this data set, so a 5×5 spatial averaging window was used during the calculation of the Stokes vector parameters which were used to calculate ψ, χ, and m.

3. Split the ψ-χ space into bins. Here we use square bins that are one degree in

each dimension.

4. For each bin, find the set of pixels in the ocean image chip that have ψ and χ values corresponding to that bin.

5. For this set of pixels, find the pixel with the highest decision variable value.

6. Repeat the process for every bin in the ψ-χ space. Save the max decision variable values for each bin, until the whole space is filled. For bins with no corresponding ocean pixels, save a value of zero.

7. Smooth the grid by applying a moving average using a 5×5 window.

8. Convert the max decision variable values into $\log_{10}$(FAR) values, using the $\log_{10}$(FAR) vs. threshold function.

9. Scale and offset the grid of $\log_{10}$(FAR) values such that it satisfies the maximum and minimum FAR values desired by the user, using equations 9.1 and 9.2. In the results shown here we generally use a maximum FAR of 10^{-5} and a minimum FAR of 10^{-9}, with the exception of the ship detection imaging mode, where a minimum FAR of $10^{-9.5}$ was used. When using a minimum FAR of 10^{-9} for the ship detection imaging mode, the mean FAR applied to the ocean pixels was greater than desired (above the constant FAR of 10^{-8}).

10. For each pixel in the image, the FAR corresponding to that pixel's ψ and χ values will be used. For pixels with $m < 0.3$, the ψ and χ values are not particularly useful in characterizing the nature of the backscattered wave, and are therefore ignored. We use a constant FAR for these low m pixels. Note that theoretically the m values of the ocean pixels should be high, and we therefore expect few ocean pixels to satisfy this criteria.

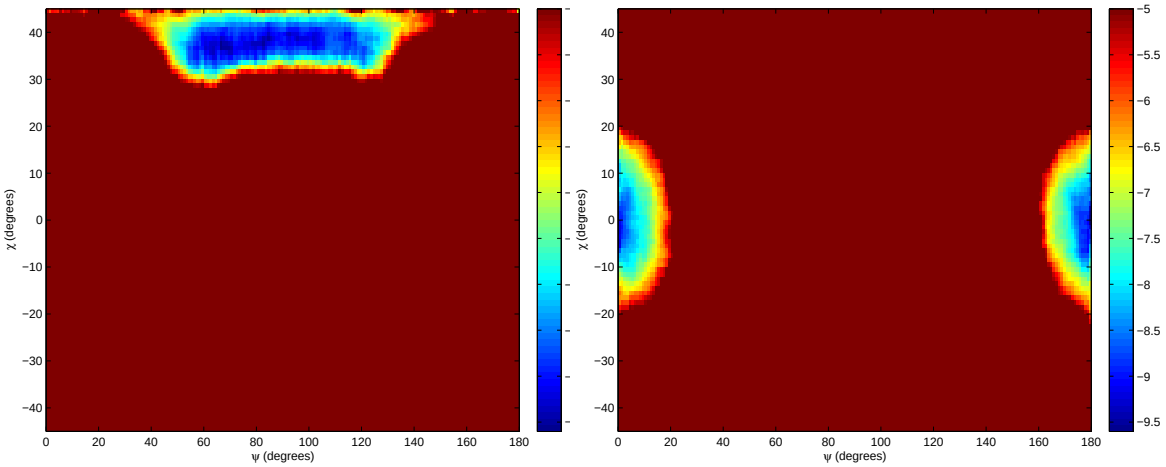

(a) RH-RV

(b) HH-HV

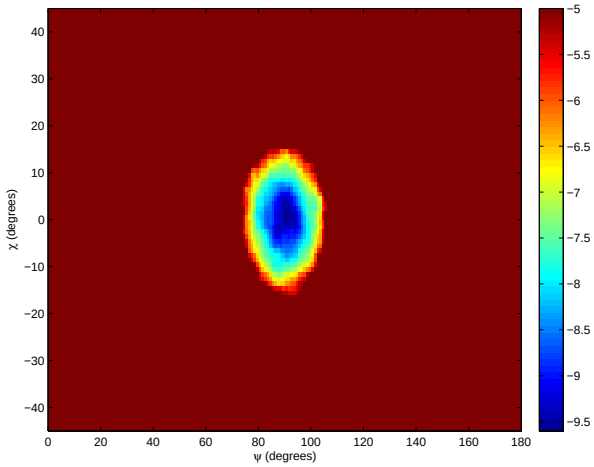

(c) VV-VH

Figure 9.1: Examples of ψ-χ FAR grids for each transmit polarization: (a) CTLR, (b) HH-HV, and (c) VV-VH. The desired value of $\log_{10}(\text{FAR})$ is shown as a function of ψ and χ for scene 0818-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

9.3 Low Resolution Imaging Mode Results

Table 9.1 shows PMD values for the low resolution imaging mode, comparing the values for a constant FAR of 10^{-8} with an adaptive FAR ranging from 10^{-5} down to 10^{-9}. For pixels with $m < 0.3$, the constant FAR of 10^{-8} was used. In the last row of the table, the mean $\log_{10}$(FAR) value applied to the ocean image chips in each scene is listed, to show the effective FAR applied to the ocean pixels by the adaptive FAR.

For brevity, only three polarimetric modes are compared: the HH-HV and VV-VH linear dual-pol modes, and the CTLR RH-RV compact polarimetry. For all three modes, use of the adaptive FAR lowers the median PMD value. For the VV-VH and CTLR modes, the mean FAR applied to the ocean was less in the adaptive case than in the constant FAR case. Despite this fact, a greater number of pixels were detected for most of the targets using the adaptive FAR, and for the CTLR mode one target was detected which was missed when using a constant FAR. None of the modes produced any false alarms, whether in the constant or adaptive case.

Table 9.1: PMD values for the low resolution imaging mode, comparing a constant FAR with an adaptive FAR. A FAR of 10^{-8} was used in the constant FAR case, with a FAR range of 10^{-5} to 10^{-9} (with low m FAR of 10^{-8}) for the adaptive case. The lowest PMD value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			HH-HV		VV-VH		CTLR	
ID (Scene)	Size	Shape	Constant	Adaptive	Constant	Adaptive	Constant	Adaptive
FQ2: $19.9° - 21.8°(1.9°)$								
R1 (0905-1-1)	S	Wedge	0.93	0.93	0.93	0.93	0.93	0.93
R2 (0905-1-2)	S	Wedge	0.81	0.79	0.86	0.79	0.79	0.79
B1 (0905-1-3)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00
H (0905-2-1)	M	Drydock	0.78	0.74	0.78	0.70	0.79	0.79
B2 (0905-2-2)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00
Median			0.93	0.93	0.93	0.93	0.93	0.93
FQ5/6/7: $23.4° - 27.6°(4.2°)$								
G1 (0829-1-1)	S	Drydock	0.55	0.55	0.73	0.69	0.55	0.48
G2 (0805-1-1)	S	Drydock	0.47	0.47	0.50	0.50	0.51	0.46
N (0805-1-2)	S	Wedge	0.51	0.51	0.54	0.55	0.57	0.49
A (0805-2-1)	M	Blocky	0.66	0.62	0.63	0.61	0.72	0.63
J (0815-2-1)	M	Pinnacle	0.51	0.53	0.58	0.60	0.36	0.34
D1 (0815-2-2)	M	Drydock	1.00	1.00	1.00	1.00	1.00	1.00
D2 (0815-2-3)	M	Drydock	0.92	0.92	0.92	0.92	0.90	0.88
M (0822-1-1)	M	Wedge	0.74	0.74	0.74	0.74	0.40	0.31
I (0822-2-1)	L	Pinnacle	0.75	0.72	0.75	0.75	0.42	0.36
P1 (0822-2-2)	L	Wedge	0.73	0.73	0.76	0.73	0.61	0.59
Median			0.69	0.67	0.73	0.71	0.56	0.49
FQ9/10/11: $28.0° - 32.0°(4.0°)$								
C1 (0815-1-1)	M	Dome	0.66	0.66	0.68	0.68	0.59	0.56
O1 (0815-1-2)	M	Wedge	1.00	1.00	1.00	1.00	1.00	1.00
P2 (0825-2-1)	M	Wedge	0.77	0.76	0.83	0.80	0.71	0.55
L (0825-1-1)	L	Tabular	1.00	1.00	1.00	1.00	0.53	0.40
Median			0.89	0.88	0.91	0.90	0.65	0.56
FQ16: $36.5° - 38.0°(1.5°)$								
K (0818-1-1)	S	Pinnacle	0.87	0.85	0.91	0.91	0.76	0.73
Q (0818-1-2)	S	Wedge	1.00	1.00	1.00	1.00	1.00	0.89
E (0818-1-3)	M	Drydock	1.00	1.00	1.00	1.00	0.84	0.71
F (0818-1-4)	-	-	0.54	0.55	0.84	0.89	0.66	0.53
O2 (0818-1-5)	M	Wedge	0.69	0.69	0.71	0.69	0.65	0.61
C2 (0818-1-6)	M	Dome	0.92	0.88	1.00	1.00	0.90	0.88
Median			0.89	0.87	0.95	0.96	0.80	0.72
Overall Median			0.78	0.76	0.84	0.80	0.72	0.63
# of Missed Targets			7	7	8	8	5	4
# of False Alarms			0	0	0	0	0	0
Mean Ocean $\log_{10}$(FAR)			-8	-7.78	-8	-8.03	-8	-8.14

9.4 Medium Resolution Imaging Mode Results

Table 9.2 shows PMD values for the medium resolution imaging mode. As with the low resolution mode, we see that use of the adaptive FAR lowers the median PMD values for the VV-VH and CTLR modes. For the HH-HV mode, the overall median PMD value stays constant, with some PMD values increasing and others decreasing. Use of the adaptive FAR does not change the number of missed targets, and all of the modes detect no false alarms with either the constant or adaptive FAR.

Regarding the mixed results for the HH-HV mode, it's worth noting that the method is theoretically more useful with a circular transmit polarization (and therefore rotational invariance) than a linear transmit polarization, where values tend to cluster around the ψ value of the transmitted waves. However, this does not explain the fact that the PMD values decreased for the VV-VH mode, and for the HH-HV mode in the low resolution case. It's possible that a larger averaging window should be used for calculation of the ψ and χ angles. The 5×5 window used for spatial averaging during the Stokes vector calculation was the same size for each imaging mode. Possibly a larger size (in pixels) should be used for the finer spatial resolutions—unfortunately this would tend to diminish the improvements that these finer spatial resolutions can provide.

Table 9.2: PMD values for the medium resolution imaging mode, comparing a constant FAR with an adaptive FAR. A FAR of 10^{-8} was used in the constant FAR case, with a FAR range of 10^{-5} to 10^{-9} (with low m FAR of 10^{-8}) for the adaptive case. The lowest PMD value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			HH-HV		VV-VH		CTLR	
ID (Scene)	Size	Shape	Constant	Adaptive	Constant	Adaptive	Constant	Adaptive
FQ2: $19.9° - 21.8°(1.9°)$								
R1 (0905-1-1)	S	Wedge	0.89	0.89	0.89	0.87	0.85	0.83
R2 (0905-1-2)	S	Wedge	0.90	0.90	0.90	0.90	0.87	0.85
B1 (0905-1-3)	M	Blocky	1.00	1.00	1.00	1.00	1.00	1.00
H (0905-2-1)	M	Drydock	0.83	0.82	0.84	0.79	0.83	0.81
B2 (0905-2-2)	M	Blocky	1.00	1.00	1.00	0.93	0.85	0.85
Median			0.90	0.90	0.90	0.90	0.85	0.85
FQ5/6/7: $23.4° - 27.6°(4.2°)$								
G1 (0829-1-1)	S	Drydock	0.62	0.60	0.73	0.68	0.50	0.44
G2 (0805-1-1)	S	Drydock	0.50	0.50	0.54	0.54	0.54	0.50
N (0805-1-2)	S	Wedge	0.50	0.49	0.53	0.53	0.58	0.50
A (0805-2-1)	M	Blocky	0.64	0.62	0.62	0.60	0.66	0.63
J (0815-2-1)	M	Pinnacle	0.48	0.47	0.54	0.53	0.32	0.28
D1 (0815-2-2)	M	Drydock	0.74	0.74	0.74	0.74	0.77	0.66
D2 (0815-2-3)	M	Drydock	0.89	0.88	0.91	0.92	0.74	0.69
M (0822-1-1)	M	Wedge	0.65	0.65	0.83	0.74	0.55	0.48
I (0822-2-1)	L	Pinnacle	0.61	0.61	0.80	0.75	0.41	0.34
P1 (0822-2-2)	L	Wedge	0.72	0.67	0.73	0.71	0.54	0.49
Median			0.63	0.62	0.73	0.70	0.54	0.50
FQ9/10/11: $28.0° - 32.0°(4.0°)$								
C1 (0815-1-1)	M	Dome	0.55	0.55	0.60	0.61	0.59	0.47
O1 (0815-1-2)	M	Wedge	0.82	0.82	1.00	1.00	0.81	0.78
P2 (0825-2-1)	M	Wedge	0.78	0.77	0.83	0.81	0.67	0.56
L (0825-1-1)	L	Tabular	0.71	0.67	0.67	0.67	0.48	0.38
Median			0.74	0.72	0.75	0.74	0.63	0.52
FQ16: $36.5° - 38.0°(1.5°)$								
K (0818-1-1)	S	Pinnacle	0.83	0.84	0.88	0.88	0.72	0.66
Q (0818-1-2)	S	Wedge	1.00	1.00	0.88	0.88	0.88	0.88
E (0818-1-3)	M	Drydock	0.85	0.85	0.94	0.94	0.75	0.65
F (0818-1-4)	-	-	0.66	0.69	0.86	0.86	0.62	0.52
O2 (0818-1-5)	M	Wedge	0.70	0.70	0.74	0.73	0.64	0.60
C2 (0818-1-6)	M	Dome	0.92	0.93	0.97	1.00	0.86	0.81
Median			0.84	0.85	0.88	0.88	0.74	0.66
Overall Median			0.74	0.74	0.83	0.79	0.67	0.63
# of Missed Targets			3	3	3	3	1	1
# of False Alarms			0	0	0	0	0	0
Mean Ocean $\log_{10}$(FAR)			-8	-7.94	-8	-8.09	-8	-8.19

9.5 Ship Detection Imaging Mode Results

Table 9.3 shows PMD values for the ship detection imaging mode. Here a FAR range of 10^{-5} down to $10^{-9.5}$ is used, with a slightly lower value of FAR_{low} than for the previous imaging modes. This was chosen to keep the mean ocean FAR below the constant FAR of 10^{-8} for the CTLR data.

As with the medium resolution mode, the median PMD value for the HH-HV case stays constant when using an adaptive FAR. The median PMD value of the VV-VH actually increases slightly. As before, the median PMD for the CTLR mode decreases, from 0.76 in the constant FAR case to 0.70 in the adaptive FAR case. One false alarm detected by the HH-HV mode in the constant FAR case is not detected in the adaptive FAR case—a small but notable benefit to the adaptive FAR. Use of the adaptive FAR has no affect on the number of missed targets. Again, we note that for the VV-VH and CTLR modes, the adaptive FAR applies a lower mean FAR to the ocean image chips (corresponding to a higher detection threshold) than the constant FAR case. Despite this fact, the CTLR mode detects a greater number of target pixels.

Another possibility, if this method was to be applied to the linear dual-pol mode data in the future, would be to use a different parameter in place of ψ—presumably one less affected by the orientation of the transmit polarization, which would be more useful in characterizing the difference between ocean and target pixels. It is possible that m and χ could be used to produce the FAR grid, rather than m being the lower bound for use of the adaptive FAR, but this has not yet been tested, as the focus of this research was on a method for use with the CTLR data, not the linear dual-pol.

9.6 Discussion of Results

According to these results, use of an adaptive FAR, varied based on the values of the ψ and χ parameters, improves the iceberg detection performance for this dataset. The method is

Table 9.3: PMD values for the ship detection imaging mode, categorized by polarization mode (HH-HV, VV-VH, and CTLR) and FAR type (constant or adaptive). A FAR of 10^{-8} was used in the constant FAR case, with a FAR range of 10^{-5} to $10^{-9.5}$ (with low m FAR of 10^{-8}) for the adaptive case. The lowest PMD value for each row is shaded a dark grey. If other PMD values in that row are within 0.05 of the minimum, they are shaded a lighter grey.

Iceberg Properties			HH-HV		VV-VH		CTLR	
ID (Scene)	Size	Shape	Constant	Adaptive	Constant	Adaptive	Constant	Adaptive
FQ16: $36.5° - 38.0°(1.5°)$								
K (0818-1-1)	S	Pinnacle	0.82	0.83	0.91	0.92	0.75	0.69
Q (0818-1-2)	S	Wedge	0.90	0.90	0.92	0.92	0.88	0.86
E (0818-1-3)	M	Drydock	0.88	0.88	1.00	1.00	0.76	0.67
F (0818-1-4)	-	-	0.81	0.82	0.88	0.89	0.76	0.61
O2 (0818-1-5)	M	Wedge	0.73	0.72	0.78	0.77	0.63	0.59
C2 (0818-1-6)	M	Dome	0.90	0.91	0.94	0.95	0.82	0.80
Median			0.85	0.85	0.91	0.92	0.76	0.70
# of Missed Targets			0	0	1	1	0	0
# of False Alarms			1	0	0	0	0	0
Mean Ocean $\log_{10}$(FAR)			-8	-7.76	-8	-8.19	-8	-8.33

more effective for the CTLR compact polarimetry than the linear dual-pol mode, indicated by a larger decrease in median PMD value from the constant FAR to the adaptive FAR for the CTLR data than for the linear dual-pol. This is presumably due to the rotational invariance of the CTLR (and therefore the greater usefulness of ψ as a discriminator between ocean and targets for this mode).

For the low resolution mode data the method was particularly effective, decreasing the number of missed targets for the CTLR mode data. For all three imaging modes, the adaptive FAR lowered the median PMD values of the CTLR data without increasing the number of false alarms, while using a slightly higher mean threshold (that is, a lower false alarm rate) across the ocean image chip in each scene.

A possible avenue of future work is to improve the method through the use of parameters other than ψ, χ, and m, or through the use of a 3-D function rather than a 2-D function (with a low m cutoff) to model the desired FAR behaviour.

Chapter 10

Discrimination of Ships and Icebergs

An important consideration after target detection is performed is whether those targets can then also be correctly identified and classified in the SAR imagery. One example of this issue is in the discrimination of ships and icebergs. Misclassifying ships as icebergs, or vice versa, can cause confusion and can make attempts to track the drifts of icebergs in SAR images significantly more difficult.

To test the discrimination performance of ships and icebergs, some Radarsat-2 fine-quad scenes were obtained from Defence Research & Development Canada. The data was located near the Strait of Gibraltar, and contained 125 ship targets, validated using Automatic Identification System (AIS) data. These targets were then manually masked out in the images, in the same manner as for the iceberg targets. Table 10.1 provides a summary of the ship detection scenes used for this purpose. The detection process used for these ship scenes was the same as for the iceberg detection case. After detection, a support vector machine (SVM) classifier was tested, and the discrimination accuracy for CTLR mode data (both using constant FAR, and adaptive FAR detection) and HH-HV mode data (using an adaptive FAR only, for brevity) was compared. Only the low resolution and medium resolution imaging modes were tested, due to the relative lack of data at incidence angles appropriate to the ship detection imaging mode.

10.1 The m–χ Decomposition

The m–χ decomposition is a decomposition proposed by R. K. Raney [46], designed for use with CTLR data, that attempts to separate the total backscattered power received by the SAR sensor into three components: a "red" component, R, corresponding to double-bounce

Table 10.1: An overview of the Radarsat-2 fine-quad scenes used for ship detection. Beam refers to the fine quad beam position, Date for the acquisition date, Inc. Angle for the incidence angle range of the scene (in degrees), and W.S. for the estimated wind speed for the scene.

Scene ID	Date Acquired	Beam	Inc. Angle	W.S. (m/s)
1126-1	2008-11-26	2	19.77°-21.78°	7.1
1126-2	2008-11-26	2	19.77°-21.78°	6.7
1020-1	2008-10-20	2	19.77°-21.78°	3.6
1220-1	2008-12-20	4	22.16°-24.08°	5.8
1220-12	2008-12-20	4	22.16°-24.08°	5.7
0729-1	2008-07-29	4	22.16°-24.08°	6.0
1120-1	2008-11-20	8	26.88°-28.71°	11.3
1120-2	2008-11-20	8	26.88°-28.71°	11.8
1127-1	2008-11-27	12	31.34°-33.03°	11.9
1127-2	2008-11-27	12	31.34°-33.03°	13.8
0214-1	2009-02-14	21	40.17°-41.61°	13.3
0214-2	2009-02-14	21	40.17°-41.61°	12.2
0209-1	2009-02-09	21	40.17°-41.61°	6.2
0209-2	2009-02-09	21	40.17°-41.61°	<3

backscatter (such as scattering from urban areas, or the interface between the hull of a ship and the ocean); a "green" component, G, corresponding to random, depolarized backscatter (such as from vegetation, or some freshwater ice); and a "blue" component, B, corresponding to surface or Bragg scattering (such as from the ocean).

These three parameters are calculated as follows [46]:

$$R = \sqrt{mS_0(1-\sin(2\chi))/2} \tag{10.1}$$

$$G = \sqrt{S_0(1-m)} \tag{10.2}$$

$$B = \sqrt{mS_0(1+\sin(2\chi))/2} \tag{10.3}$$

Note that the sign in front of $\sin(2\chi)$ has been adjusted from the original equations in order to be consistent with a right-circular transmit polarization, such that a positive χ value (and

therefore a positive $\sin(2\chi)$ value) represents left-handed rotation, and is therefore indicative of surface scattering.

Also note that $S_0 = R^2 + G^2 + B^2$.

Since this decomposition aims to identify the type of backscattering taking place in each pixel of the image, it seems that it has potential for target discrimination applications, where we would expect targets to differ in terms of their dominant backscattering mechanisms. We will therefore explore the use of the m–χ decomposition parameters for target discrimination, alongside other familiar parameters that have been used previously in this thesis.

10.2 Discrimination Methodology

To extract the targets for discrimination, we performed detection on the CTLR and HH-HV data using an adaptive FAR ranging from 10^{-5} down to 2.5×10^{-10}, with a constant FAR of 2.5×10^{-9} for pixels with $m < 0.3$. We also performed detection on the CTLR data using a constant FAR of 2.5×10^{-9}, approximately equal to the average ocean FAR in the adaptive case, in order to compare the discrimination results of the constant and adaptive FAR methods. For each mode, adjacent 8-connected detected pixels were clustered together, and clusters smaller than four pixels in size were discarded.

The clusters of detected pixels that corresponded to validated target locations were extracted from the original SLC imagery, and the data for those targets were saved as separate image chips. The Stokes vector and covariance matrix elements for these targets were then recalculated, using a 3×3 spatial averaging window, but also ensuring that detected pixels were not averaged with non-detected pixels and vice versa. This meant that near the edges of targets, not all pixels within the 3×3 window were used—only pixels that were part of the target cluster. For each target the mean values for the following parameters were calculated: σ^0 (of both channels), S_0, S_1, S_2, S_3, m, ψ, χ, as well as the three parameters of the m–χ decomposition [46], as described above.

Various combinations of these features were used as input data for a support vector machine (SVM) classifier with a Gaussian radial basis function (RBF) kernel (with $\gamma = 1$), a classifier widely used in the literature (see, for example, the work of Scholkopf et al. [47] or Vapnik [48]). Previous work by Howell et al. [7] used a maximum likelihood classifier (MLC) to discriminate ships and icebergs, but to use the MLC, the probability distribution functions of the ships and icebergs must either be known, or assumed (e.g., assumed to be Gaussian). The SVM has an advantage in that knowledge or assumptions about the PDF of the targets is not required. Rather than attempting to calculate the log-likelihood for each target class, the SVM attempts to divide the feature space using a hyperplane such that each class will reside entirely on its own side of the plane. The SVM is also fairly computationally efficient, provided the amount of training data and the number of input parameters is kept relatively small. As well, through the use of different kernel functions, the SVM can be adapted to a wide variety of problems. Here we use the Gaussian RBF kernel, which has the following form [47]:

$$K(\vec{x}, \vec{x_i}) = \exp\left(-\gamma ||\vec{x} - \vec{x_i}||^2\right) \tag{10.4}$$

The Gaussian RBF was chosen due to its wide use in the literature and also due the fact that its value is based on the squared Euclidean distance between the input vector of the target, $\vec{x}$, and the vector corresponding to some centre position $\vec{x_i}$ (such as the centre of a class, or the origin). This makes sense when attempting to discriminate between ships and icebergs. Generally we would expect ships to have higher amounts of backscattered power than icebergs, such that when using power-based parameters as input to the SVM, the ship targets would be farther from the origin of the feature space than iceberg targets.

All input parameters were scaled to have zero mean and unit variance before training.

Classification accuracy was measured using stratified 10-fold cross-validation (see, for example, Kohavi [49] or Salzberg and Fayyad [50]), wherein the data were randomly split into ten different groups, or folds, each containing an approximately equal number of ships

and icebergs. The SVM was run ten times, using each fold once for testing, and the remaining nine folds for training. The reported accuracies in the next section are the average accuracies across the ten testing folds. Similarly, the reported confusion matrices are the sums of the ten testing confusion matrices.

At first θ, the incidence angle, was used as another input parameter to the SVM. Figure 10.1 shows 2-D SVM training for the medium resolution imaging mode, using CTLR m–χ G and θ as features. However, there were two concerns with using the incidence angles directly: (i) since the ships and icebergs were from different scenes, their incidence angles would not match exactly, leading to potential overfitting; and (ii) there were no icebergs with incidence angles above 40° in the data set, causing the model to think that any high incidence angle target should automatically be classified as a ship. To solve this problem, θ was rounded to the nearest ten degrees before using it as input to the SVM. While this reduced the potential classification accuracy, it avoided any potential overfitting, and compensated for the lack of iceberg targets with $\theta > 40°$ by rounding all targets with $35° < \theta < 45°$ into a single 40° incidence angle category. When θ is used as input to the SVM in the next section, it is always this rounded θ that is used rather than the actual value.

10.3 Discrimination Results

The discrimination process was performed a number of different times for each imaging and polarization mode (and, for the CTLR, both adaptive and constant FAR), each time using different features as input to the SVM, in an attempt to determine the best overall feature space.

Once the targets were classified, the classification accuracy was quantified in the form of a confusion matrix, also known as a contingency table or error matrix, which takes the form of a table where each row represents the number of actual objects in each class, and each column represents the number of objects classified as each class. The classification accuracy

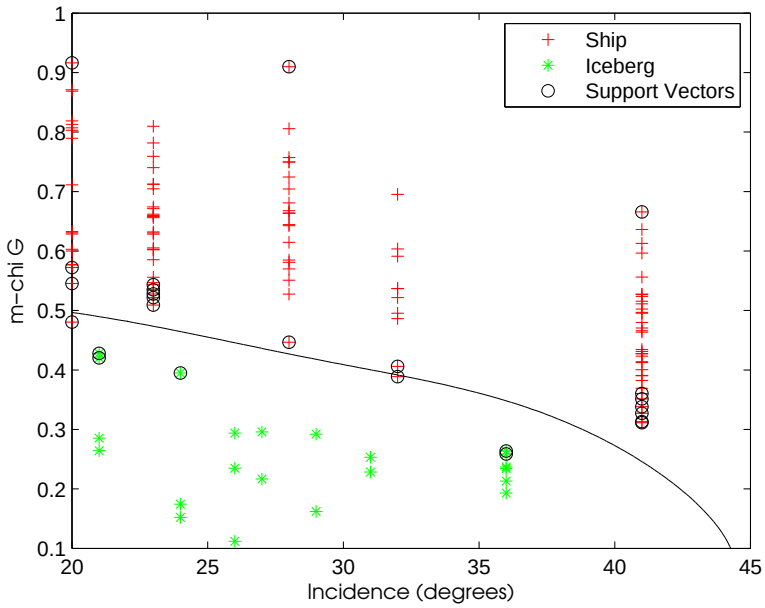

Figure 10.1: 2-D SVM training using m–χ G (equation 10.2) and the incidence angle as features. The decision boundary is shown as the black line.

can be summarized using a number of accuracy measures which are common in the literature (see a review by Congalton for reference [51]):

Overall Accuracy

The number of correctly classified targets divided by the total number of targets. From a confusion matrix, can be calculated by taking the sum of the diagonal elements in the matrix, then dividing that by the sum of all elements in the matrix [51]. This is a simple measure that tells the percentage of targets that were correctly classified.

User's Accuracy

The number of correctly classified targets in a given class, divided by the number of targets classified as that class. Can be calculated by looking at a single column of a confusion matrix, and dividing the correctly classified element in that column by the sum of the elements in that column [51]. This

measure is called the user's accuracy because it is a measure of accuracy from the user's perspective—of the number of targets they have predicted to be within a certain class, how many of those targets are actually of that class?

Producer's Accuracy

The number of correctly classified targets in a given class, divided by the number of actual targets of that class. Can be calculated by looking at a single row of a confusion matrix, and dividing the the correctly classified element in that row by the sum of the elements in that row [51]. This measure is called the producer's accuracy because it is a measure of accuracy from the perspective of the producer of the classifier—of the number of targets that are actually within a certain class, how many of those targets will then be correctly classified?

Kappa Coefficient

Also referred to as Cohen's kappa, this measure adjusts the overall accuracy to account for the random chance agreement. A value of zero indicates that the classifier is no better than random chance, while a positive value of one indicates perfect agreement. Can be calculated using the equation [51]:

$$\kappa = \frac{N \sum_{i=1}^{r} x_{ii} - \sum_{i=1}^{r} (x_{i+} x_{+i})}{N^2 - \sum_{i=1}^{r} (x_{i+} x_{+i})} \tag{10.5}$$

x_{ii} is the number of targets in row i and column i, x_{i+} and x_{+i} are the sums of row i and column i, respectively, r is the number of rows in the confusion matrix, and N is the total number of targets.

10.3.1 Low Resolution Imaging Mode Results

The summarized discrimination accuracies for the low resolution imaging mode are shown in Table 10.2. The CTLR constant FAR data is shown in the first section, with the CTLR

adaptive FAR in the second section, and the HH-HV adaptive FAR in the third section. The best feature set for each section is shaded grey.

For the CTLR data, the highest overall accuracy was yielded by a feature space of the three m–χ parameters, plus the incidence angle θ (rounded to the nearest ten degrees as described above). This feature space yielded an accuracy of 95.0% for the constant FAR detection, and a Kappa coefficient of 0.81. Unfortunately, it seems that use of the adaptive FAR slightly decreased the discrimination accuracy, such that the same feature space yielded an accuracy of 94.2% for the adaptive FAR detection.

In comparison, the HH-HV mode yielded a maximum accuracy of 89.8% for a feature space of m–χ R and G.

The confusion matrices for the highest accuracy feature set for each polarimetric mode and FAR are shown as Tables 10.3 (for the CTLR constant FAR), 10.4 (for the CTLR adaptive FAR), and 10.5 (for the HH-HV adaptive FAR). For the adaptive FAR CTLR, six ships were incorrectly classified as icebergs, and two icebergs were incorrectly classified as ships. In comparison, for the HH-HV data, nine ships were classified as icebergs, and four icebergs were classified as ships. This is in addition to the fact that the HH-HV mode detected less targets than the CTLR in the first place, so that it had less targets to classify. The CTLR data had 119 ships and 20 icebergs to classify, while the HH-HV only had 111 ships and 17 icebergs.

10.3.2 Medium Resolution Imaging Mode Results

Table 10.6 shows the results for the medium resolution imaging mode. An overall accuracy of 100% for the medium resolution data was achieved using the CTLR constant FAR mode and a feature space of the m–χ red and green (double-bounce and random) parameters alongside θ. For the adaptive FAR CTLR data, the highest accuracy was 99.3%, using the same feature space as in the constant FAR case.

The most accurate results for the HH-HV polarization were achieved using a feature

space of HH σ^0, HV σ^0, S_2, S_3, and θ, with an overall accuracy of 96.5%. While the overall accuracies are high, due to the much larger number of ships than icebergs in the data set the user's accuracy for icebergs will be lower. For the CTLR adaptive FAR data, the user's accuracy for icebergs is 96%, while for the HH-HV adaptive FAR, this accuracy is reduced to 90.5%.

The confusion matrices for the highest accuracy features sets are shown in Tables 10.7 (for the constant FAR CTLR), 10.8 (for the adaptive FAR CTLR), and 10.9 (for the adaptive FAR HH-HV). In the adaptive FAR CTLR case one ship target was incorrectly classified as an iceberg target, while all other targets were classified correctly. For the HH-HV case, two ship targets were classified as icebergs, and three icebergs were classified as ships. In addition to this, as with the low resolution imaging mode data, the CTLR data was working with a larger set of targets due to its superior detection performance. The CTLR data needed to classify 123 ships and 24 icebergs, while the HH-HV data only needed to classify 121 ships and 22 icebergs.

Misclassifications were less, as expected, than in the low resolution imaging mode. The increase in classification accuracy when going from the low resolution imaging mode to the medium resolution imaging mode is about 5% in the CTLR case, and about 6% in the HH-HV case. The CTLR compact polarimetry produces more accurate discrimination results than the HH-HV for all of the feature sets tested, for both the medium and low resolution imaging modes. The m–χ decomposition seems to be a strong choice for use as input to the SVM, yielding the most accurate discrimination performance of the CTLR data for both the medium and low resolution imaging modes.

Unfortunately, as in the low resolution imaging mode, use of the adaptive FAR results in a small decrease in discrimination accuracy for the CTLR mode data. This is a sign that detecting a greater number of target pixels can be detrimental to the discrimination performance, which is a direction that could be investigated further—rather than averaging

the polarimetric parameters across the whole of each target cluster, the parameters could be averaged over the centre of the targets, or near the peak of their backscattered power. The CTLR compact polarimetry data seems to be a strong choice for this application compared to the linear dual-pol HH-HV, with higher classification accuracies across all of the data tested.

Table 10.2: Discrimination accuracies for the RCM low resolution imaging mode. Each row represents a different set of features used as input to the SVM, with the best feature set for each polarization shaded in grey. The top section of the table shows results using CTLR compact polarimetry (adaptive and constant FAR), and the bottom section shows results using HH-HV linear dual-pol (adaptive FAR only).

Features	Overall Accuracy	User's Accuracies (Ships, Icebergs)	Producer's Accuracies (Ships, Icebergs)	Kappa Coef.
CTLR Mode (Constant FAR)				
m–χ R, m–χ G, θ	0.899	0.973, 0.607	0.908, 0.850	0.650
m–χ R, m–χ G	0.907	0.973, 0.630	0.916, 0.850	0.669
m–χ R, m–χ G, m–χ B, θ	0.950	0.983, 0.783	0.958, 0.900	0.808
S_0, S_1, S_2, S_3, θ	0.921	0.991, 0.655	0.916, 0.950	0.729
RH σ^0, RV σ^0, θ	0.928	0.982, 0.692	0.933, 0.900	0.740
RH σ^0, RV σ^0, S_2, S_3	0.899	0.982, 0.600	0.899, 0.900	0.662
RH σ^0, RV σ^0, S_2, S_3, θ	0.942	0.991, 0.731	0.941, 0.950	0.792
CTLR Mode (Adaptive FAR)				
m–χ R, m–χ G, θ	0.885	0.964, 0.571	0.899, 0.800	0.599
m–χ R, m–χ G	0.921	0.982, 0.667	0.924, 0.900	0.720
m–χ R, m–χ G, m–χ B, θ	0.942	0.983, 0.750	0.950, 0.900	0.784
S_0, S_1, S_2, S_3, θ	0.935	0.991, 0.704	0.933, 0.950	0.771
RH σ^0, RV σ^0, θ	0.921	0.991, 0.655	0.916, 0.950	0.729
RH σ^0, RV σ^0, S_2, S_3	0.907	0.982, 0.621	0.908, 0.900	0.680
RH σ^0, RV σ^0, S_2, S_3, θ	0.921	0.991, 0.655	0.916, 0.950	0.729
HH-HV Mode (Adaptive FAR)				
m–χ R, m–χ G, θ	0.875	0.970, 0.519	0.883, 0.824	0.567
m–χ R, m–χ G	0.898	0.962, 0.591	0.919, 0.765	0.608
m–χ R, m–χ G, m–χ B, θ	0.891	0.971, 0.560	0.901, 0.824	0.604
S_0, S_1, S_2, S_3, θ	0.805	0.939, 0.367	0.829, 0.647	0.360
HH σ^0, HV σ^0, θ	0.867	0.980, 0.500	0.865, 0.882	0.564
HH σ^0, HV σ^0, S_2, S_3	0.867	0.961, 0.500	0.883, 0.765	0.529
HH σ^0, HV σ^0, S_2, S_3, θ	0.875	0.961, 0.520	0.892, 0.765	0.548

Table 10.3: Confusion matrix for the low resolution imaging mode, using CTLR data (constant FAR) and a feature space of m–χ R, m–χ G, m–χ B, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	114	5
Actual Iceberg	2	18

Table 10.4: Confusion matrix for the low resolution imaging mode, using CTLR data (adaptive FAR) and a feature space of m–χ R, m–χ G, m–χ B, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	113	6
Actual Iceberg	2	18

Table 10.5: Confusion matrix for the low resolution imaging mode, using HH-HV data (adaptive FAR) and a feature space of m–χ R, m–χ G, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	102	9
Actual Iceberg	4	13

Table 10.6: Discrimination accuracies for the RCM medium resolution imaging mode. Each row represents a different set of features used as input to the SVM, with the best feature set for each polarization shaded in grey. The top section of the table shows results using CTLR compact polarimetry (adaptive and constant FAR), and the bottom section shows results using HH-HV linear dual-pol (adaptive FAR only).

Features	Overall Accuracy	User's Accuracies (Ships, Icebergs)	Producer's Accuracies (Ships, Icebergs)	Kappa Coef.
CTLR Mode (Constant FAR)				
m–χ R, m–χ G, θ	1.000	1.000, 1.000	1.000, 1.000	1.000
m–χ R, m–χ G	0.973	1.000, 0.857	0.968, 1.000	0.907
m–χ R, m–χ G, m–χ B, θ	0.993	0.992, 1.000	1.000, 0.958	0.975
S_0, S_1, S_2, S_3, θ	0.952	0.975, 0.840	0.968, 0.875	0.829
RH σ^0, RV σ^0, θ	0.966	0.992, 0.852	0.968, 0.958	0.882
RH σ^0, RV σ^0, S_2, S_3	0.946	0.992, 0.767	0.943, 0.958	0.819
RH σ^0, RV σ^0, S_2, S_3, θ	0.973	0.992, 0.885	0.976, 0.958	0.904
CTLR Mode (Adaptive FAR)				
m–χ R, m–χ G, θ	0.993	1.000, 0.960	0.992, 1.000	0.976
m–χ R, m–χ G	0.959	1.000, 0.800	0.951, 1.000	0.864
m–χ R, m–χ G, m–χ B, θ	0.986	0.992, 0.958	0.992, 0.958	0.950
S_0, S_1, S_2, S_3, θ	0.959	0.984, 0.846	0.968, 0.917	0.856
RH σ^0, RV σ^0, θ	0.966	0.984, 0.880	0.976, 0.917	0.878
RH σ^0, RV σ^0, S_2, S_3	0.946	0.992, 0.767	0.943, 0.958	0.946
RH σ^0, RV σ^0, S_2, S_3, θ	0.980	0.992, 0.920	0.984, 0.958	0.927
HH-HV Mode				
m–χ R, m–χ G, θ	0.916	0.966, 0.692	0.934, 0.818	0.700
m–χ R, m–χ G	0.867	0.964, 0.546	0.876, 0.818	0.576
m–χ R, m–χ G, m–χ B, θ	0.937	0.975, 0.760	0.950, 0.864	0.771
S_0, S_1, S_2, S_3, θ	0.930	0.983, 0.714	0.934, 0.909	0.758
HH σ^0, HV σ^0, θ	0.923	0.974, 0.704	0.934, 0.864	0.730
HH σ^0, HV σ^0, S_2, S_3	0.888	0.965, 0.600	0.900, 0.818	0.626
HH σ^0, HV σ^0, S_2, S_3, θ	0.965	0.975, 0.905	0.984, 0.864	0.863

Table 10.7: Confusion matrix for the medium resolution imaging mode, using CTLR data (constant FAR) and a feature space of m–χ R, m–χ G, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	123	0
Actual Iceberg	0	24

Table 10.8: Confusion matrix for the medium resolution imaging mode, using CTLR data (adaptive FAR) and a feature space of m–χ R, m–χ G, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	122	1
Actual Iceberg	0	24

Table 10.9: Confusion matrix for the medium resolution imaging mode, using HH-HV data (adaptive FAR) and a feature space of HH σ^0, HV σ^0, S_2, S_3, and θ.

	Predicted Ship	Predicted Iceberg
Actual Ship	119	2
Actual Iceberg	3	19

Chapter 11

Conclusions and Future Work

This thesis has explored methods for the automated detection of icebergs in compact polarimetric SAR imagery, with minimal user input required. By using Radarsat-2 fine-quad imagery to simulate some of the wide swath imaging modes that will be available once the Radarsat Constellation is launched, a first look at the expected iceberg detection performance of this data has been provided. Overall, the results are very positive for the compact polarimetry. For the three imaging modes tested, the compact polarimetry consistently detected a greater number of iceberg targets than the linear dual-pol modes. The compact polarimetry also detected a greater number of pixels within each target than the linear dual-pol data. As well, experiments on the discrimination performance of the compact polarimetry and linear HH-HV showed that the compact polarimetry data also produces higher classification accuracy when discriminating between ship and iceberg targets.

A summary of the author's main contributions is as follows, in roughly the same order the topics were presented in this thesis:

- The iceberg detection performance of compact polarimetric SAR compared to linear dual-pol data was investigated. This is the first such investigation in the literature.

- A model of N vs. incidence angle, suitable for pseudo quad-pol reconstruction of ocean imagery, was proposed. Unfortunately, the signal level of the reconstructed pseudo-HV is below the noise floor of the planned Radarsat Constellation imaging modes.

- The values of the Stokes vector elements and a number of derived parameters were investigated for both ocean and iceberg targets.

- A modified form of the Liu et al. [30, 31] likelihood ratio test detector was proposed, based on the skew-normal probability distribution function of Azzalini [40] rather than the multivariate normal (Gaussian) distribution.

- A new method was proposed through which the orientation and ellipticity angles of the received polarization ellipse were used to vary the false alarm rate (i.e., the detection threshold) from pixel-to-pixel across the SAR image.

- The ability of compact polarimetric and dual-pol HH-HV data to discriminate between ships and icebergs was tested using a support vector machine classifier. This is the first investigation of ship and iceberg discrimination using compact polarimetric SAR data in the literature.

There were a number of challenges this thesis wished to address involving the use of compact polarimetric SAR data. First was the fact that compact polarimetry can be used in a wide variety of ways, such as through pseudo quad-pol reconstruction, or in the derivation of the Stokes parameters or polarization ellipse angles. As well, since the use of compact polarimetry for remote sensing of the Earth is relatively new, less is understood than for linear dual-pol data in terms of how to best utilize the compact polarimetric information.

The detection performance of pseudo quad-pol data was assessed, and unfortunately it was found that for the Radarsat Constellation data, the results produced by the pseudo quad-pol data is inferior to that of the native compact polarimetry channels. A likely reason for this is due to the higher noise floor of the Radarsat Constellation data. When working with Radarsat-2 data, the detection performance of the pseudo quad-pol data was quite favourable, particularly the use of the reconstructed pseudo-HV parameter [10]. Unfortunately, since this parameter is a cross-polarized parameter, its signal-to-noise ratio tends to be quite low, and will almost always lie below the noise floor of the Radarsat Constellation for ocean pixels. This will make the use of cross-pol parameters, and therefore the pseudo quad-pol data in general, problematic for data collected by this SAR sensor. However, it

is possible that the pseudo-HV could be used for target discrimination applications (where the backscatter from targets would be much greater than the backscatter from the ocean), or with other SAR sensors that feature a lower noise floor (closer to that of the original Radarsat-2 data). More work could be done in this area, to test the expected values of the pseudo-HV for various types of targets (e.g., icebergs, ships, and sea ice).

Use of an LRT decision variable derived from the skew-normal distribution, rather than the normal Gaussian distribution, was tested for use with Stokes vector data as input. The skew-normal LRT improved detection performance of the Stokes vector data for the low resolution imaging mode, but results were mixed for the medium resolution and ship detection imaging modes.

A method was proposed wherein the properties of the received polarization ellipse were used to adjust the detection threshold used for each pixel. This method reduced the number of targets missed by the CTLR data for the low resolution mode, and reduced the overall median PMD values for the other imaging modes. This method attempts to account for the fact that in general, many studies in the literature treat the ocean clutter as a single variable PDF, independent of any other parameters. But in fact, this thesis showed that the vast majority of the ocean pixels with high decision variable values (that is, potential false alarm pixels) tend to have predictable polarization ellipse shapes which are often different than the polarization ellipse behaviour from actual targets. By taking advantage of the fact that with compact polarimetry, we receive a polarization ellipse unaffected by the preferred orientation of the transmit polarization, the ability to detect icebergs can be improved while using a lower overall false alarm rate over the ocean regions of the image.

