ICESat full waveform signal analysis
for the classification of land cover types
over the cryosphere

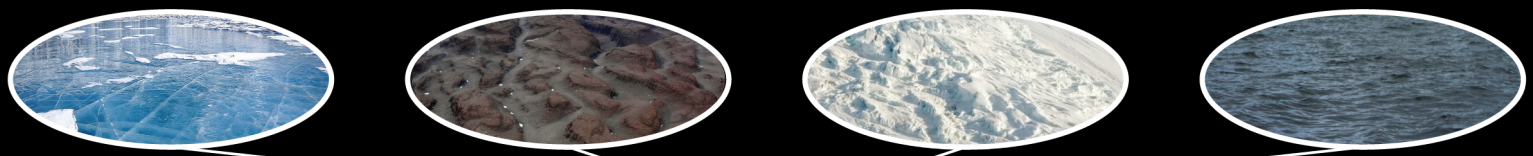

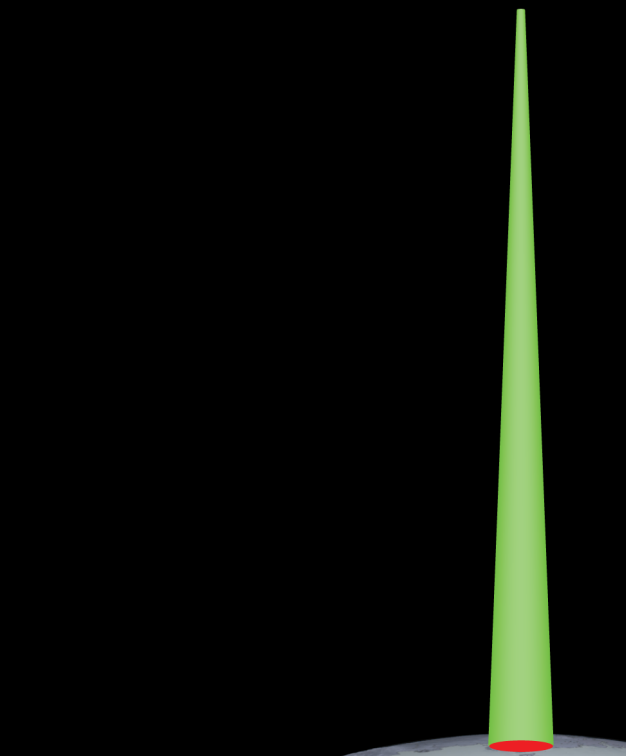

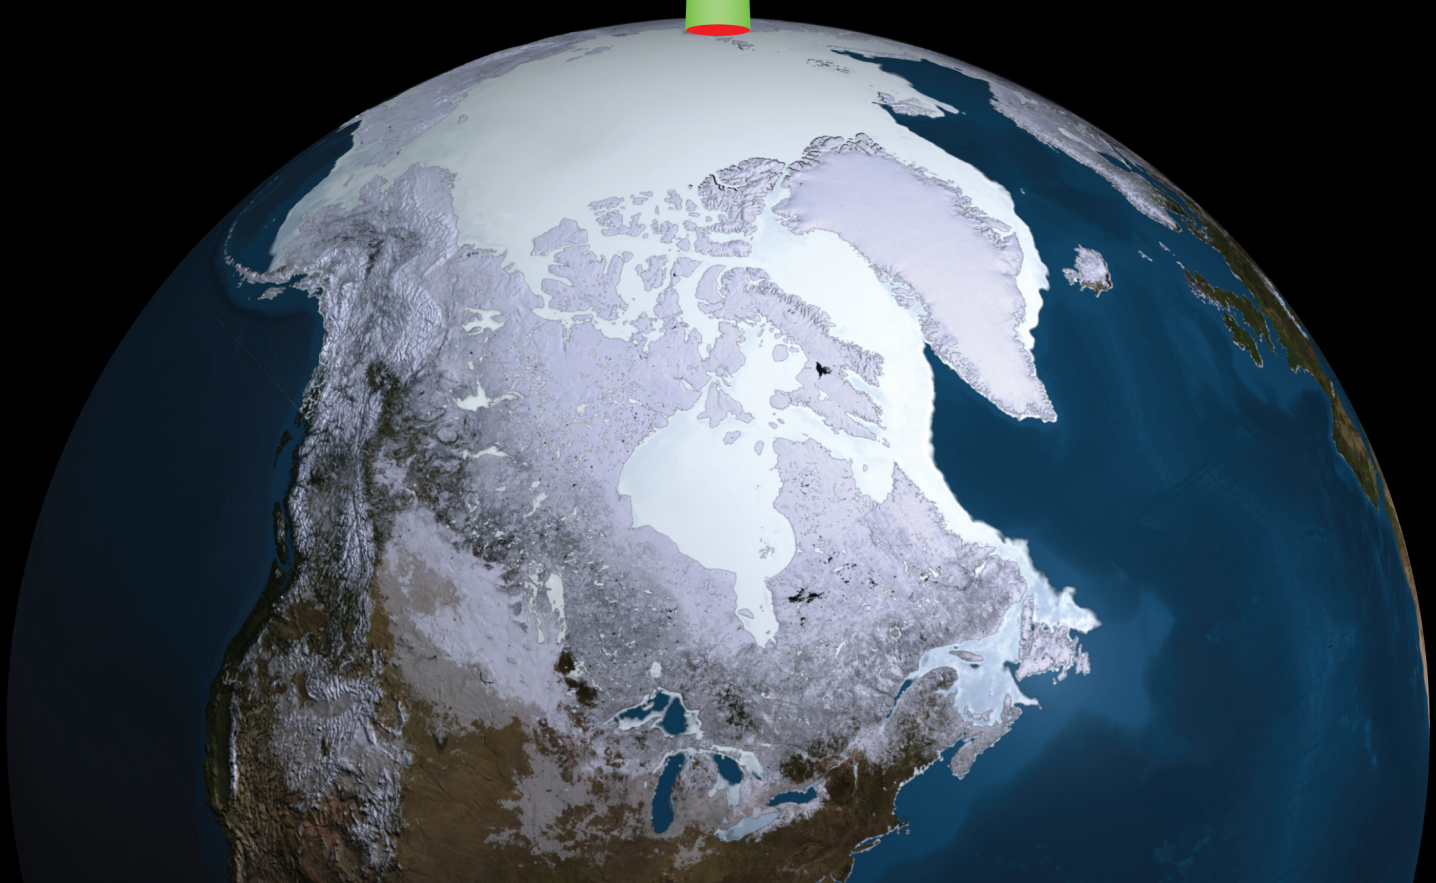

Cover: Illustration of a laser beam illuminating a polar region and the classification of ice, rock, snow or water. Upper images taken and adapted from National Science Foundation, from left to right: courtesy Joe Mastroianni; courtesy Peter Rejcek, National Science Foundation; courtesy Chris Kannen; courtesy Zee Evans. Lower image courtesy NASA Goddard's Scientific Visualization Studio.