Finally, some scenes containing ship targets were obtained, and the ability of the CTLR and HH-HV modes to discriminate between ships and icebergs was assessed and compared. It was found that the CTLR data also yields superior discrimination performance in addition to detecting a greater number of targets in the first place. The parameters of the m–χ

decomposition proposed by R. K. Raney [46] were found to be particularly well suited to this application. The study was somewhat limited due to the relatively small number of iceberg targets available compared to ship targets. Future work could be done using a greater amount of iceberg data for training and testing of the classifier. Collection of iceberg data from different seasons, since the properties of the ice (and therefore its backscattering behaviour) can vary significantly from year to year, could also be helpful. As well, while the results of ship and iceberg discrimination have been presented, it is also important to test discrimination methods between icebergs and sea ice, which could be a challenging problem in the case of thick, multi-year sea ice. Unfortunately, not enough sea ice was present in the available data to investigate this issue, but this could be an area of future work as well.

Overall, this thesis has presented a broad investigation of the use of compact polarimetry for iceberg detection applications—one of the few such investigations into iceberg detection using SAR data present in the open literature, and currently the only investigation using compact polarimetry data. While the best algorithms and methods for use with the compact polarimetry data will no doubt be improved as the availability and use of such data increases, the strength of the compact polarimetry for this application, and for related maritime applications such as ship detection and target discrimination, is clear. The use of compact polarimetry yields a much greater amount of available polarimetric information than linear dual-pol SAR, without the swath width restrictions of fully polarimetric data. This makes it a strong choice for use in maritime surveillance applications.

Bibliography

[1] Wikipedia, "Polarization (waves)." http://en.wikipedia.org/wiki/Polarization_(waves), 2014. [Online; accessed March, 2014].

[2] National Snow & Ice Data Center, "Thicker on top, more down under." http://nsidc.org/arcticseaicenews/2014/02/thicker-on-top-higher-down-under/, 2014. [Online; accessed February, 2014].

[3] European Parliament, "European Parliament resolution of 9 October 2008 on Arctic governance," 2008.

[4] Canadian Geographic Magazine, "In-depth: Singing icebergs." http://www.canadiangeographic.ca/magazine/ma06/indepth/. [Online; accessed March, 2014].

[5] E. Lewis, B. Currie, and S. Haykin, *Detection and Classification of Ice*. John Wiley and Sons, Inc., 1987.

[6] C. Willis, J. Macklin, K. Partington, K. Teleki, W. Rees, and R. Williams, "Iceberg detection using ERS-1 synthetic aperture radar," *International Journal of Remote Sensing*, vol. 17, no. 9, pp. 1777–1795, 1996.

[7] C. Howell, J. Youden, K. Lane, D. Power, C. Randell, and D. Flett, "Iceberg and ship discrimination with ENVISAT multi-polarization ASAR," in *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, pp. 113–116, IEEE, 2004.

[8] C. Howell, J. Mills, D. Power, J. Youden, K. Dodge, C. Randell, S. Churchill, and D. Flett, "A multivariate approach to iceberg and ship classification in HH/HV ASAR data," in *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 3583–3586, IEEE, 2006.

[9] C. Howell, D. Power, M. Lynch, K. Dodge, P. Bobby, C. Randell, P. Vachon, and G. Staples, "Dual polarization detection of ships and icebergs – recent results with ENVISAT ASAR and data simulations of Radarsat-2," in *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 206–209, 2008.

[10] M. Denbina and M. J. Collins, "Iceberg detection using pseudo quad-pol reconstruction of compact polarimetric SAR," *Atmosphere-Ocean*, vol. 50, no. 4, pp. 437–446, 2012.

[11] J.-S. Lee and E. Pottier, *Polarimetric Radar Imaging: From Basics to Applications*. Boca Raton, FL: CRC Press, 2009.

[12] B. J. O'Connell, "Marine radar for improved ice detection." `http://www.ccg-gcc.gc.ca/e0014403`, 2013. [Online; accessed November, 2014].

[13] K. Lane, D. Power, J. Youden, C. Randell, and D. Flett, "Validation of synthetic aperture radar for iceberg detection in sea ice," in *Geoscience and Remote Sensing Symposium, 2004. IGARSS '04. Proceedings. 2004 IEEE International*, vol. 1, pp. – 128, Sept 2004.

[14] C. Wesche and W. Dierking, "Iceberg signatures and detection in sar images in two test regions of the weddell sea, antarctica," *Journal of Glaciology*, vol. 58, pp. 325–339, Jan. 2012.

[15] J.-W. Kim, D. jin Kim, S.-H. Kim, and B.-J. Hwang, "Iceberg detection using full-polarimetric radarsat-2 sar data in west antarctica," in *Synthetic Aperture Radar (APSAR), 2011 3rd International Asia-Pacific Conference on*, pp. 1–4, Sept 2011.

[16] W. Dierking and C. Wesche, "C-band radar polarimetry—useful for detection of icebergs in sea ice?," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 52, pp. 25–37, Jan 2014.

[17] J.-C. Souyris, P. Imbo, R. Fjortoft, S. Mingot, and J.-S. Lee, "Compact polarimetry based on symmetry properties of geophysical media: The $\pi/4$ mode," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 3, pp. 634–646, 2005.

[18] N. Stacy and M. Preiss, "Compact polarimetric analysis of X-band SAR data," in *Proc. European Conference on Synthetic Aperture Radar (EUSAR)*, 2006.

[19] R. Raney, "Hybrid-polarity SAR architecture," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 11, pp. 3397–3404, 2007.

[20] M. Nord, T. Ainsworth, J.-S. Lee, and N. Stacy, "Comparison of compact polarimetric synthetic aperture radar modes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 1, pp. 174–188, 2009.

[21] L. Carter, D. Campbell, and B. Campbell, "Geologic studies of planetary surfaces using radar polarimetric imaging," *Proceedings of the IEEE*, vol. 99, no. 5, pp. 770–782, 2011.

[22] Canadian Space Agency, "RADARSAT Constellation website." `http://www.asc-csa.gc.ca/eng/satellites/radarsat`. [Online; accessed February, 2014].

[23] M. Denbina and M. J. Collins, "Iceberg detection using simulated dual-polarized Radarsat Constellation data," *Canadian Journal of Remote Sensing*, vol. 40, no. 3, pp. 165–178, 2014.

[24] M. J. Collins, M. Denbina, and G. E. Atteia, "On the reconstruction of quad-pol SAR data from compact polarimetry data for ocean target detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 1, pp. 591–600, 2013.

[25] M. Denbina and M. J. Collins, "Iceberg detection using analysis of the received polarization ellipse in compact polarimetry," in *Proc. IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2014.

[26] P. Vachon and J. Wolfe, "C-band cross-polarization wind speed retrieval," *IEEE Geoscience and Remote Sensing Letters*, vol. 8, no. 3, pp. 456–459, 2011.

[27] National Snow and Ice Data Center, "International ice patrol (iip) iceberg sightings database," 1995, updated annually.

[28] J. Inglada, C. Henry, and J.-C. Souyris, "ASAR multi-polarization images phase difference: Assessment in the framework of persistent scatterers interferometry," in *Proceedings of the Envisat & ERS Symposium*, 2004.

[29] C. Oliver and S. Quegan, *Understanding Synthetic Aperture Radar Images*. SciTech Publishing, 2004.

[30] C. Liu and A. Meek, "Likelihood ratio test polarimetric SAR ship detection application." Defence Research and Development Canada, Technical Memorandum, TM 2005-243, 2005.

[31] C. Liu, P. W. Vachon, R. A. English, and N. Sandirasegaram, "Ship detection using radarsat-2 fine quad mode and simulated compact polarimetry data." Defence Research and Development Canada, Technical Memorandum, TM 2009-285, 2010.

[32] N. Redding, "Estimating the parameters of the k distribution in the intensity domain." DSTO Technical Report DSTO-TR-0839, 1999.

[33] G. G. Stokes, "On the composition and resolution of streams of polarized light from different sources," *Trans. Camb. Philosoph. Soc.*, vol. 9, pp. 399–416, 1852.

[34] E. Collett, *Field Guide to Polarization*. SPIE Press, Bellingham, WA, 2005.

[35] R. Raney, "Dual-polarized SAR and stokes parameters," *IEEE Geoscience and Remote Sensing Letters*, vol. 3, no. 3, pp. 317–319, 2006.

[36] W. Cameron, N. Youssef, and L. Leung, "Simulated polarimetric signatures of primitive geometrical shapes," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 34, no. 3, pp. 793–803, 1996.

[37] R. Touzi, F. Charbonneau, R. K. Hawkins, and P. W. Vachon, "Ship detection and characterization using polarimetric SAR," *Canadian Journal of Remote Sensing*, vol. 30, no. 3, pp. 552–559, 2004.

[38] N. Mascarenhas, "An overview of speckle noise filtering in SAR images," in *Image Processing Techniques, First Latino-American Seminar on Radar Remote Sensing*, p. 71, 1996.

[39] S. Shitole, Y. Rao, B. Mohan, A. Bhattacharya, and A. Das, "Region growing based improved SAR speckle filter for polarimetric data," in *Synthetic Aperture Radar (APSAR), 2013 Asia-Pacific Conference on*, pp. 517–520, Sept 2013.

[40] A. Azzalini and A. Capitanio, "Statistical applications of the multivariate skew-normal distribution," *Journal of the Royal Statistical Society*, vol. B, no. 61, pp. 579–602, 1999.

[41] L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Prentice Hall, 1990.

[42] X. Xing, Z. Chen, H. Zou, and S. Zhou, "A fast algorithm based on two-stage CFAR for detecting ships in SAR images," in *Proceedings of the 2nd Asian-Pacific Conference on Synthetic Aperture Radar*, pp. 506–509, 2009.

[43] J. Ai, X. Qi, W. Yu, F. Liu, and L. Shi, "A new CFAR ship detection algorithm based on 2-D joint log-normal distribution in SAR images," *IEEE Geoscience and Remote Sensing Letters*, vol. 7, no. 4, pp. 806–810, 2010.

[44] D. Staglianò, A. Lupidi, and F. Berizzi, "Ship detection from SAR images based on CFAR and wavelet transform," in *Proceedings of the 2012 Tyrrhenian Workshop on*

Advances in Radar and Remote Sensing, 2012.

[45] R. Raney, "Desdyni adopts hybrid polarity sar architecture," in *IEEE Radar Conference – Radar09*, pp. 1–4, 2009.

[46] R. K. Raney, J. T. S. Cahill, G. W. Patterson, and D. B. J. Bussey, "The $m - chi$ decomposition of hybrid dual-polarimetric radar data with application to lunar craters," *Journal of Geophysical Research*, vol. 117, 2012.

[47] B. Scholkopf, K.-K. Sung, C. J. C. Burges, F. Girosi, P. Niyogi, T. Poggio, and V. Vapnik, "Comparing support vector machines with gaussian kernels to radial basis function classifiers," *Signal Processing, IEEE Transactions on*, vol. 45, pp. 2758–2765, Nov 1997.

[48] V. Vapnik, *The Nature of Statistical Learning Theory, 2nd Ed.* Springer Science & Business Media, 2000.

[49] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *Proceedings of 14th International Joint Conference on Artificial Intelligence*, pp. 1137–1143, 1995.

[50] S. L. Salzberg and U. Fayyad, "On comparing classifiers: Pitfalls to avoid and a recommended approach," *Data Mining and Knowledge Discovery*, pp. 317–328, 1997.

[51] R. G. Congalton, "A review of assessing the accuracy of classifications of remotely sensed data," *Remote Sensing of Environment*, pp. 35–46, 1991.

Appendix A

Stokes Vector Density Plots of Ocean Pixels

A.1 Low Resolution Mode (CTLR)

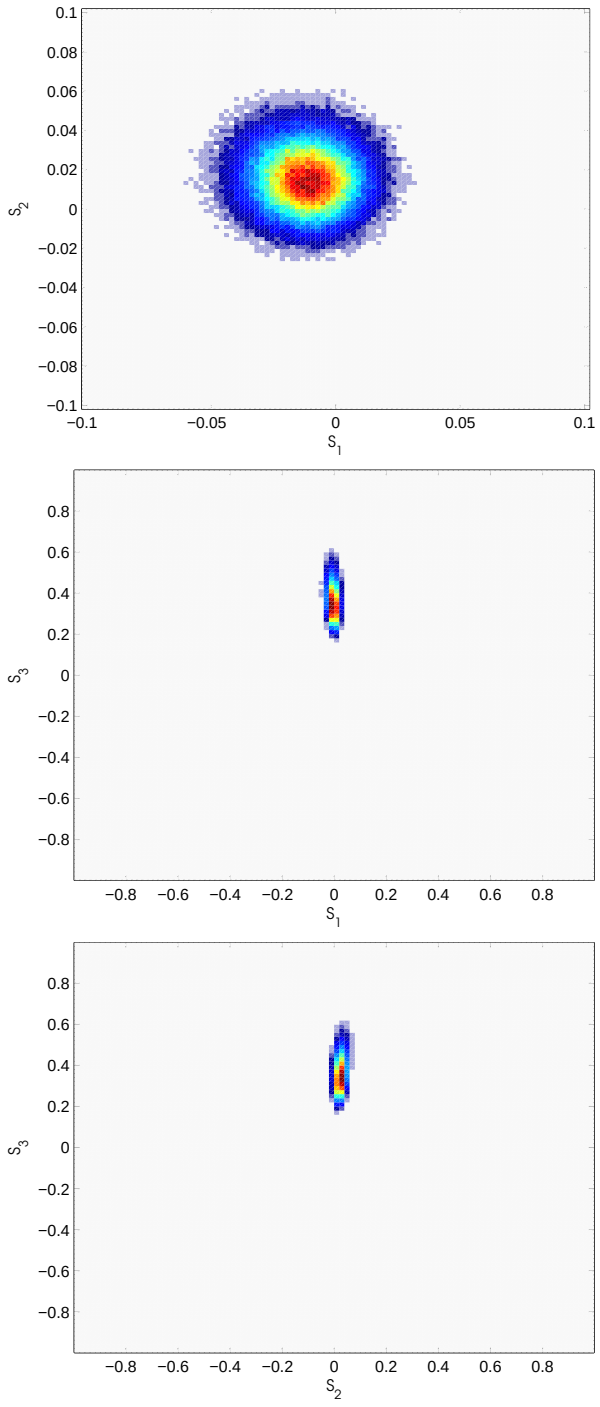

Figure A.1: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-1, low resolution mode, for right-circular transmit polarization.

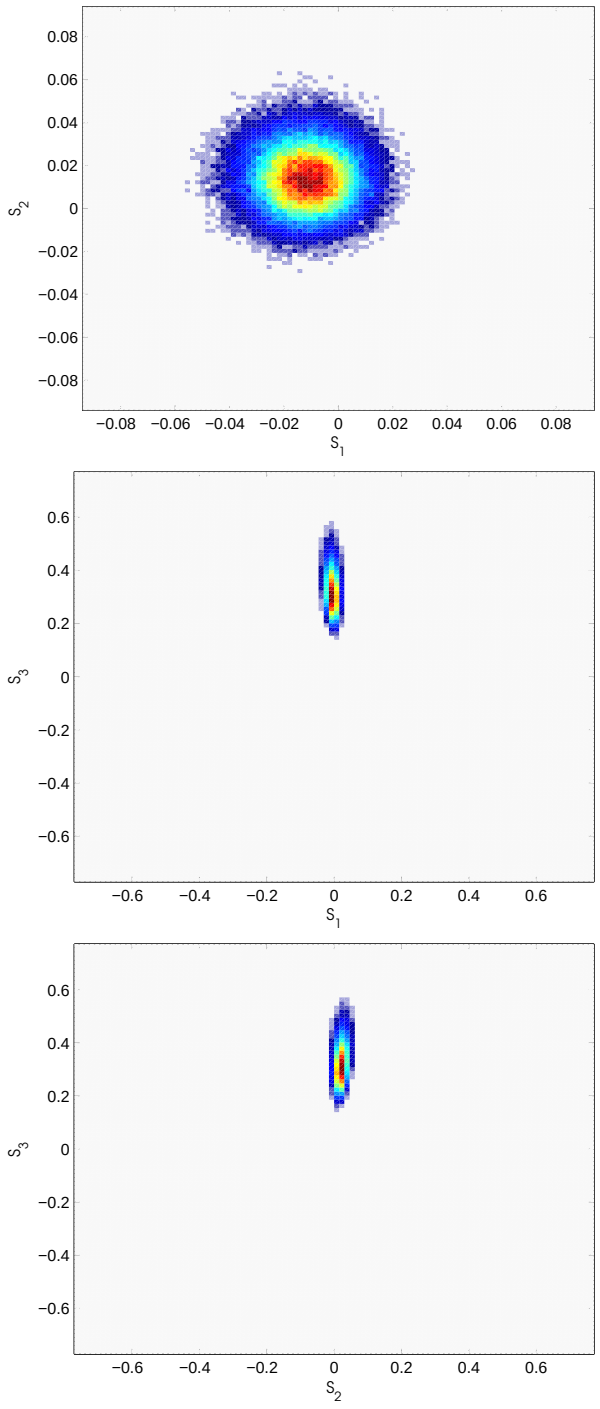

Figure A.2: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-2, low resolution mode, for right-circular transmit polarization.

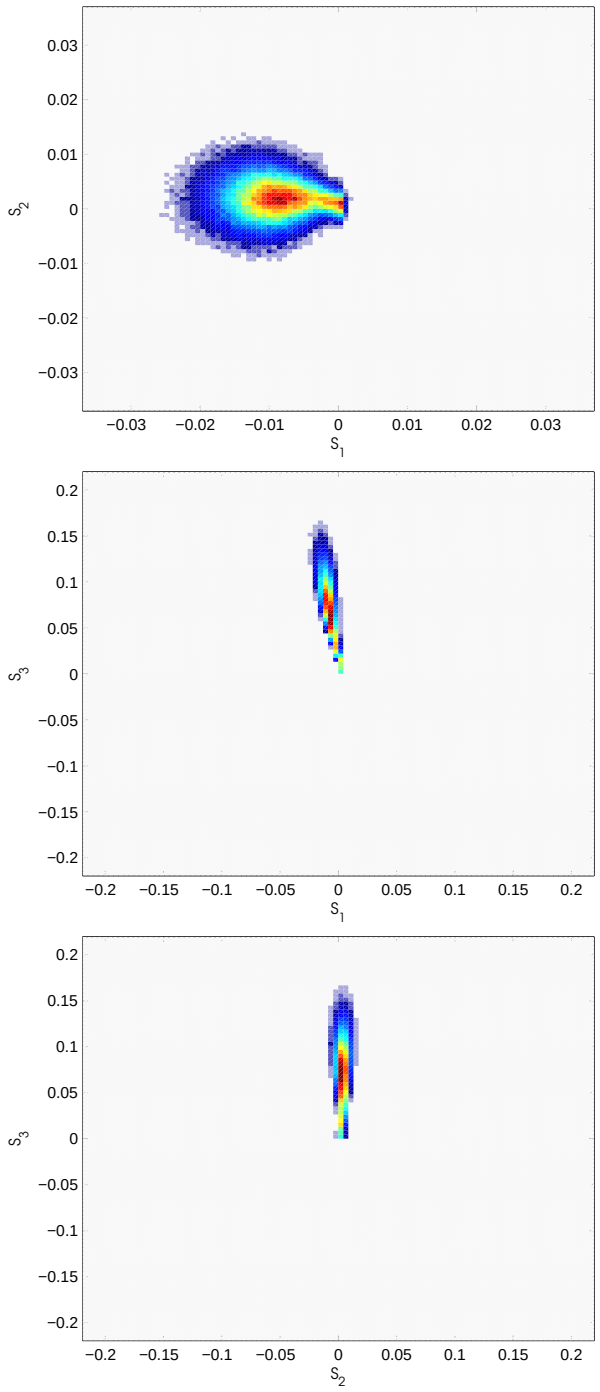

Figure A.3: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0829-1, low resolution mode, for right-circular transmit polarization.

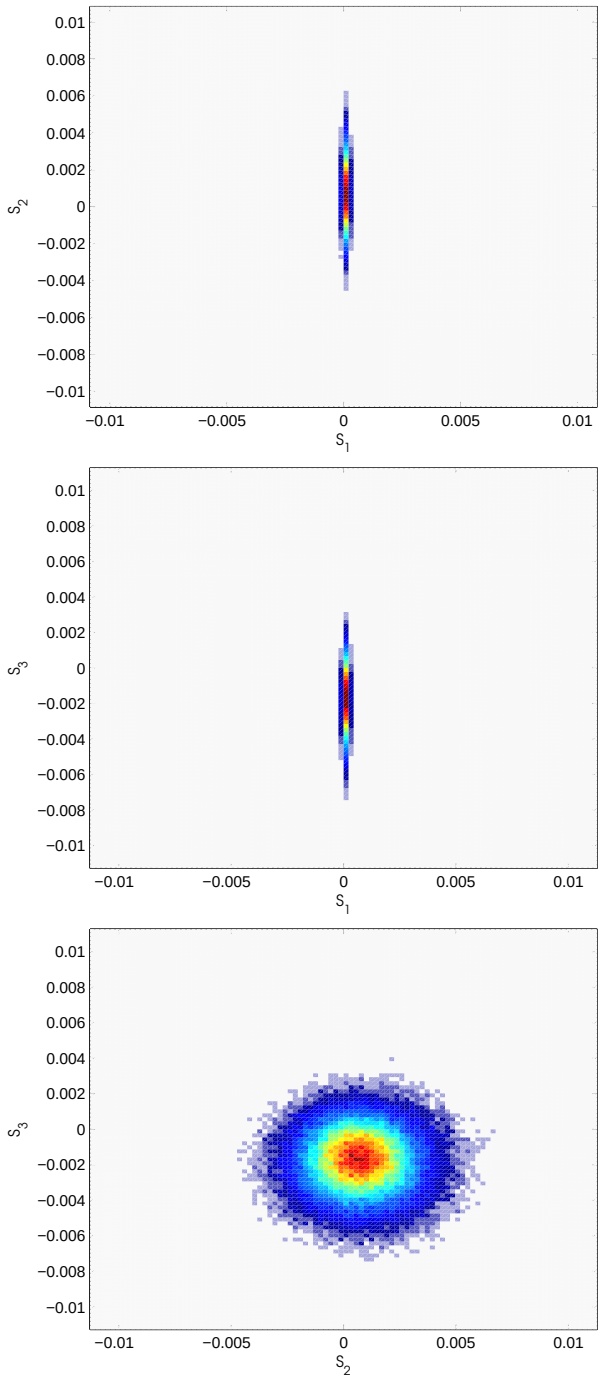

Figure A.4: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-1, low resolution mode, for right-circular transmit polarization.

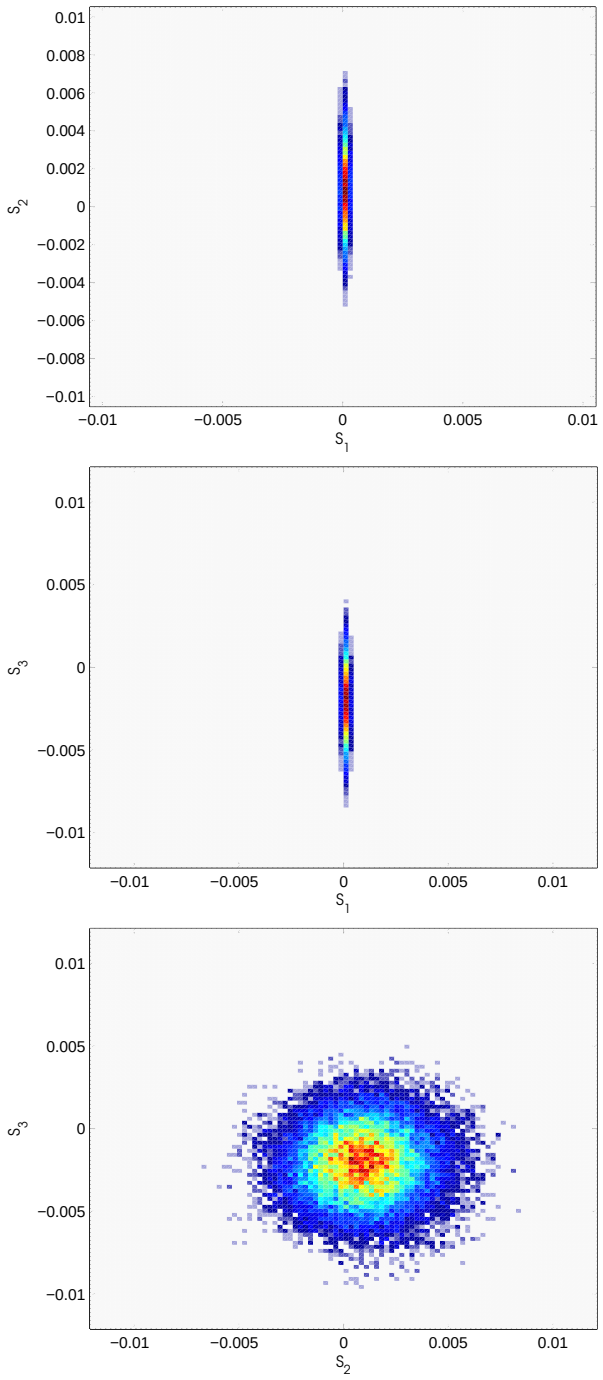

Figure A.5: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-2, low resolution mode, for right-circular transmit polarization.

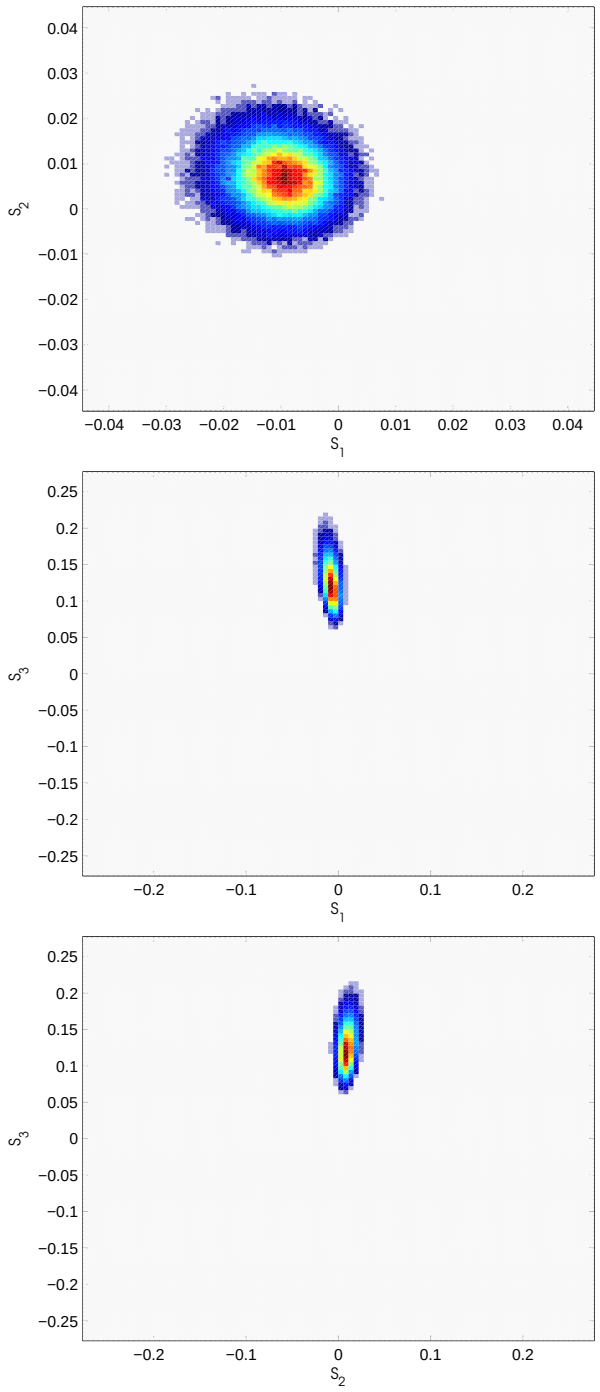

Figure A.6: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-2, low resolution mode, for right-circular transmit polarization.

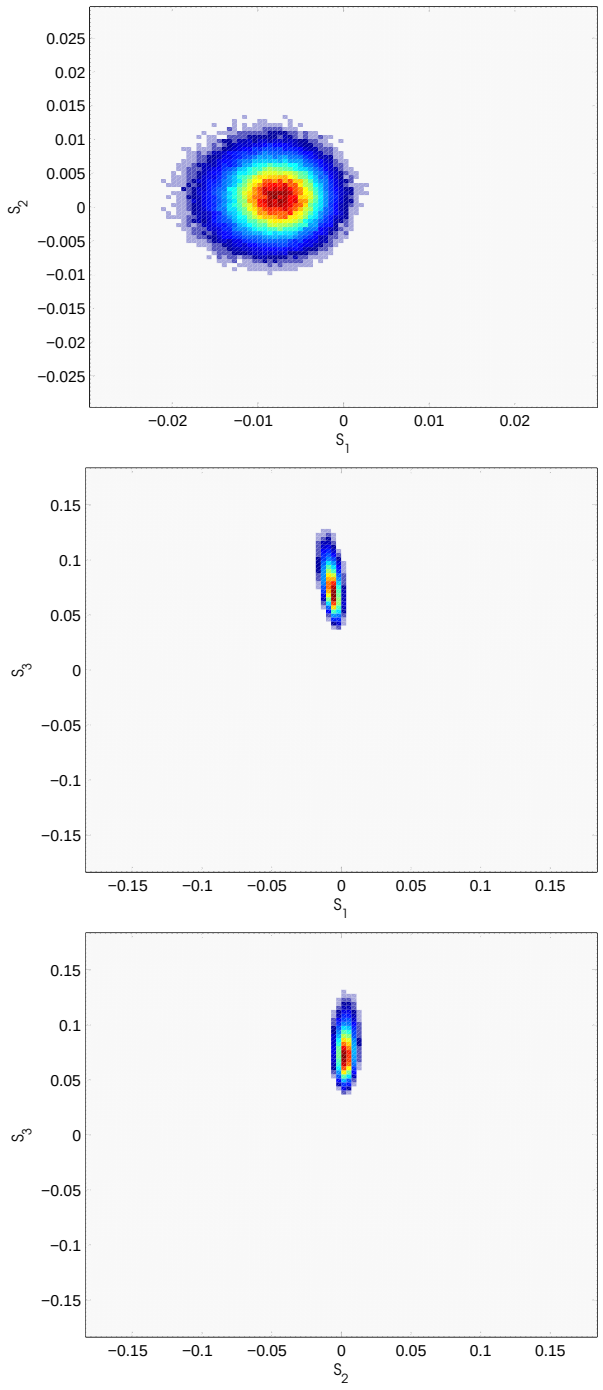

Figure A.7: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-1, low resolution mode, for right-circular transmit polarization.

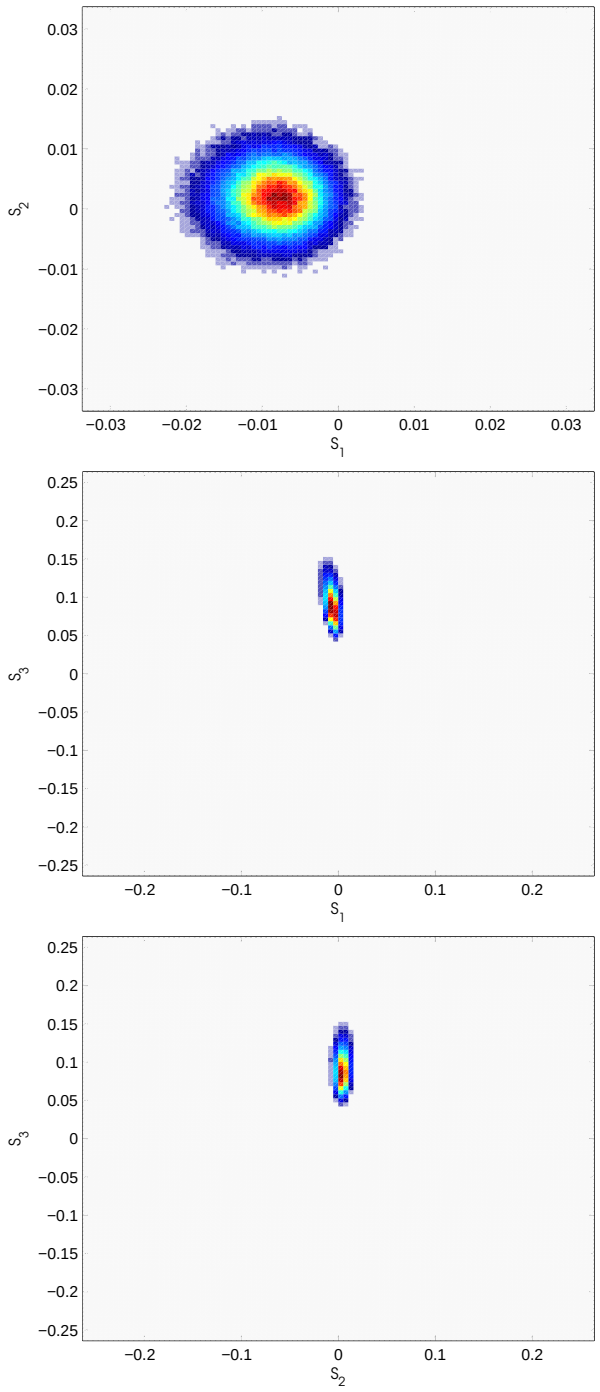

Figure A.8: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-2, low resolution mode, for right-circular transmit polarization.

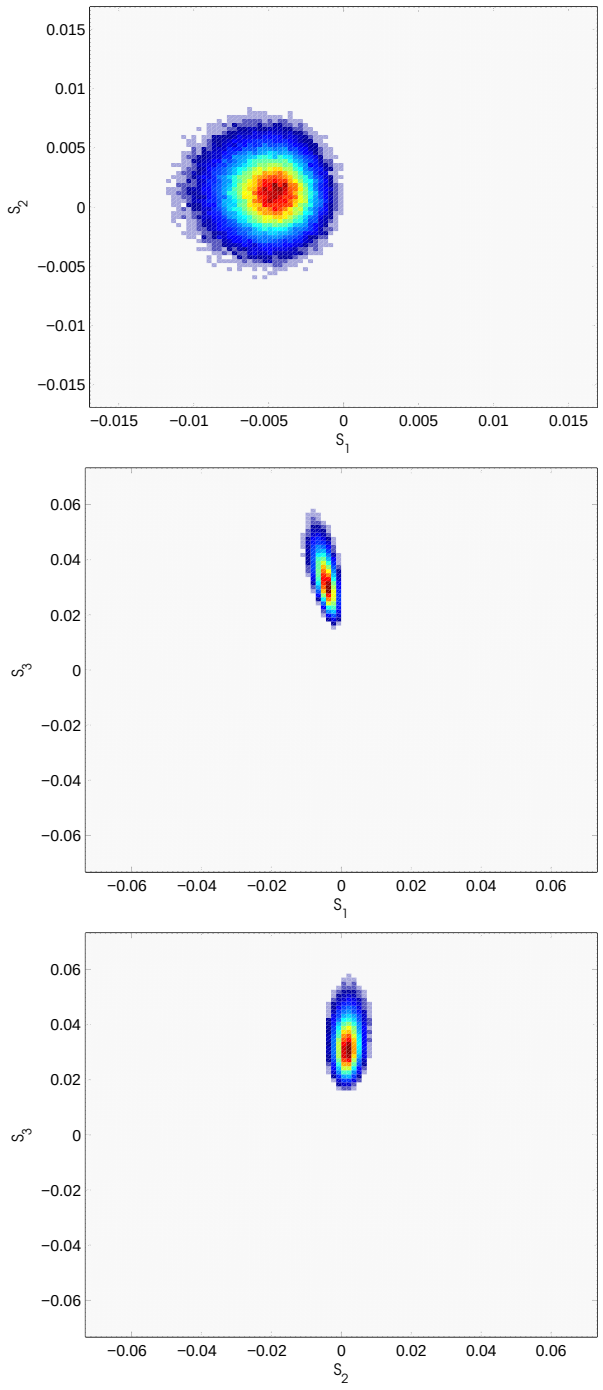

Figure A.9: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-1, low resolution mode, for right-circular transmit polarization.

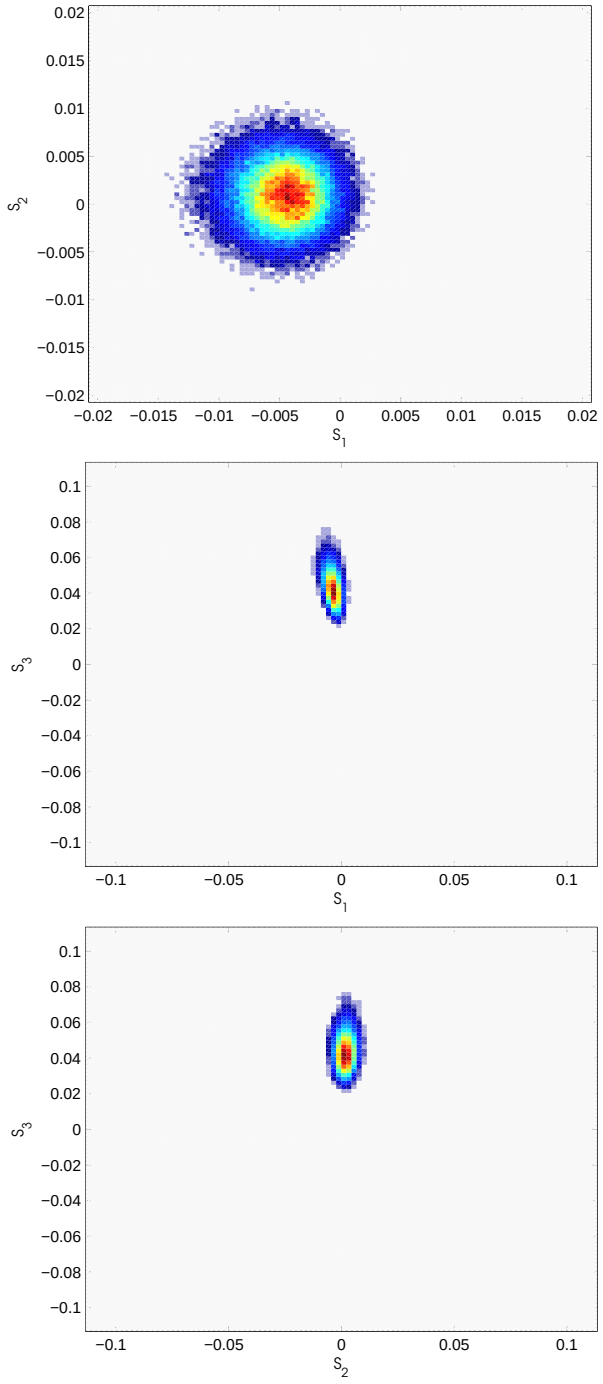

Figure A.10: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-2, low resolution mode, for right-circular transmit polarization.

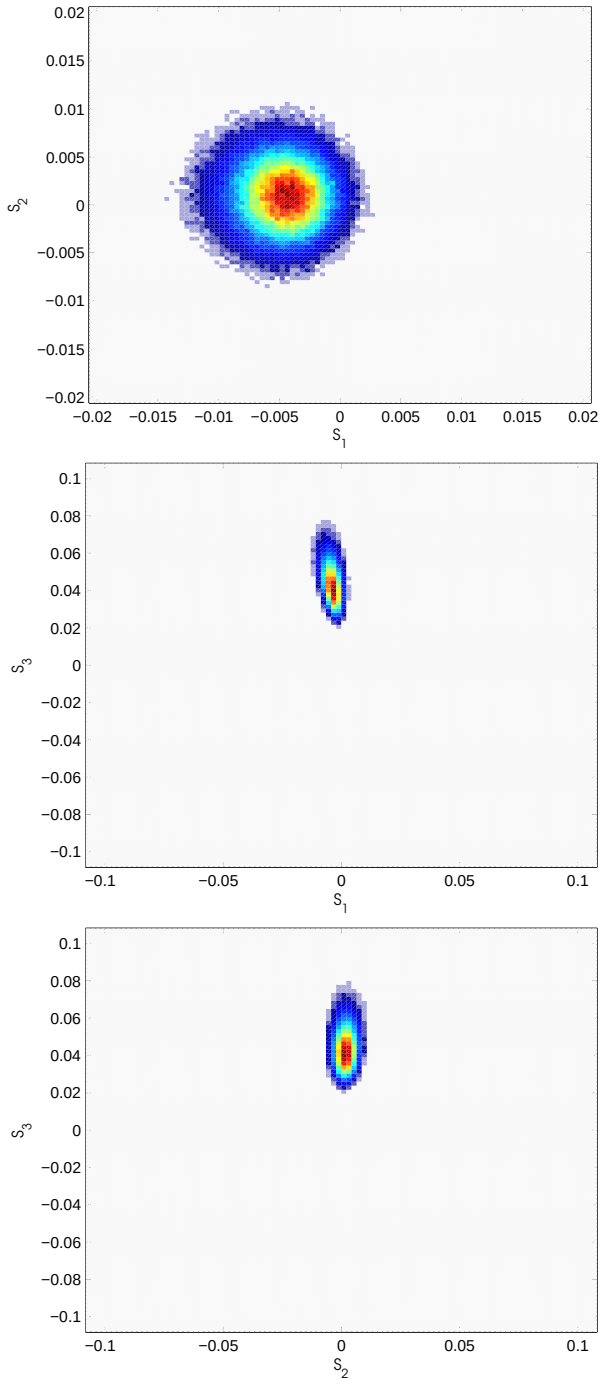

Figure A.11: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-1, low resolution mode, for right-circular transmit polarization.

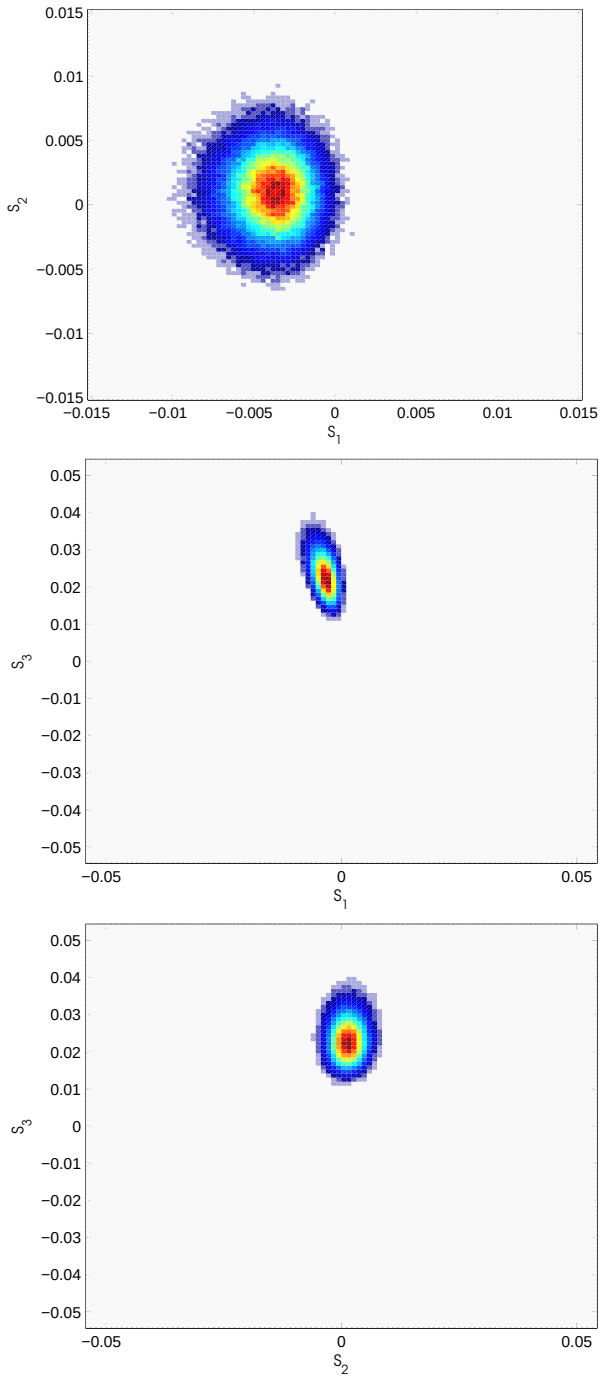

Figure A.12: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, low resolution mode, for right-circular transmit polarization.

A.2 Low Resolution Mode (Linear)

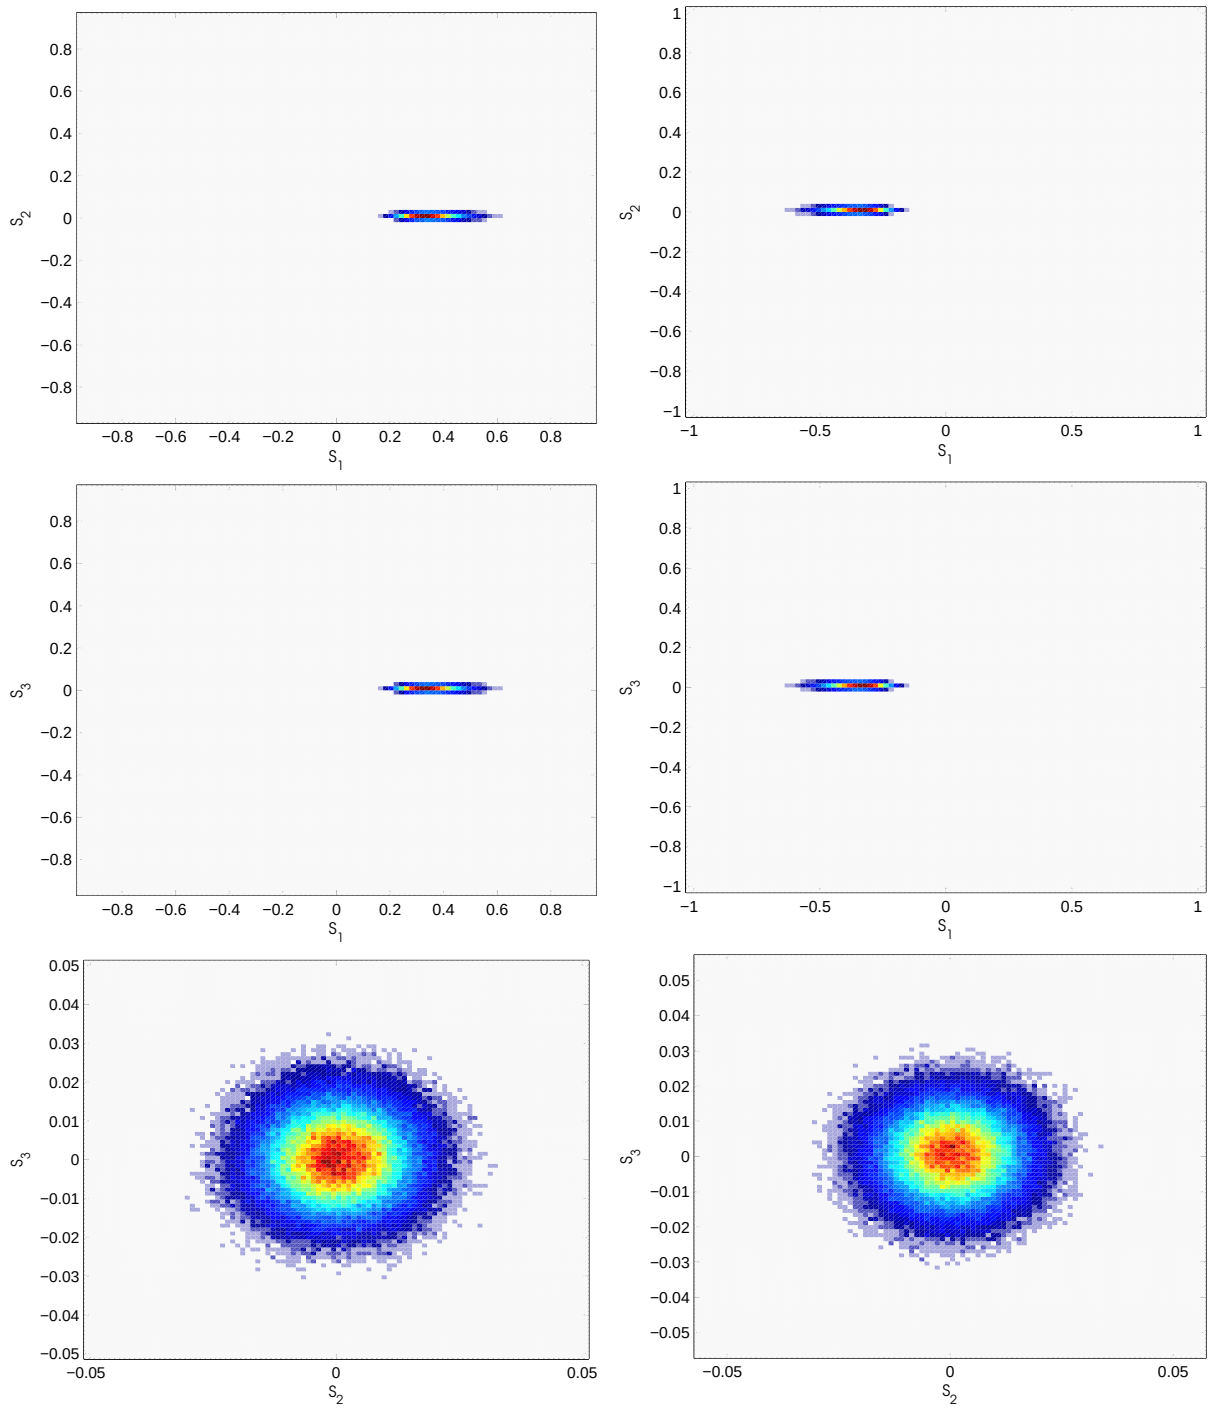

(a) HH-HV (b) VV-VH

Figure A.13: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

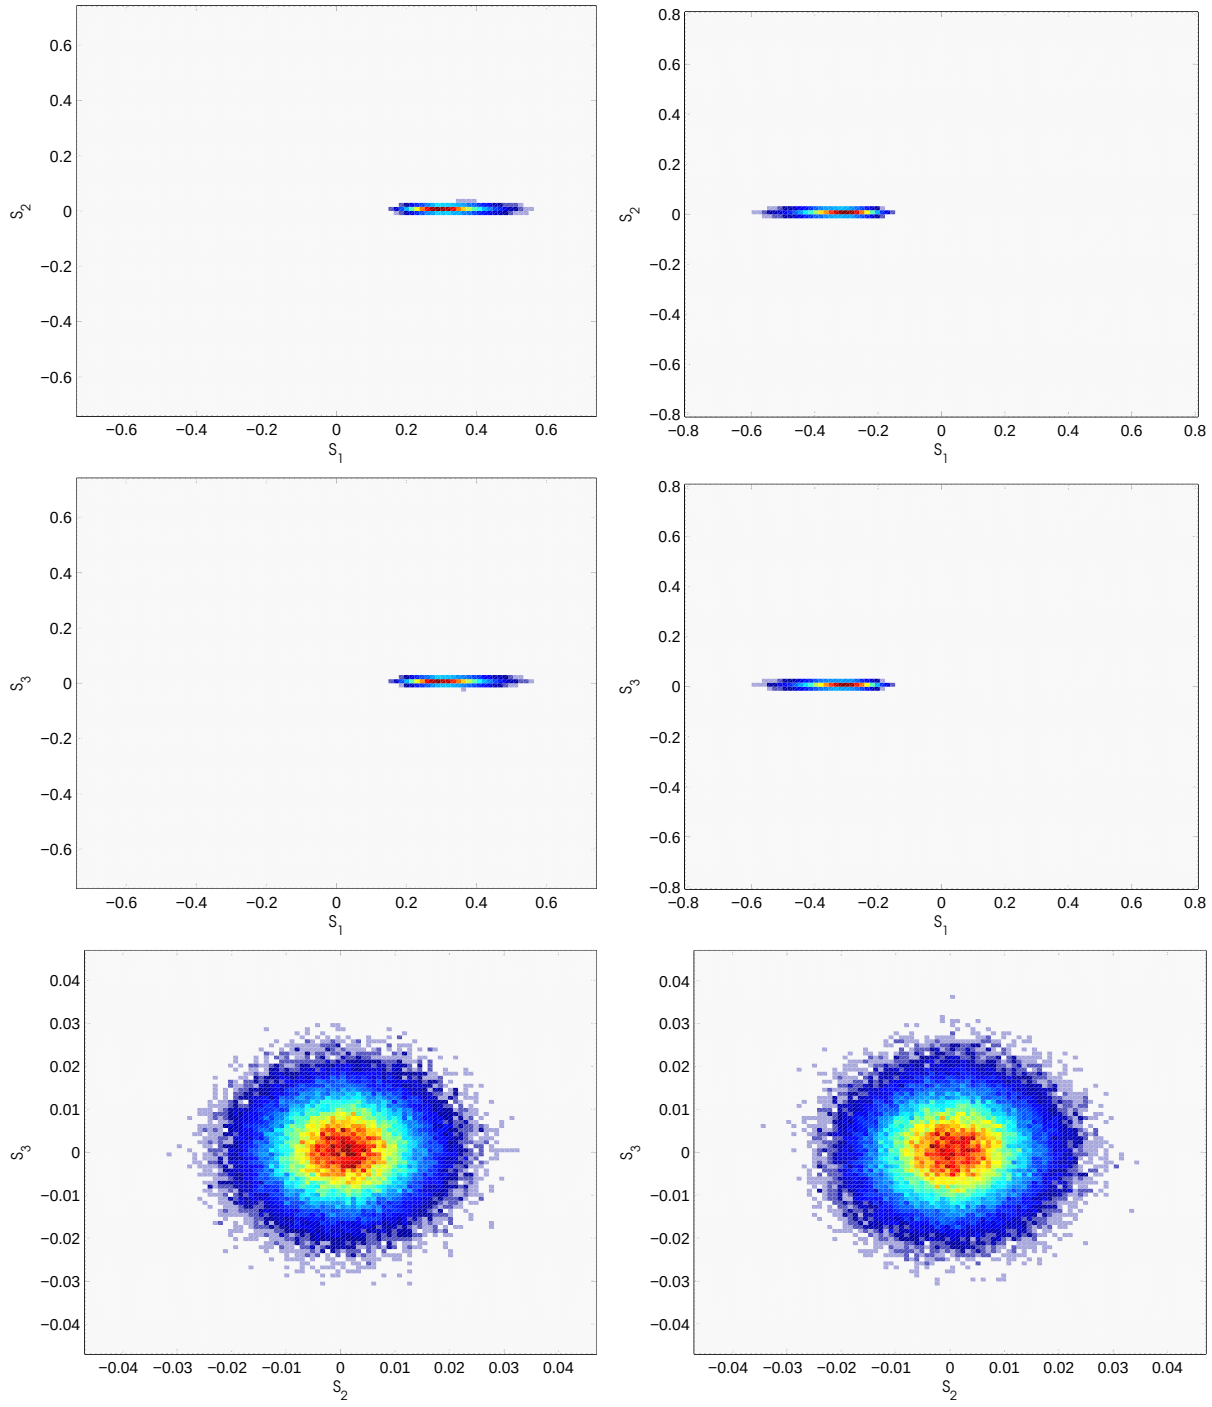

(a) HH-HV (b) VV-VH

Figure A.14: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-2, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

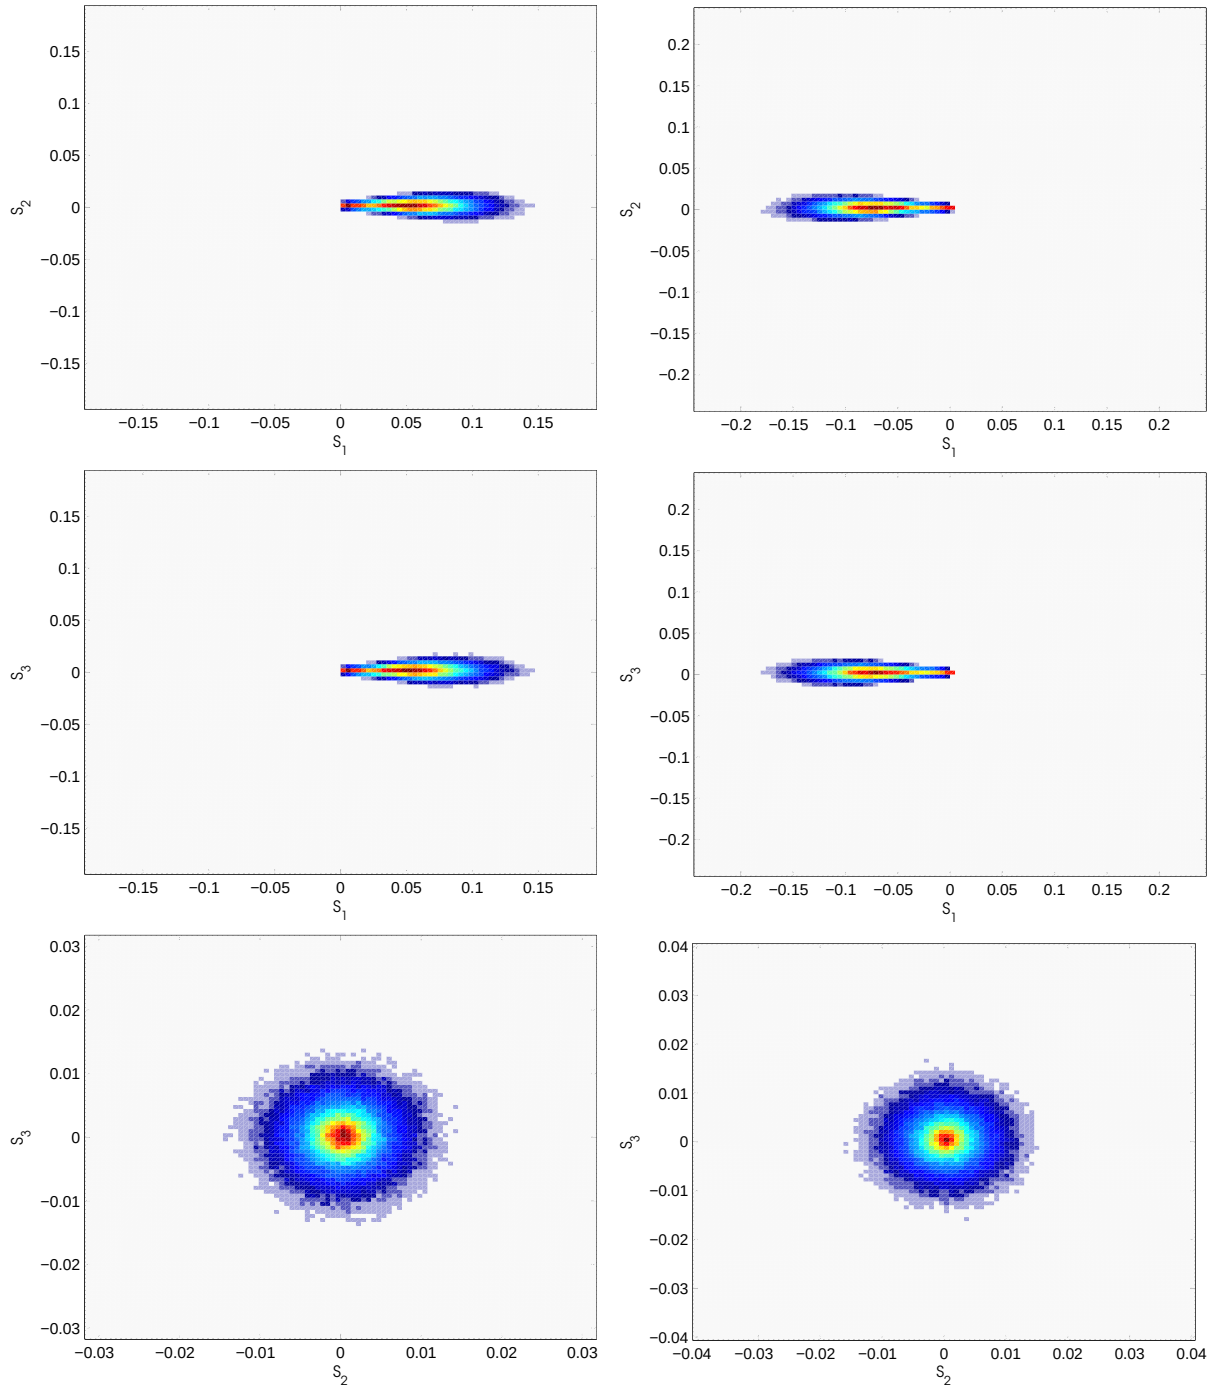

(a) HH-HV (b) VV-VH

Figure A.15: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0829-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

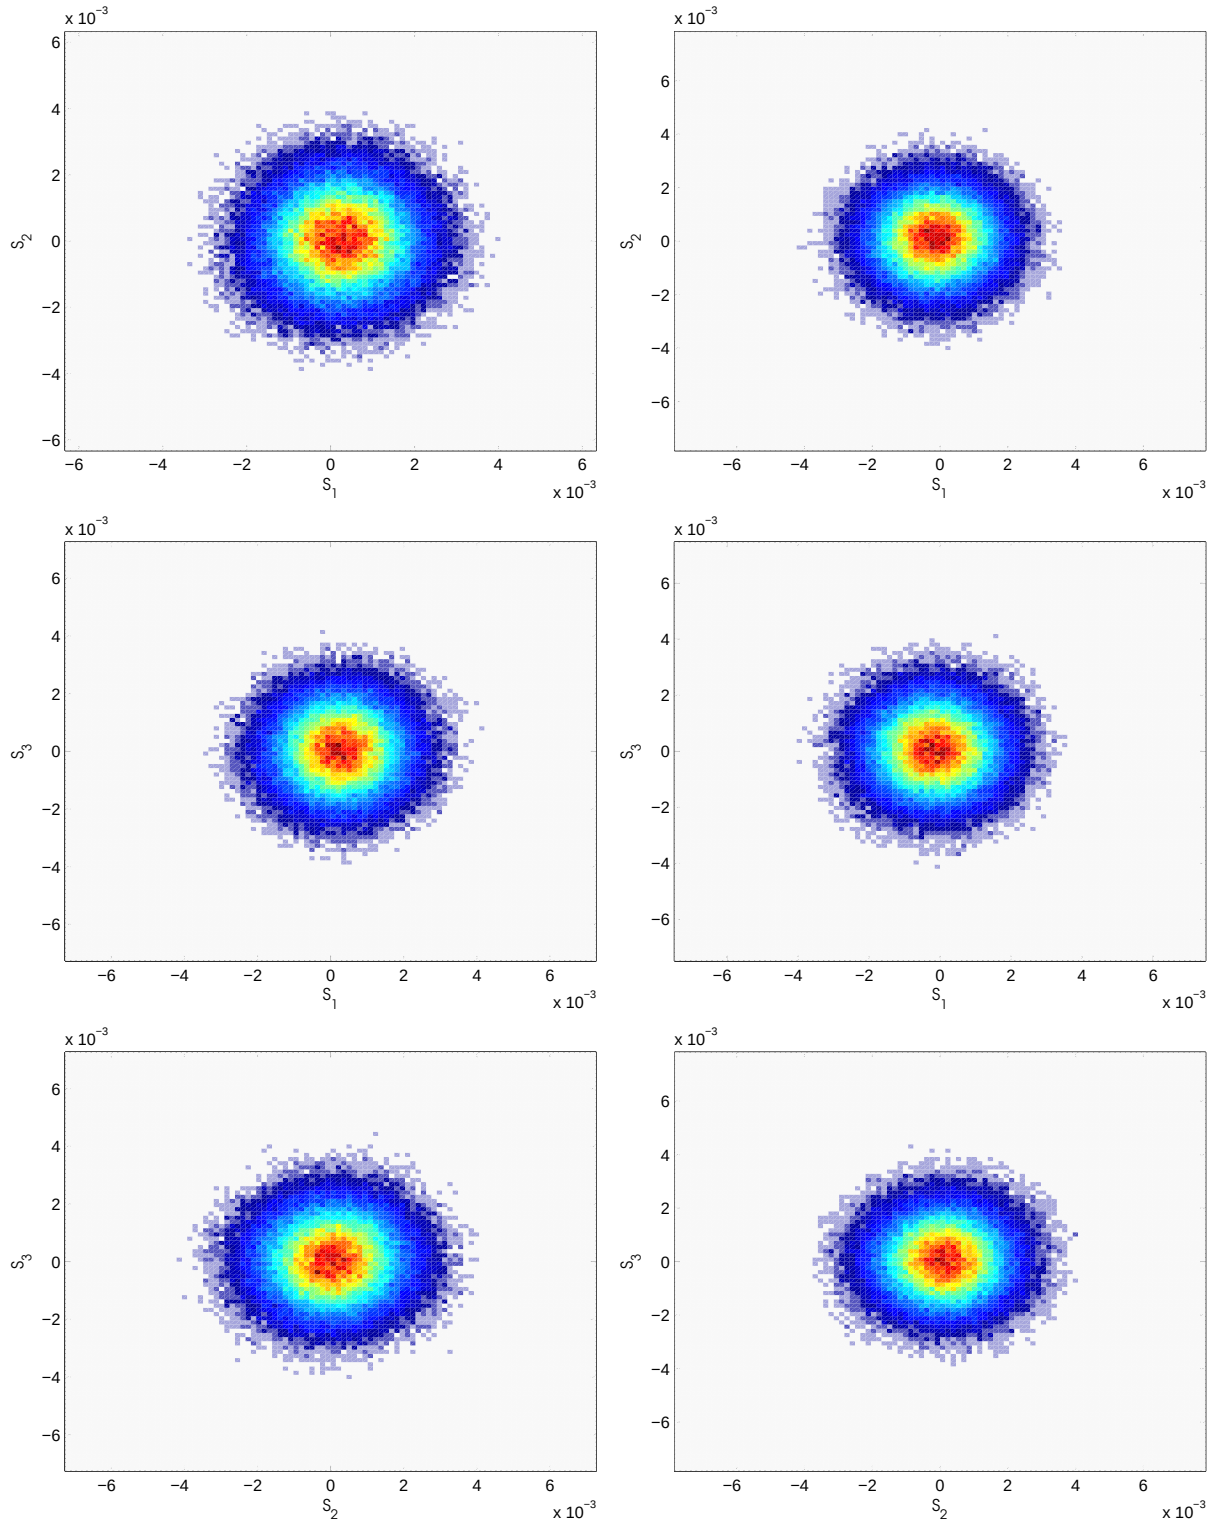

(a) HH-HV

(b) VV-VH

Figure A.16: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

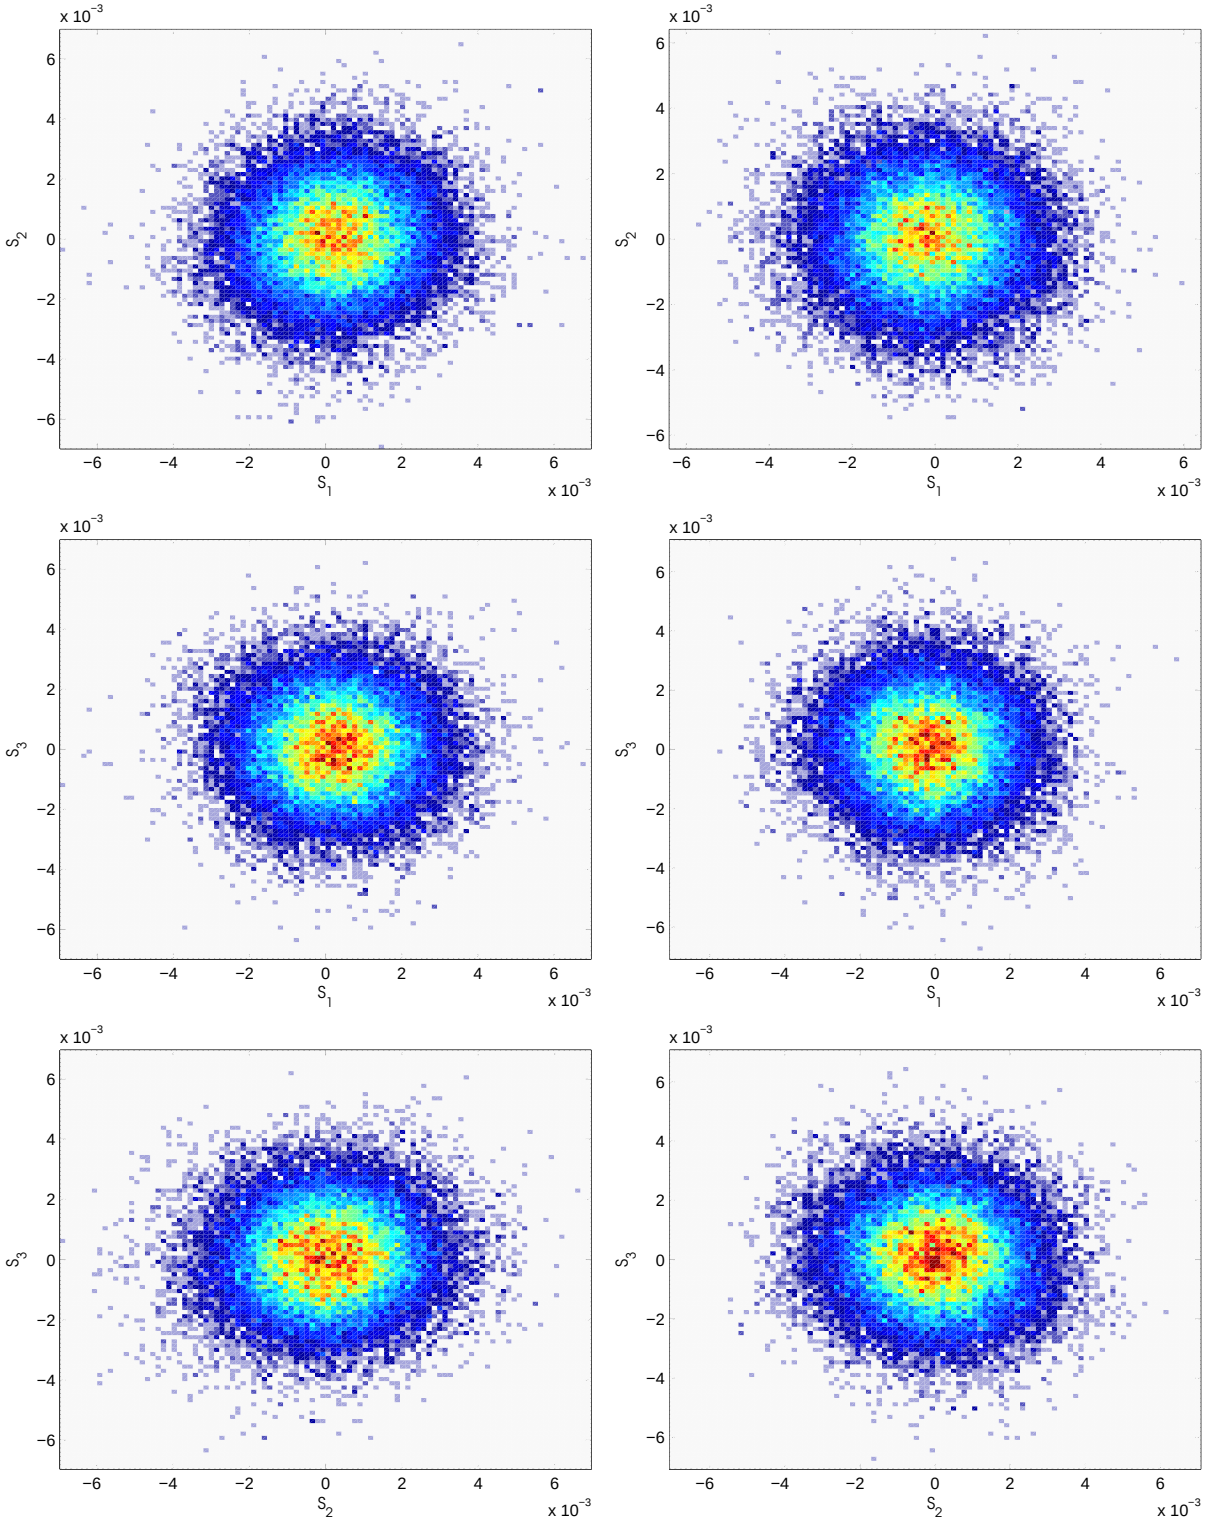

(a) HH-HV (b) VV-VH

Figure A.17: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-2, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

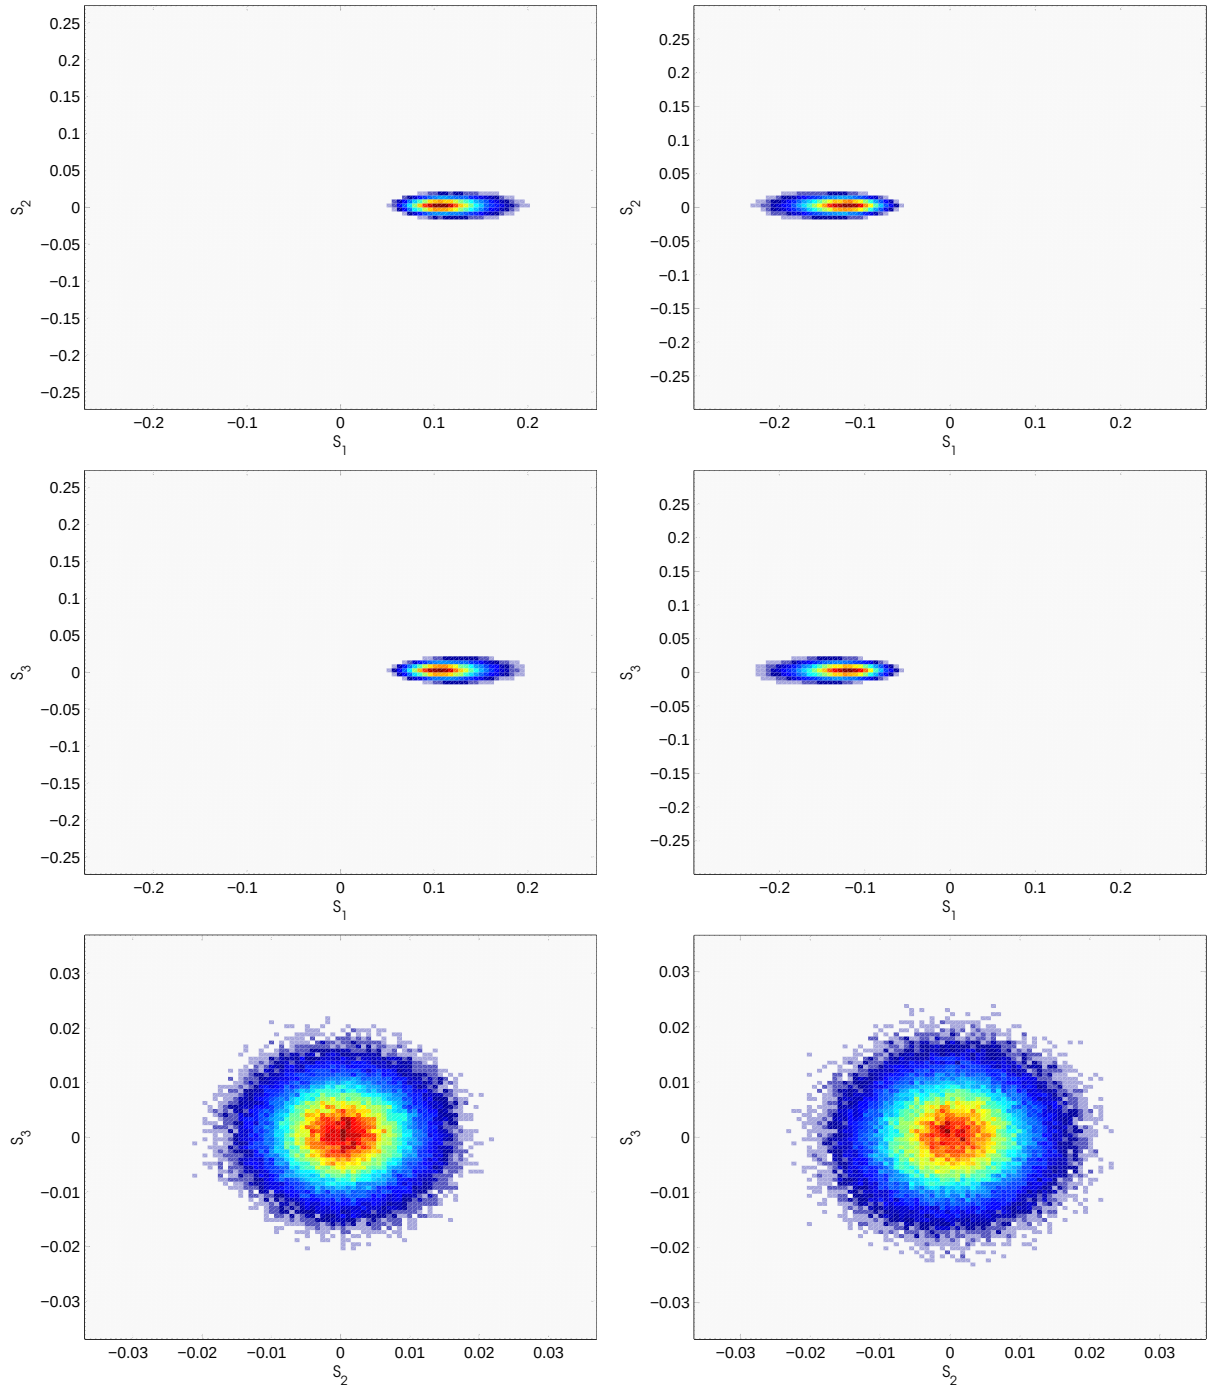

(a) HH-HV (b) VV-VH

Figure A.18: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-2, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

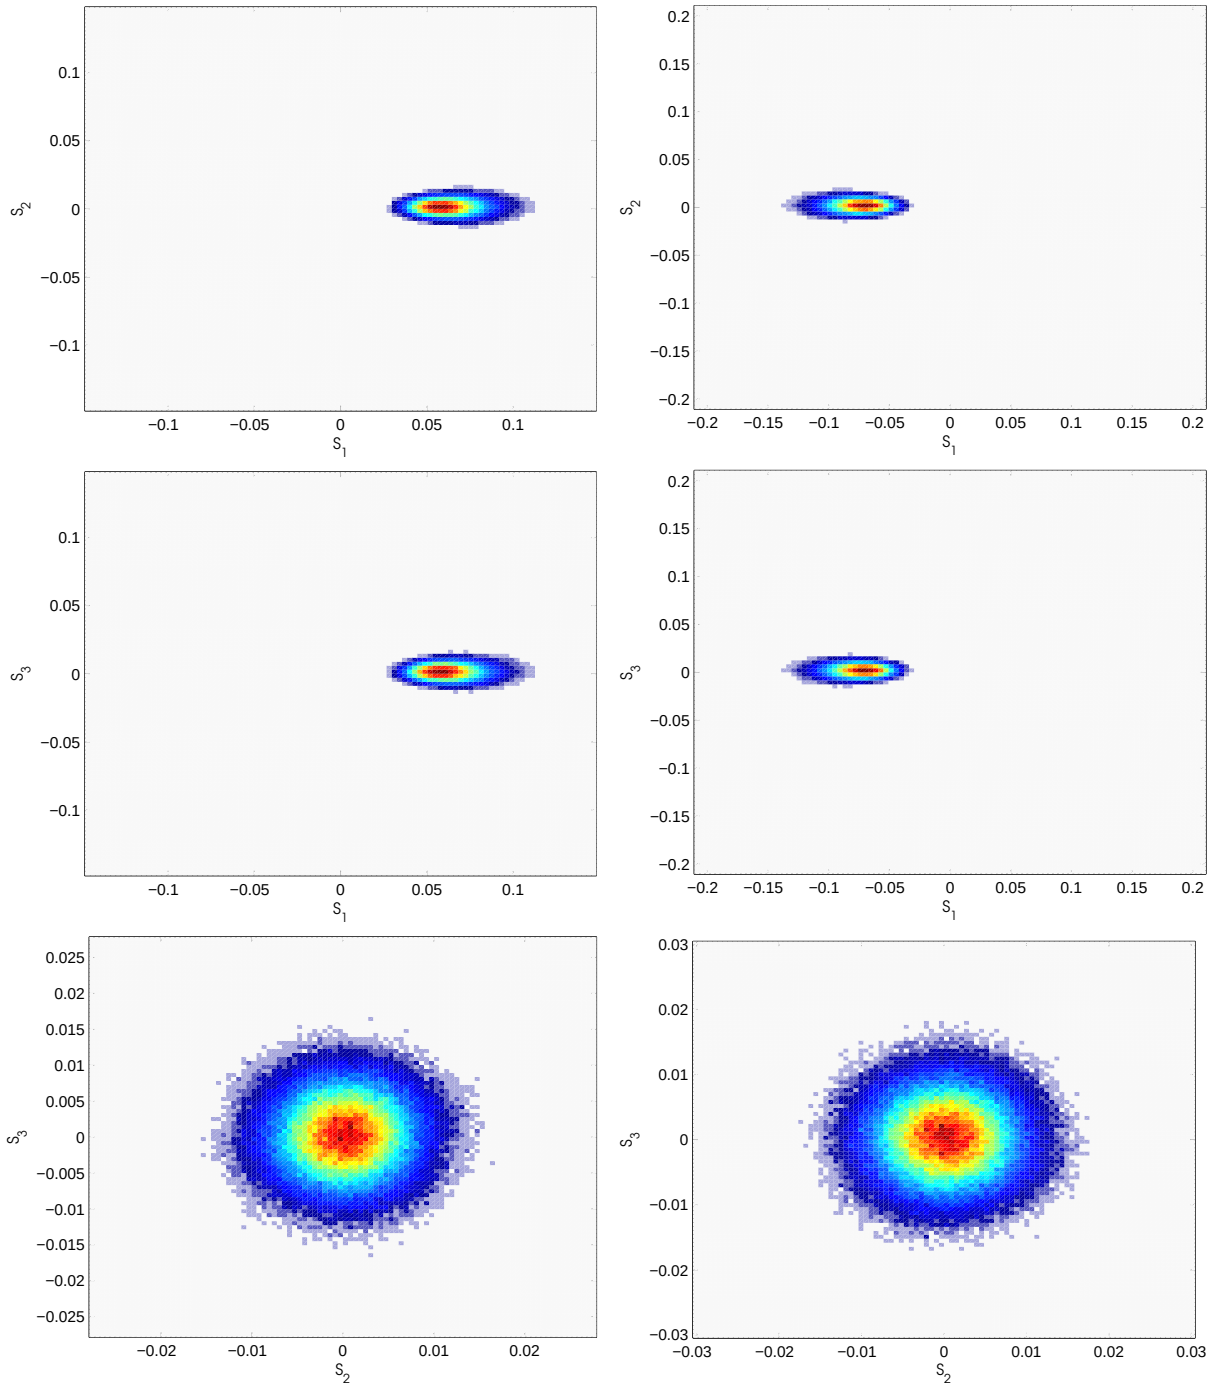

(a) HH-HV (b) VV-VH

Figure A.19: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

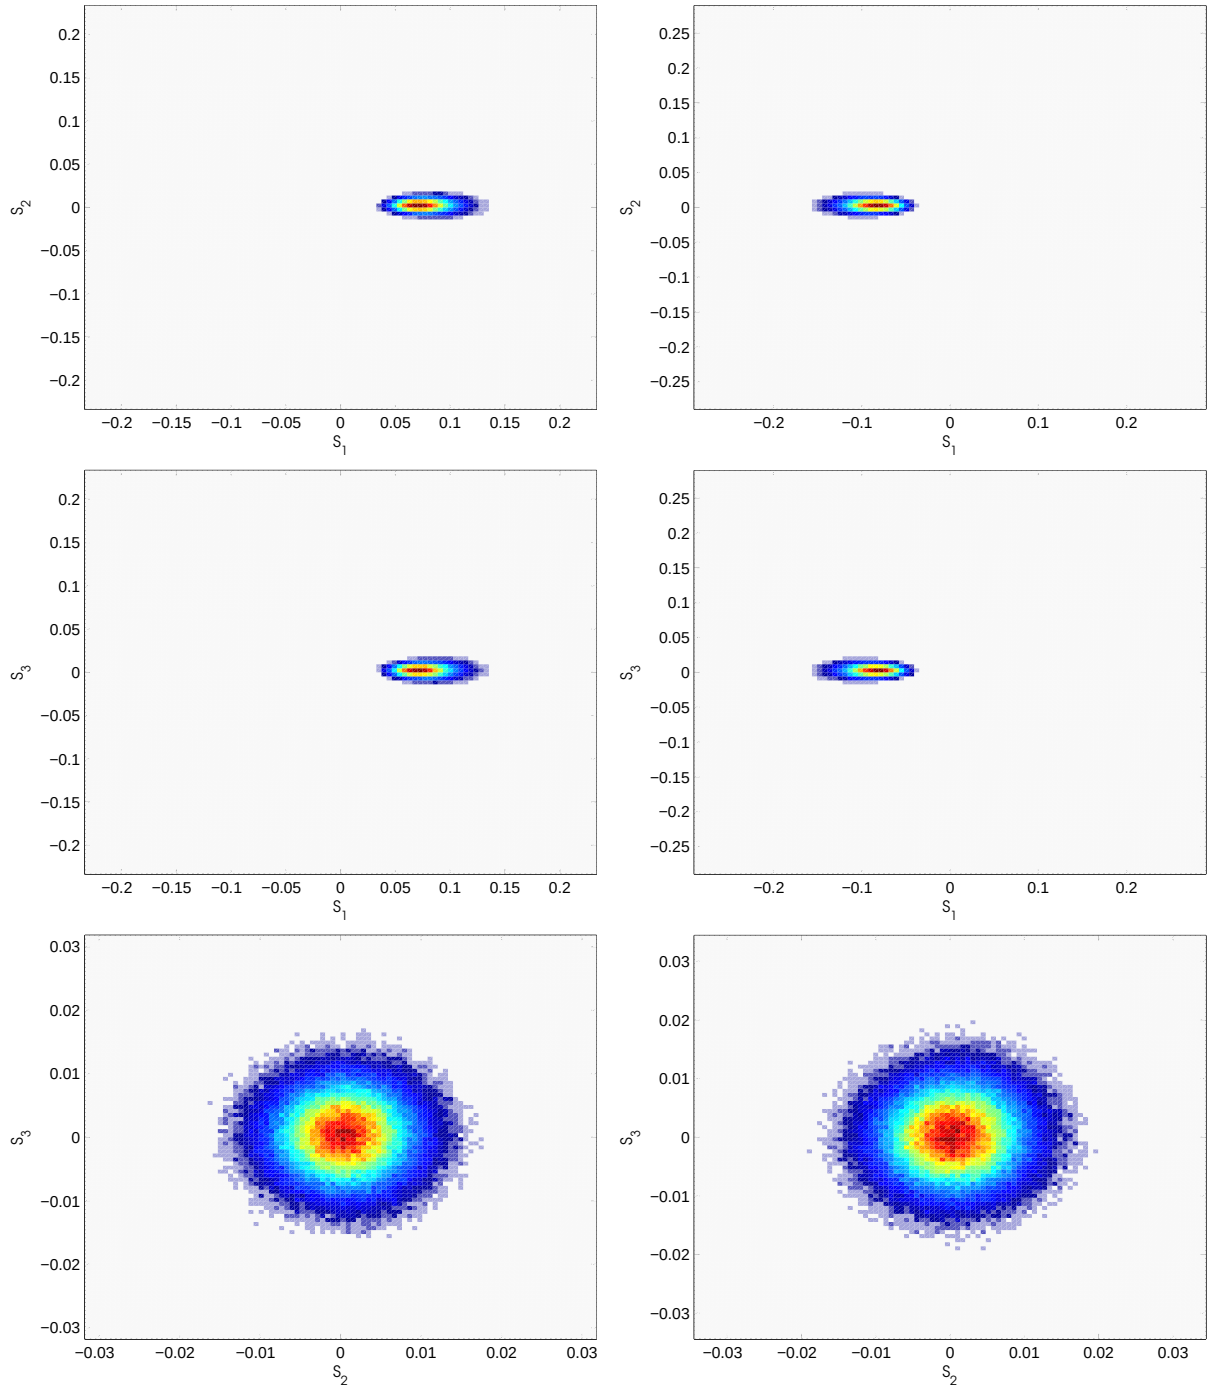

(a) HH-HV (b) VV-VH

Figure A.20: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-2, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