MSc Geomatics Graduation Thesis

ICESat full waveform signal analysis for the classification of land cover types over the cryosphere

Ramses Amadeus Molijn, 1097199
June 2009
Optical and Laser Remote Sensing & Physical and Space Geodesy
Department of Earth Observation and Space Systems
Faculty of Aerospace Engineering
Delft University of Technology (DUT)
The Netherlands
Center for Space Research (CSR)
The University of Texas at Austin
United States of America

Graduation Professor (DUT): Prof. dr. ir. R. Klees
Supervisors (DUT): Dr. B.C. Gunter & Dr. R.C. Lindenbergh
Internship supervisor (CSR): Dr. T.J. Urban
Co-reader (DUT): V.H. Duong
Contact email address: ramolyn@gmail.com

Abstract

The Earth needs attention. That is the bottom line of today's discussion on the impact of humans on the environment and their contribution to global warming. With the rapid sea level rise and the melting ice caps, it is becoming increasingly important to understand these processes. Mass balance estimates of the cryosphere (e.g. Greenland and Antarctica) are one of the keys to unlocking the secrets of today's concerns and predicting tomorrow's problems. The primary objective of this research is to construct an automatic classification algorithm to that will distinguish between the polar land cover types: ice, rock, snow and water. The algorithm is based on full waveform laser altimetry measurements gathered by the Ice, Cloud and Elevation Satellite (ICESat). The results of this research have the potential to play an important role in improving the current methods for computing mass balance estimates and monitoring seasonal changes in land cover type of the cryosphere.

Mass balance estimates are an important technique for monitoring the total water content of the polar regions. Analyzing temporal differences in ice mass change offers an insight into the melting of the ice sheets and its effects on global water cycle, such as ocean currents. The more accurately mass balances can be estimated, the better the impacts on the Earth's processes can be determined.

One of the techniques to compute mass balance estimates involves the use of surface altimetry data. One of the main Earth observing satellites for surface altimetry is ICESat, launched in 2003 by the National Aeronautics and Space Administration (NASA). This satellite measures the Earth's surface elevations with a technique called laser ranging: 40 times per second, the satellite sends out a laser pulse to the Earth's surface and records how long it takes until each reflected pulse returns. The time duration multiplied by the speed of light yields the range between satellite and the ground surface. With the location of the satellite at the time of measurement, the Earth's surface can be modeled.

This method holds true to its simplicity only if the transmitted laser pulses would be infinitely short and would only encounter a flat ground surface. However, the received pulses are broadened (in terms of energy over time) by the finite transmitted pulse length and, more importantly, by all objects along the laser's path line-of-sight. The waveforms that result from these influences form the most important piece of information for the

classification algorithm of this research study.

The research basically analyzes the effects of ice, rock, snow and water on the returned waveforms. To verify the land cover types, validation data sets have been used from various Earth observing satellite missions, airborne altimetry missions, and terrestrial weather stations. Aided by these validation data sets, a classification algorithm was constructed based on ICESat's waveforms. The most important waveform attributes for classification are the waveform returned energy, width, kurtosis and saturation energy.

The returned energy, from which the reflectivity of the laser pulse can be computed, is the most important attribute to classify snow, since snow generally gives higher reflectivity than other land cover types. The waveform width is used to classify rock, since rock generally makes the waveform wider due to elevation differences within the laser's surface footprint. In addition, small width values are used for classifying ice. The waveform kurtosis is a measure for the peakedness of the waveform and indicates water, since flat water surfaces makes the waveform sharply peaked rather than bluntly peaked. Finally, simply put, the waveform saturation energy is a measure for the unexpectedness of returned energy with respect to the returned energy of the previous measurements. For the classification algorithm, this attribute is used to indicate a change in land cover type.

The classification algorithm has been tested against validation data to analyze its accuracy. This resulted in the most important conclusion of the research, namely that classification based on ICESat's full waveform signals and attributes works well, with a classification accuracy of 74%. The algorithm has been applied to other test areas as well, showing seasonal changes in land cover types.

Two major benefits of the classification algorithm can be identified. First, rough classification maps of the polar regions per season improve the mass balance estimates by indicating which areas to exclude. Since spatial interpolation between ICESat's measurements is required for computing the mass accumulation or loss, an indication of rock surfaces can help to make the estimate more accurate. In addition, ice, snow and water have different mass densities, which means that one meter accumulation in height does not result in equal mass gain for each land cover type. The second major benefit is the potential monitoring of changing land cover types over time. For example, glacier retreat and shrinking snow extent over time could be observed. The implementation of the classification algorithm gets us closer to answering the question: what is the impact of climate change on the cryosphere?

Preface

Climate change is one of the most challenging issues that we currently face and will face in the future if we do not take significant action to reduce our impact on the environment. As the Climate Change 2007 Synthesis Report of the Intergovernmental Panel on Climate Change (IPCC) points out: "Partial loss of ice sheets on polar land could imply meters of sea level rise, major changes in coastlines and inundation of low-lying areas, with greatest effects in river deltas and low-lying islands." That statement concerns only the possible sea level rise due to melting of the ice caps, but does not even indicate the direct wide-spread implications to mankind and our fellow inhabitants. It may sound popular, but right now could very well be the defining moment to get a grip on the situation and significantly reduce or even avert potential disastrous effects on today's society.

Contributing to finding solutions to today's problems was one of the main reasons for my interest in this research study. Although I was hoping to see already the direct effects of this study on a better understanding of the Earth's problems, I am optimistic about the implications it might have. If the results of this study will be indeed used by the scientific community, I will be satisfied knowing that I was able to have a (small) contribution to the research on today's problems.