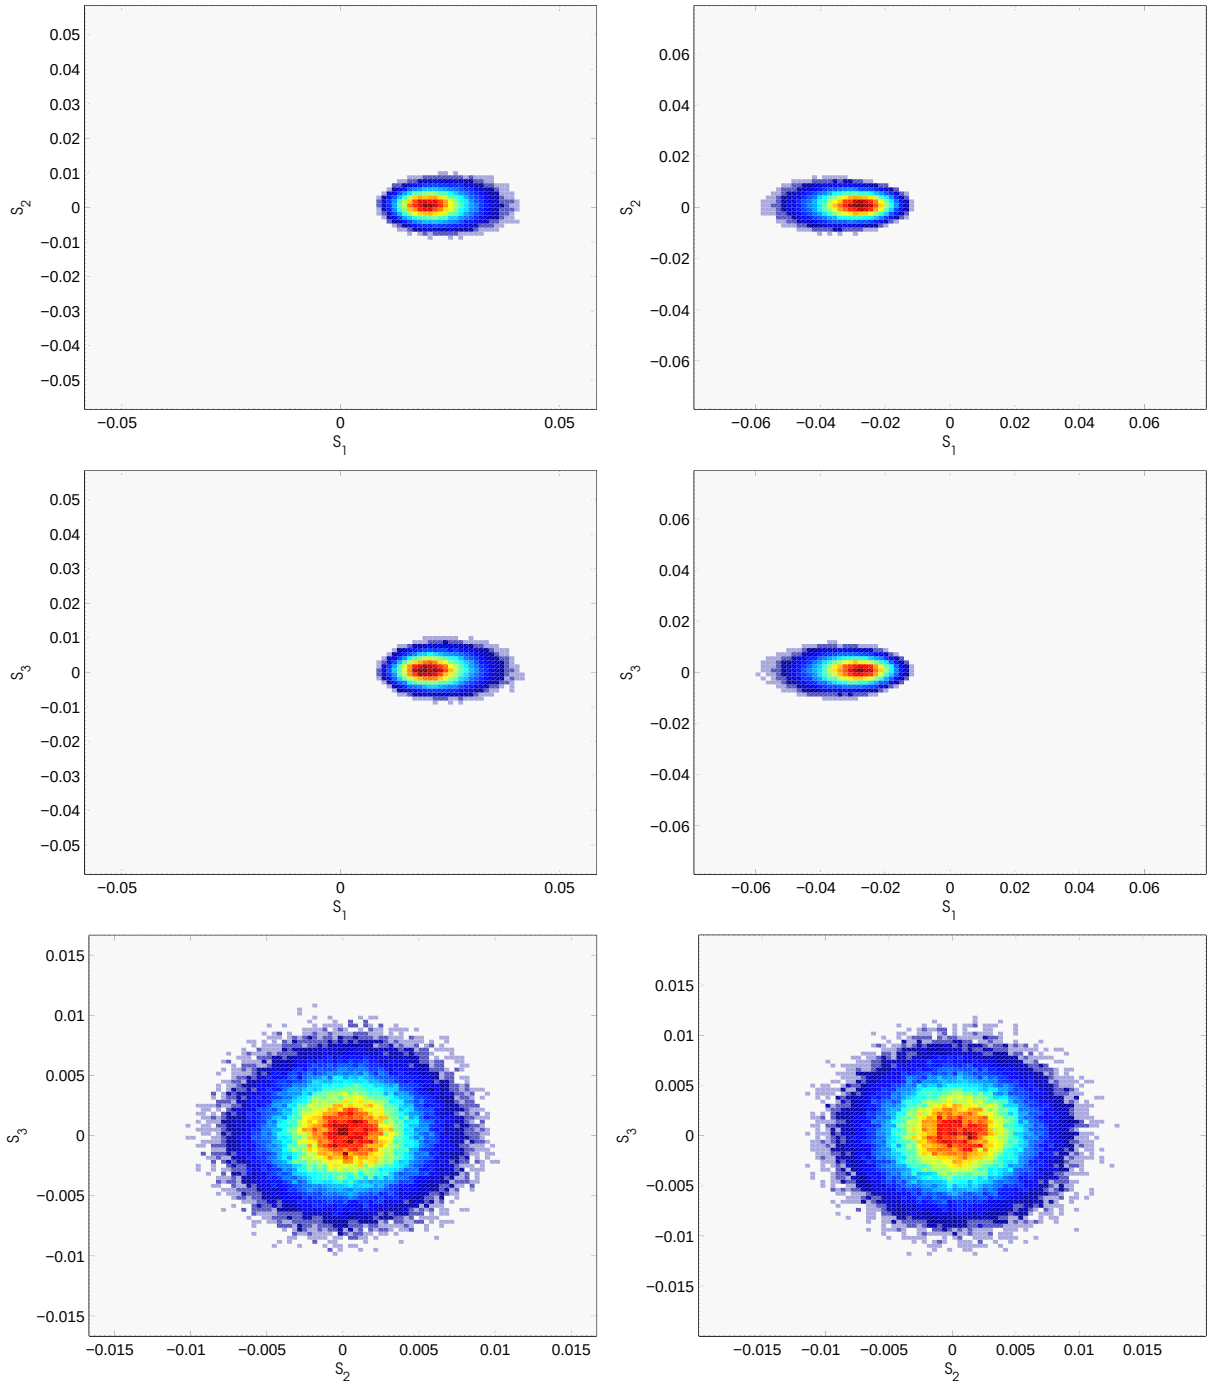

(a) HH-HV

(b) VV-VH

Figure A.21: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

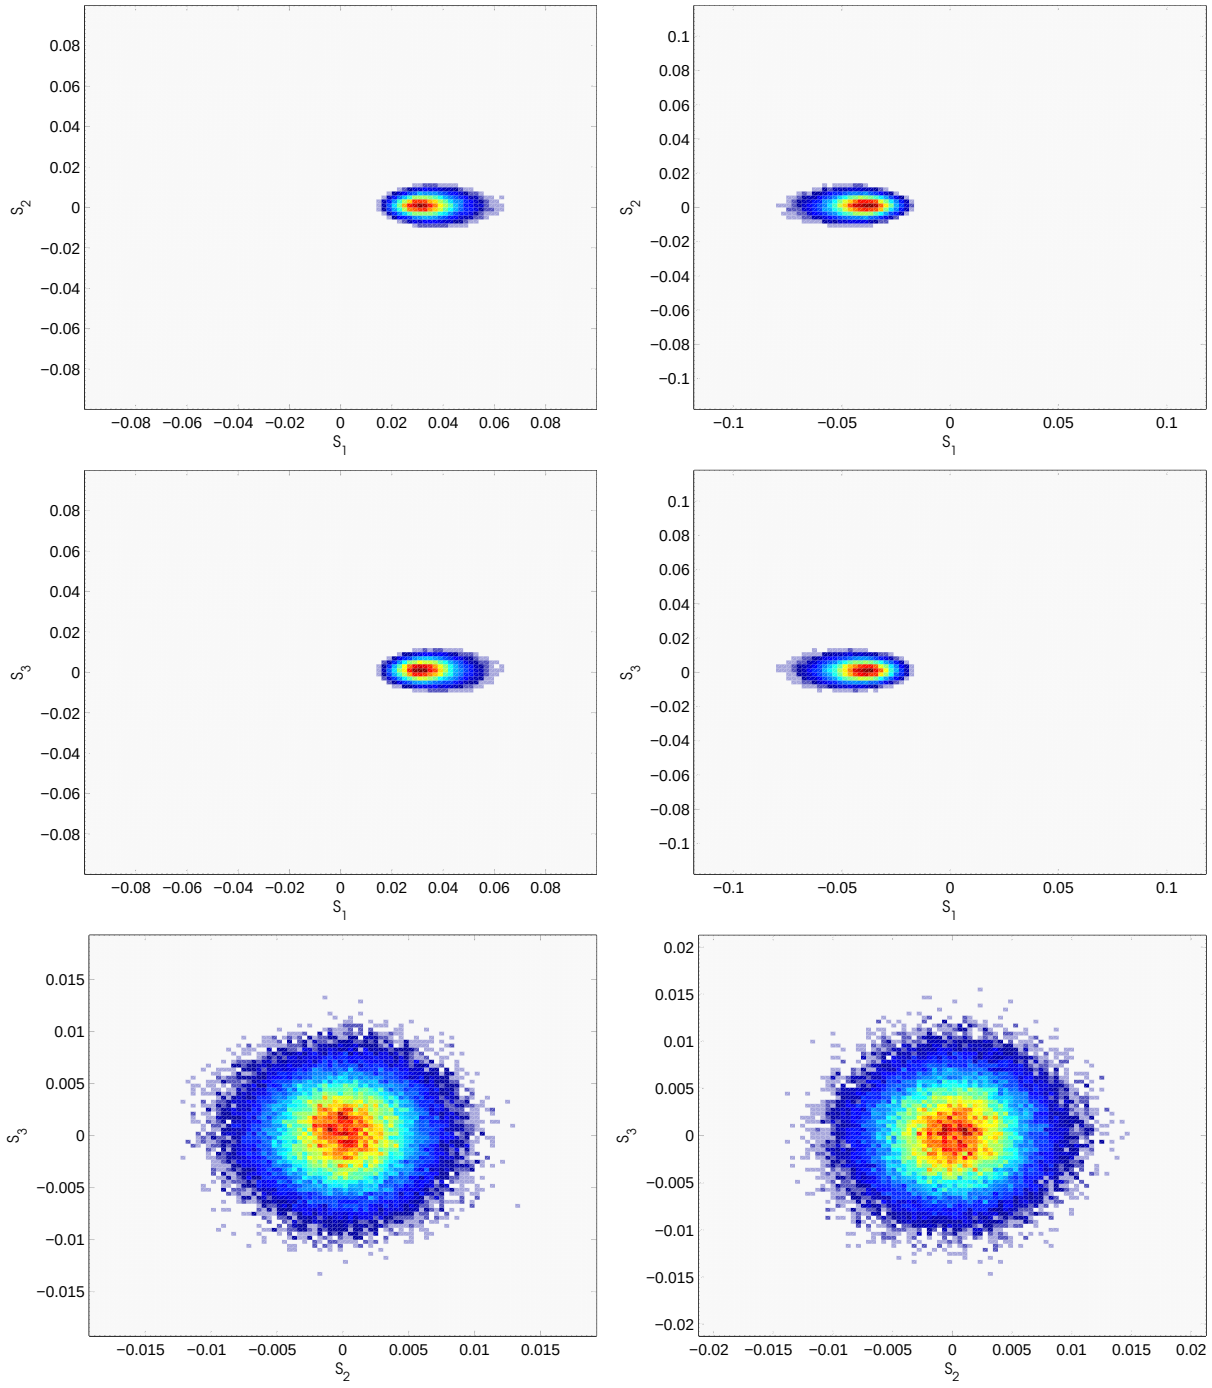

(a) HH-HV (b) VV-VH

Figure A.22: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-2, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

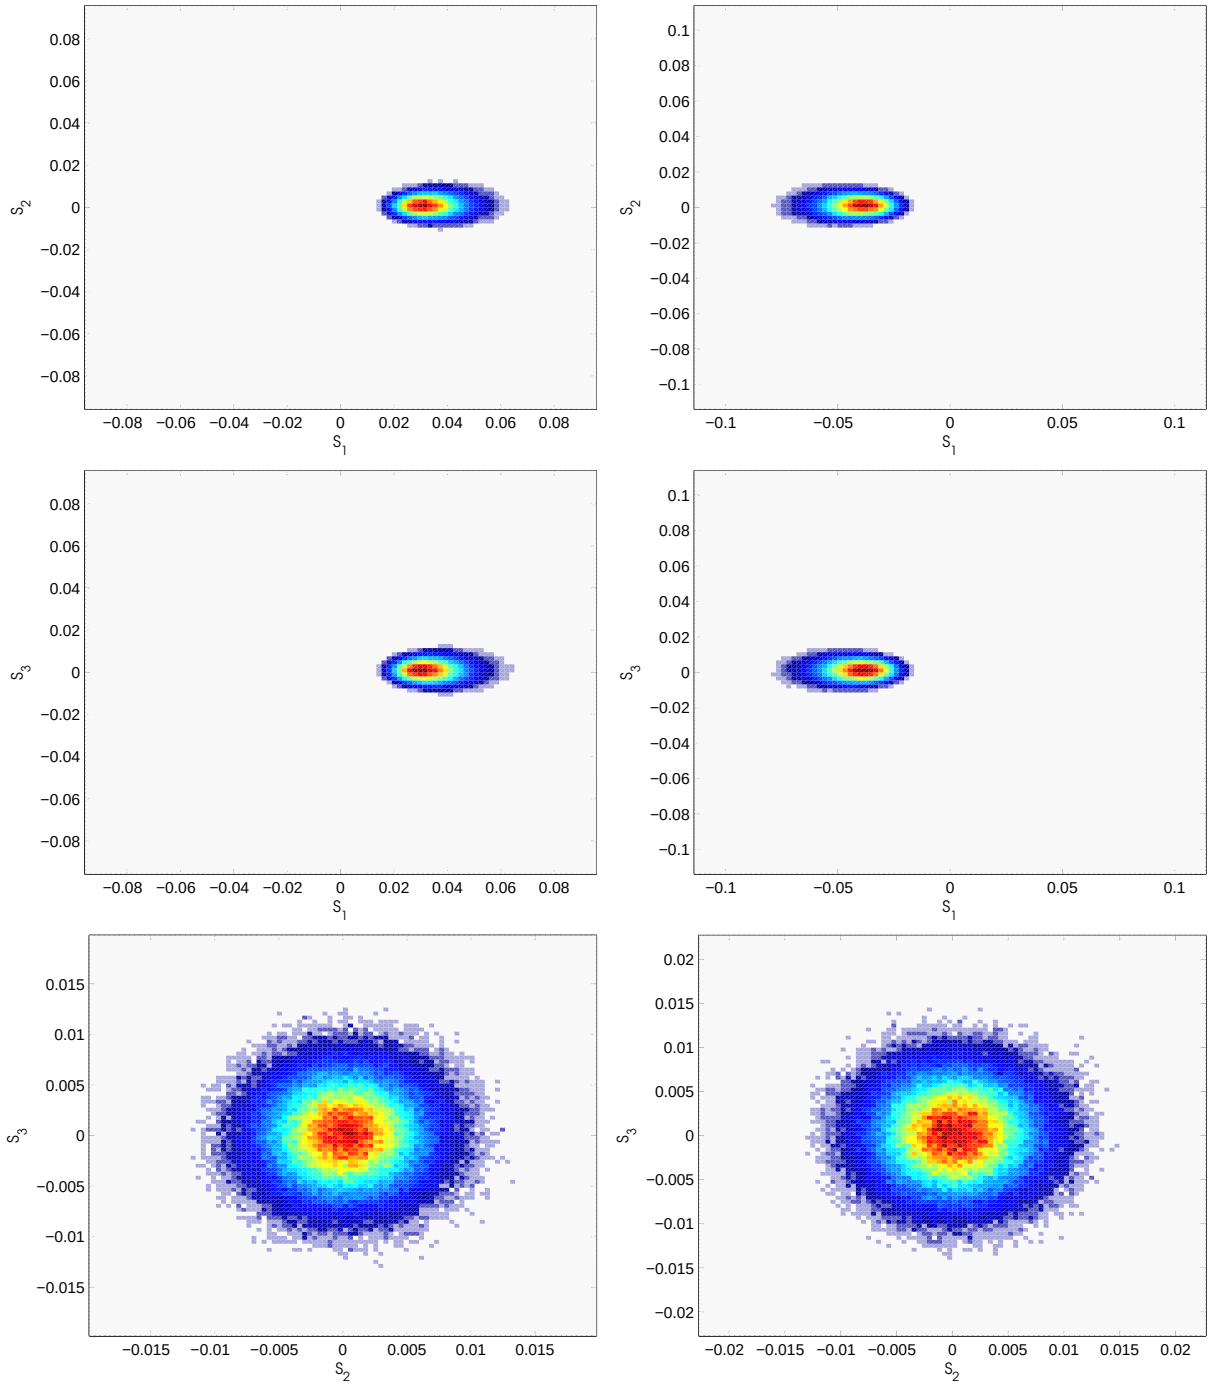

(a) HH-HV

(b) VV-VH

Figure A.23: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

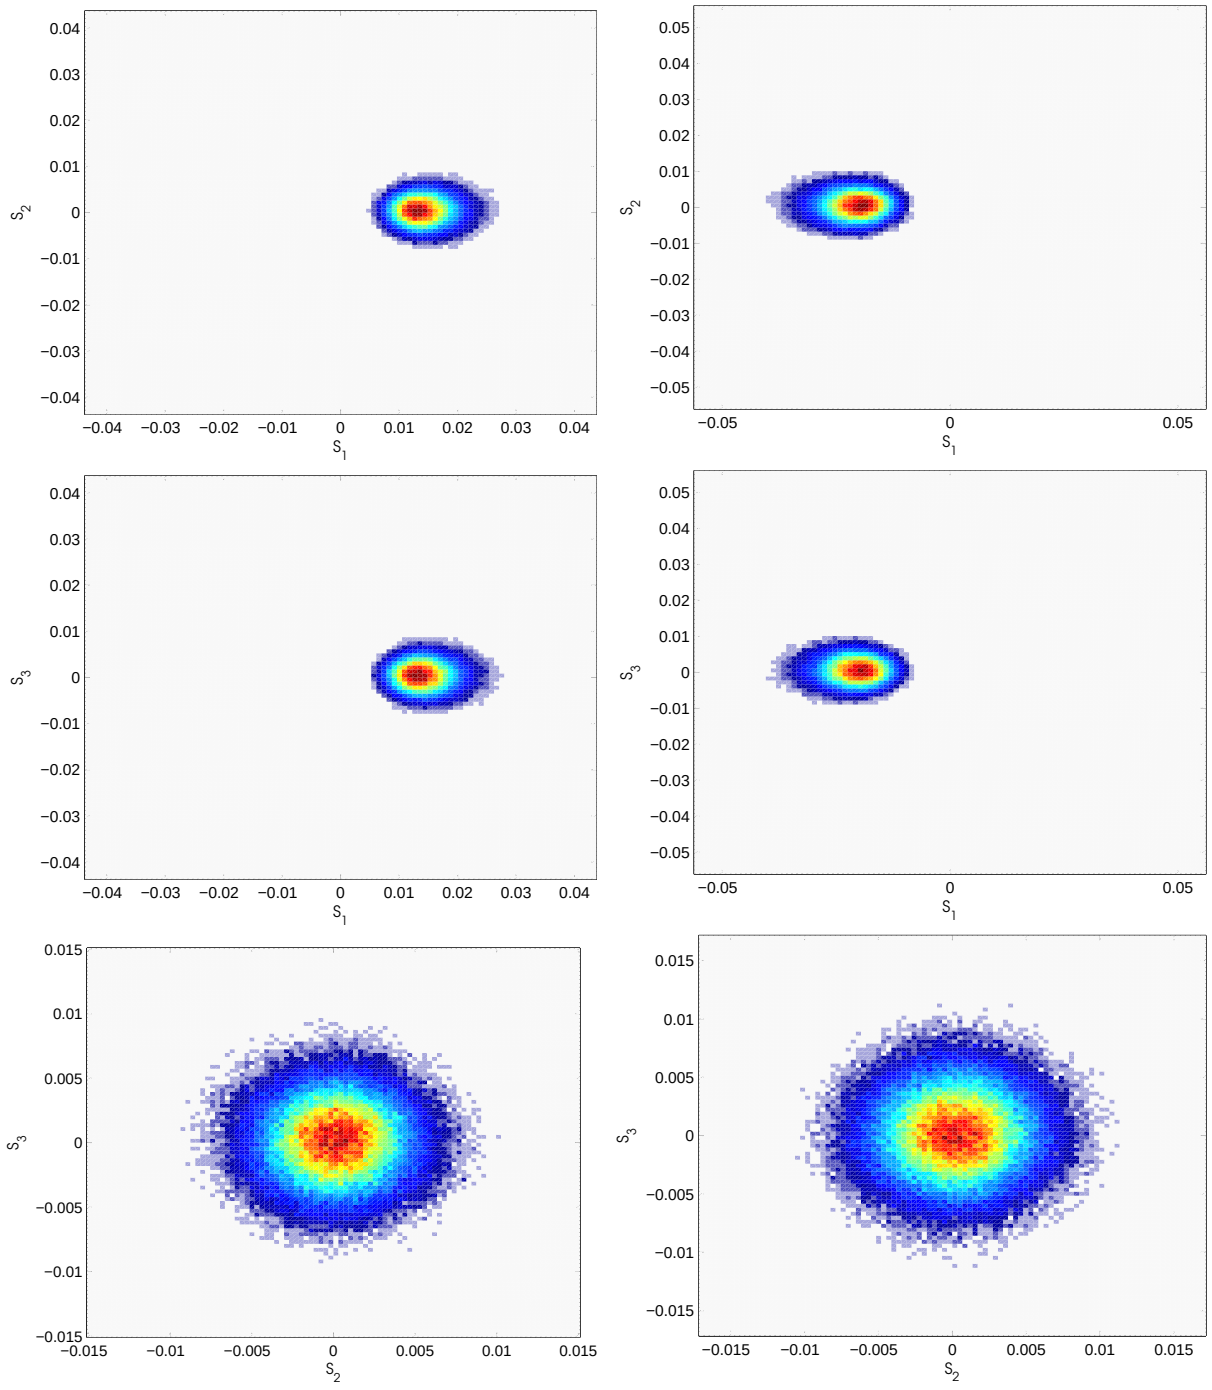

(a) HH-HV

(b) VV-VH

Figure A.24: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, low resolution mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

A.3 Medium Resolution (CTLR)

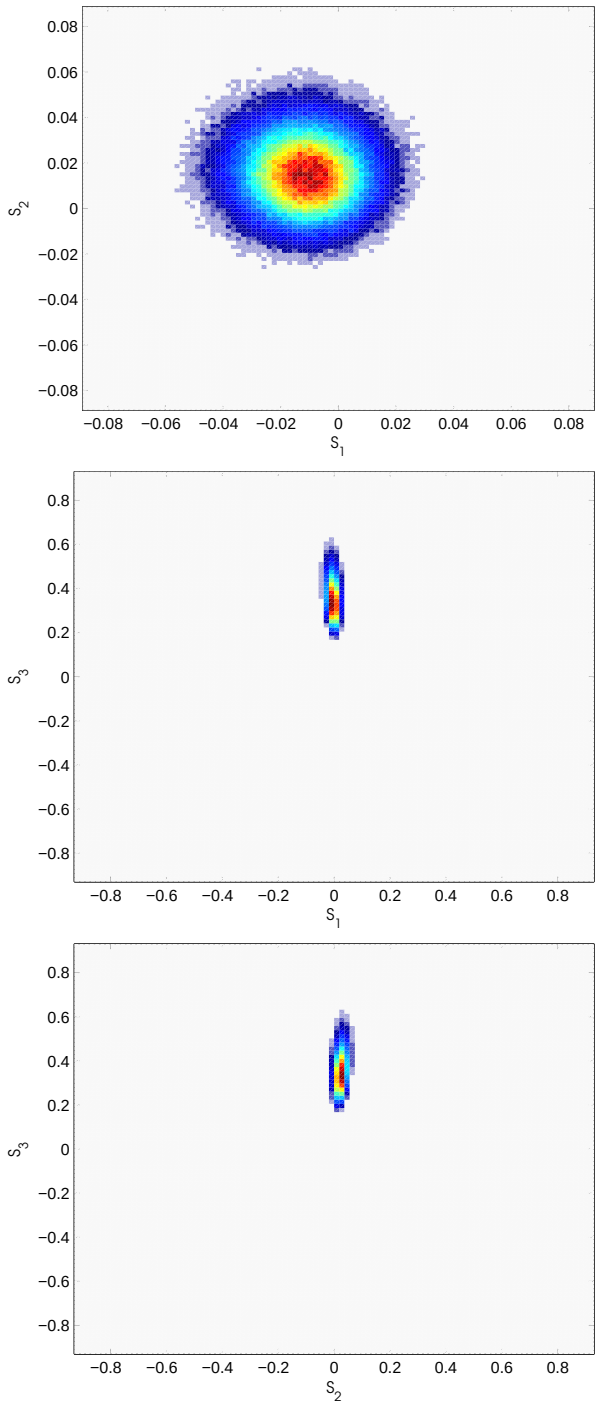

Figure A.25: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-1, medium resolution mode, for right-circular transmit polarization.

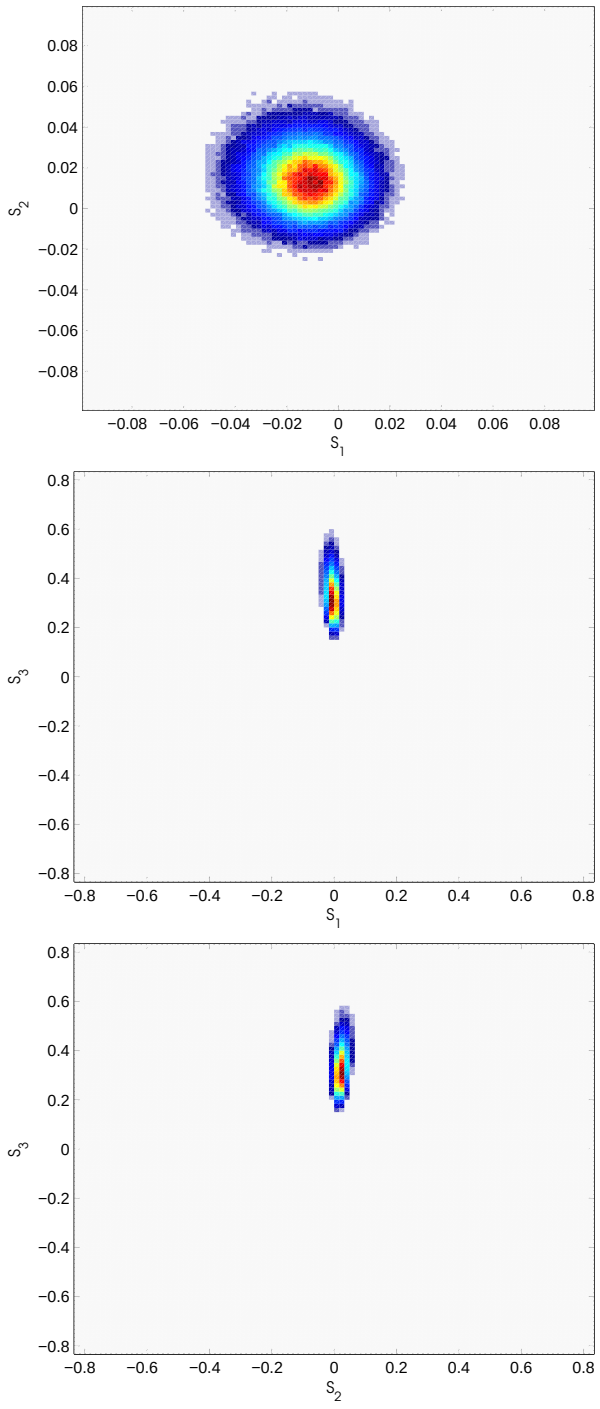

Figure A.26: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0905-2, medium resolution mode, for right-circular transmit polarization.

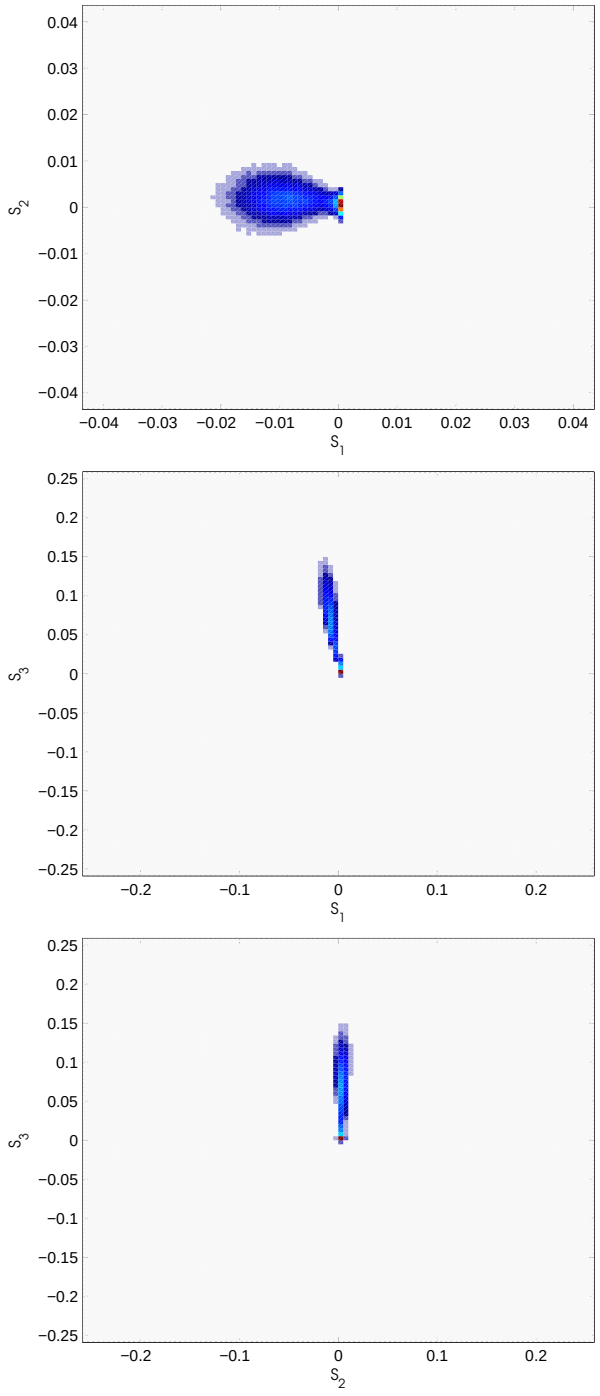

Figure A.27: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0829-1, medium resolution mode, for right-circular transmit polarization.

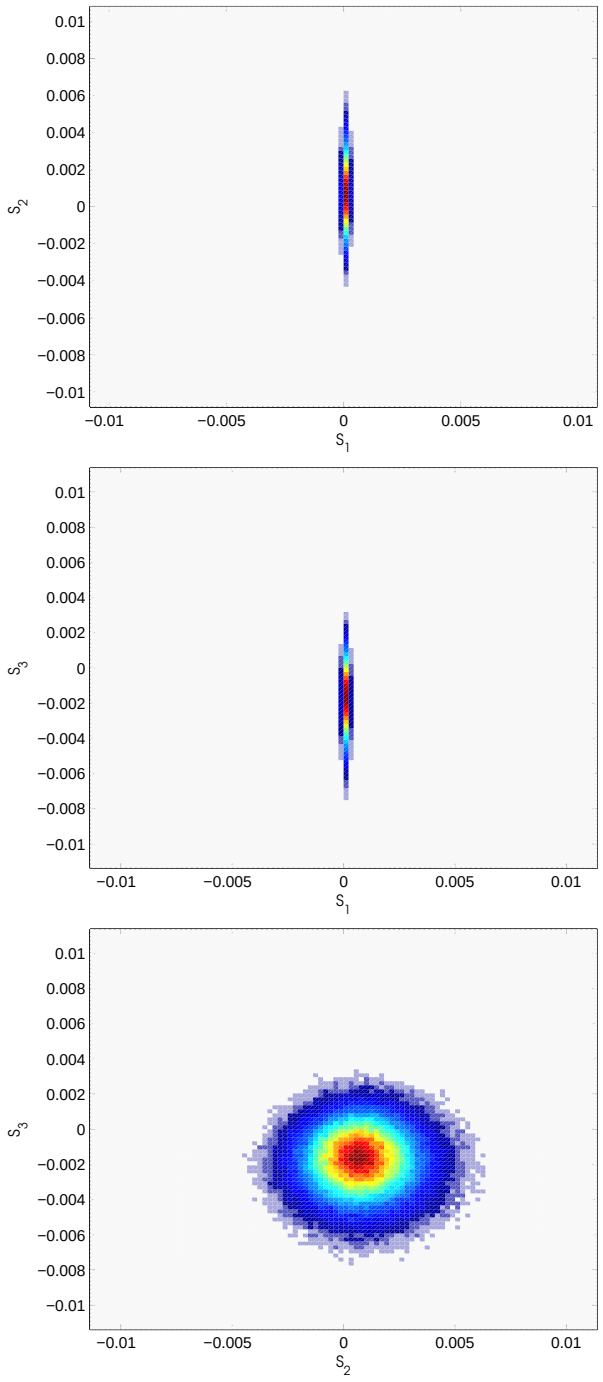

Figure A.28: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-1, medium resolution mode, for right-circular transmit polarization.

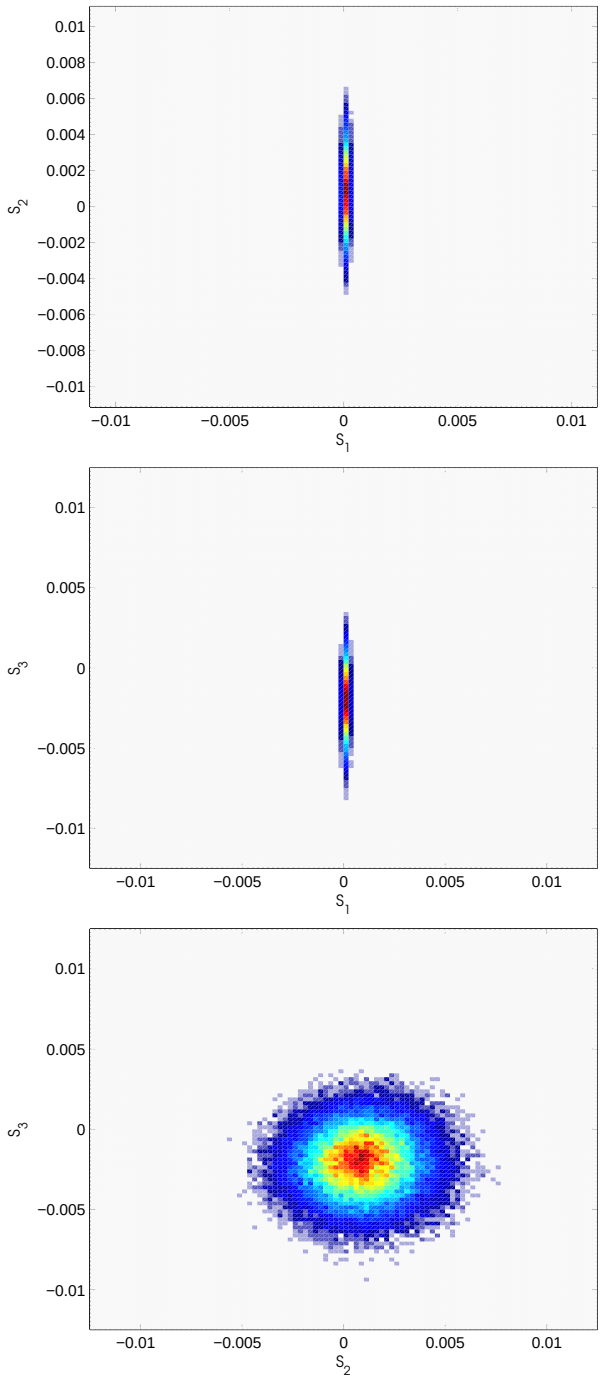

Figure A.29: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0805-2, medium resolution mode, for right-circular transmit polarization.

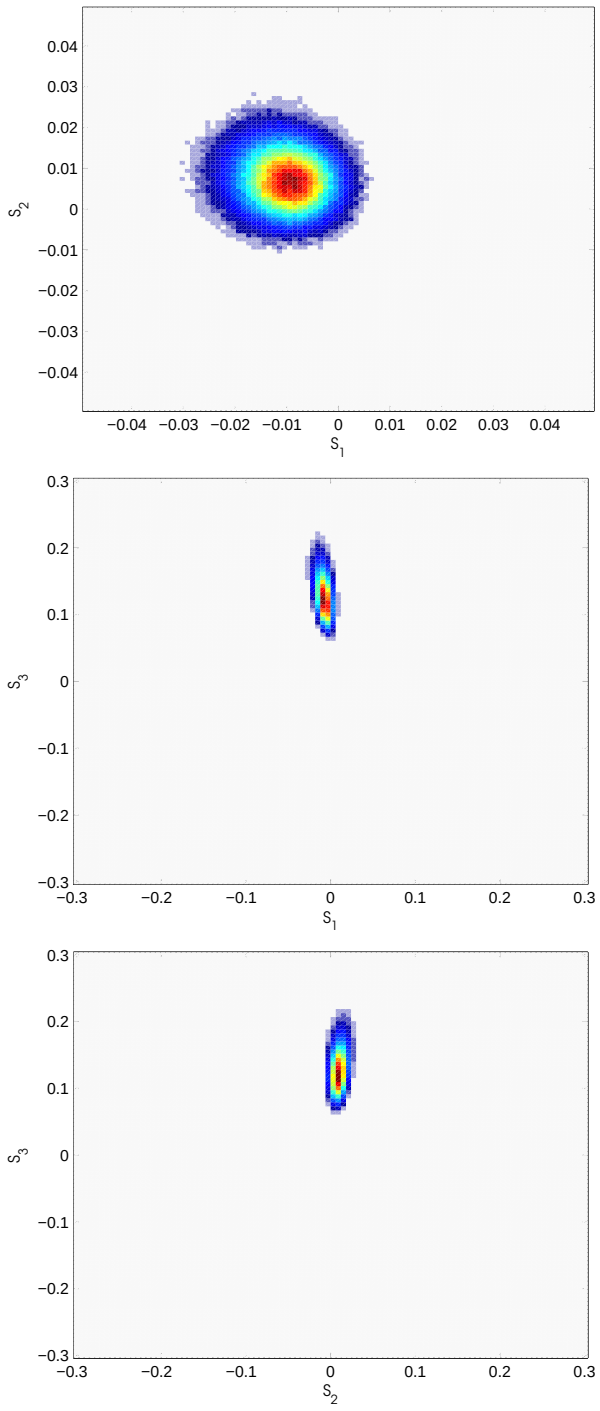

Figure A.30: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-2, medium resolution mode, for right-circular transmit polarization.

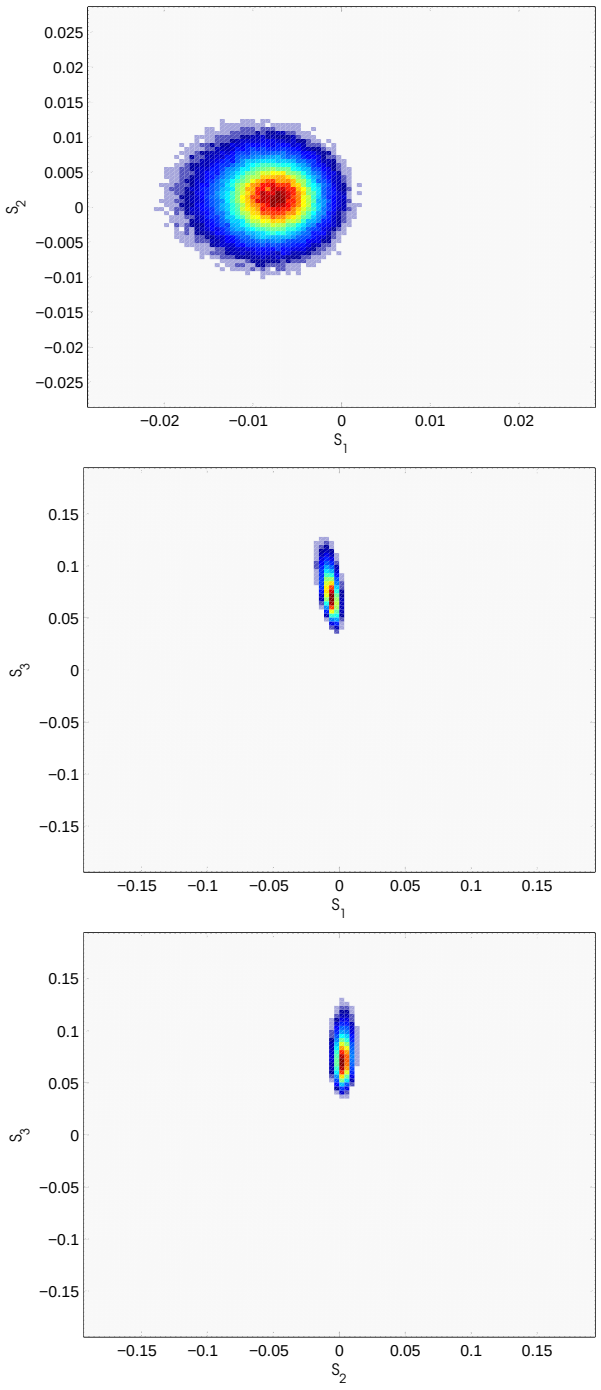

Figure A.31: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-1, medium resolution mode, for right-circular transmit polarization.

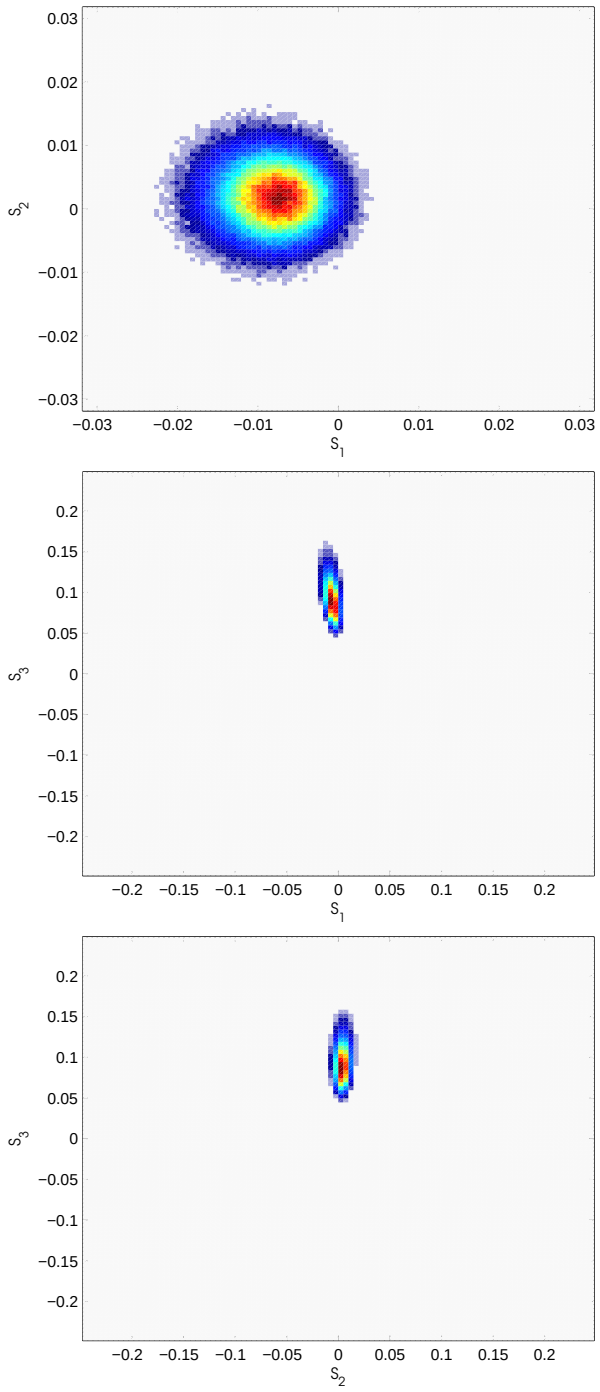

Figure A.32: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0822-2, medium resolution mode, for right-circular transmit polarization.

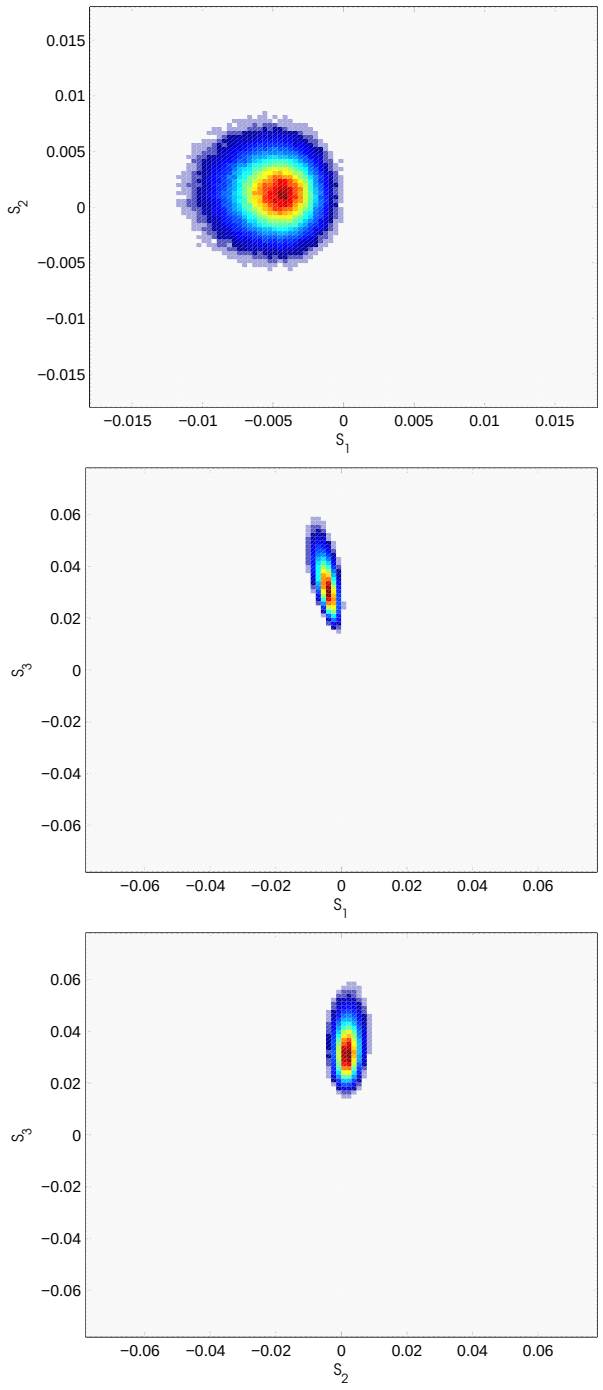

Figure A.33: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0815-1, medium resolution mode, for right-circular transmit polarization.

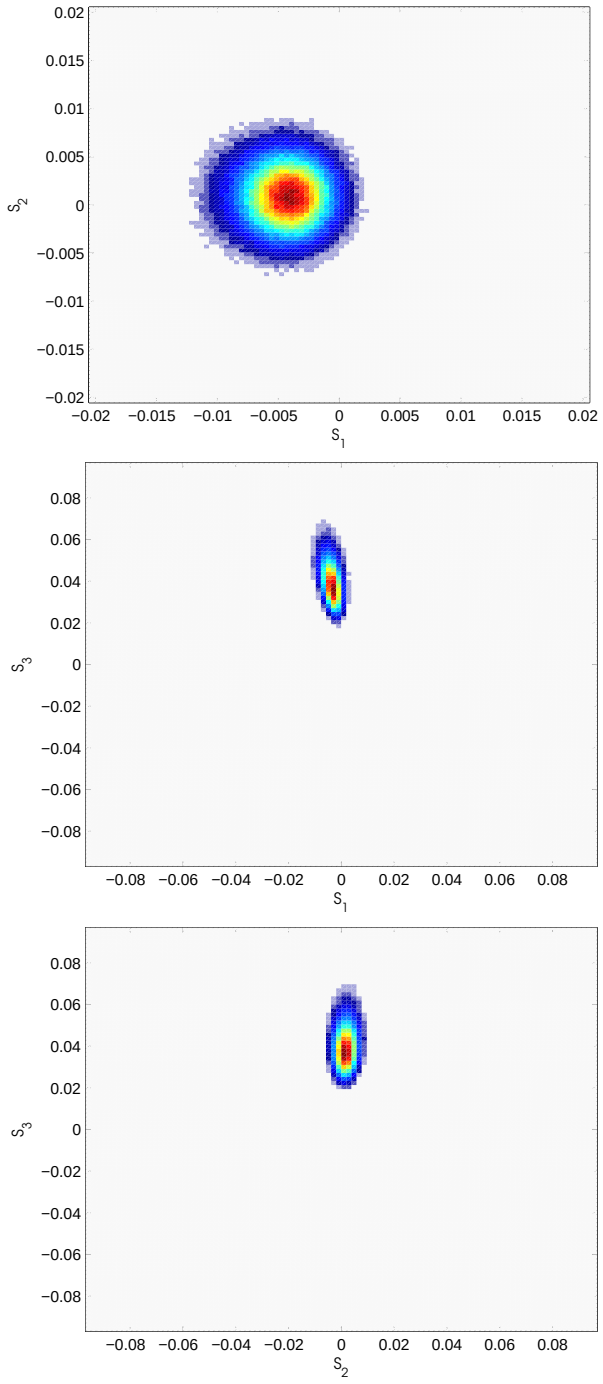

Figure A.34: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-2, medium resolution mode, for right-circular transmit polarization.

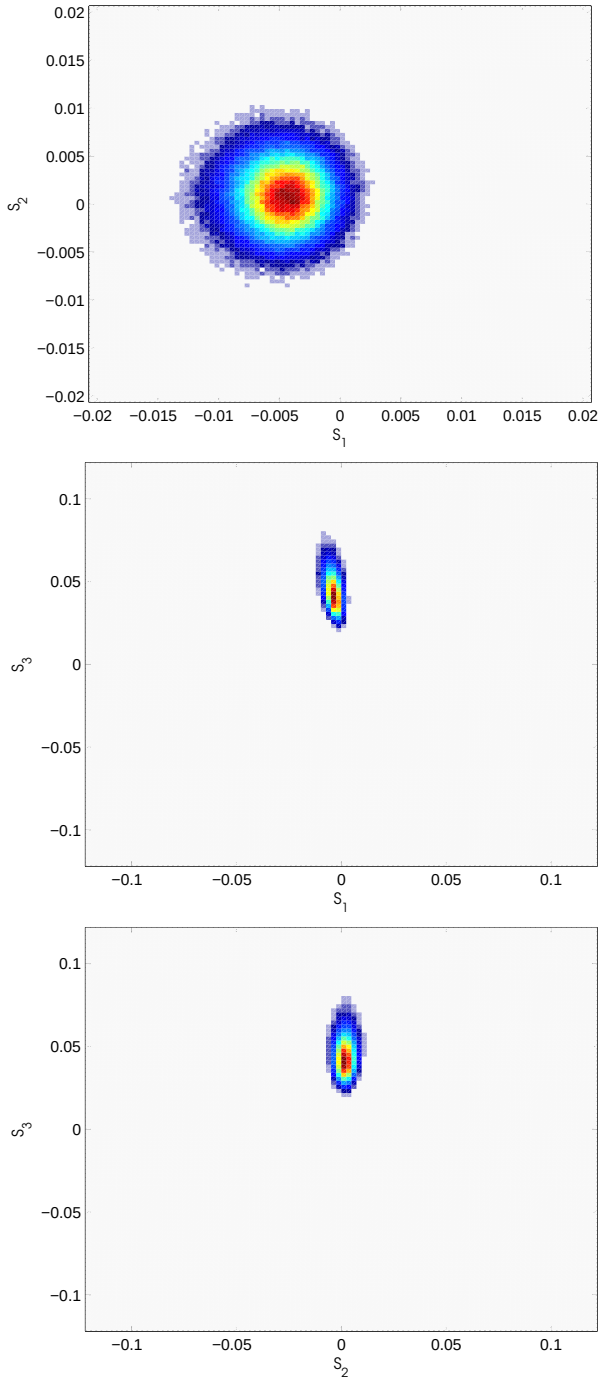

Figure A.35: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0825-1, medium resolution mode, for right-circular transmit polarization.

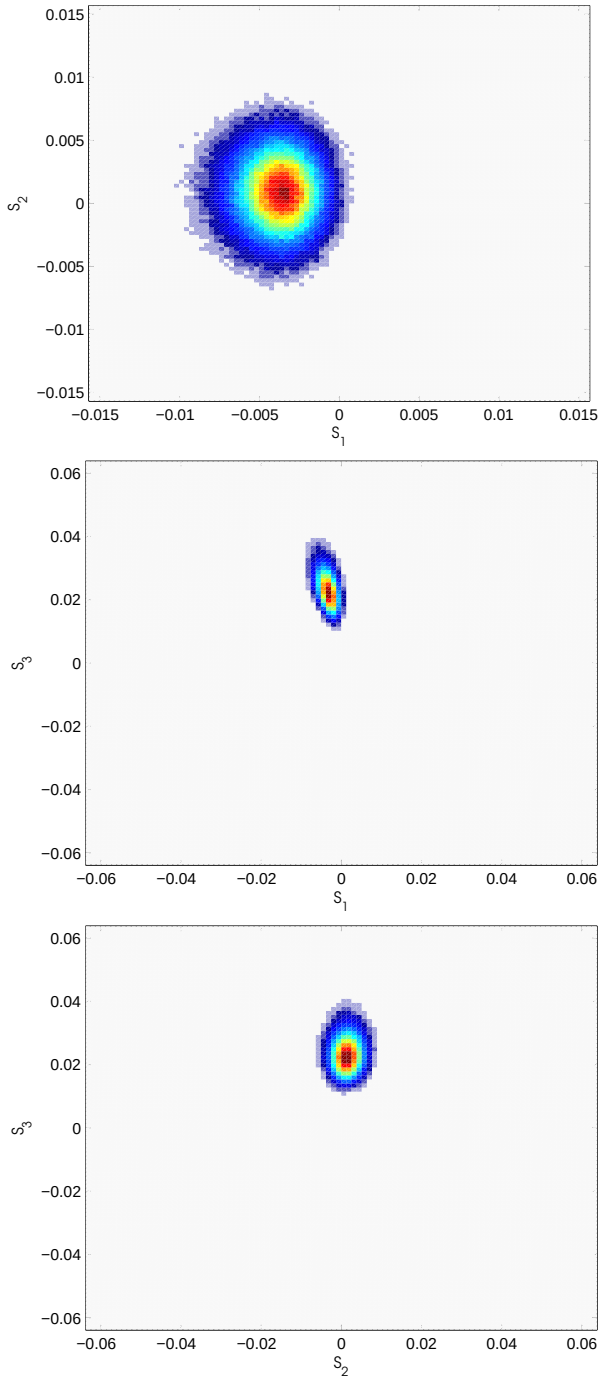

Figure A.36: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, medium resolution mode, for right-circular transmit polarization.

A.4 Ship Detection Mode

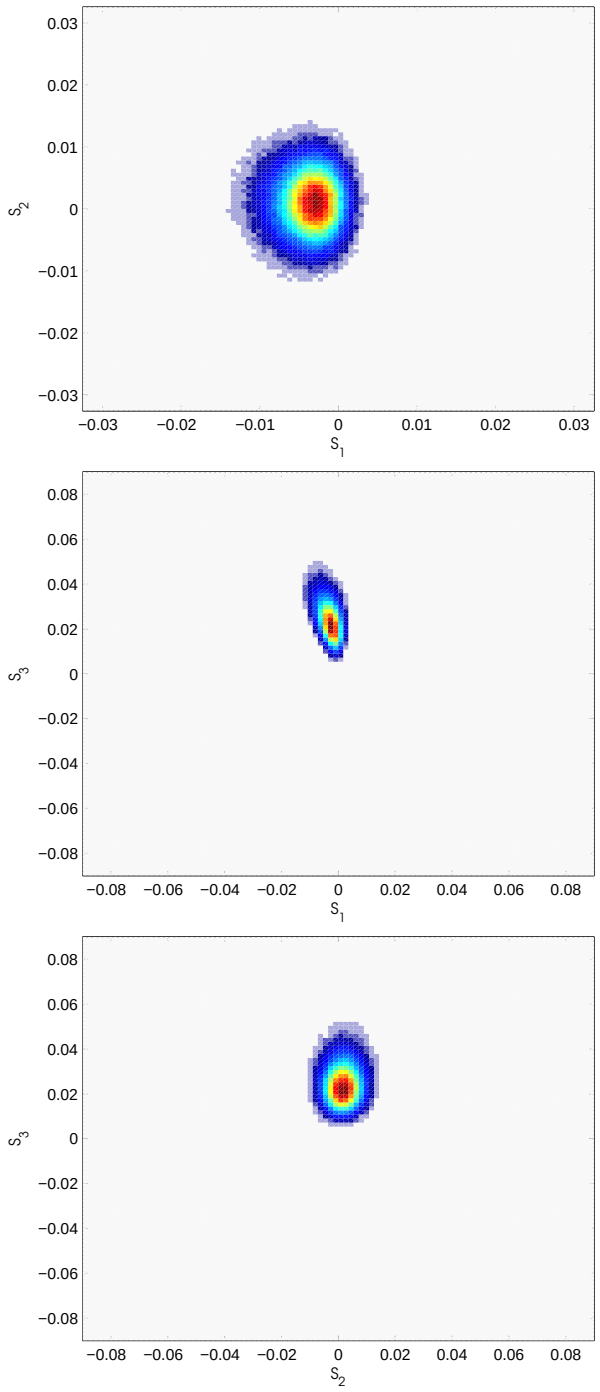

Figure A.37: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, ship detection mode, for right-circular transmit polarization.

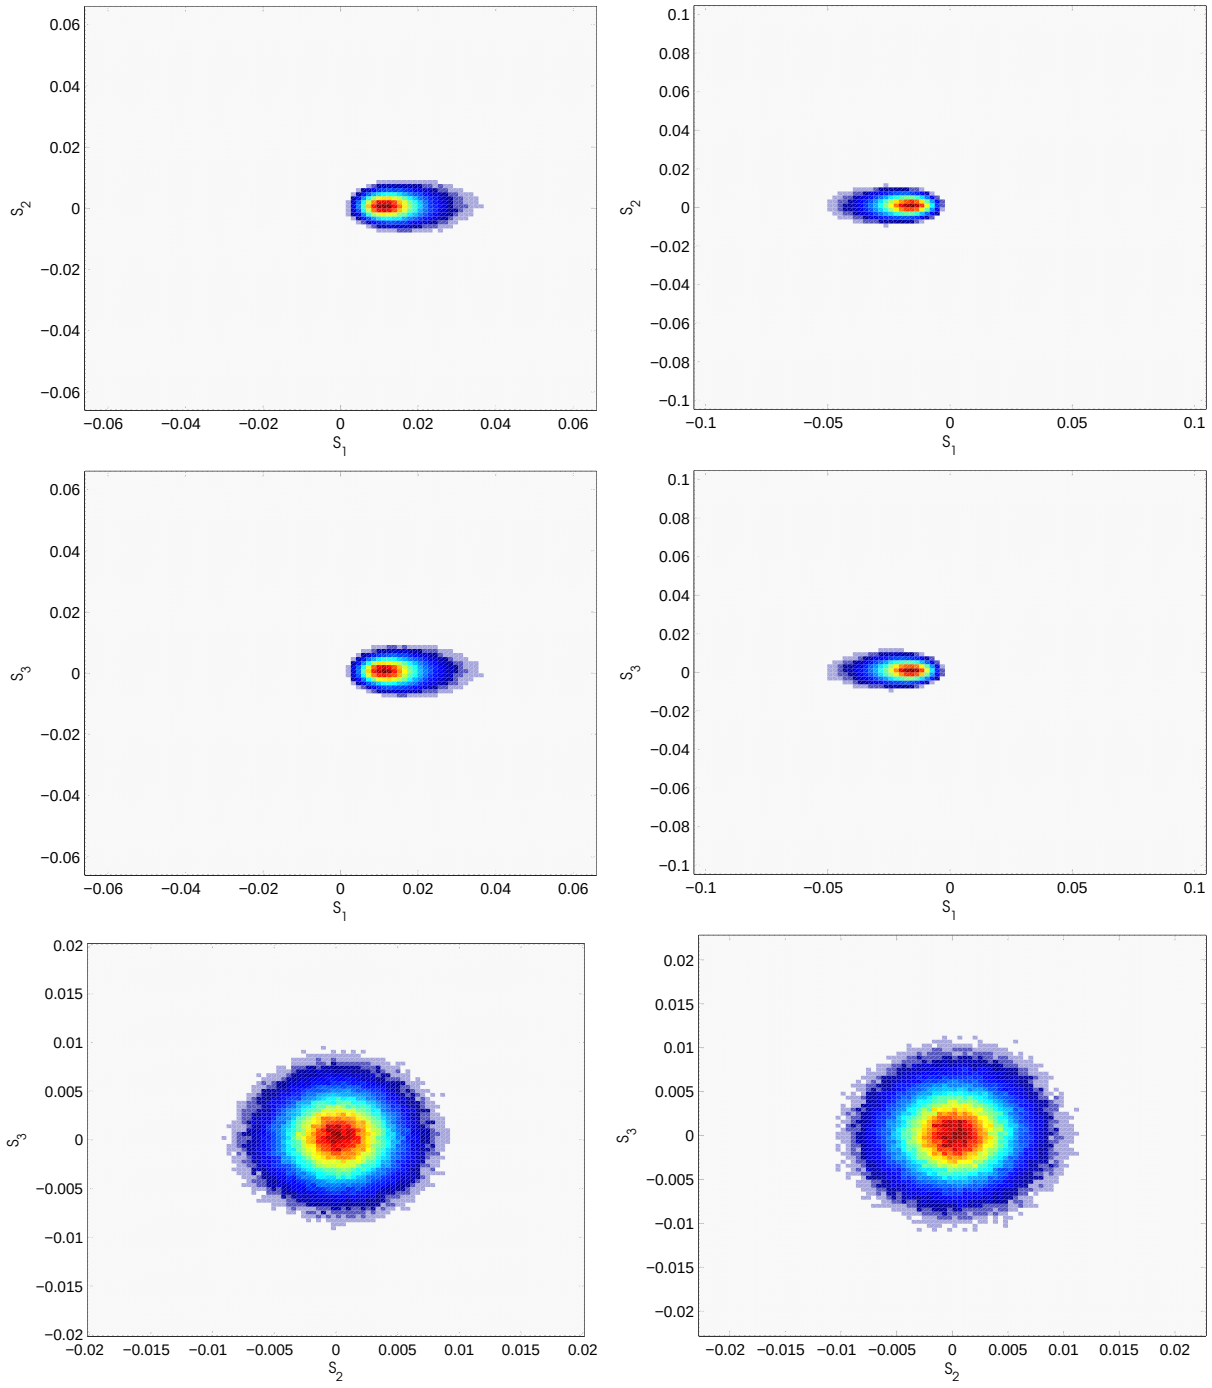

(a) HH-HV

(b) VV-VH

Figure A.38: Stokes vector density plots showing the most common (red) and least common (blue) values for an ocean image chip taken from scene 0818-1, ship detection mode, for horizontal transmit polarization (left column) and vertical transmit polarization (right column).

Appendix B

Plots of ψ, χ, and m for Iceberg and Ocean Pixels

B.1 Low Resolution Mode

Ocean pixels are shown as a density plot (2-D histogram, with blue representing the least common values and red representing the most common values). Iceberg pixels are plotted with green symbols and ship pixels are plotted with black symbols.

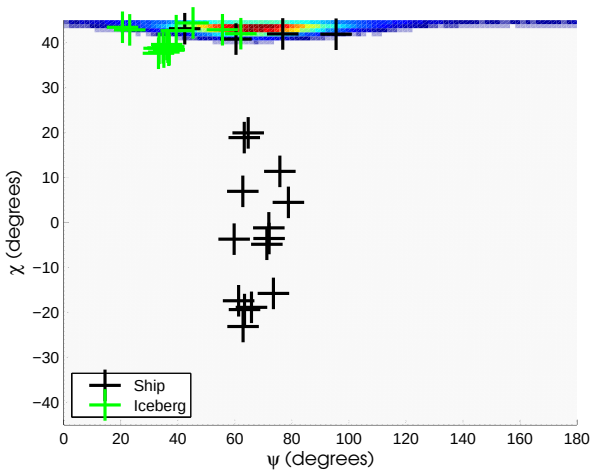

(a) RH-RV

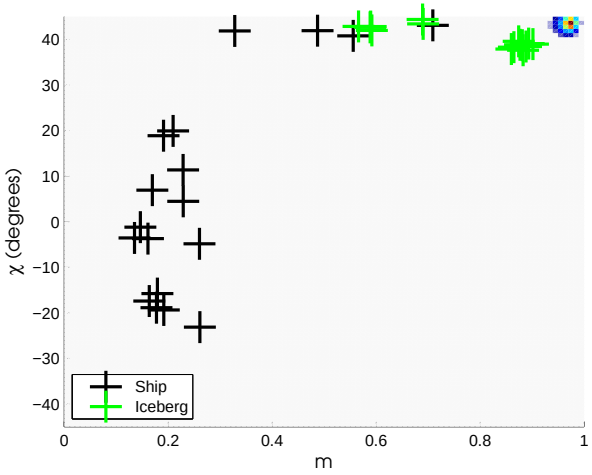

(b) RH-RV

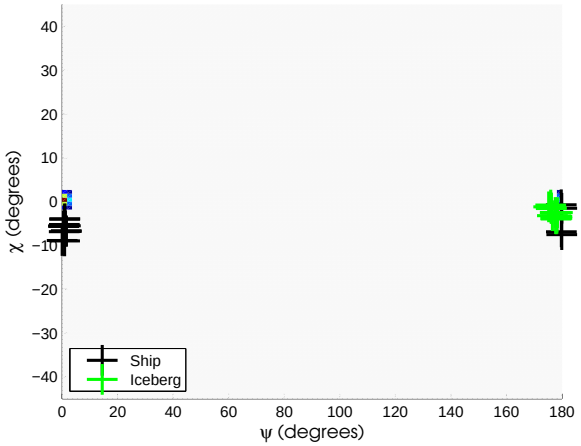

(c) HH-HV

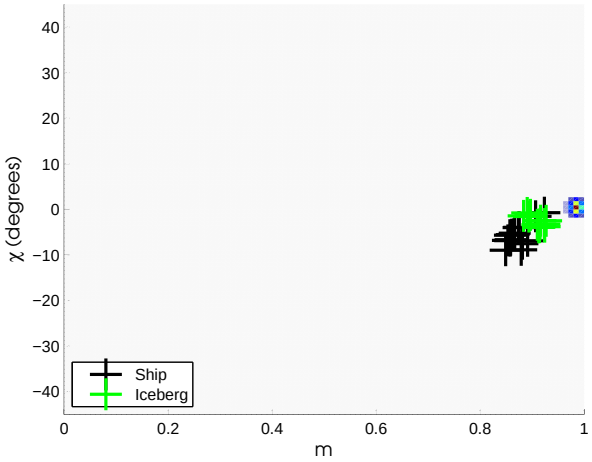

(d) HH-HV

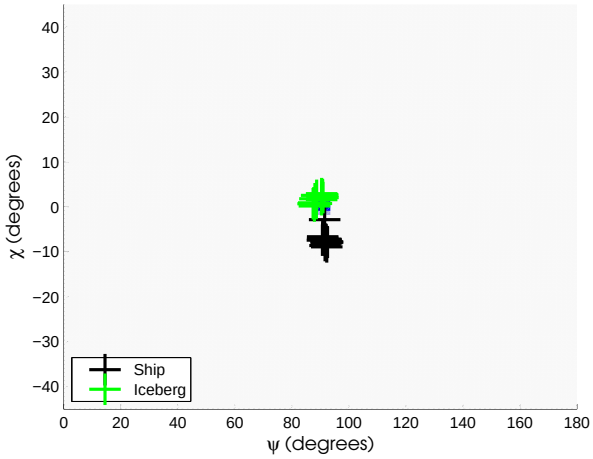

(e) VV-VH

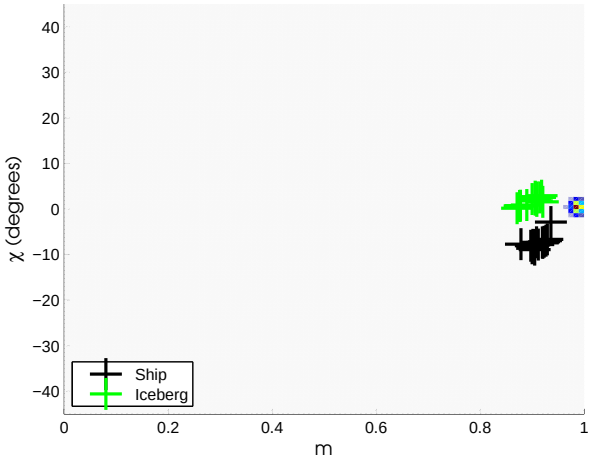

(f) VV-VH

Figure B.1: Plot of χ vs. ψ (left column) and χ vs. m (right column) for scene 0905-1, low resolution mode.

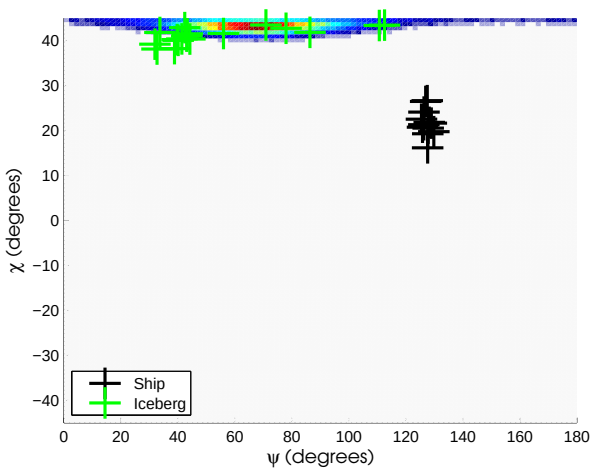

(a) RH-RV

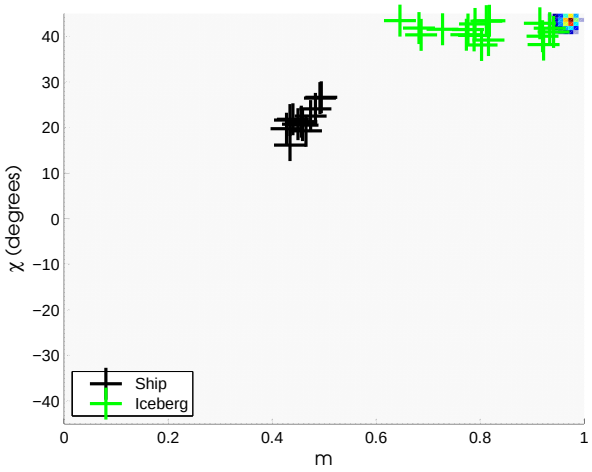

(b) RH-RV

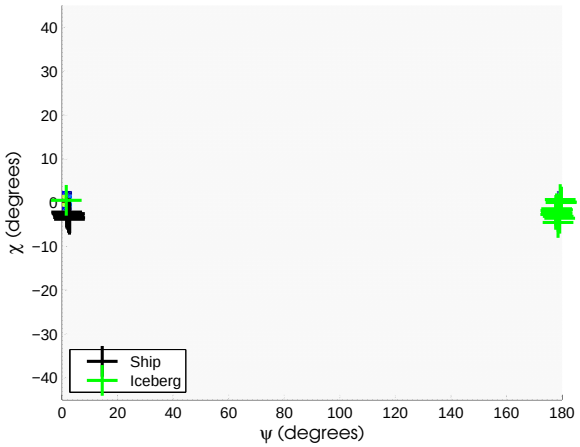

(c) HH-HV

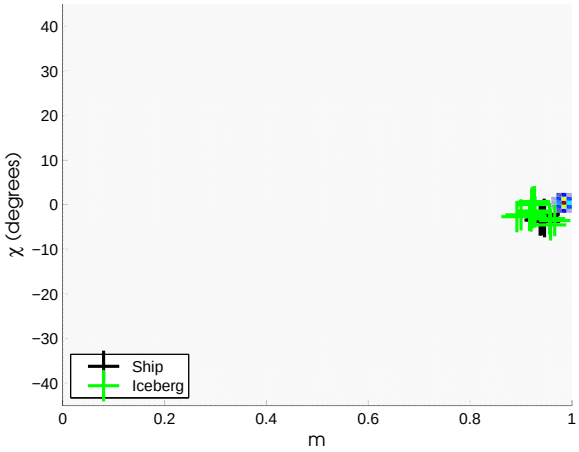

(d) HH-HV

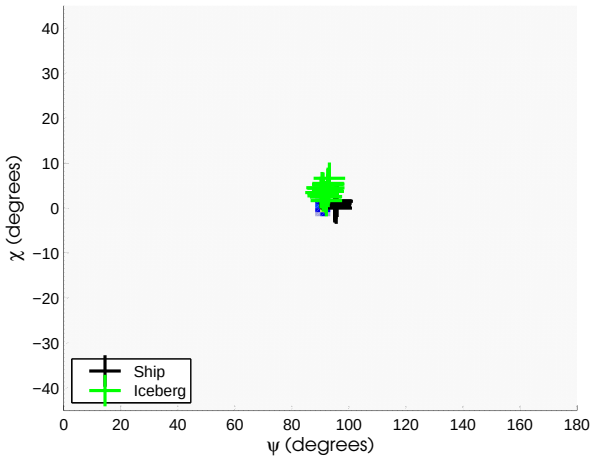

(e) VV-VH

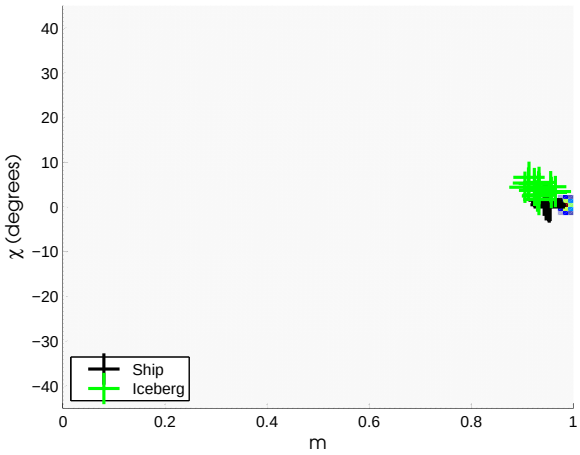

(f) VV-VH

Figure B.2: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0905-2, low resolution mode.

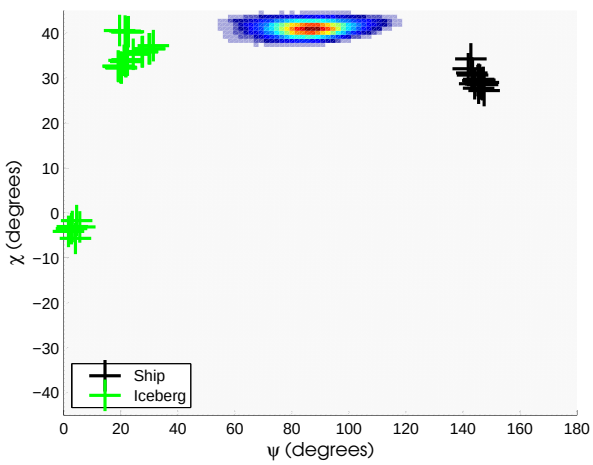

(a) RH-RV

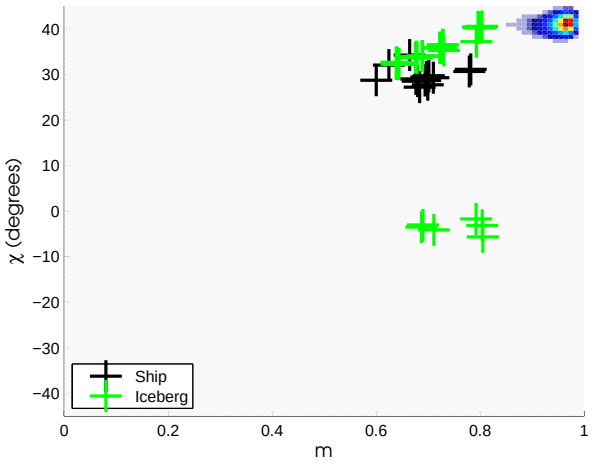

(b) RH-RV

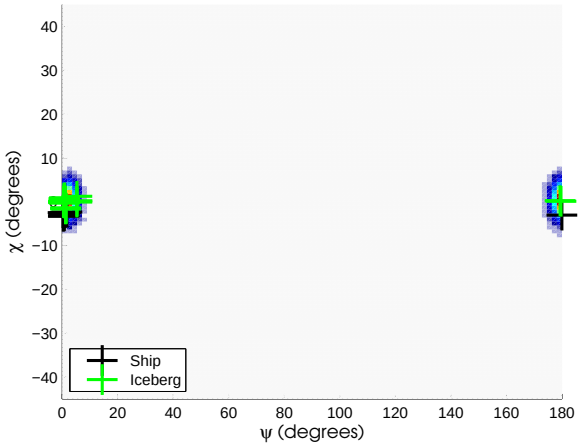

(c) HH-HV

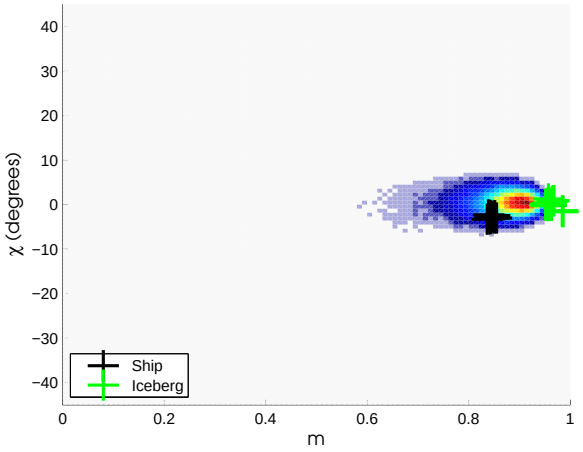

(d) HH-HV

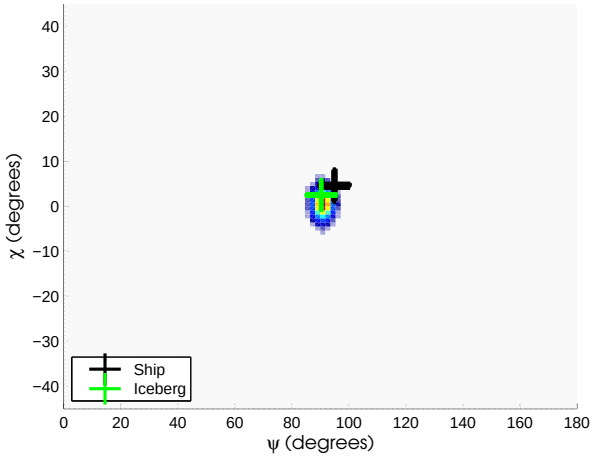

(e) VV-VH

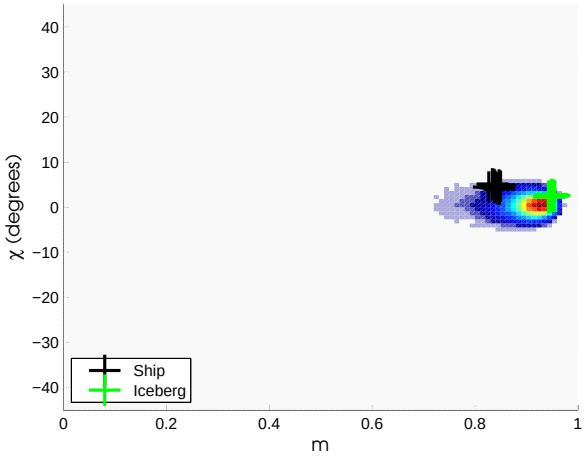

(f) VV-VH

Figure B.3: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0829-1, low resolution mode.

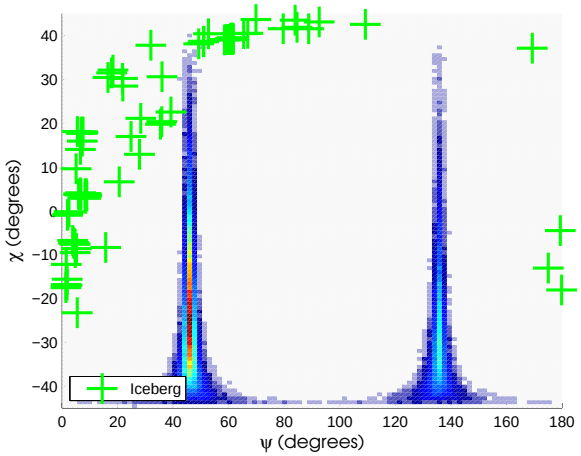

(a) RH-RV

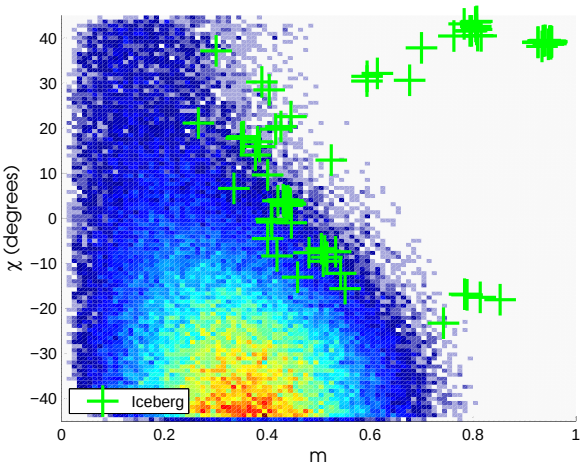

(b) RH-RV

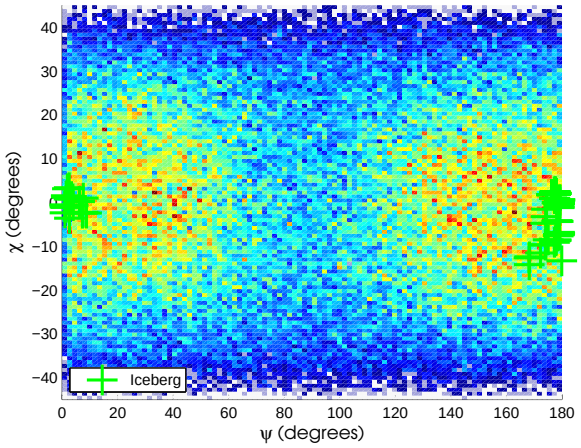

(c) HH-HV

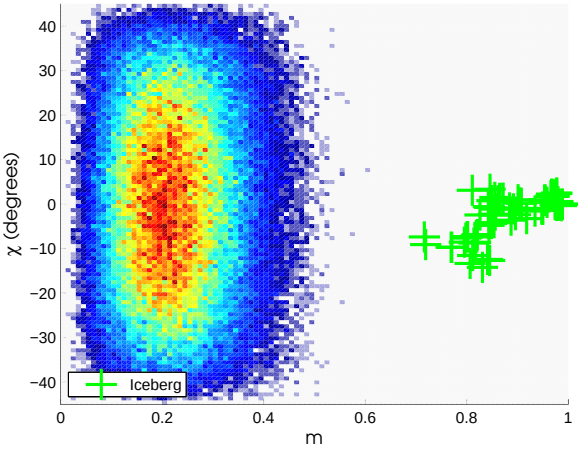

(d) HH-HV

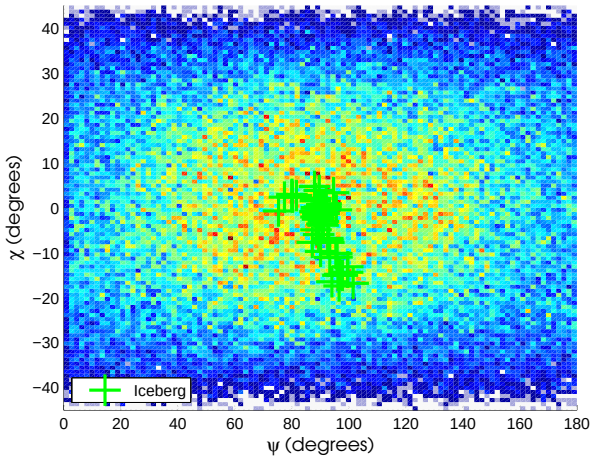

(e) VV-VH

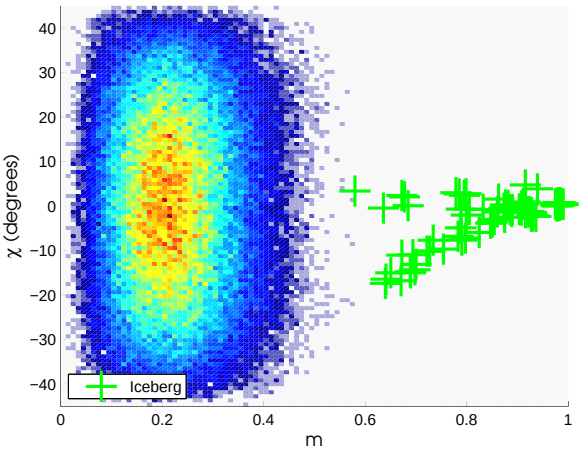

(f) VV-VH

Figure B.4: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0805-1, low resolution mode.

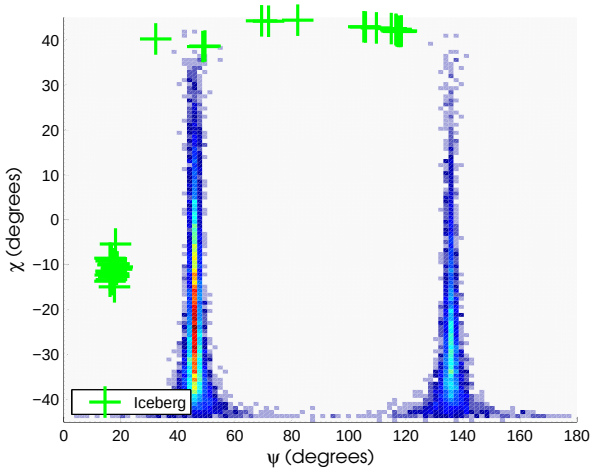

(a) RH-RV

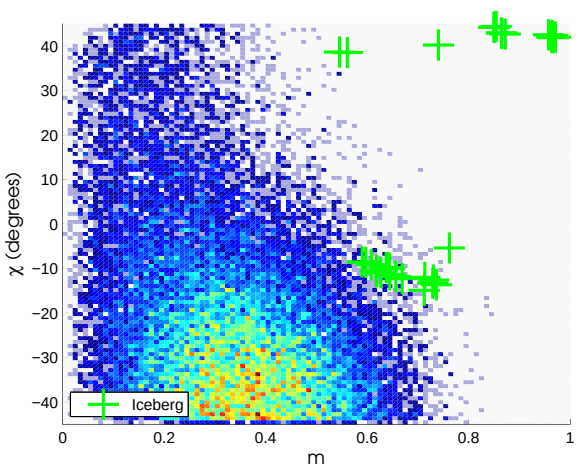

(b) RH-RV

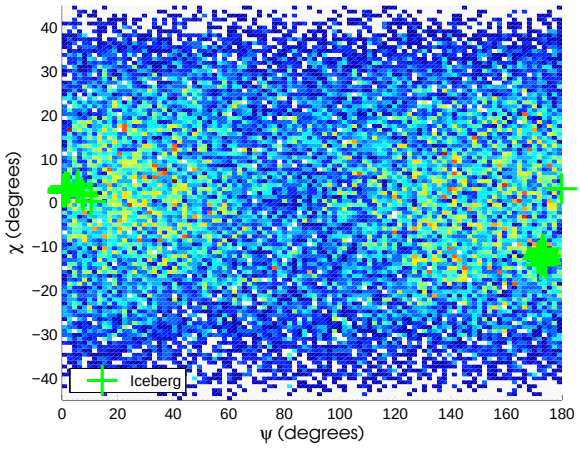

(c) HH-HV

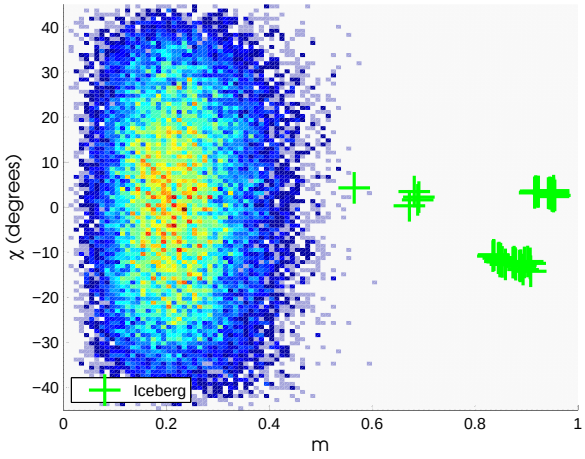

(d) HH-HV

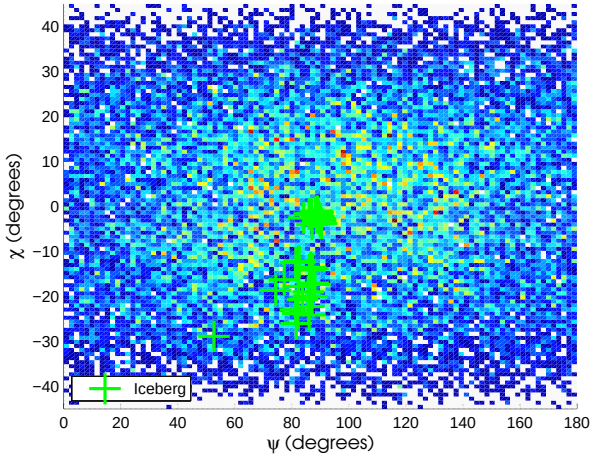

(e) VV-VH

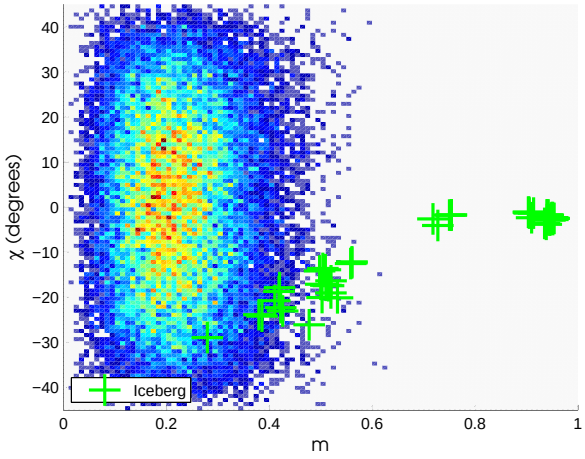

(f) VV-VH

Figure B.5: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0805-2, low resolution mode.