This research study is the final project of the MSc. Geomatics of Delft University of Technology (DUT) in The Netherlands. The research has been conducted partly at DUT and partly at the Center for Space Research (CSR) at the University of Texas at Austin. In the first three months, a literature study and some basic processing of ICESat data was conducted at DUT. In addition, preparations were made for the four months internship at the CSR. In Austin, the study focused on data filtering and correction techniques, the waveform interpretation and the integration of validation data with ICESat data. In addition to the benefits of consultation with CSR colleagues for this research study, I intended to provide the CSR with helpful feedback on the data and assistance where ever I could. At the end of the internship, I gave a seminar to the all colleagues of the CSR that were interested in my research.

After the internship, the research focused more on the surface topography analysis and the classification algorithm. Along the way, more and more analysis was put into finding

more optimal attributes for the classification algorithm, such as the atmospheric signal, the gain and the saturation energy correction. Finally, after fifteen months, the work is finished and a classification algorithm can be presented.

Looking back at the research, several difficulties and lessons learned can be summarized. First of all, the interpretation of all of the flags, and the understanding of the correction and filtering techniques was surprisingly difficult. Since ICESat suffered from some setbacks during its lifetime, it was difficult to determine which flags and corrections were still valid and how to apply these properly. Fortunately, the colleagues at the CSR were able to assist me in this process. Another difficulty was the somewhat ambiguous interpretation of the validation data. Since the primary validation data set was acquired before ICESat was operational, special care had to be taken in comparing this data to ICESat data. Finally, a review of the gain and saturation of the signals after thirteen months delayed the research study once more, because some unexpected signal characteristics were discovered that had to be taken into account for proper classification. The most important lesson learned regarding how to tackle the research question is that special attention should be paid to the relevancy of subprojects with respect to the end goal of the research.

In addition, several experiences and gained skills can be identified as well. First of all, to become acquainted with the academic scene of research was helpful and I got a good impression of the working environment. Still to the last day, I am amazed by the researchers' dedication to the matter of research that is so important for current discussions. On a personal level, I became more skilled in a technical and social way. First of all, my abilities of scripting in MATLAB improved significantly over the research study. In addition, I learned new computing languages such as Fortran and IDL. Even though programming is not my hobby in particular, I learned to appreciate that it forms a powerful tool to get sensible and interpretable results from abstract data. Without a doubt, the freedom in programming results in endless and unlimited possibilities for the sake of research. However, interpretation of computer-made results remains dependent on human skills and their personal judgment. Finally, I realized the importance of clear and straightforwards communication towards colleagues, not only face to face but via the electronic highway as well. Especially, the communication with DUT, the CSR, the Goddard Space flight Center (GSFC), and the National Snow and Ice Data Center (NSIDC) was essential for my research and I hope I was able to provide these parties with sensible feedback as well.

Besides these institutes and centers as a whole, I would like to thank several individuals as well. First of all, I would like to express my gratitude towards my girlfriend Sanne. I am a lucky man with the support she gave me by accompanying me to Austin, by giving me the freedom to work hard on this research study and by reading my lengthy chapters. Of extremely high importance was the support of my supervisors as well: Brian Gunter and Roderik Lindenbergh at DUT and Tim Urban at the CSR. The weekly and two-weekly meetings in Delft with Brian and Roderik gave me the extra drive to get the most out of this research study; their support via email during my time in Austin is highly appreciated as well. In Austin, I received great support from Tim on direct ICESat related issues.

Especially the explanations on the data products and the interpretation of their attributes were essential to this research. I hope I compensated the CSR's time and effort by giving constructive feedback on ICESat's data. Next to the direct players in the research study, I also would like to thank the CSR employees Roberto Gutierrez for the information on the White Sands data, Amy Neuenschwander for the waveform interpretation, Robert Harpold for the help with Fortran, Charles Webb for the data product processing, and Dr. Schutz for giving me the opportunity to conduct research at the CSR. In addition, I would like to thank Boris Jutzi (Research Institute for Optronics and Pattern Recognition) for sharing his experience with the deconvolution technique, Stephen Palm (GSFC) for helping me with extracting ICESat's atmospheric signals, Hieu Duong (DUT) for sharing his own findings with ICESat data, and David Korn (NSIDC) for the accurate ICESat data delivery. Furthermore, I would like to thank US Geological Survey (USGS), Portland State University, and Ohio State University for validation data sets and their answers to my data related questions. Finally, I would like to thank my graduation professor Prof. dr. ir. Klees for his interest during my research, and last but not least my parents for their patience and ever lasting support.

Contents

List of Figures

List of Tables

List of Abbreviations

ATBD	Algorithm Theoretical Basis Document
ATM	Airborne Topographic Mapper
CSR	Center for Space Research
DEM	digital elevation model
DUT	Delft University of Technology
EOS	Earth Observation System
EROS	Earth Resources Observation and Science
fft	fast-Fourier transform
FWHM	full width at half maximum
GLAS	Geoscience Laser Altimeter System
GSFC	Goddard Space Flight Center
ICESat	Ice, Cloud and Elevation Satellite
ifft	inverse fast-Fourier transform
ITRF	International Terrestial Reference Frame
LIMA	LANDSAT Image Mosaic of Antarctica
LPA	laser profiling array
LTER	Long Term Ecological Research
MOLA	Mars Orbiter Laser Altimeter
NASA	National Aeronautics and Space Administration
NSIDC	National Snow and Ice Data Center
rms	root mean square
ROI	region of interest
SDS	science data system
SI	International System of Units
SLC	Scan Line Corrector
SNR	signal-to-noise ratio
SRTM	Shuttle Radar Topography Mission
std	standard deviation
USGS	US Geological Survey

Chapter 1

Introduction

This research study focuses on analyzing full waveform signals of the Ice, Cloud and Elevation Satellite (ICESat) for the classification of land cover types over the cryosphere. The title of this research can be decomposed into three important pieces of information: the full waveform signals of ICESat, the classification of land cover types, and the cryosphere. Together, these indicate the limits of the playing field of this research study: the first part limits the means of the research to signals that are produced by ICESat's laser ranging; the second part limits the goal of the research to the characteristics of the Earth's surface; the third part limits the domain of the research and implies the classification of ice, rock, snow and water.

In this chapter, the structure of the report and research study will be introduced. First of all, a brief explanation on ICESat's full waveform signals will be given in Section 1.1. In Section 1.2 the relevance of the research project will be explained. Next, the research objectives and research question will be given in Section 1.3. Finally, in Section 1.4 and Section 1.5 the report outline and research outline will be discussed with text and flow diagrams.