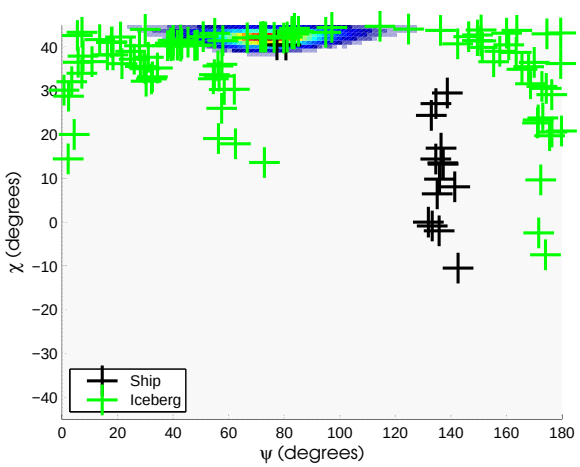

(a) RH-RV

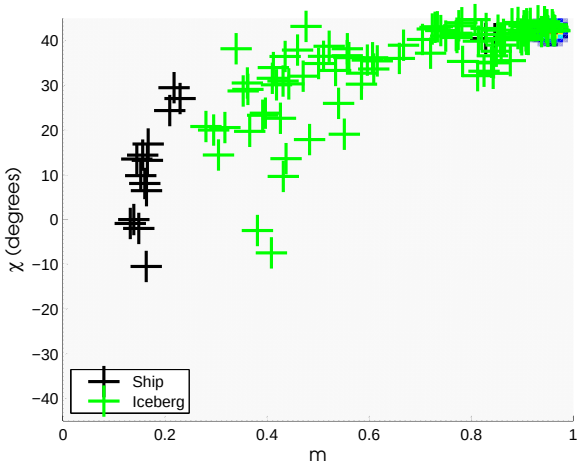

(b) RH-RV

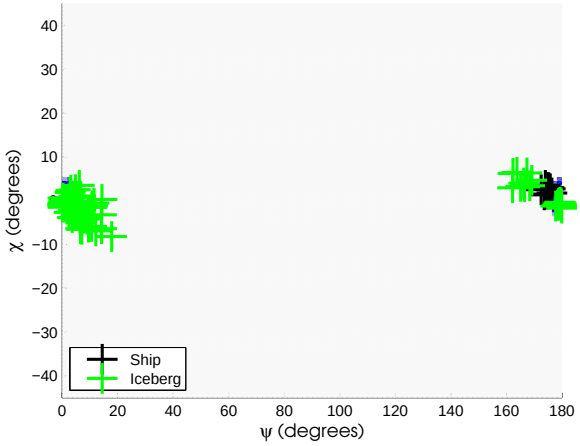

(c) HH-HV

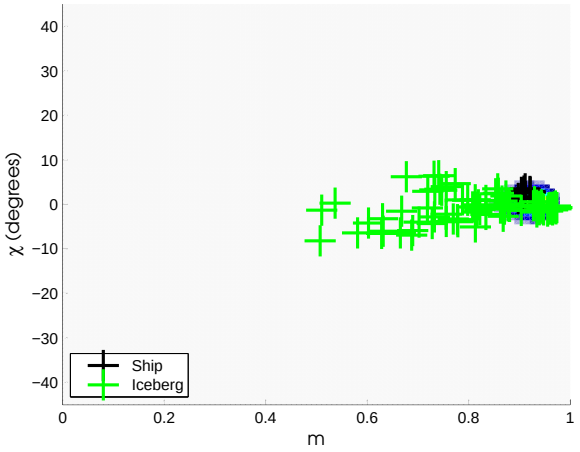

(d) HH-HV

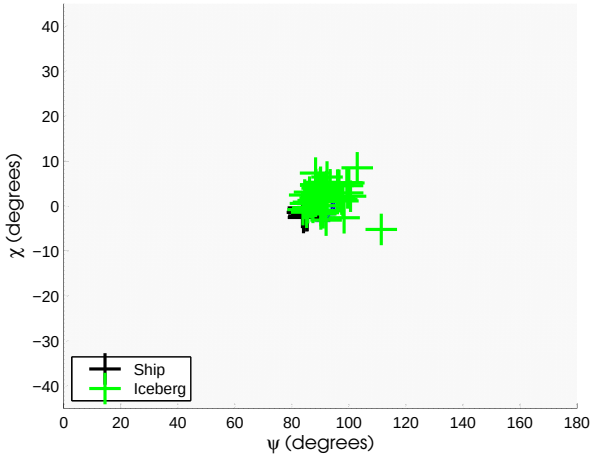

(e) VV-VH

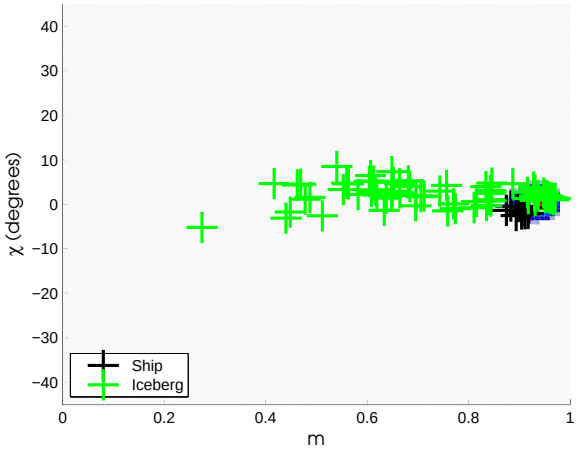

(f) VV-VH

Figure B.6: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0815-2, low resolution mode.

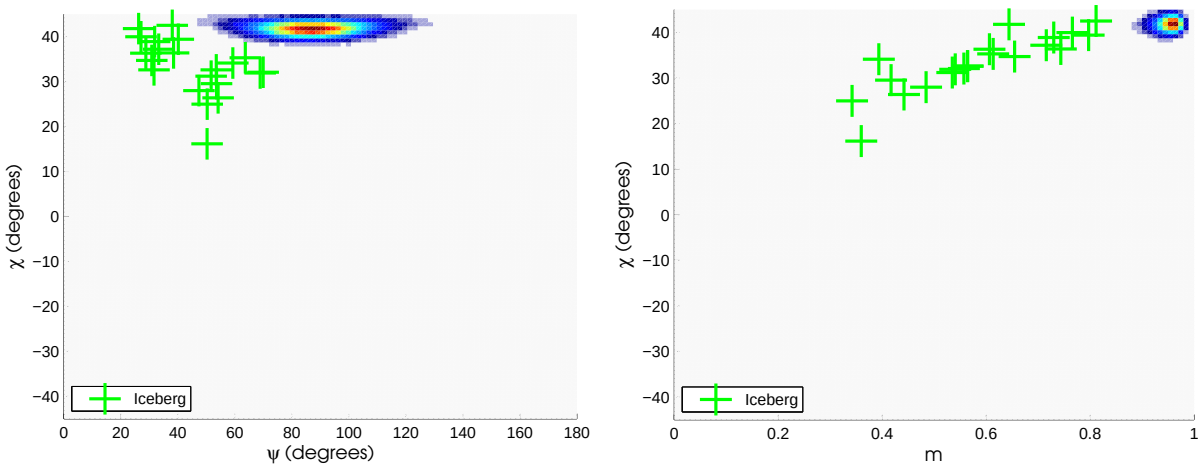

(a) RH-RV

(b) RH-RV

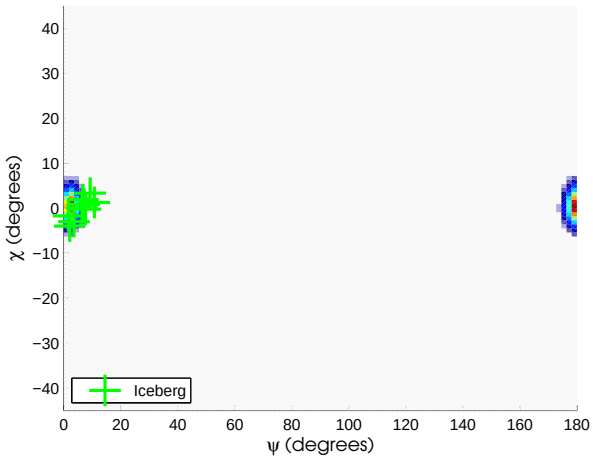

(c) HH-HV

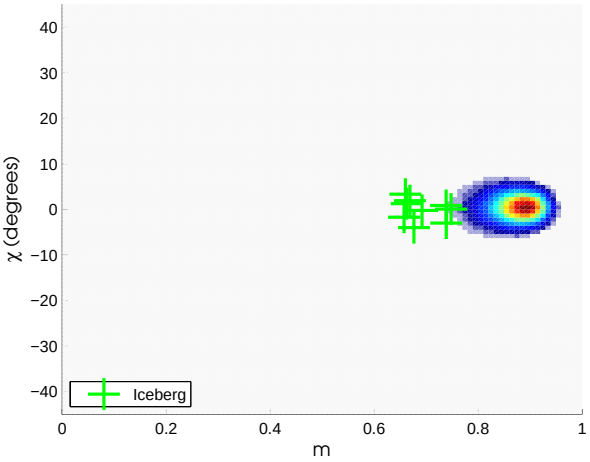

(d) HH-HV

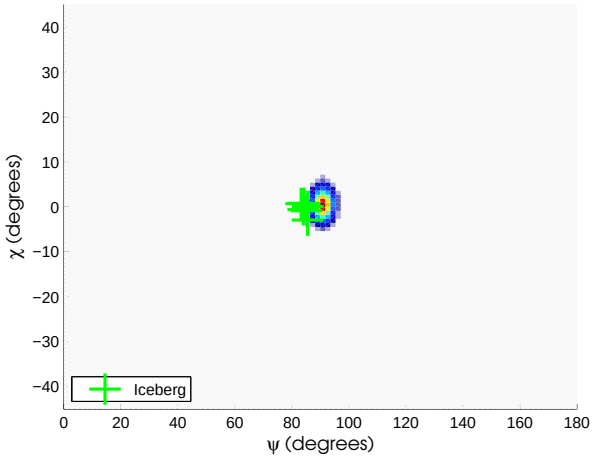

(e) VV-VH

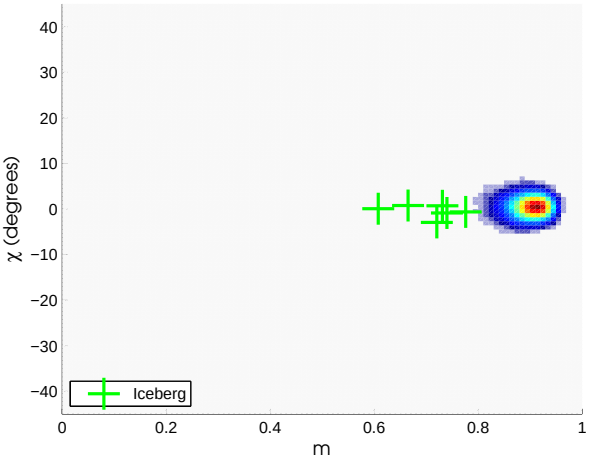

(f) VV-VH

Figure B.7: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0822-1, low resolution mode.

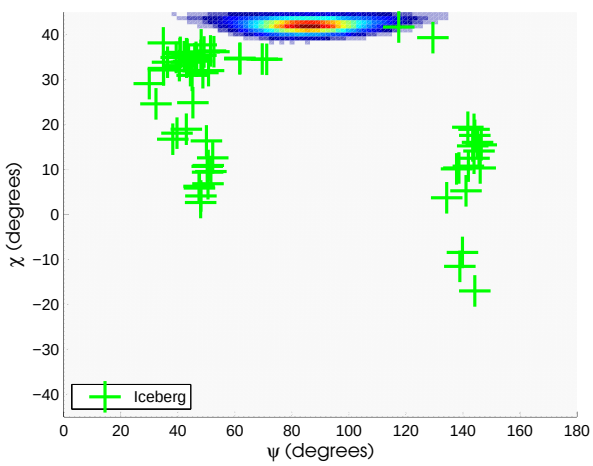

(a) RH-RV

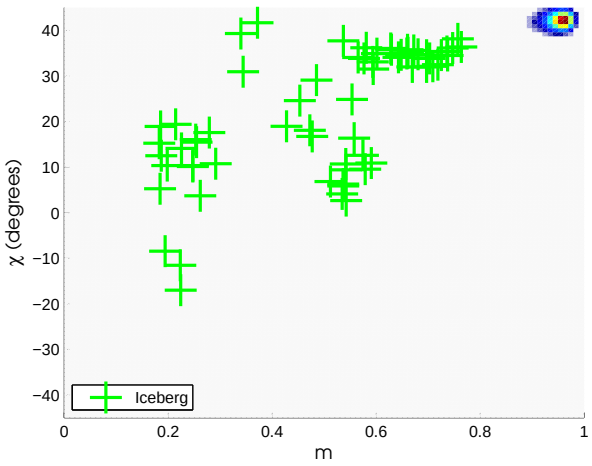

(b) RH-RV

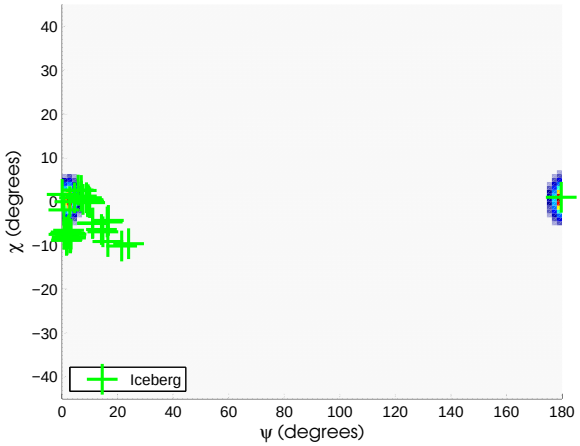

(c) HH-HV

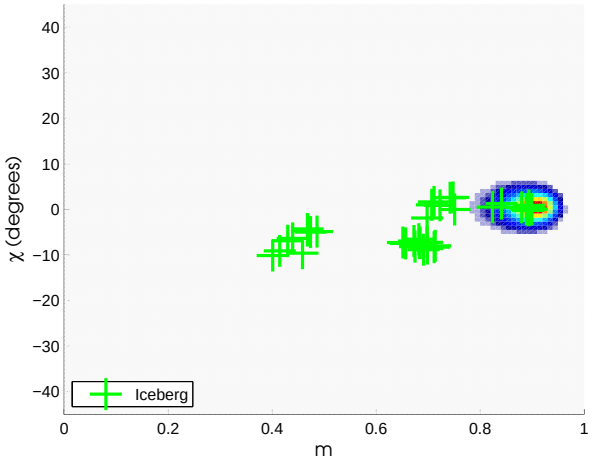

(d) HH-HV

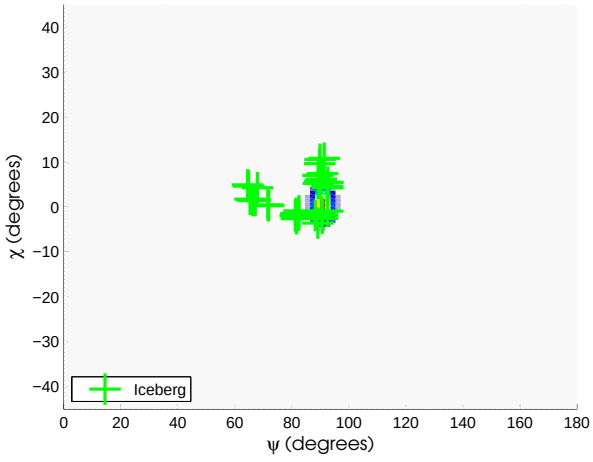

(e) VV-VH

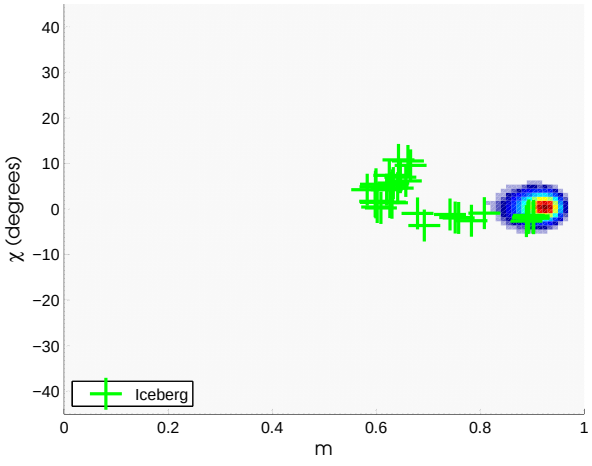

(f) VV-VH

Figure B.8: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0822-2, low resolution mode.

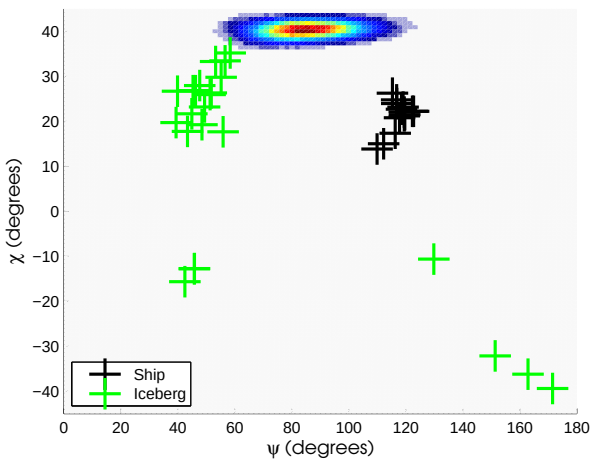

(a) RH-RV

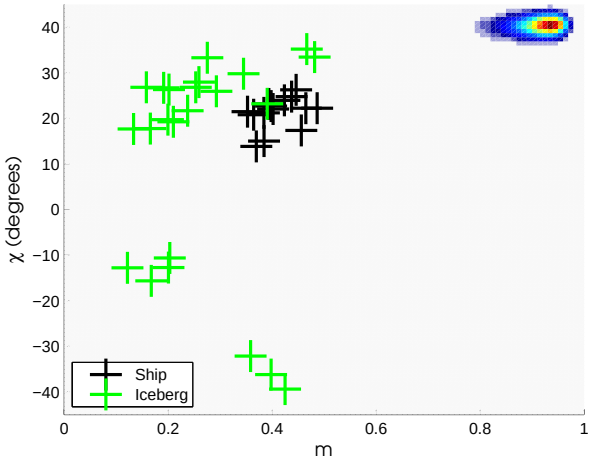

(b) RH-RV

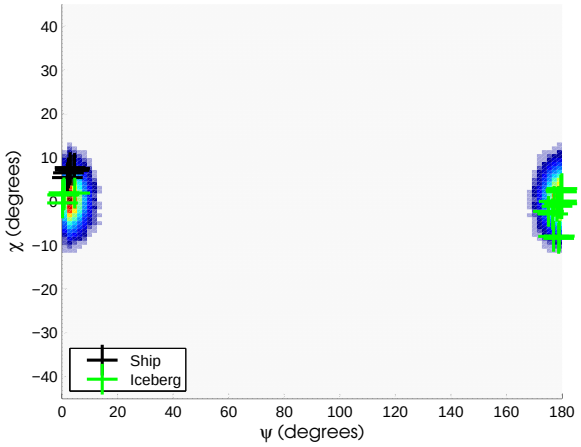

(c) HH-HV

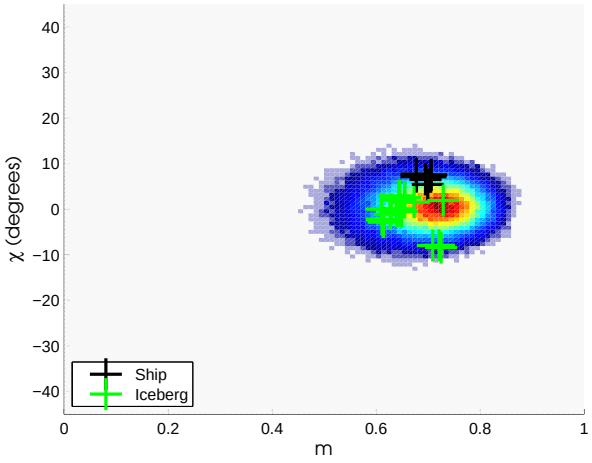

(d) HH-HV

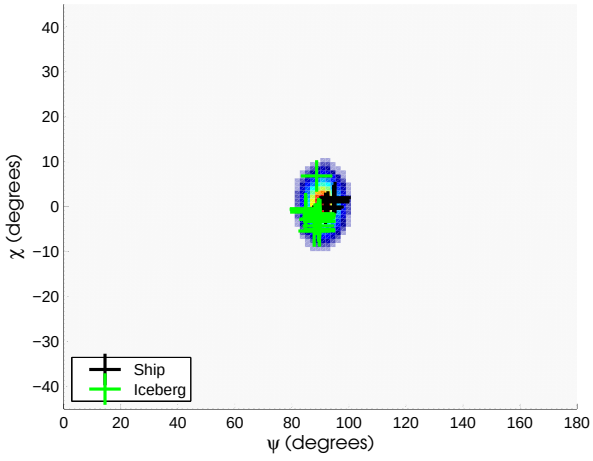

(e) VV-VH

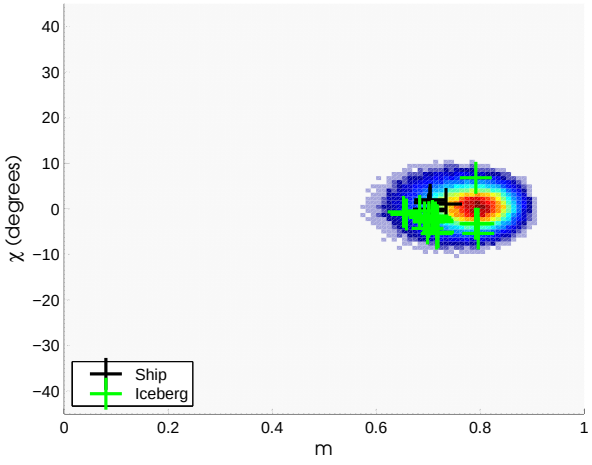

(f) VV-VH

Figure B.9: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0815-1, low resolution mode.

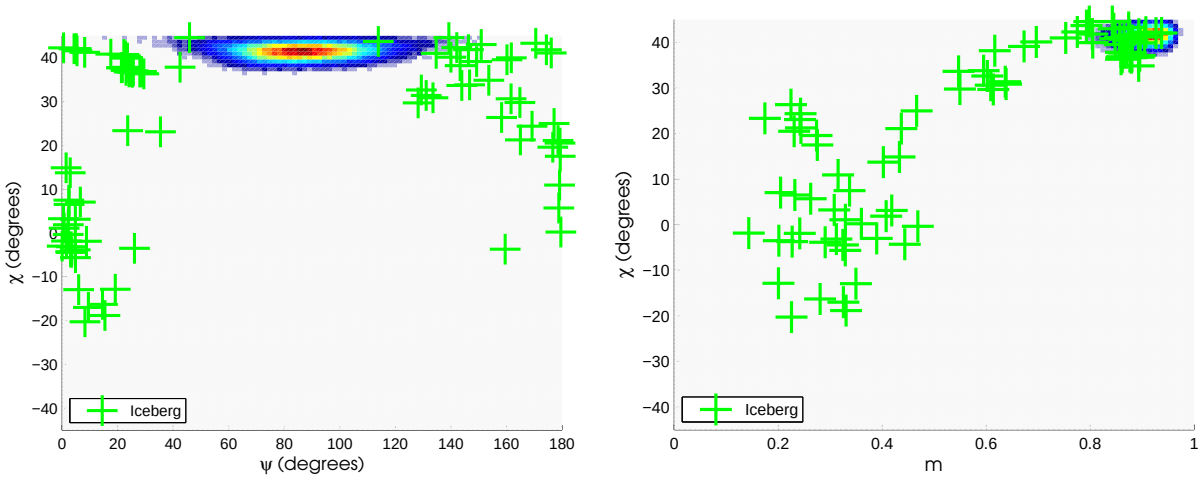

(a) RH-RV

(b) RH-RV

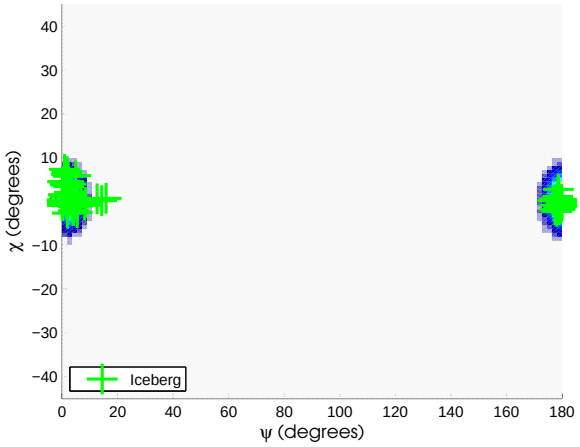

(c) HH-HV

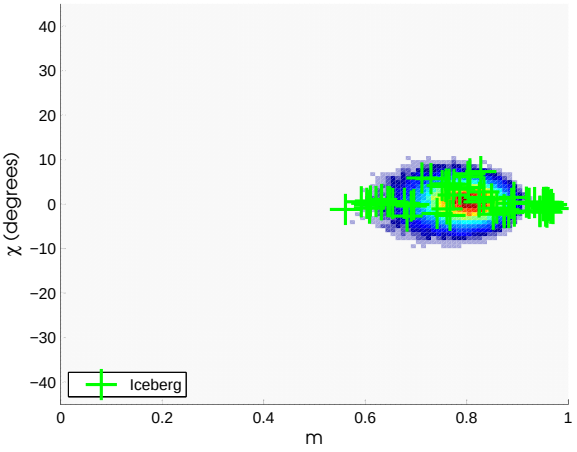

(d) HH-HV

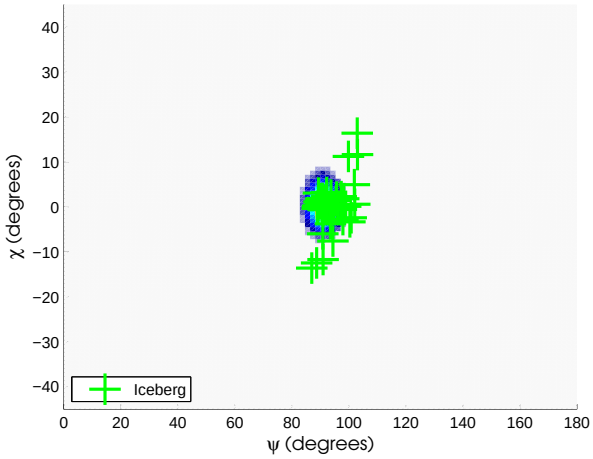

(e) VV-VH

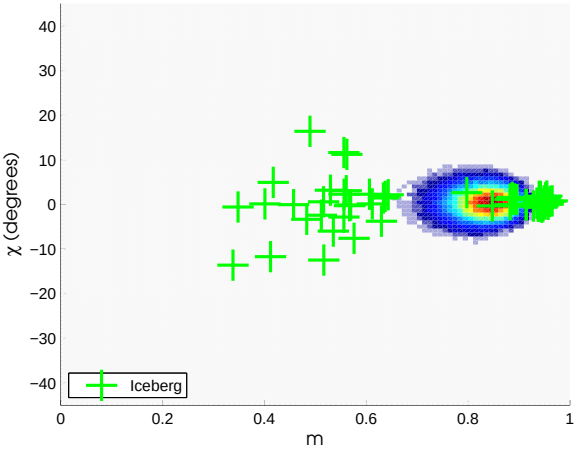

(f) VV-VH

Figure B.10: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0825-2, low resolution mode.

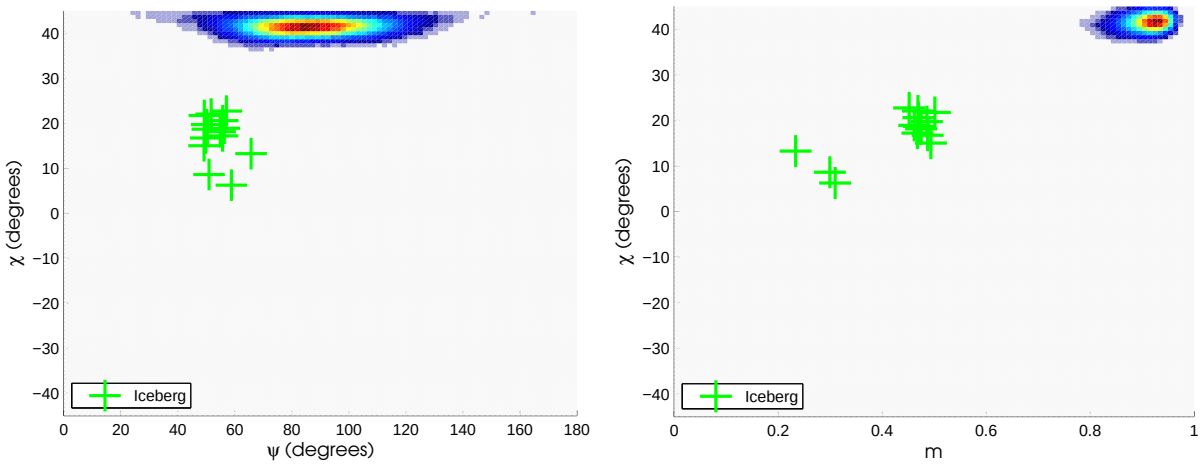

(a) RH-RV

(b) RH-RV

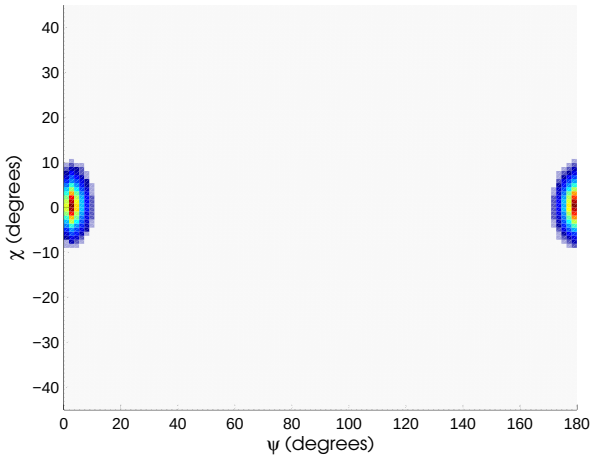

(c) HH-HV

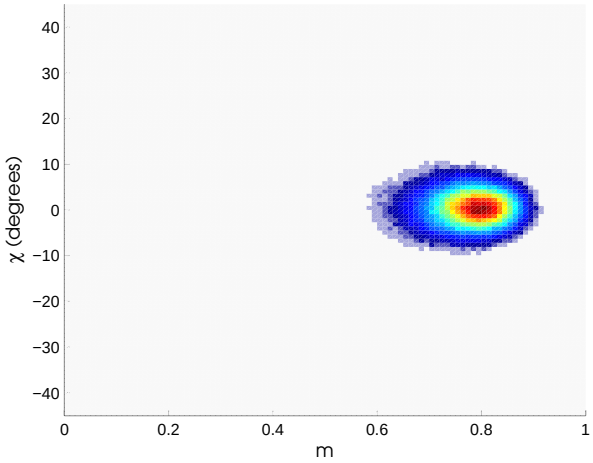

(d) HH-HV

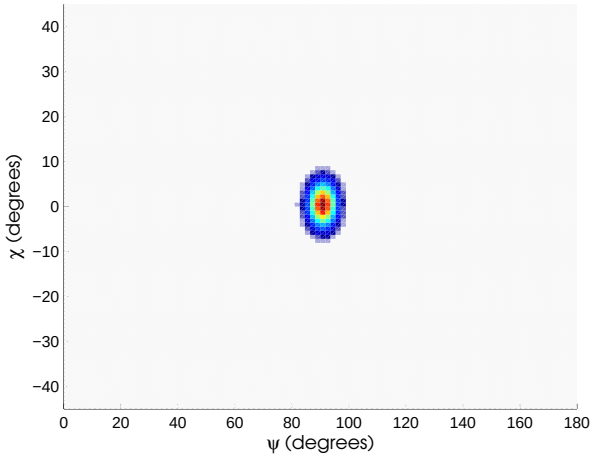

(e) VV-VH

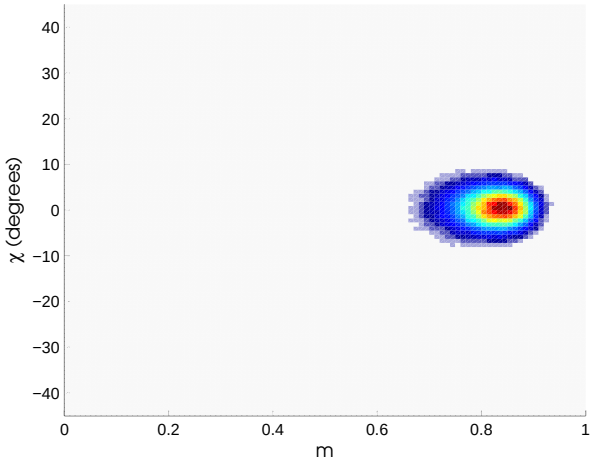

(f) VV-VH

Figure B.11: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0825-1, low resolution mode.

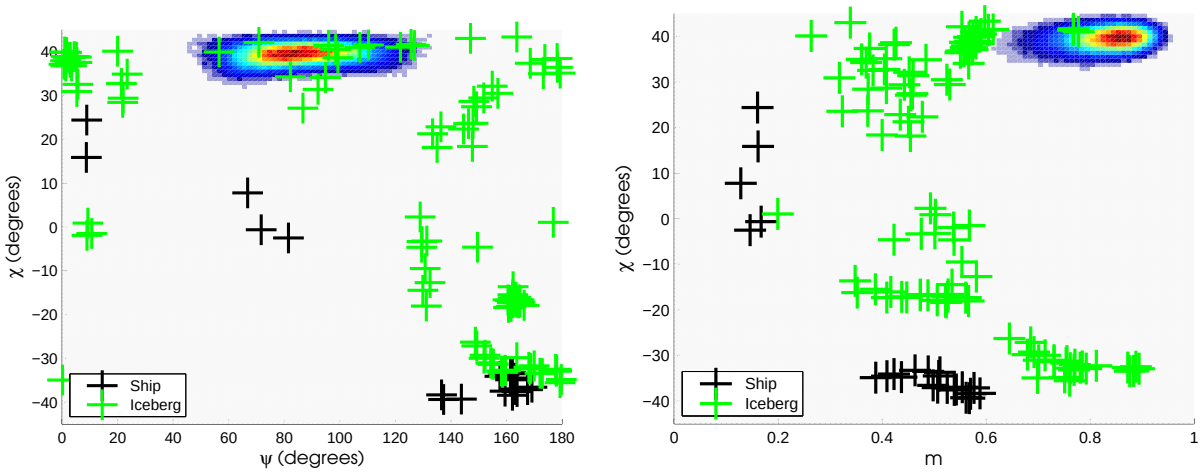

(a) RH-RV

(b) RH-RV

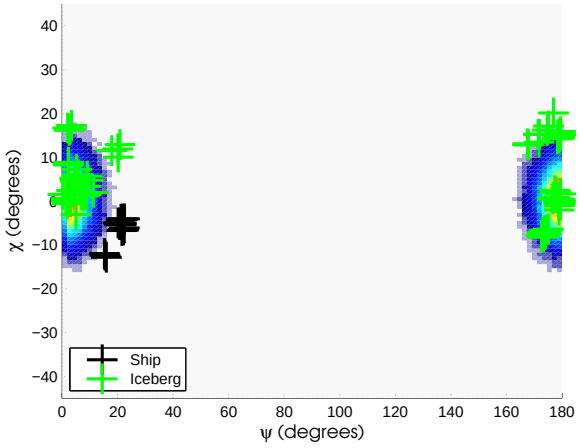

(c) HH-HV

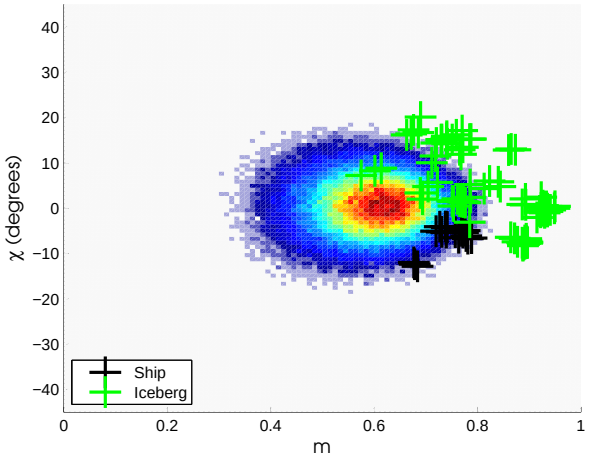

(d) HH-HV

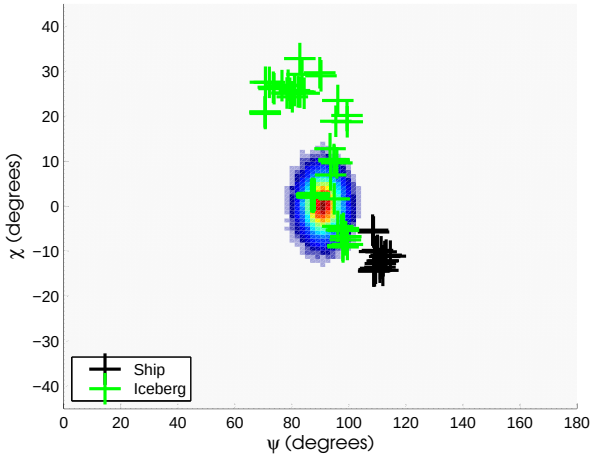

(e) VV-VH

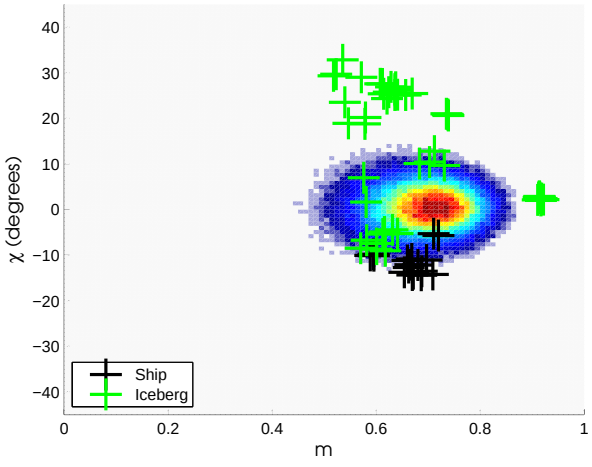

(f) VV-VH

Figure B.12: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0818-1, low resolution mode.

B.2 Medium Resolution Mode

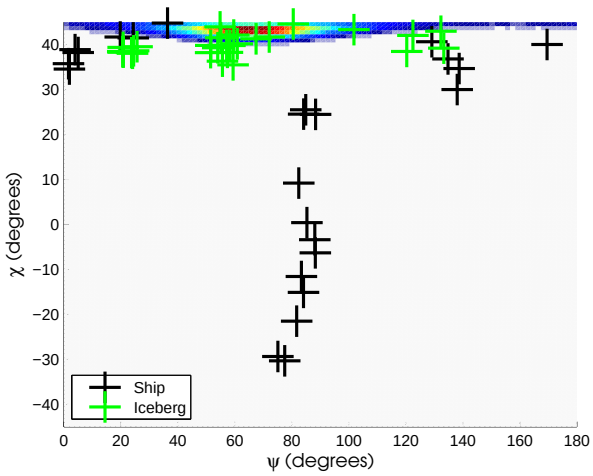

(a) RH-RV

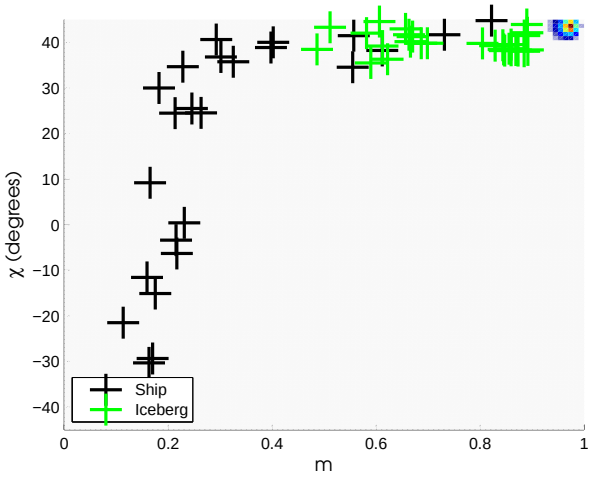

(b) RH-RV

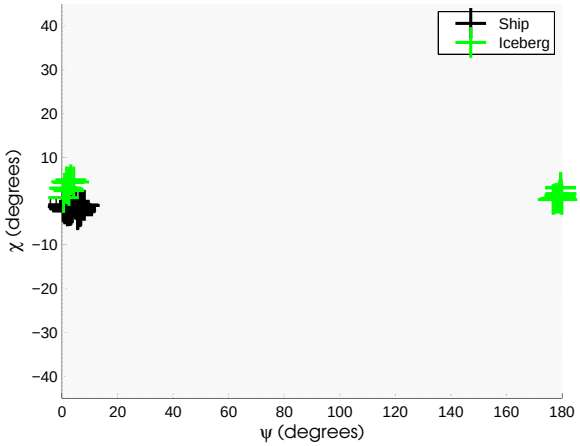

(c) HH-HV

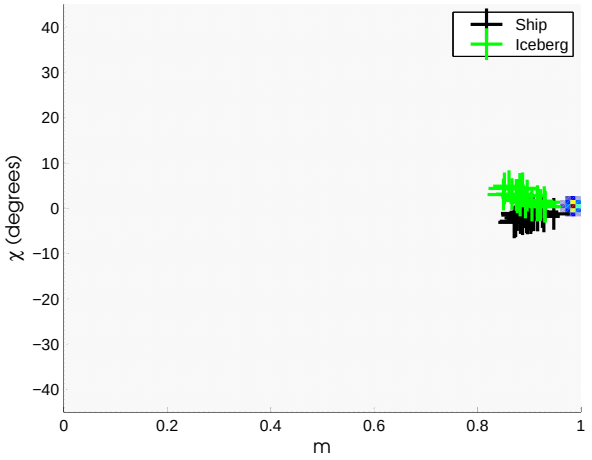

(d) HH-HV

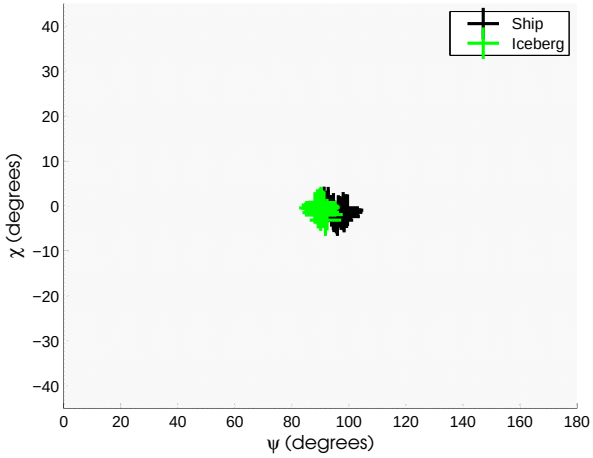

(e) VV-VH

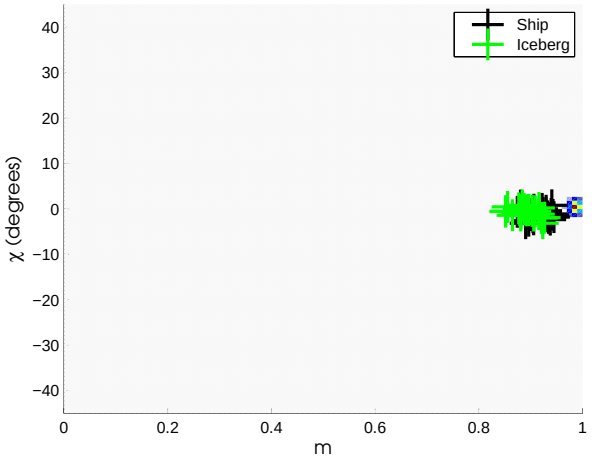

(f) VV-VH

Figure B.13: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0905-1, medium resolution mode.

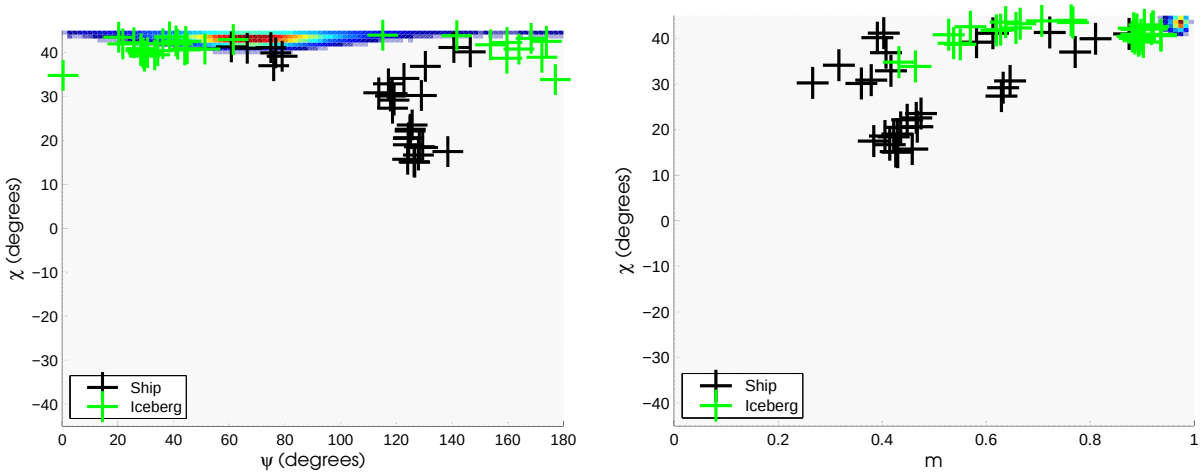

(a) RH-RV

(b) RH-RV

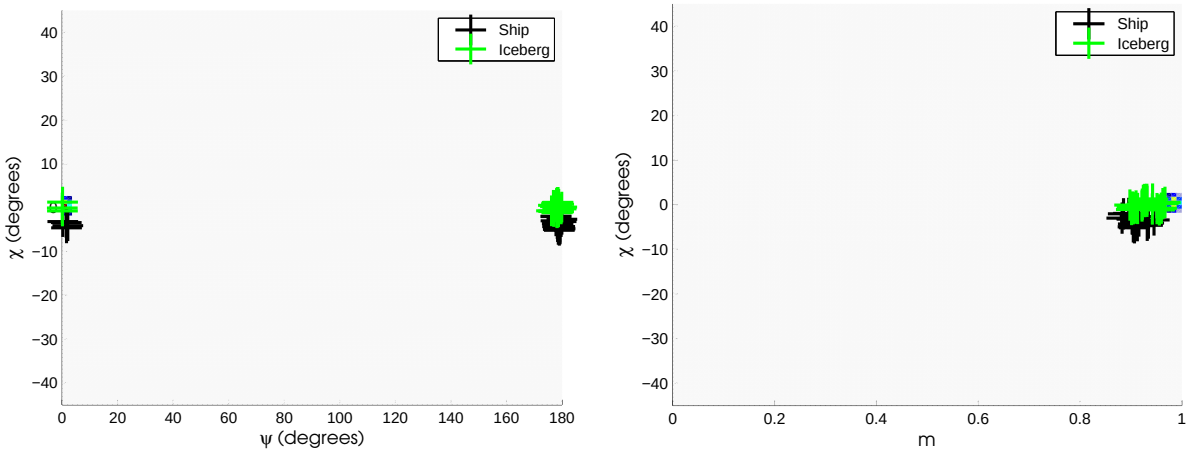

(c) HH-HV

(d) HH-HV

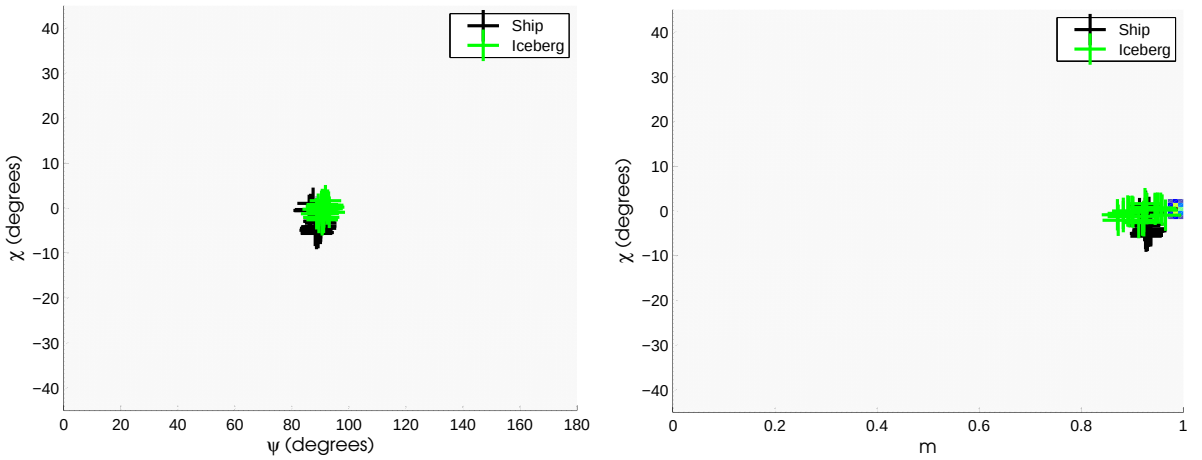

(e) VV-VH

(f) VV-VH

Figure B.14: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0905-2, medium resolution mode.