1.1 Brief introduction to full waveform analysis

Full waveform signals play a central role in this research. The basic principle of ICESat's measurements is simple: the satellite sends out a laser pulse to the Earth's surface and records how long it takes until the reflected pulse returns. The time duration multiplied by the speed of light yields the range between satellite and the ground surface. With the location of the satellite, the Earth's surface can be modeled in a pre-defined reference frame. Yet, this is only the beginning.

This method holds true to its simplicity only if the transmitted pulse would be infinitely small and not disturbed by any other factor than a flat ground surface. However, the re-

ceived pulse is broadened in terms of energy over time by the transmitted pulse length and, more importantly, by all objects along the laser's path line-of-sight. The laser beam cross-section on the Earth's surface is about $120m$ in diameter and because the elevation differences within this footprint may differ, the received waveform (i.e. the distribution of returned laser energy over time) becomes complex. Furthermore, the reflectance of land cover types and atmospheric disturbances like aerosols and clouds have their effect on the waveform as well. In addition, also the satellite's waveform recorder has its effects on the waveform's shape.

Based on the signals that are recorded by ICESat and additional information computed on-ground, data products are created by the National Snow and Ice Data Center (NSIDC) [35]. These data products contain hundreds of attributes, from raw satellite parameters to computed elevation estimates and corrections. For example, these include the raw waveforms; pre-computed signal characteristics, such as signal reflectivity and signal peak location; footprint geometry values; data quality flags and corrections. For this research study, the attributes are essential for analysis of the signals and serve as the framework for constructing the classification algorithm.

After this brief explanation of what ICESat's waveforms are and which factors influence these, one might wonder if these waveforms could be used not only for estimation of surface elevation, but for other applications as well. One of these applications is the classification of land cover types. This would make ICESat data even more successful than it already is.

1.2 Relevance of research

This research study basically comes down to analyzing an alternative application of ICESat data other than the primary goal of the mission. The relevance directly relates to the improvement of current techniques for mass balance estimates of the polar regions.

Currently, most of the mass balance estimates from space are computed either with gravity field determination or surface altimetry (e.g. ICESat), or a combination of the two. To estimate the water content from a mass estimate (for example the amount of snow on Greenland), information is required on the composition of the Earth's surface. This composition changes over time and consequently the water content too. For these approaches, assumptions are made concerning the land cover types.

Since ICESat takes measurements along flight direction and these tracks cross each other once in a while, cross-over measurements can be used for analysis. Several techniques to compute mass balance estimates use these cross-overs to determine the elevation difference over time. However, because the cross-overs are not regularly spaced, in particular away from the Earth's poles, the height change observations at the cross-overs cover

individual points rather than an extended region. To estimate the height change over a complete region, the observations at the cross-overs should be interpolated to get a smooth grid of estimations. This smoothing and interpolation results inevitably in height change estimation errors. The main contributing factor to these errors is the uncertainty between observations, i.e. the difference between the actual height change and the interpolated height change at a certain point. In addition, to estimate the mass gain or loss from height changes, information is required of the land cover type. Here, the relevance of this research kicks in.

Consider two cross-overs on Greenland both showing an surface elevation increase of $1m$ at time *t2* with respect to the year before (*t1*). The elevation increase is presumably caused by snow accumulation and the two cross-overs are $100km$ apart. Interpolation would probably result in an estimated elevation increase of $1m$ right at the center of the two points. Whether this is indeed true, is unknown and consequently this creates uncertainty in the analysis. Even though there is not a cross-over at the central point, there still might be a single observation taken by ICESat at time *t2*. If this observation could be used to determine whether there was bare rock, then it might be concluded that no snow accumulation at this central point occurred. Consequently the interpolation can be corrected for. Thus, knowing the locations of rock areas, the ice sheet mass balance estimates could be improved by filtering out these areas.

In addition to the benefit of identifying areas of bare rock at the polar regions, another important advantage of classification is related to mass densities. Snow, ice and water have different mass densities and information where these land cover types occur could make the estimation of total water content of ice sheets more accurate. Furthermore, the classification could be used for monitoring ice sheet landform change over time. For example, glacier retreat and advance over time and changing snow extent could be monitored. Finally, the classification as a new application can be seen as an extension of the ICESat's possibilities and makes the mission even more valuable than it already is.

1.3 Objective and research question

Having explained the relevance of this research, the overall objective and research question can be stated. This will make the goal of the research more concrete and tangible. The research question will subsequently be divided into subquestions.

The main objective of the research is to improve the overall accuracy of the ICESat-derived elevations and elevation changes over the cryosphere (in particular Antarctica and Greenland). This is a broad objective and could be approached in many ways. To narrow down the scope of the research, the focus will be on two approaches. The first approach relates to finding a methodology to classify land cover types. For this methodology, the

relevant ICESat attributes have to be identified. This is the most important approach and forms the red thread throughout the research. The second approach involves the analysis of surface topography and the ICESat measurement system for recording the signals. The results of this approach will contribute to construct the land cover type classification algorithm.

From the objective and approaches, the main research question and underlying question can be identified. The main research question is:

Can ICESat's full waveform signals be used to discriminate between land cover types over the cryosphere?

The research question can be subdivided into two underlying questions, accompanied by brief explanations:

- *Which attributes can help to distinguish between the land cover types?* The ICESat's full waveform signals consists of several attributes, from waveform width to signal reflectivity. the relevant attributes will be used to construct a classification algorithm. The land cover types that are distinguished for this research are ice, rock, snow and water. To verify these land cover types, validation areas will be identified, together with validation data sets.
- *What effect do surface topography and the ICESat measurement system have on ICESat's full waveform signals?* Surface topography influences the full waveform signals, because of the elevation changes within the footprint. It could assist in identifying classification criteria for the discrimination between land cover types. In addition, the measurement system of ICESat should be taken into account as well, because it influences the recording of the waveform signals.

The priority of this research study is finding a classification algorithm based on full waveform signal analysis. Consequently, most attention will be paid to answering the first subquestion. The second subquestion is mainly answered in the process of finding the classification algorithm.

1.4 Report outline

The report outline gives a rough overview of the research structure and the relationship between the chapters. It is intended to construct the report such that the chapters have minimal overlap and are as decorrelated as possible. The chapters have been structured in the report outline in parallel order as well as in sequential order.

Since the research process is quite complex, the report outline can help to introduce the general structure of the research study. It is the first step in explaining how the research

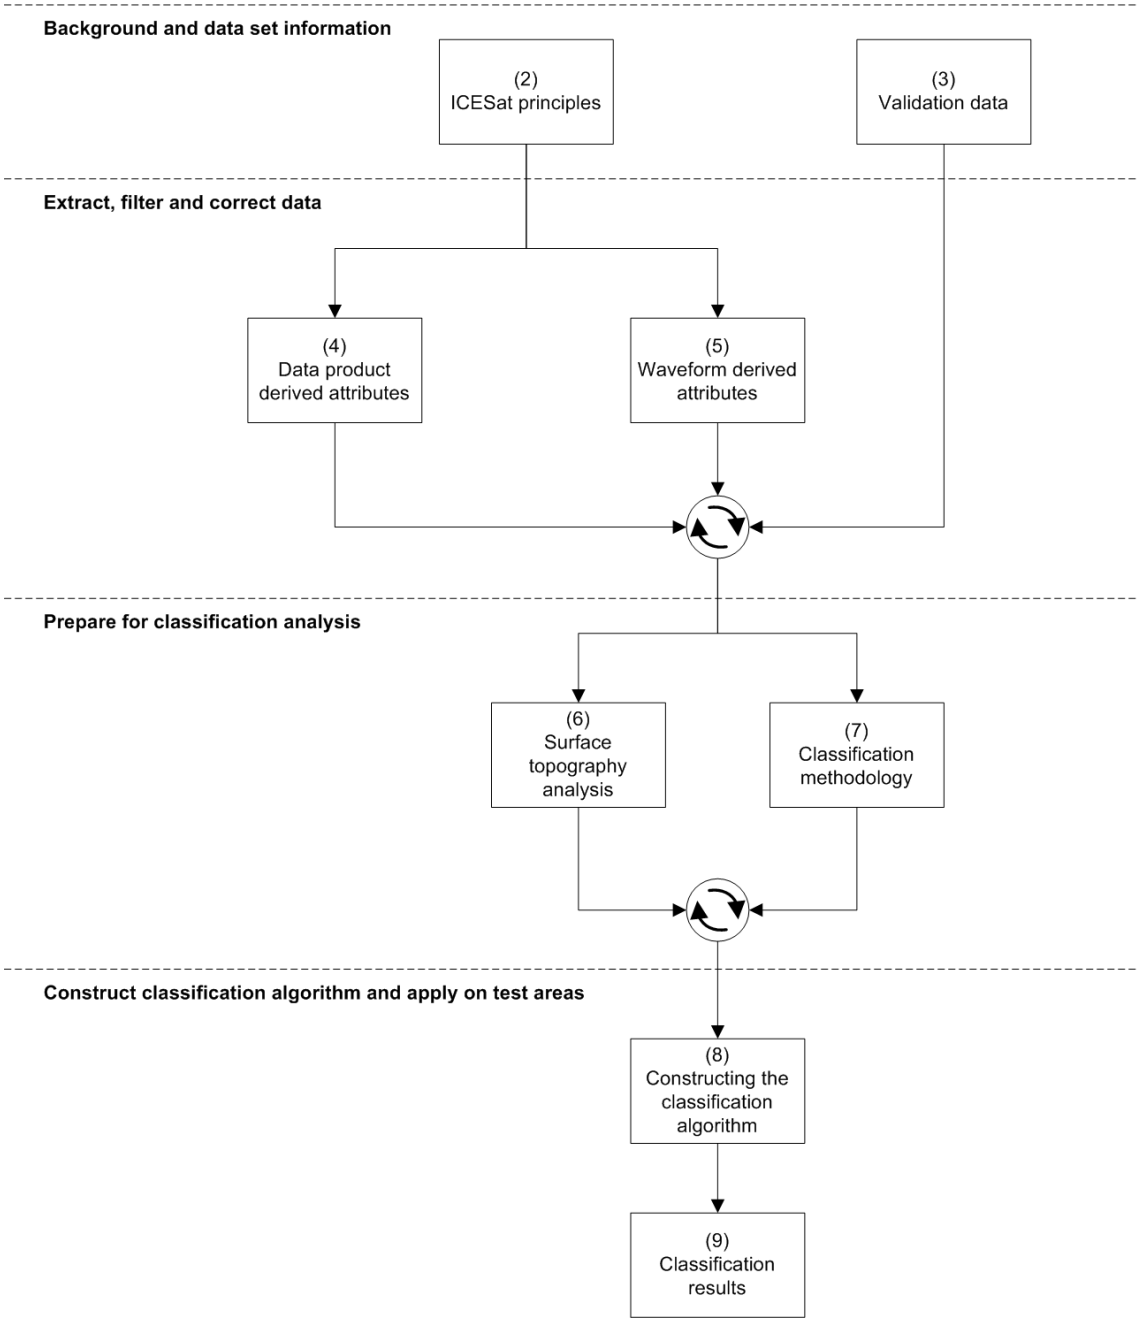

Figure 1.1: The outline of the report, indicating the four segments and the corresponding chapters. The chapters are indicated by square boxes. The corresponding chapter numbers are indicated between parentheses.

question is approached and what it takes to find the answer. The report can be divided into four segments, each one covering multiple chapters, see Figure 1.1. The first segment

addresses the background and data set information. In Chapter 2, the ICESat principles are analyzed, which includes the mission's objectives, the mission's operations and the data products. In Chapter 3 the validation areas and data sets will be discussed. A distinction is made between the validation data for surface topography analysis and the validation data for classification analysis.

The second segment can be summarized as the data extraction, filtering and correction process. In Chapter 4 an overview will be given of the data product derived attributes that have no relationship with the full waveform signals. Here, the characteristics of these attributes are explained, together with the filtering and correction techniques that are applied to them. Chapter 5 gives an overview of all waveform derived attributes. It discusses the exact meaning of waveforms and how these are affected. In addition, the filtering techniques that are applied on the waveforms will be discussed. After these two chapters, the data is ready to be used for more advanced processing and the data can be combined with the validation data sets.

In the third segment, ICESat data and validation data sets are combined to prepare for classification analysis. In Chapter 6 the ICESat waveform signals are analyzed for effects of surface topography. First of all, the theory behind slope and roughness analysis is discussed, followed by a description about an angle-related correction technique that is applied to waveforms. Finally, the results of the surface topography analysis are discussed. In Chapter 7 an explanation will be given on the methodology of applying ICESat data on the classification validation data sets. In addition, the methodology of interpreting the resulting data will be discussed.