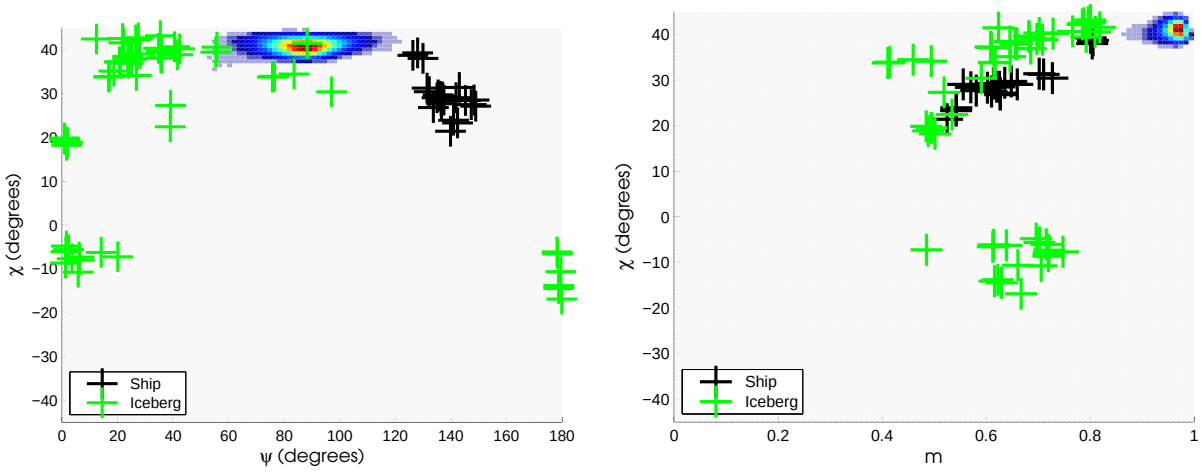

(a) RH-RV

(b) RH-RV

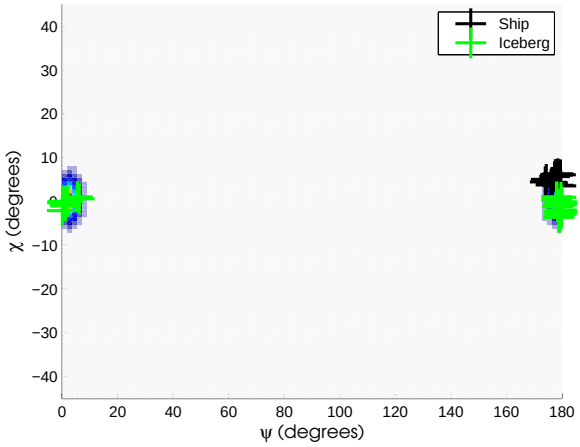

(c) HH-HV

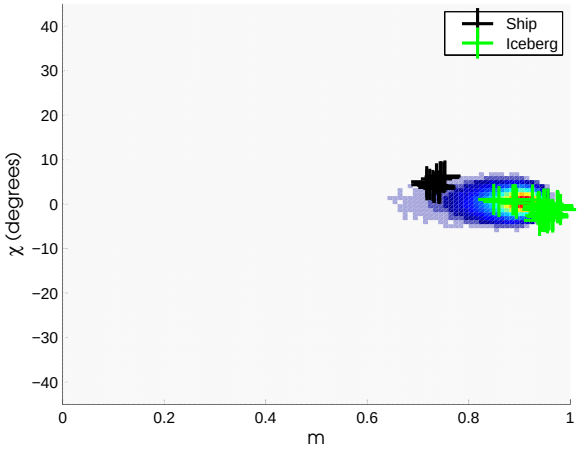

(d) HH-HV

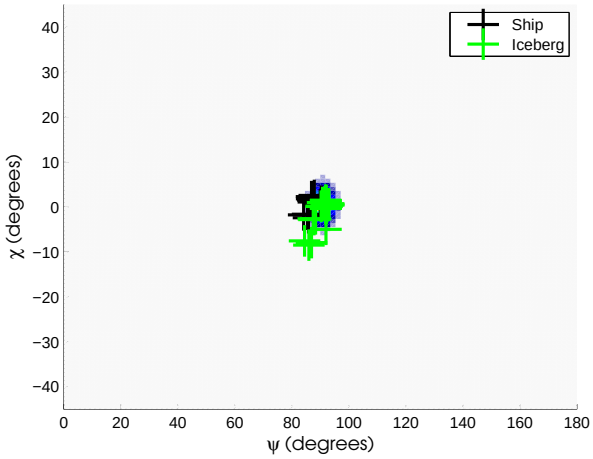

(e) VV-VH

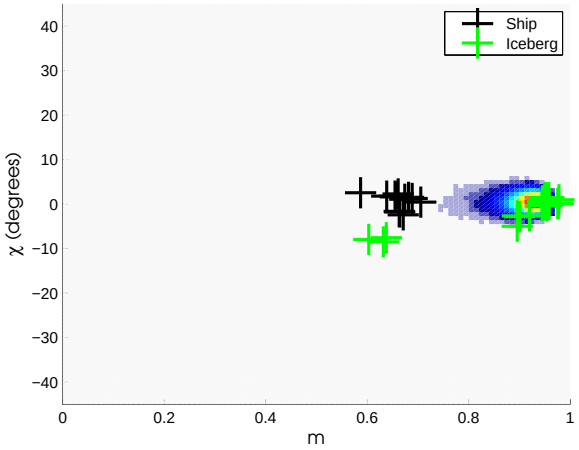

(f) VV-VH

Figure B.15: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0829-1, medium resolution mode.

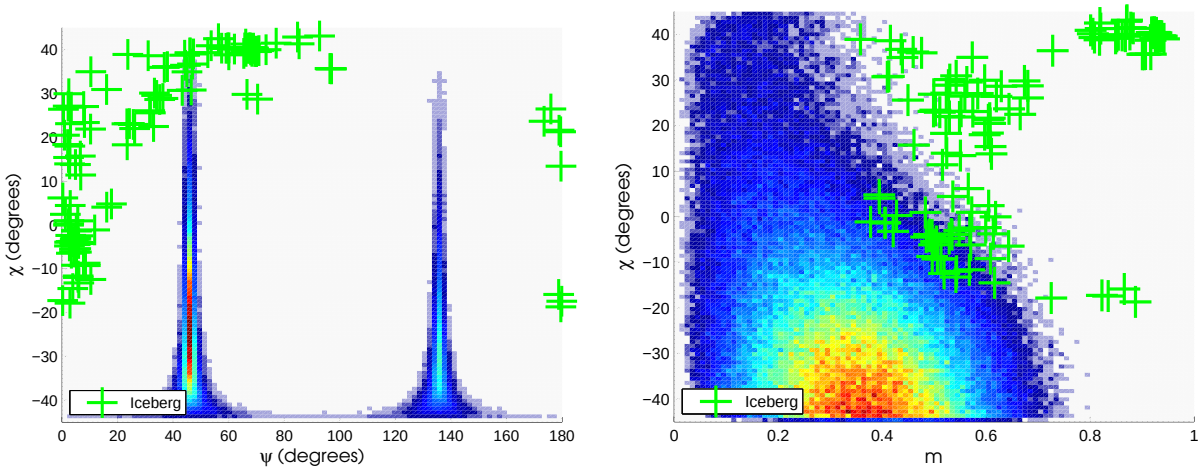

(a) RH-RV

(b) RH-RV

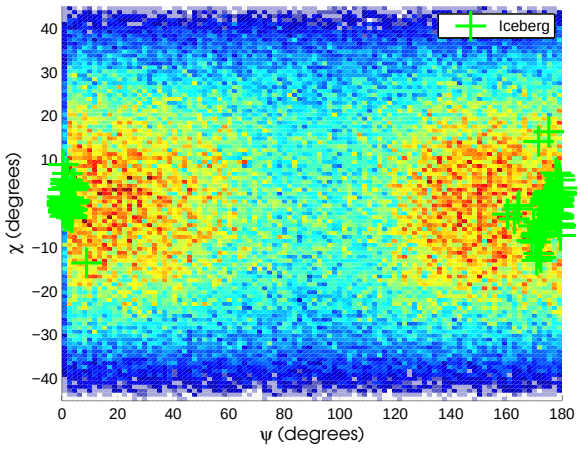

(c) HH-HV

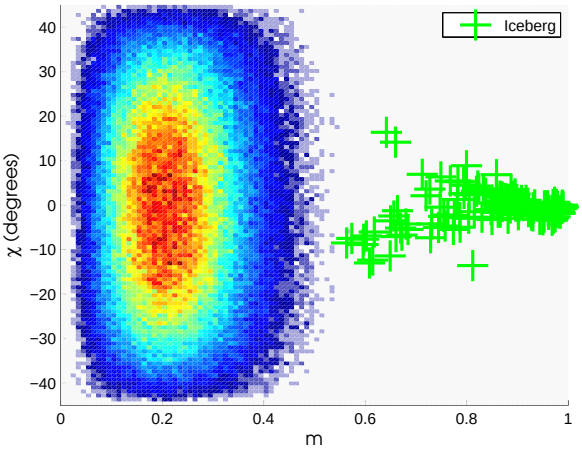

(d) HH-HV

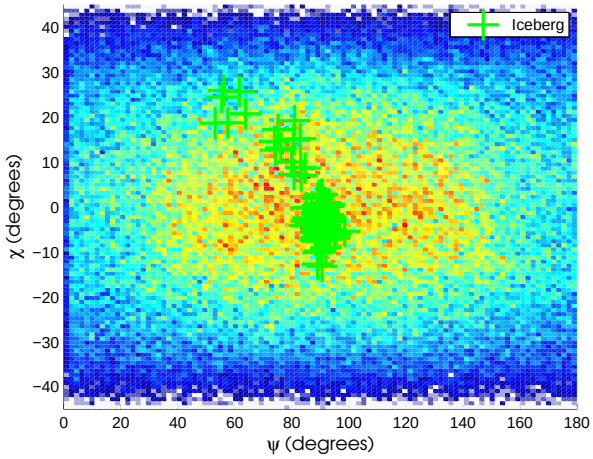

(e) VV-VH

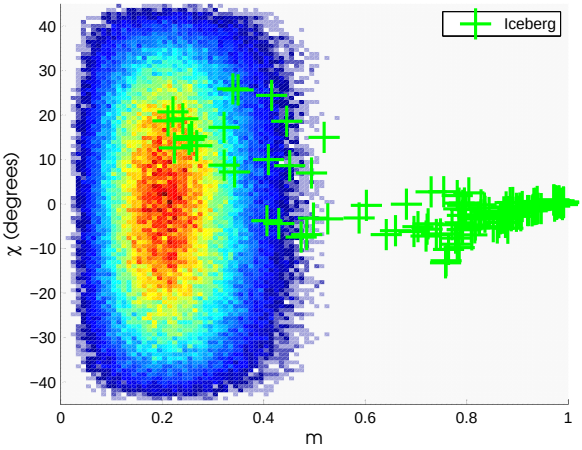

(f) VV-VH

Figure B.16: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0805-1, medium resolution mode.

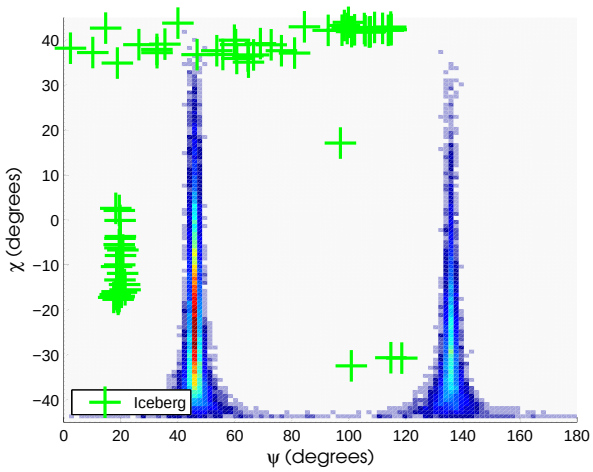

(a) RH-RV

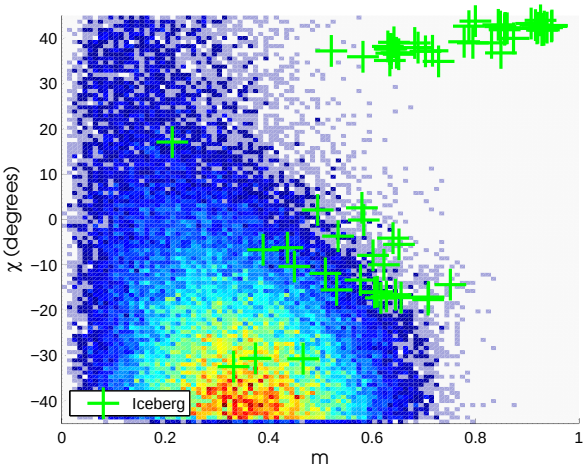

(b) RH-RV

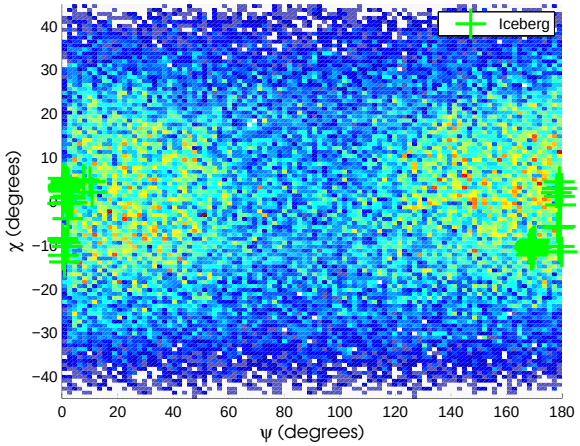

(c) HH-HV

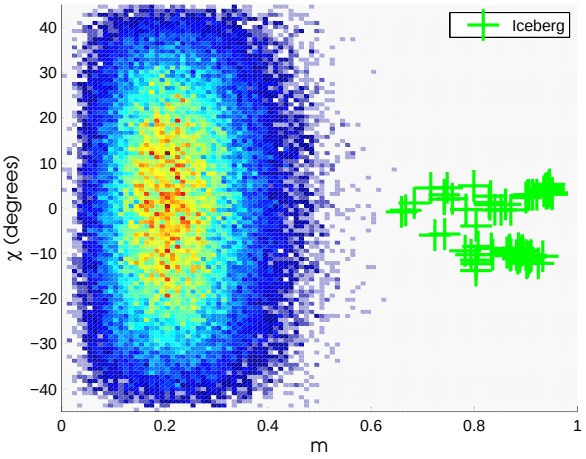

(d) HH-HV

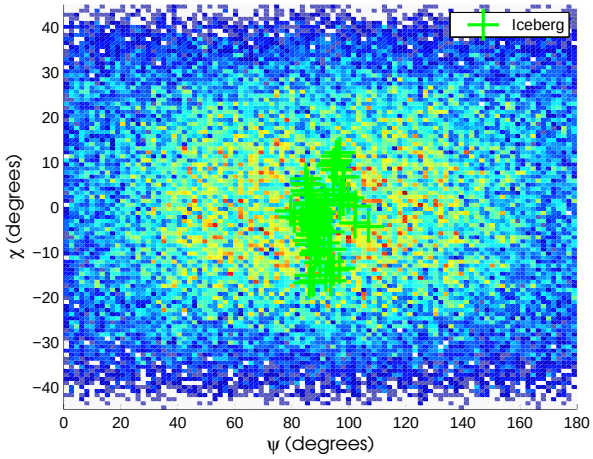

(e) VV-VH

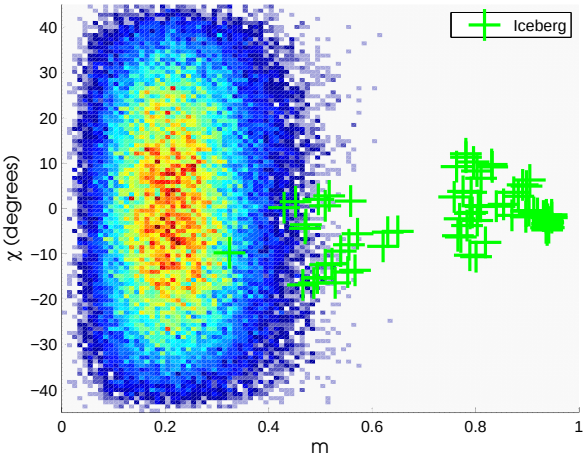

(f) VV-VH

Figure B.17: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0805-2, medium resolution mode.

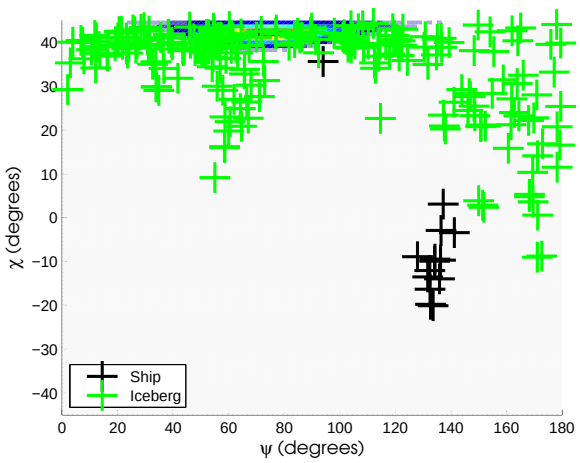

(a) RH-RV

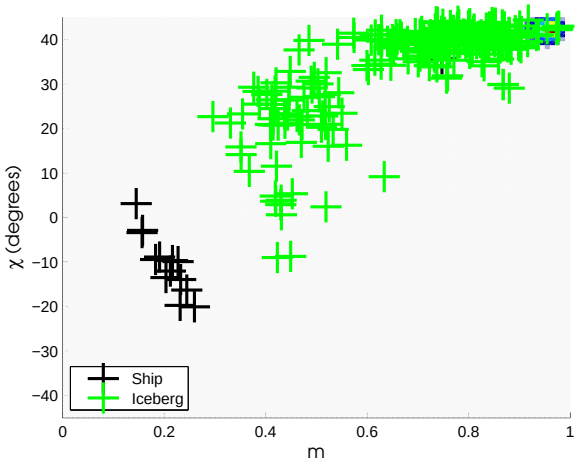

(b) RH-RV

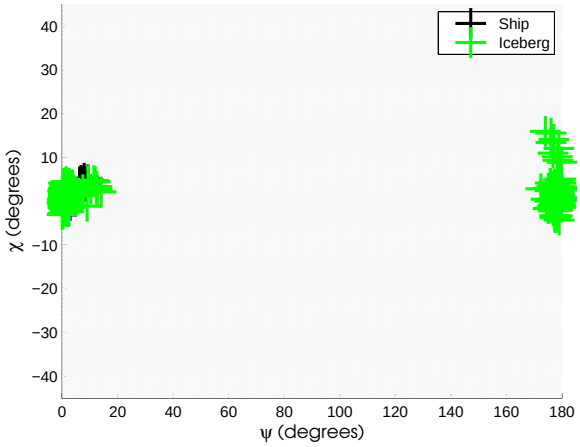

(c) HH-HV

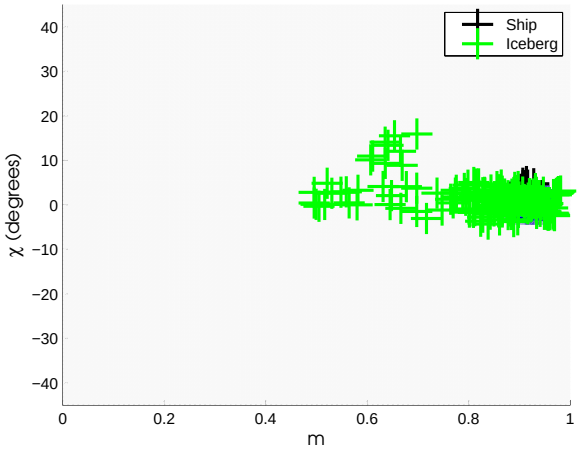

(d) HH-HV

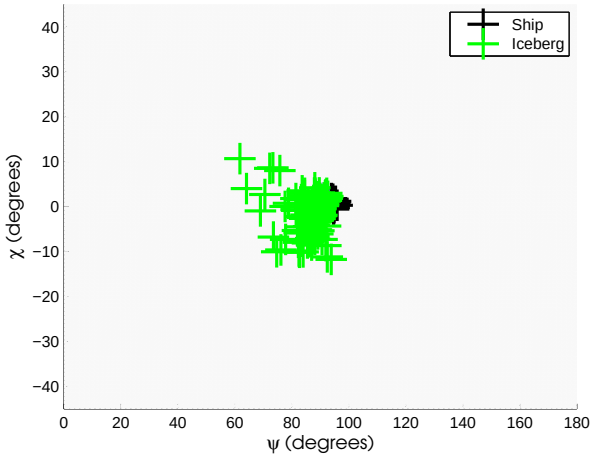

(e) VV-VH

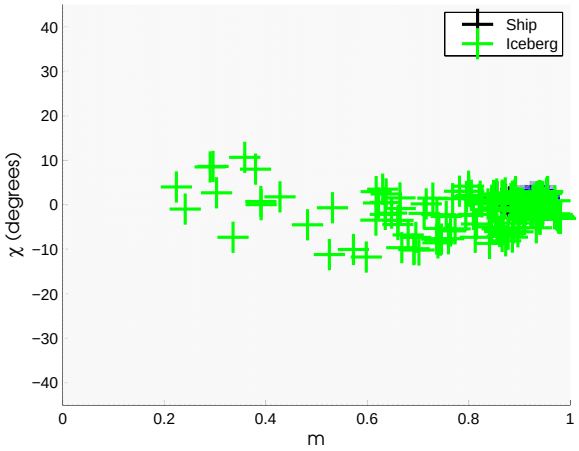

(f) VV-VH

Figure B.18: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0815-2, medium resolution mode.

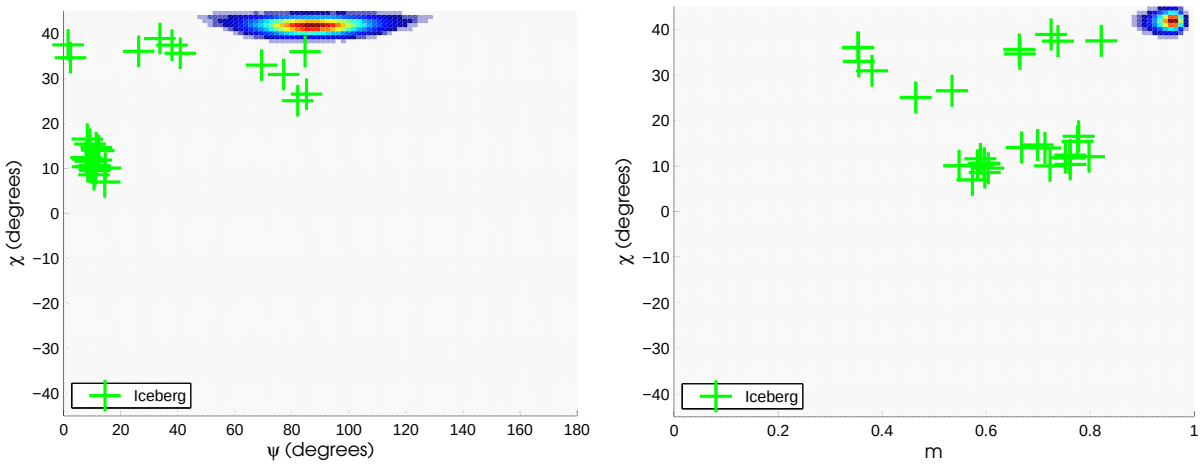

(a) RH-RV

(b) RH-RV

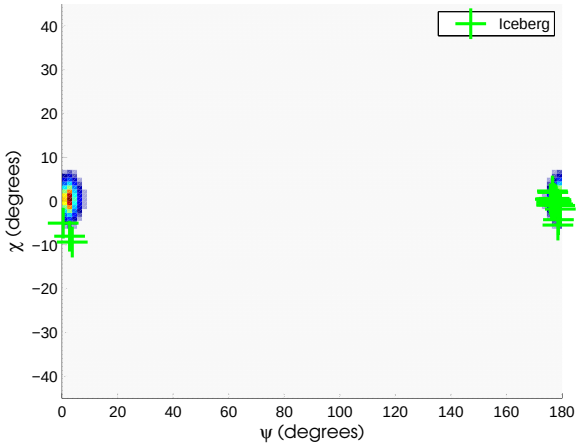

(c) HH-HV

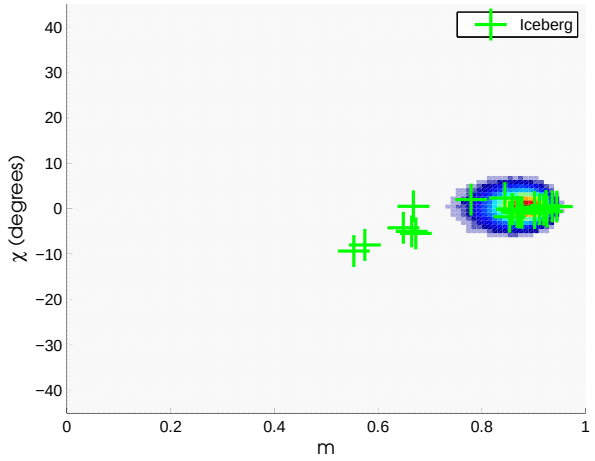

(d) HH-HV

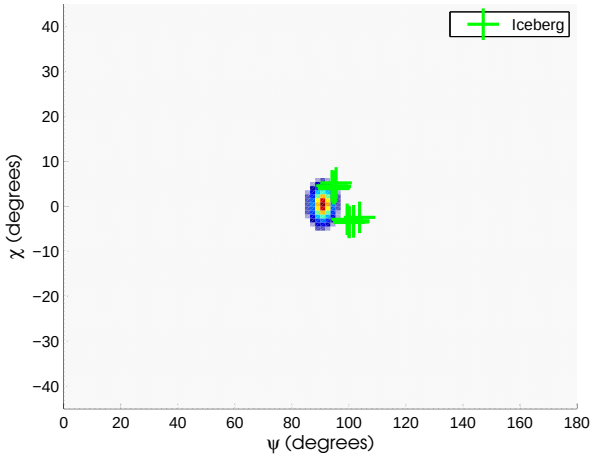

(e) VV-VH

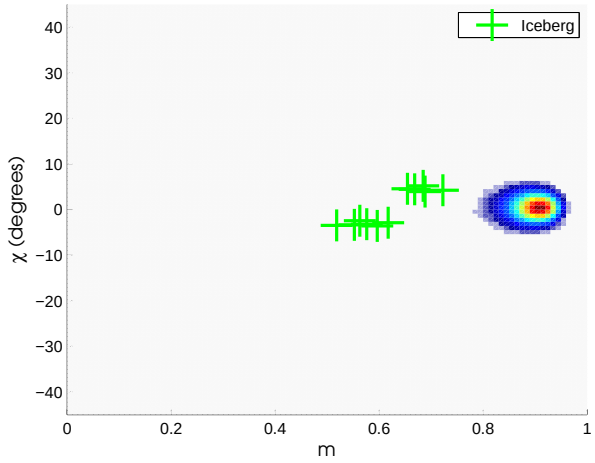

(f) VV-VH

Figure B.19: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0822-1, medium resolution mode.

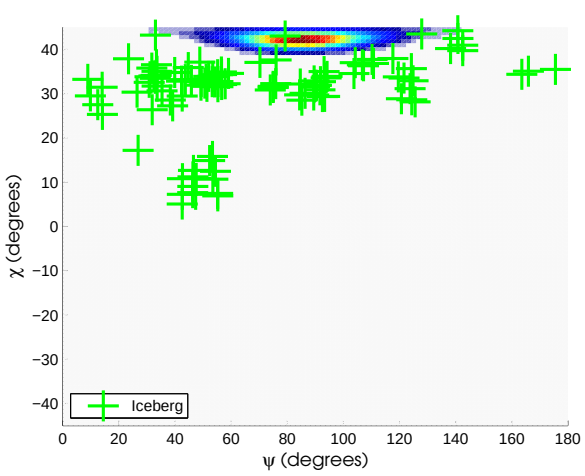

(a) RH-RV

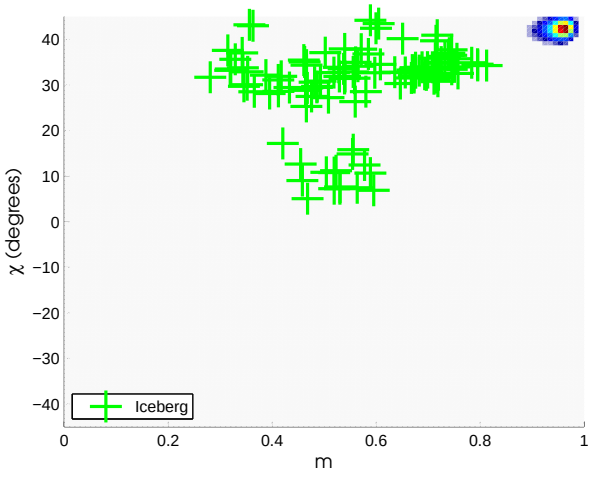

(b) RH-RV

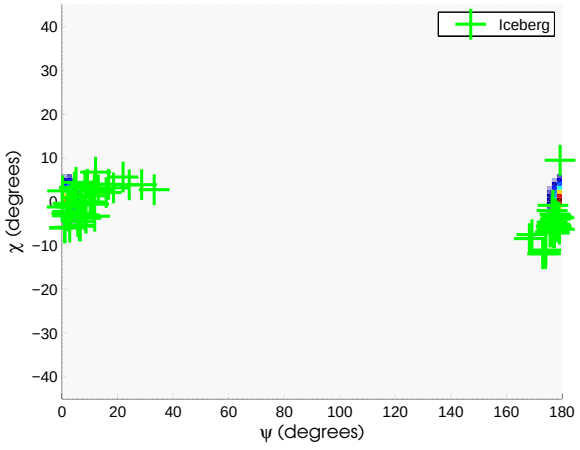

(c) HH-HV

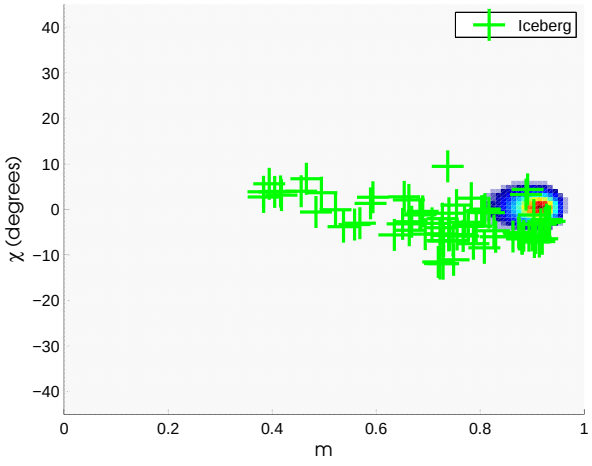

(d) HH-HV

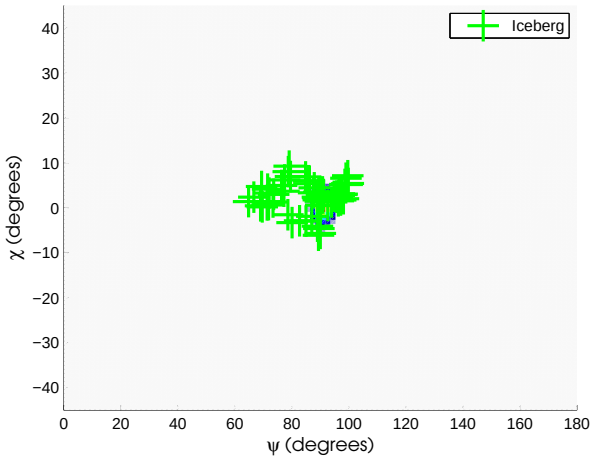

(e) VV-VH

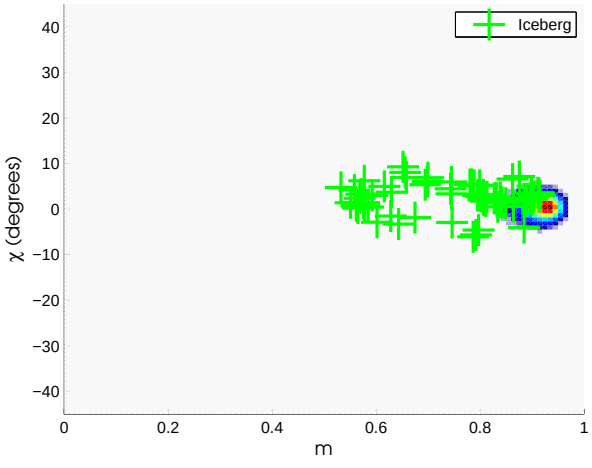

(f) VV-VH

Figure B.20: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0822-2, medium resolution mode.

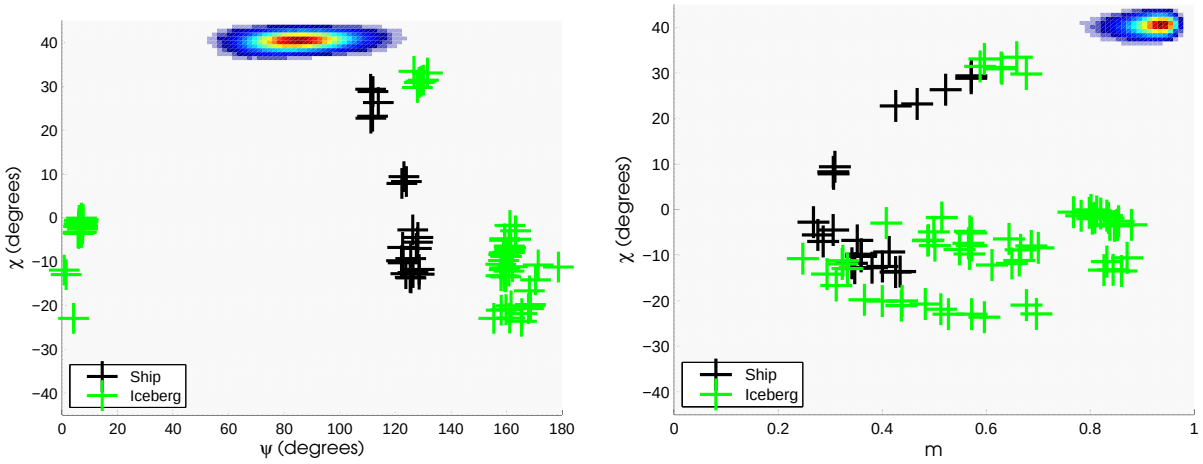

(a) RH-RV

(b) RH-RV

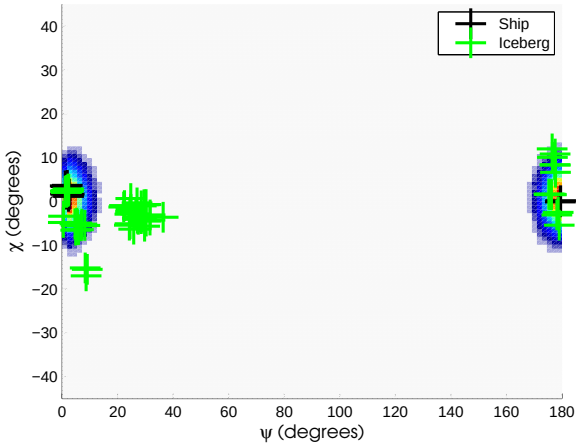

(c) HH-HV

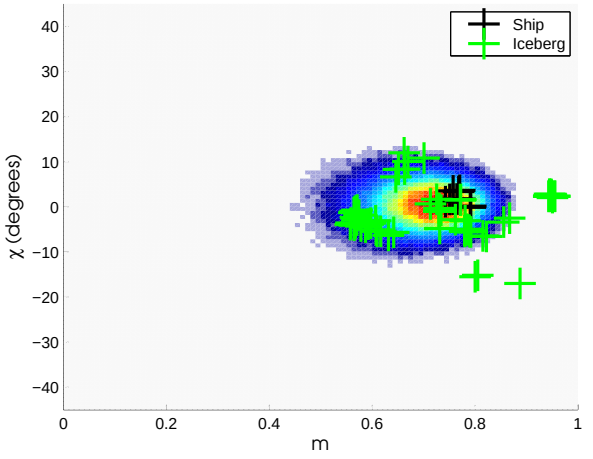

(d) HH-HV

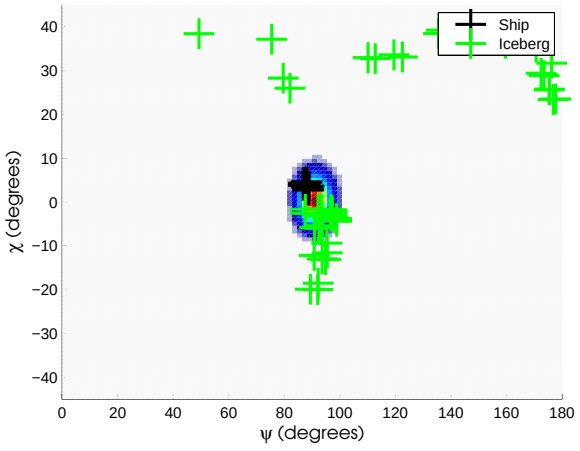

(e) VV-VH

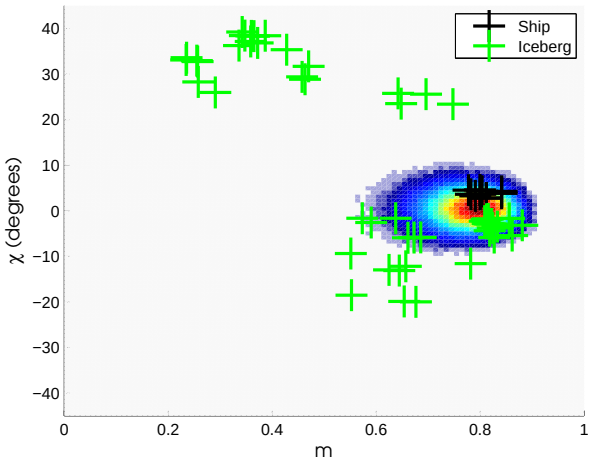

(f) VV-VH

Figure B.21: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0815-1, medium resolution mode.

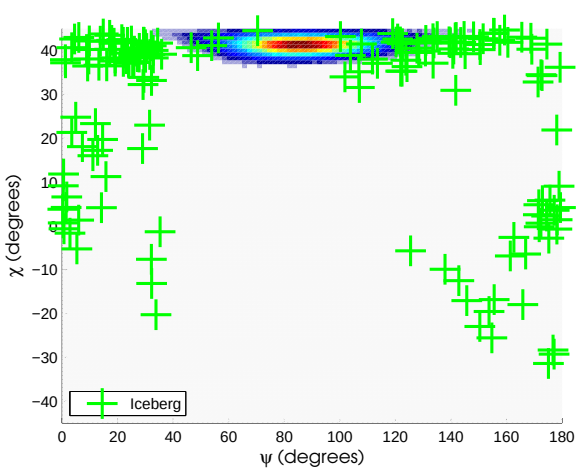

(a) RH-RV

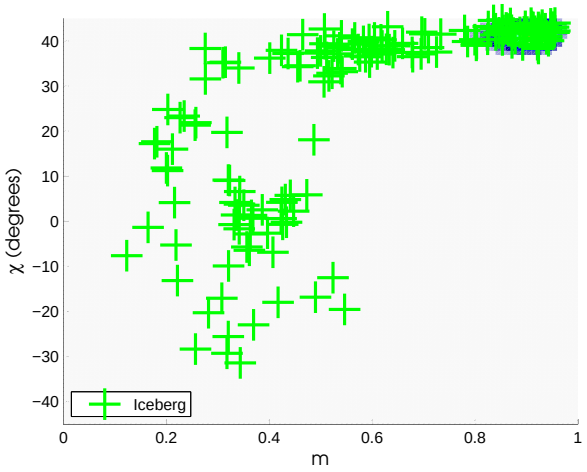

(b) RH-RV

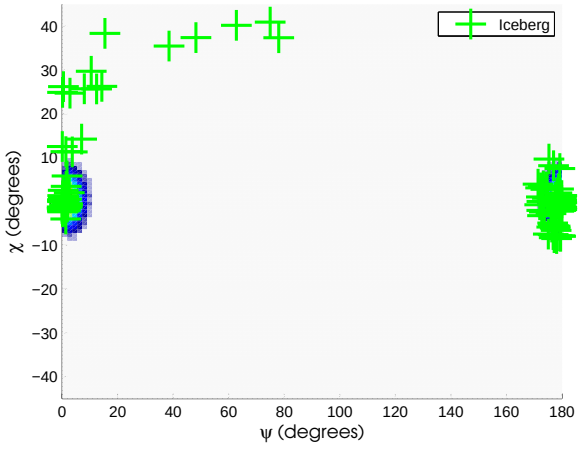

(c) HH-HV

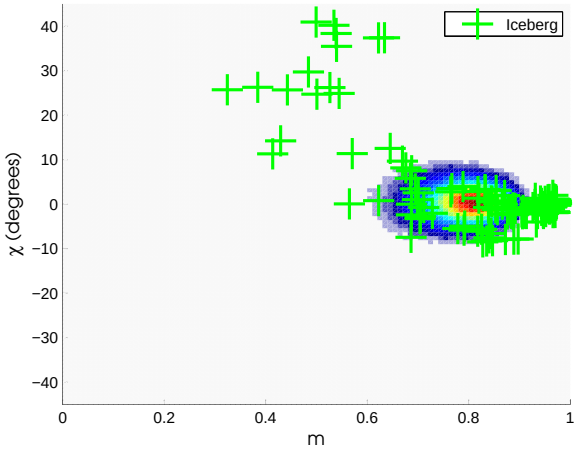

(d) HH-HV

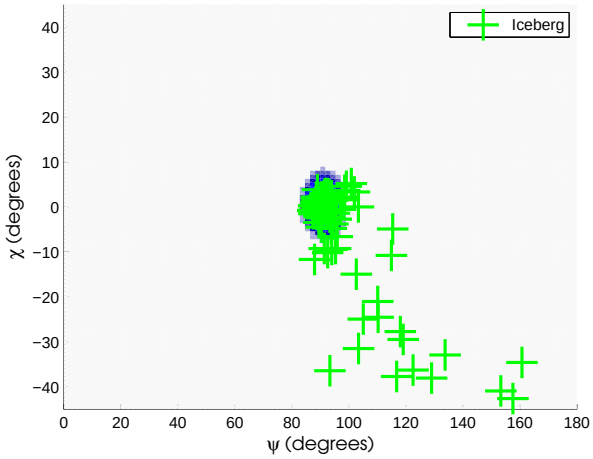

(e) VV-VH

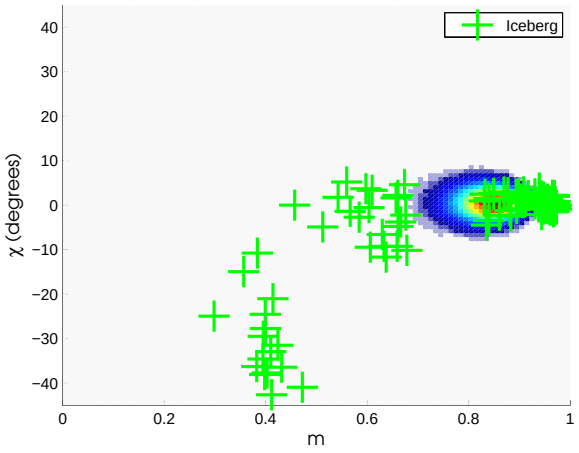

(f) VV-VH

Figure B.22: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0825-2, medium resolution mode.