In the fourth and final segment the classification algorithm is constructed and applied on test areas. In Chapter 8, the procedure of constructing the algorithm is explained with several case studies, a detailed analysis of data attributes and a step-by-step breakdown of the algorithm. Finally, in Chapter 9 the algorithm is applied on several test areas in Antarctica and Greenland.

1.5 Research methodology and outline

The report outline gave a simple overview of the chapter structure of the report. However, since the structure of the research is more complex, a detailed overview is required to understand the processes from raw data to classification results. This section provides the research outline and is useful for reference while reading the report.

Before reading the rest of the report it is advised to briefly study the research outline of Figure 1.2. It shows the detailed methodology of this research study in a two page flow diagram. All chapters except 'ICESat Principles', 'Data product derived attributes'

and 'Waveform derived attributes' (Chapter 2, Chapter 4 and Chapter 5, respectively) are included in the flow diagram. The chapters that have been left out are included in Figure 2.3 in Chapter 2. These three chapters are combined in the research outline of Figure 1.2 as the yellow colored box 'Pre-filtered and corrected ICESat footprints and waveforms'.

The research outline can be used as a guide for the report in order to better understand how the different parts of analysis contribute to the overall objective of the research.

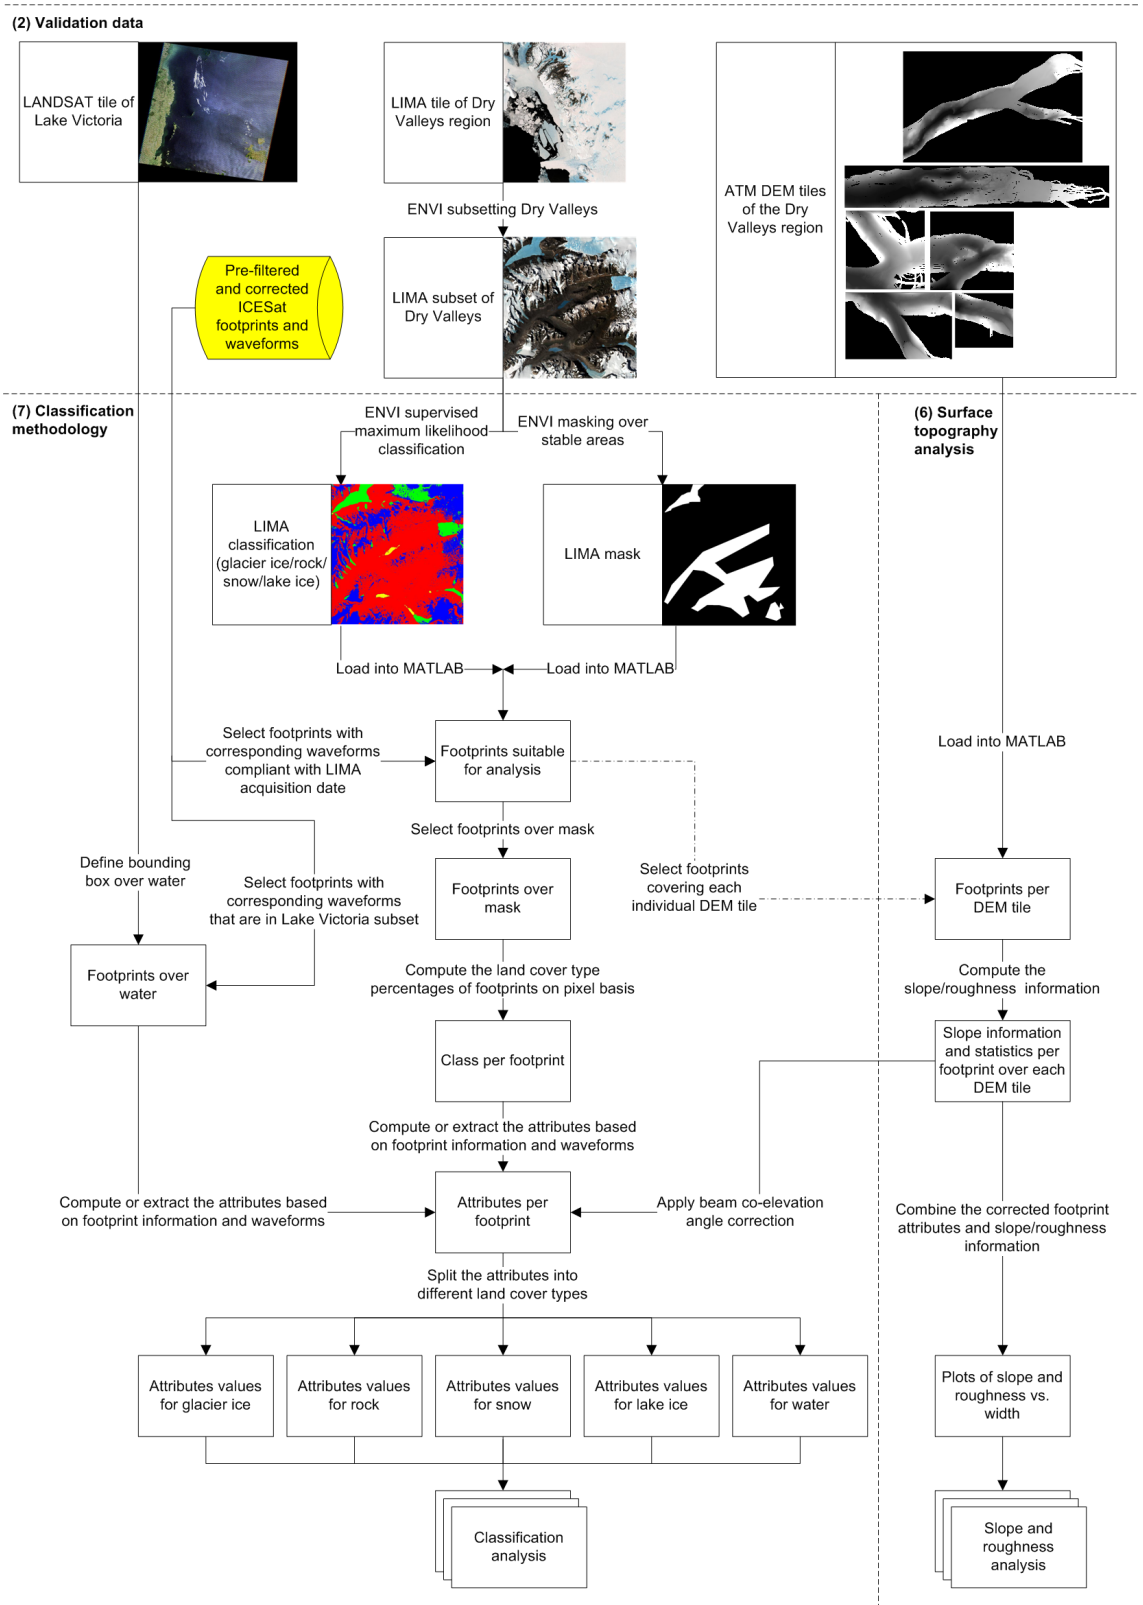
(2) Validation data
LANDSAT tile of Lake Victoria
LIMA tile of Dry Valleys region
ENVI subsetting Dry Valleys
LIMA subset of Dry Valleys
ATM DEM tiles of the Dry Valleys region
Pre-filtered and corrected ICESat footprints and waveforms
(7) Classification methodology
ENVI supervised maximum likelihood classification
ENVI masking over stable areas
LIMA classification (glacier ice/rock/ snow/lake ice)
LIMA mask
Load into MATLAB
Load into MATLAB
Select footprints with corresponding waveforms compliant with LIMA acquisition date
Footprints suitable for analysis
Select footprints over mask
Define bounding box over water
Select footprints with corresponding waveforms that are in Lake Victoria subset
Footprints over mask
Select footprints covering each individual DEM tile
Footprints over water
Compute the land cover type percentages of footprints on pixel basis
Class per footprint
Compute or extract the attributes based on footprint information and waveforms
Compute or extract the attributes based on footprint information and waveforms
Attributes per footprint
Apply beam co-elevation angle correction
Split the attributes into different land cover types
Attributes values for glacier ice
Attributes values for rock
Attributes values for snow
Attributes values for lake ice
Attributes values for water
Classification analysis
(6) Surface topography analysis
Load into MATLAB
Footprints per DEM tile
Compute the slope/roughness information
Slope information and statistics per footprint over each DEM tile
Combine the corrected footprint attributes and slope/roughness information
Plots of slope and roughness vs. width
Slope and roughness analysis

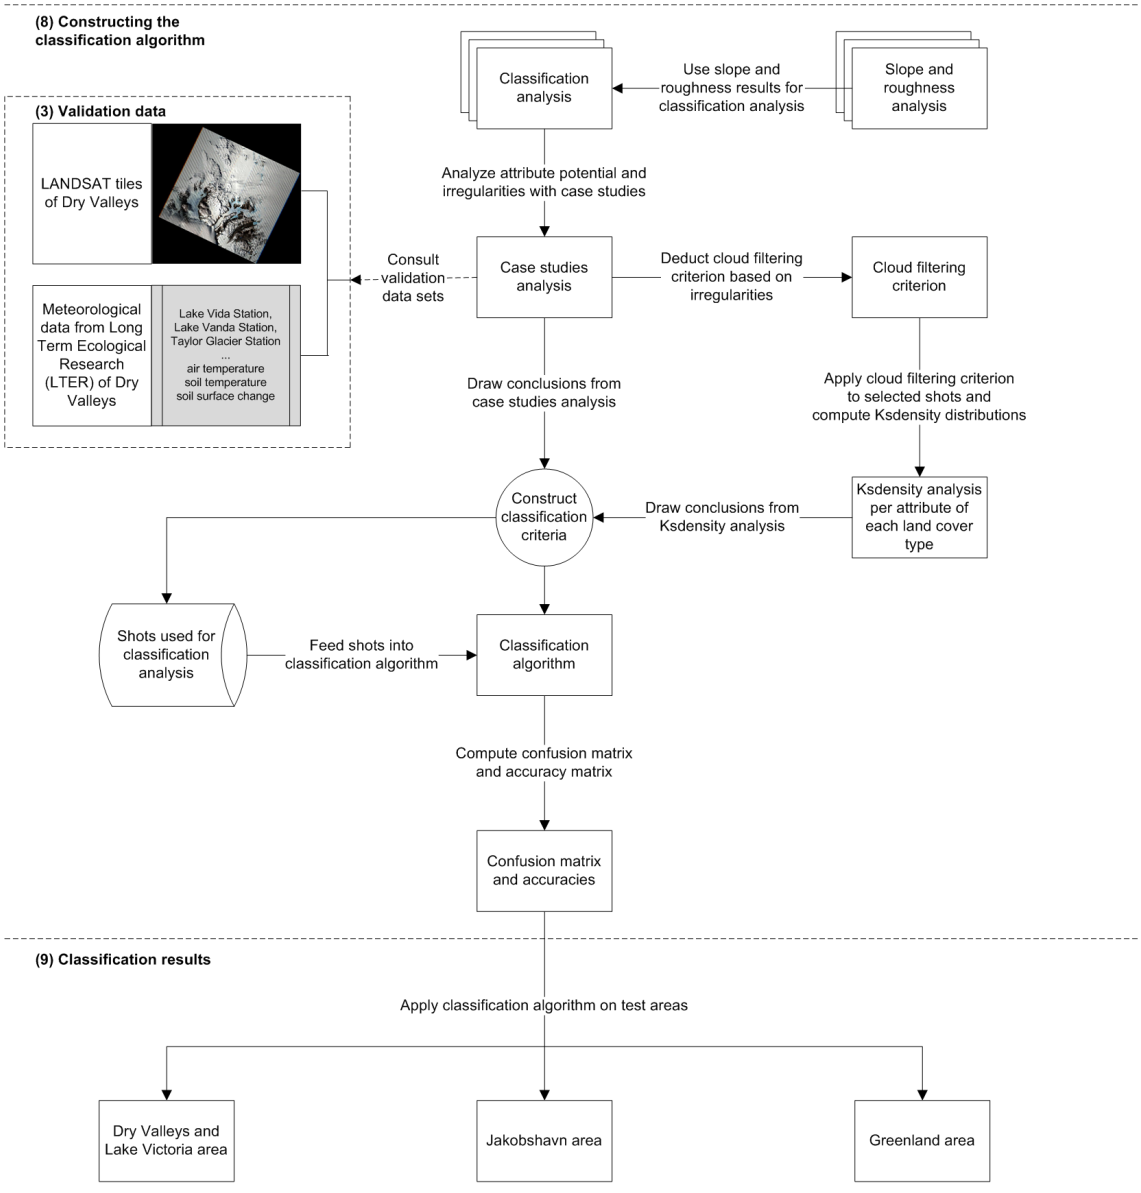

Figure 1.2: The outline of the research, every segment refers to one of the chapters. The corresponding chapter numbers are indicated between parentheses. The yellow box combines three chapters, see Figure 2.3.