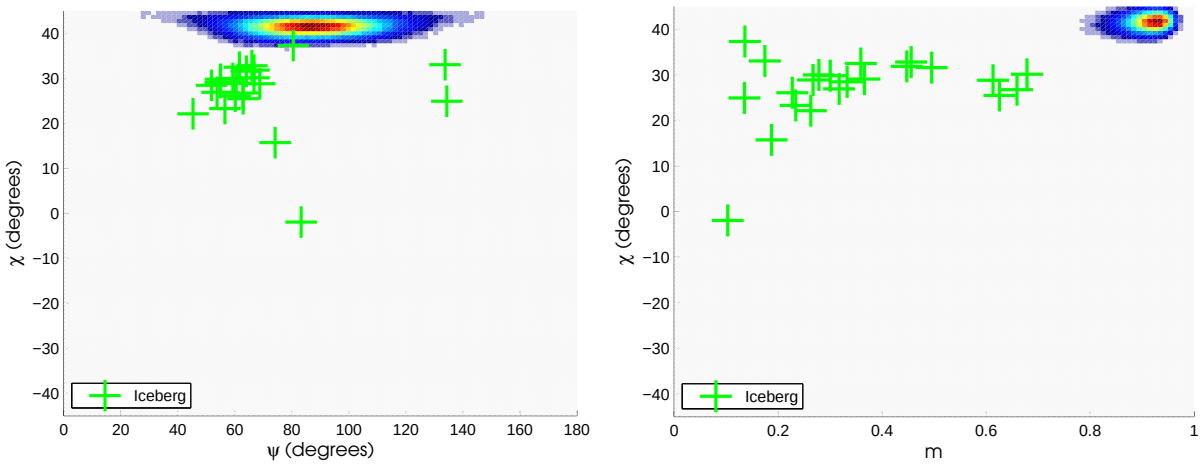

(a) RH-RV

(b) RH-RV

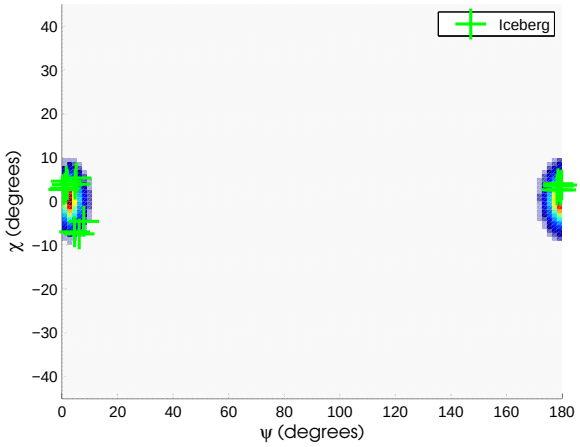

(c) HH-HV

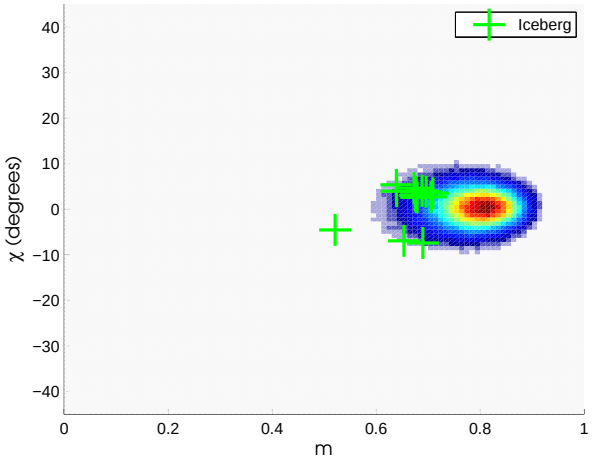

(d) HH-HV

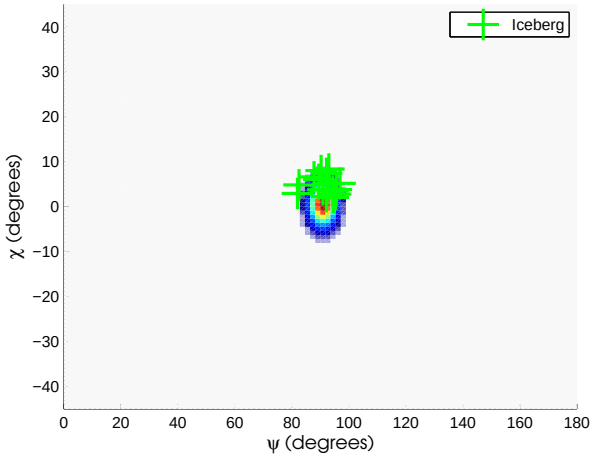

(e) VV-VH

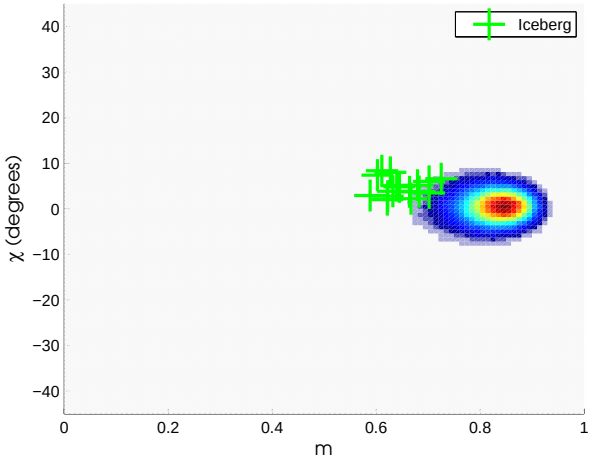

(f) VV-VH

Figure B.23: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0825-1, medium resolution mode.

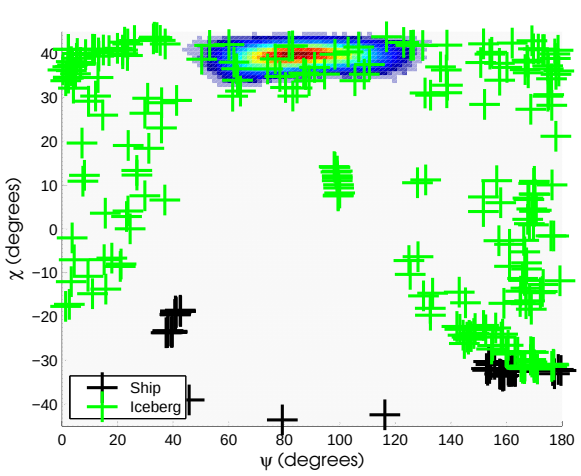

(a) RH-RV

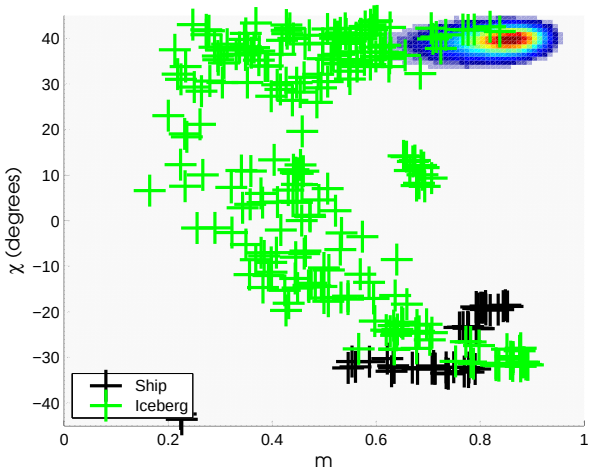

(b) RH-RV

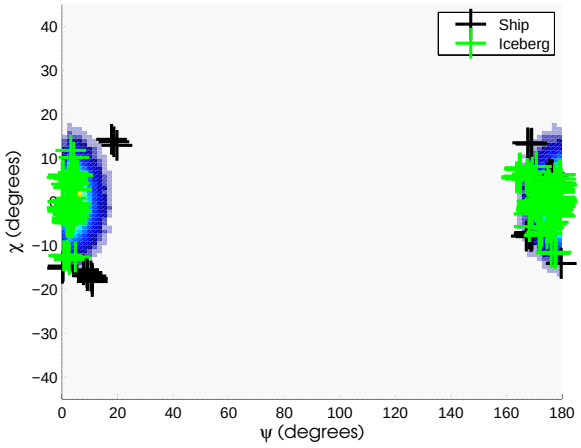

(c) HH-HV

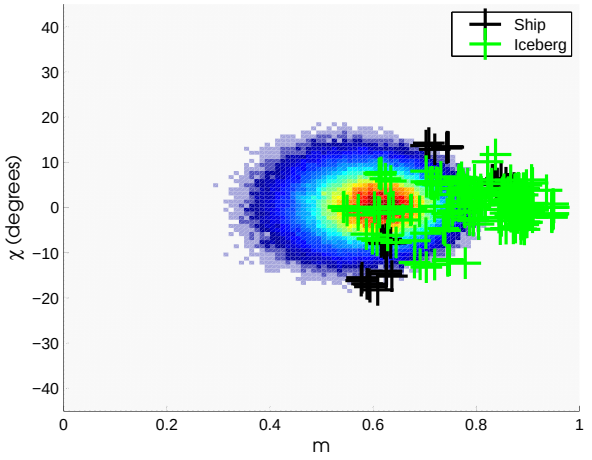

(d) HH-HV

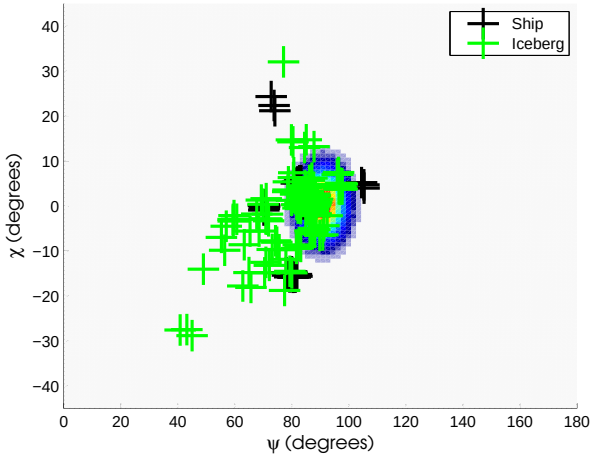

(e) VV-VH

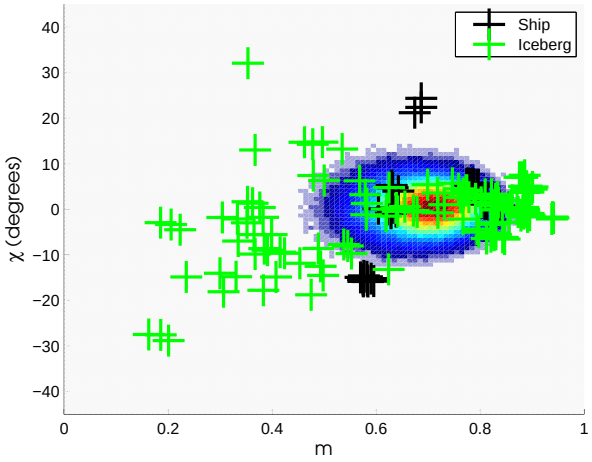

(f) VV-VH

Figure B.24: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0818-1, medium resolution mode.

B.3 Ship Detection Mode

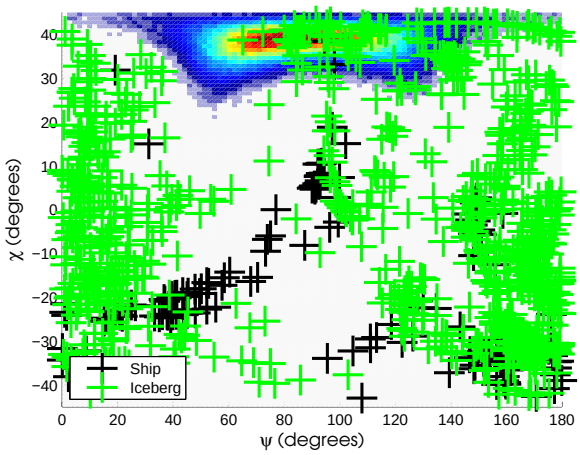

(a) RH-RV

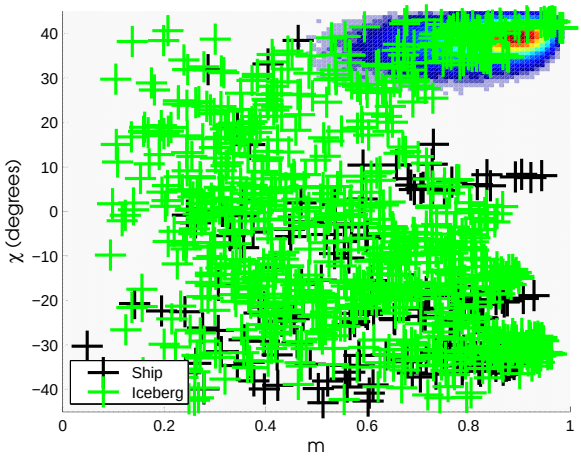

(b) RH-RV

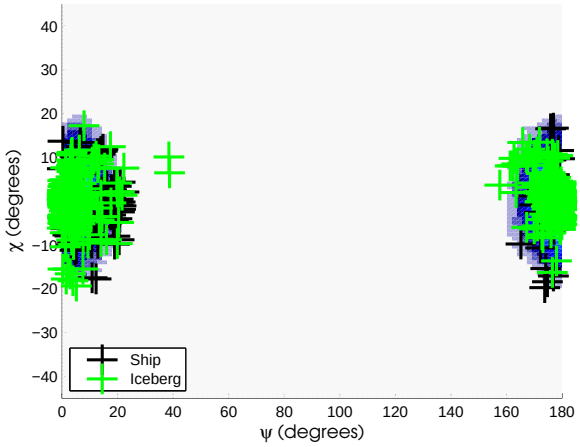

(c) HH-HV

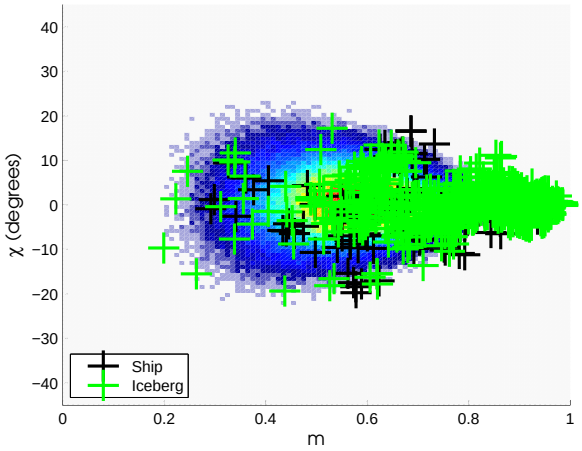

(d) HH-HV

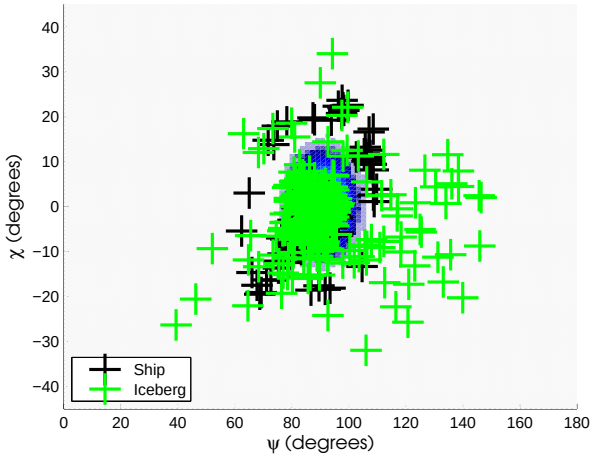

(e) VV-VH

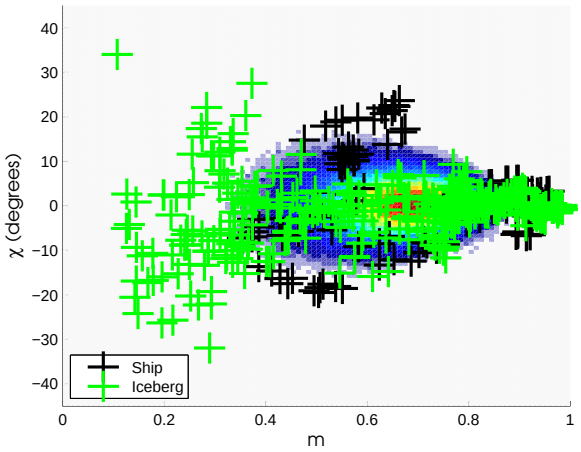

(f) VV-VH

Figure B.25: Plot of χ vs. ψ (top row) and χ vs. m (bottom row) for scene 0818-1, ship detection mode.

Appendix C

Plots of log_{10}(FAR) as a Function of ψ and χ

C.1 Low Resolution Mode

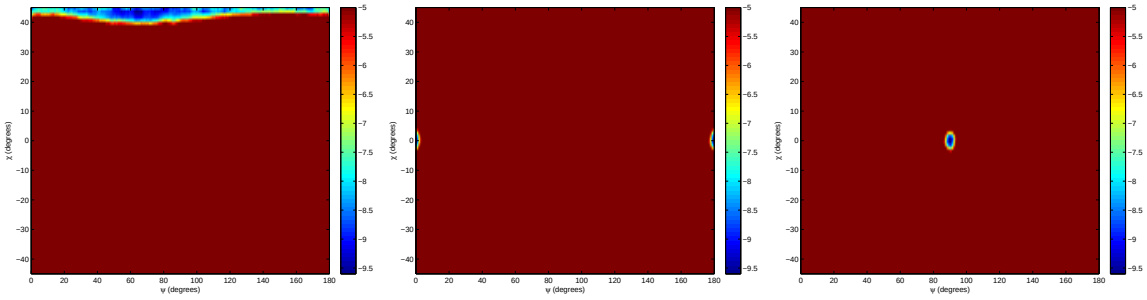

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.1: $\log_{10}$(FAR) as a function of ψ and χ for scene 0905-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

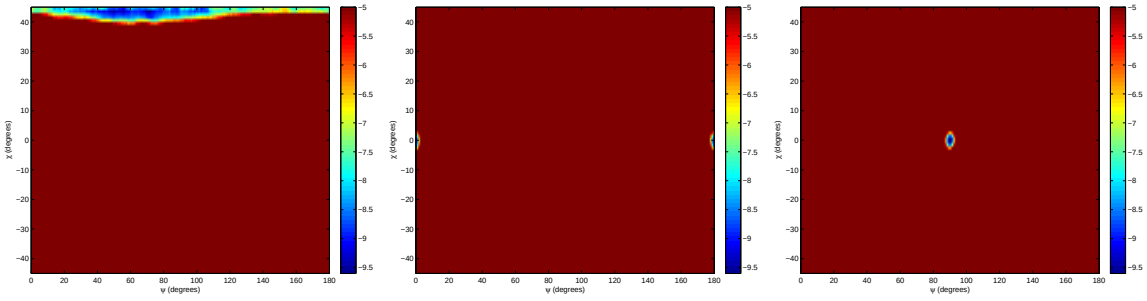

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.2: $\log_{10}$(FAR) as a function of ψ and χ for scene 0905-2, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

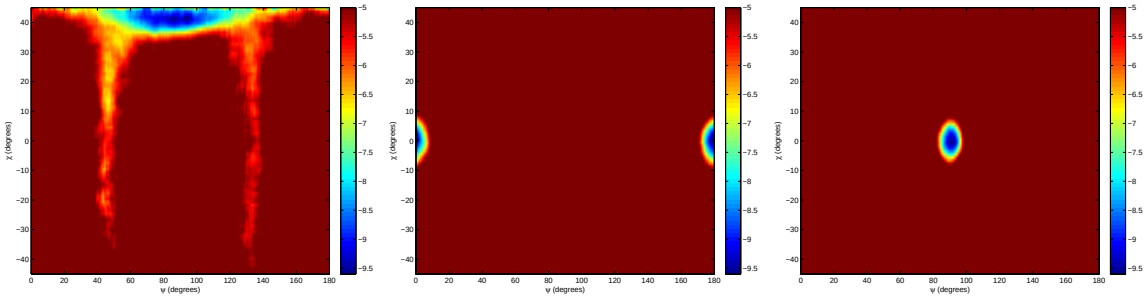

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.3: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0829-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

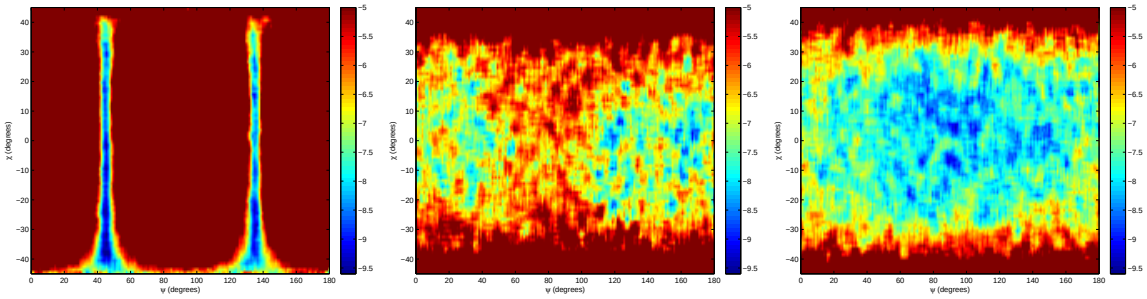

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.4: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0805-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

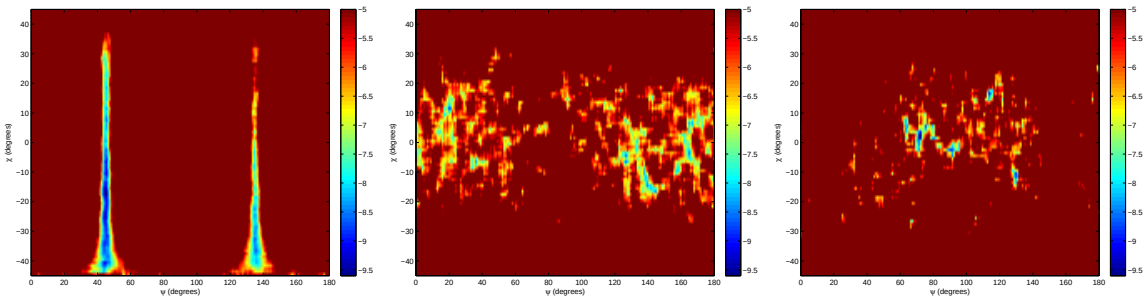

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.5: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0805-2, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

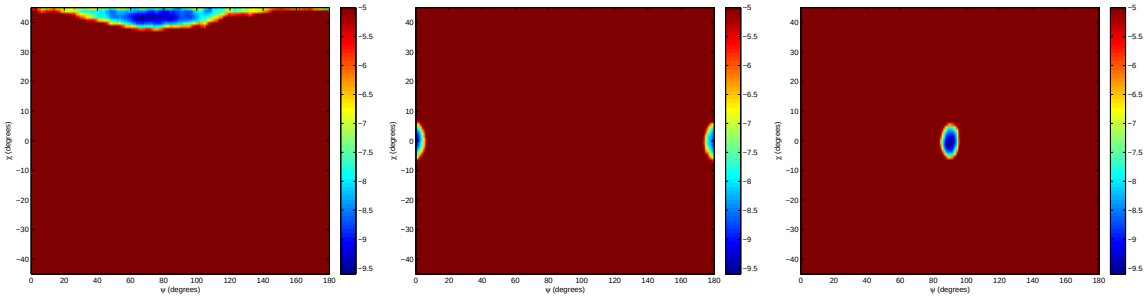

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.6: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0815-2, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

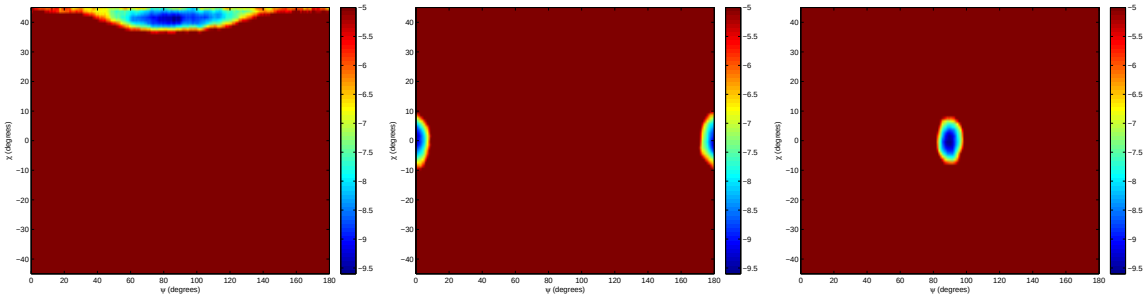

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.7: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0822-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

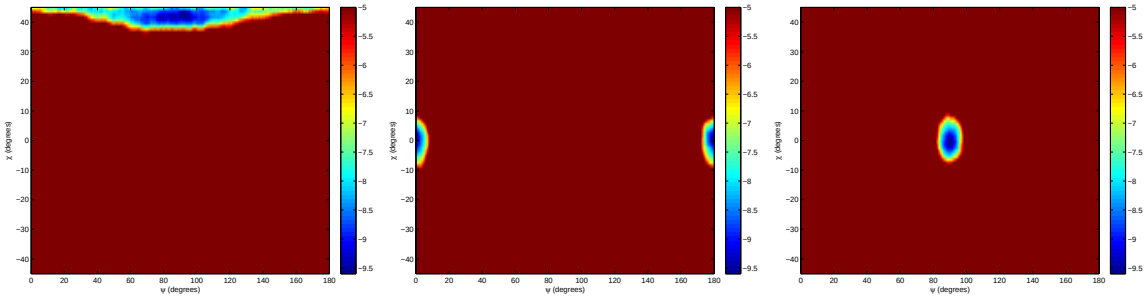

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.8: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0822-2, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

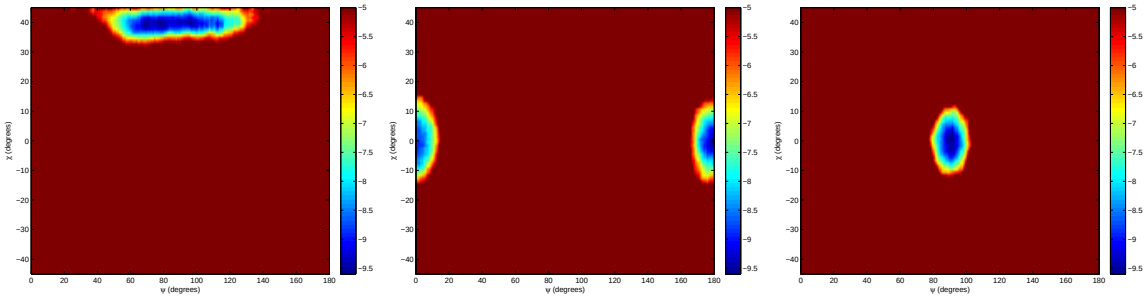

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.9: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0815-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

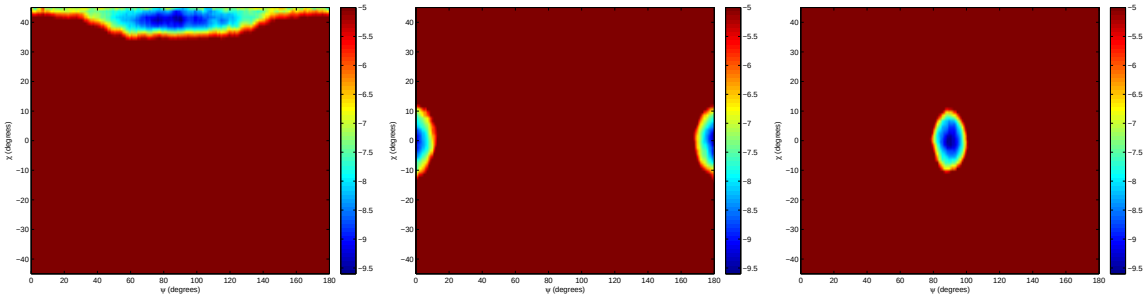

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.10: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0825-2, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

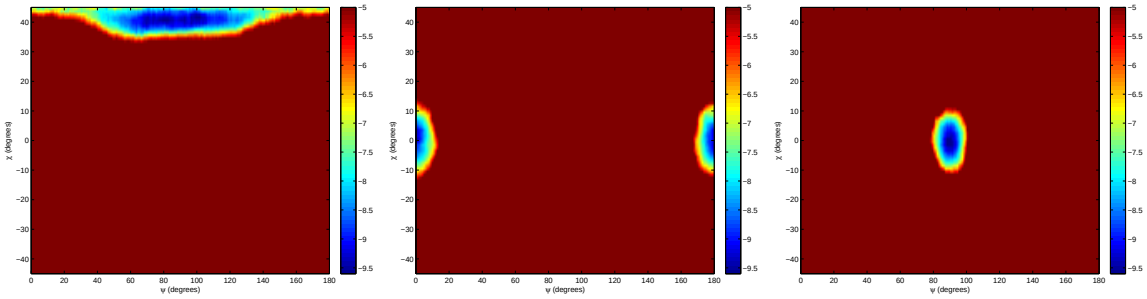

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.11: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0825-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

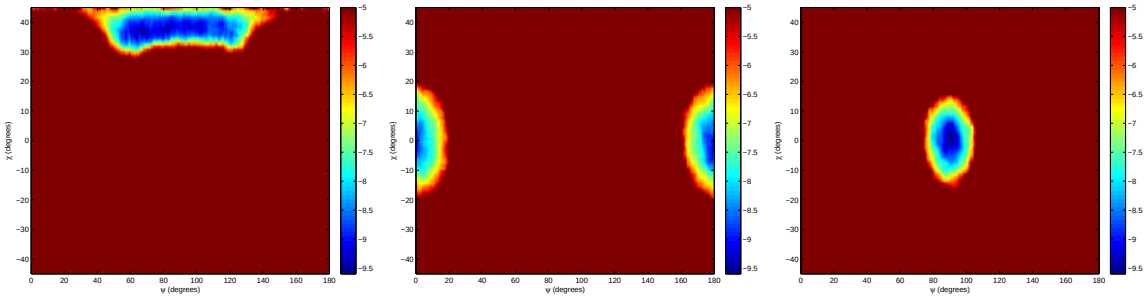

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.12: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0818-1, low resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

C.2 Medium Resolution Mode

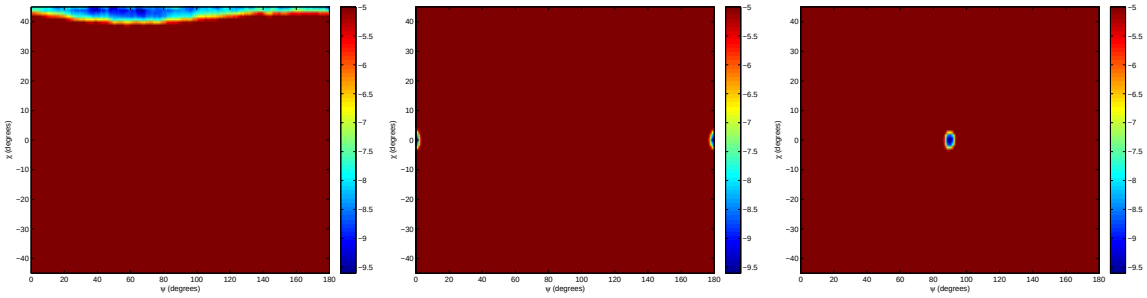

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.13: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0905-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

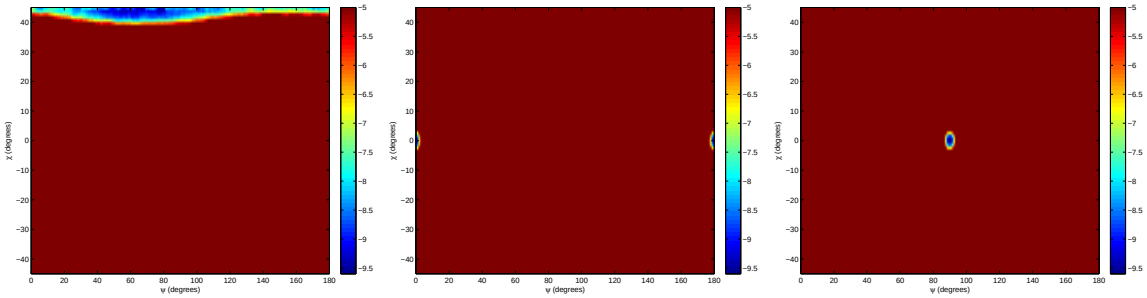

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.14: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0905-2, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

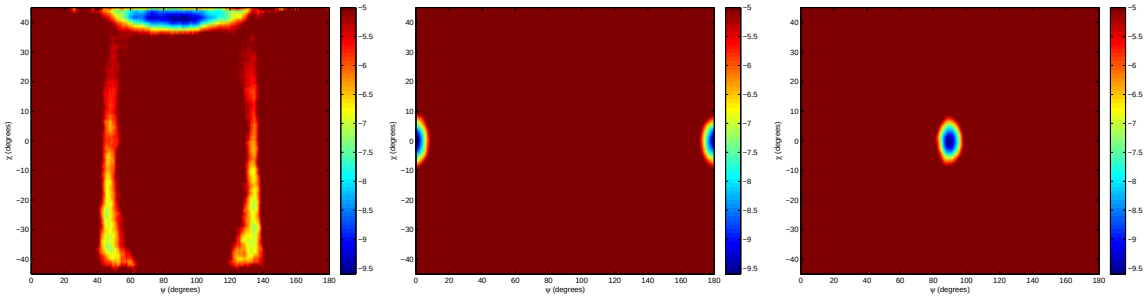

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.15: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0829-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

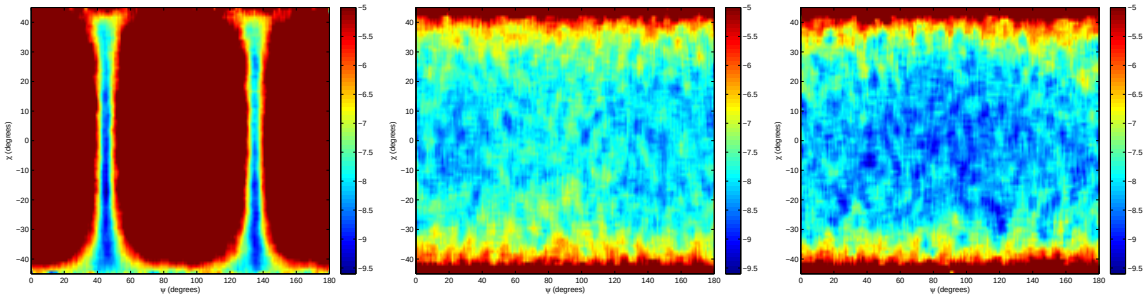

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.16: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0805-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

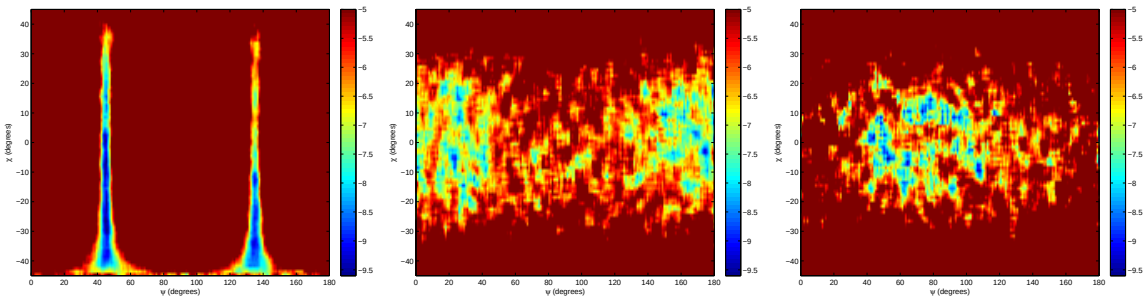

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.17: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0805-2, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

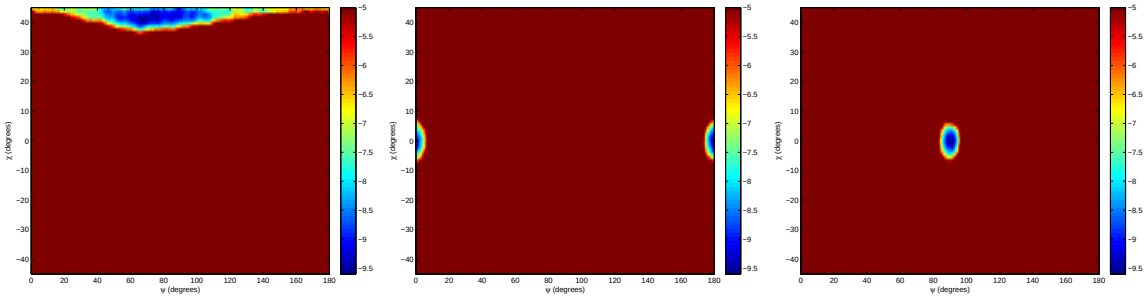

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.18: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0815-2, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

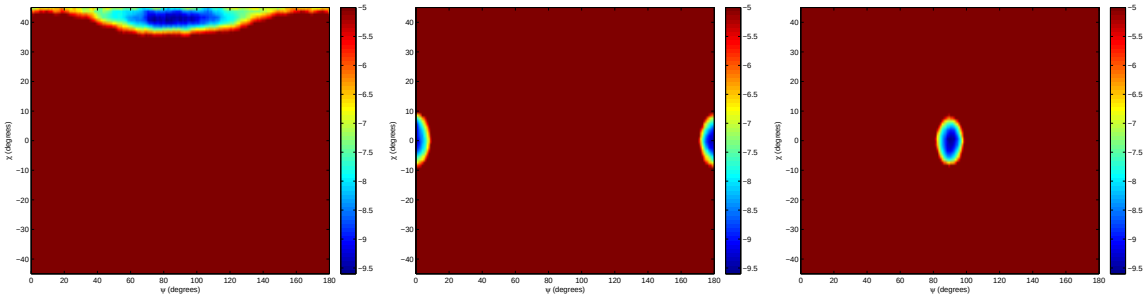

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.19: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0822-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

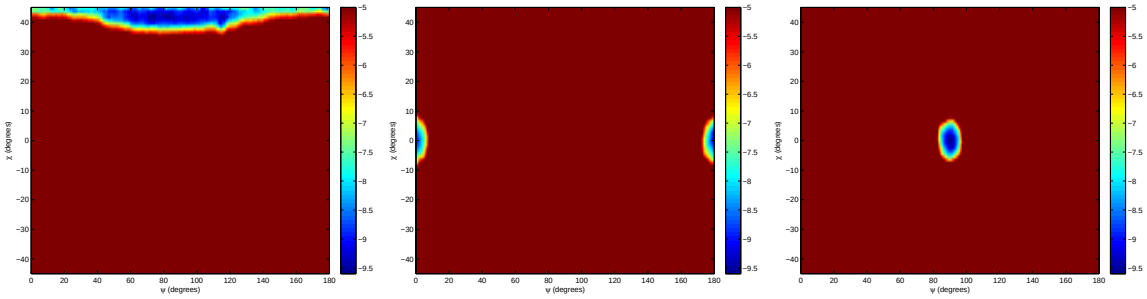

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.20: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0822-2, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

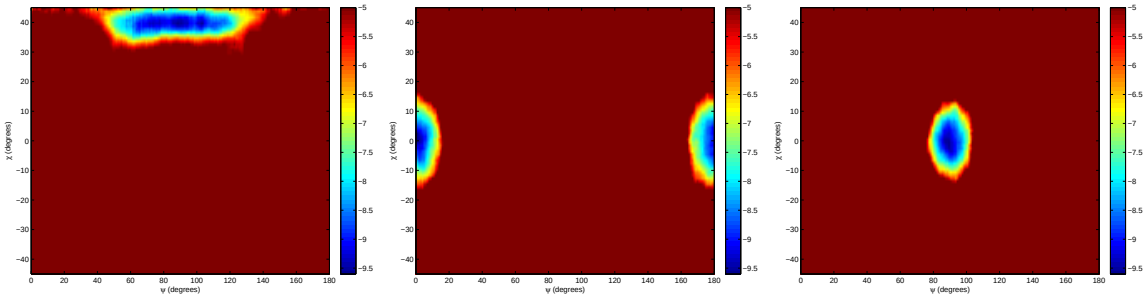

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.21: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0815-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

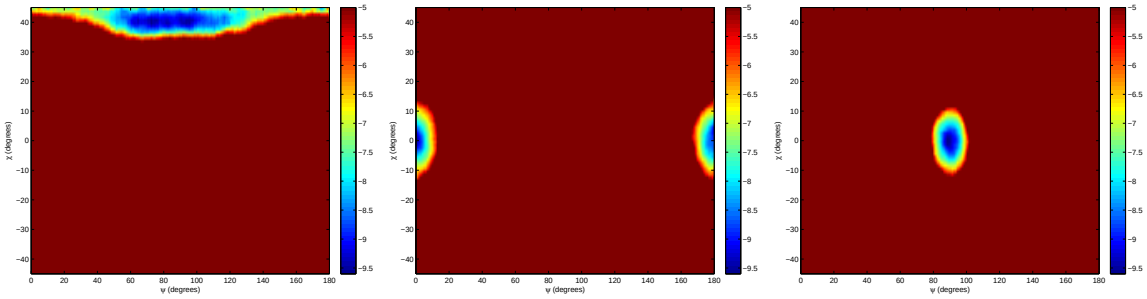

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.22: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0825-2, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

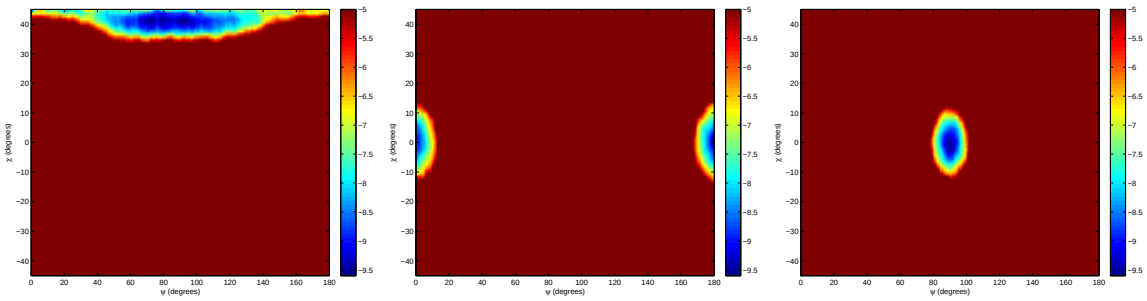

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.23: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0825-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

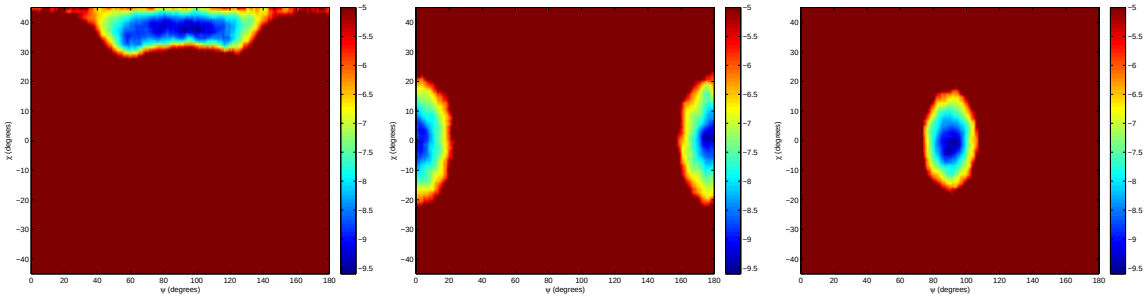

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.24: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0818-1, medium resolution mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9}$.

C.3 Ship Detection Mode

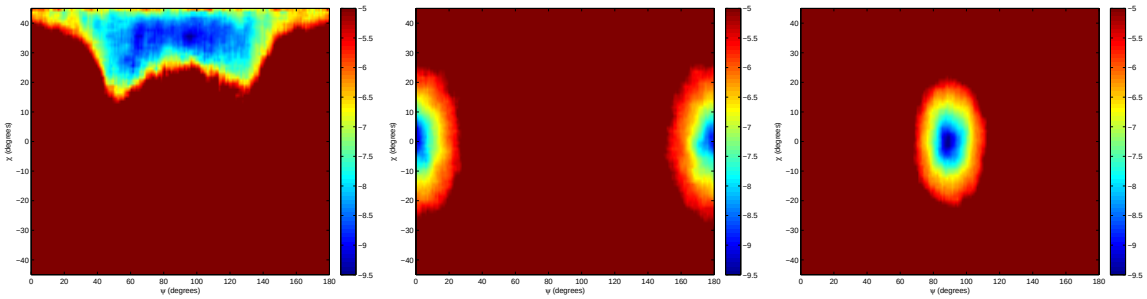

(a) RH-RV (b) HH-HV (c) VV-VH

Figure C.25: $\log_{10}(\text{FAR})$ as a function of ψ and χ for scene 0818-1, ship detection mode, using $\text{FAR}_{high} = 10^{-5}$ and $\text{FAR}_{low} = 10^{-9.5}$.