Chapter 2

ICESat principles

In this chapter some background information about the ICESat mission will be described, in order to provide the reader an insight into the mission's principles and the data products that are used for this research. In Section 2.1 the objectives, purpose and relevance of the ICESat mission are described. Section 2.2 explains the characteristics of the ICESat mission, including a detailed overview of the main measurement instrument. Finally, Section 2.3 addresses the essentials of the ICESat data in general and a detailed analysis of the available ICESat data used for this research.

2.1 ICESat objectives

The ICESat mission was launched on the 12th of January 2003 and became operational in February 2003. The mission is the benchmark Earth Observation System (EOS) mission of the National Aeronautics and Space Administration (NASA). EOS is one of NASA's programs that comprises of a series of Earth observing satellites designed to monitor and understand climate related processes [25]. The main purpose of the ICESat mission is to measure the mass balance of ice sheets, observe cloud and aerosol heights, and determine land topography and vegetation characteristics [23]. The main measurement principle is laser ranging to determine surface topography elevations and carry out atmospheric profiling. The satellite is the first Earth orbiting satellite that takes measurements with a laser altimeter system. The mission is especially relevant considering the recent concern about the effects of global warming on ice sheet mass balances.

Before ICESat was launched, the science team set up the most important science goals for the mission. These goals represent the very essence of the mission. In the end, all components in the mission contribute to realizing these goals. Below, the most important science goals are listed per discipline: cryosphere, land processes and atmospheric science. The goals are extracted from the Geoscience Laser Altimeter System (GLAS) Science Re-

quirements [5].

Cryosphere: *"to measure long-term changes in the volumes (and mass) of the Greenland and Antarctic ice sheets to sufficient accuracy to assess their impact on global sea level, and to measure seasonal and inter-annual variability of the surface elevation in sufficient spatial and temporal detail to permit identification of long-term trends and to help explain those trends. A further goal is to provide a precise elevation topography of these ice sheets and describe the nature of surface characteristics (e.g., roughness), including sea ice."*

Land processes: *"to conduct topographic measurements of the Earth's land surface on a global basis in order to contribute to a global grid of ground control points for georeferencing of topographic maps and digital elevation models. The secondary land processes science goal is to detect topographic change at the meter per year level or better in selected regions of limited spatial extent."*

Atmospheric science: *"to determine the radiative forcing and vertically resolved atmospheric heating rate due to cloud and aerosol by directly observing the vertical structure and magnitude of cloud and aerosol parameters that are important for the radiative balance of the earth-atmosphere system, but which are ambiguous or impossible to obtain from existing or planned passive remote sensors. A further goal is to directly measure the height of atmospheric transition layers (inversions) which are important for dynamics and mixing, the planetary boundary layer and lifting condensation level."*

The document also notes that, where conflicts occur with the higher-level requirements between the disciplines, the cryosphere requirements should have primary importance. Following from the science goals, the two main measurement types are surface elevation estimation and atmospheric profiling, both using laser principles. At the time of writing, ICESat still provides elevation estimates with unprecedented accuracy based on a technique never used before in space for Earth applications.

A quick glance at the higher level requirements gives an idea about the order of accuracy and resolution that is associated with the measurements. The ice sheet elevation changes are to be estimated with an accuracy of $1.5cm$ per year over areas of $100km$ by $100km$ for low surface sloped areas ($<0.6°$). The vertical sampling resolution of the atmospheric signal is $75m$, while the vertical sampling resolution of the elevation signal is $15cm$ [13]. In Section 2.2, more specifications of the measurements are provided.

Laser ranging as a measurement technique for ice sheet monitoring is in itself not unique; however, it is for Earth observing systems from space. Existing laser knowledge applied to airborne lidar systems and the Mars Orbiter Laser Altimeter (MOLA) mission served as a basis for the design and operation of the space operating laser system of ICESat [23]. Bringing this technique into space for Earth observation has two major advantages over airborne laser ranging, namely the significant improvement of temporal and spatial coverage. This holds especially for the polar regions, which are generally less accessible and

more expensive for airplanes to observe. ICESat provides a combination of unprecedented temporal resolution, spatial coverage and spatial accuracy, that is of particular relevance with today's discussion about the impact of global warming on the ice sheets and sea level rise. Even though ICESat's spatial resolution and vertical accuracy cannot beat airborne laser ranging resolution and accuracy, the combination with temporal and spatial coverage makes this space system truly unique.

2.2 ICESat mission

The satellite should be considered as only one aspect of the ICESat mission. The mission as a whole includes the space segment, the data and information support segment, the mission control segment and the science segment. In the end, all segments have one common goal, which is satisfying the primary objectives and science requirements of ICESat.

The space segment is predominantly the satellite. The satellite has a near-circular orbit at an approximate altitude of $600km$, with an inclination of $94°$. During the first 120 days after launch, the satellite was in a repeat orbit of 8 days, after which the orbit was changed to a 183 day repeat orbit [8]. The satellite was launched from Vandenberg Air Force Base, California, with a nominal lifetime scheduled for 3 years, with a 5 year goal.

The GLAS instrument is the primary payload on ICESat. The main purpose of the instrument is to send out and receive the laser pulse and measure the time it takes for the transmitted pulse to come back to the satellite [8]. Figure 2.1 illustrates the GLAS instrument with annotations to the vital parts. The instrument includes three identical and redundant lasers, each sending out pulses at $1064nm$ wavelength (in the near-infrared) for measuring the surface topography and at $532nm$ wavelength (in the green) for measuring atmospheric characteristics [33]. When operating in nominal mode, the laser continuously fires pulses at a rate of $40Hz$, each with a pulse width of $5ns$. Each of the three solid state Nd:YAG lasers has an expected lifetime of about 2.5 billion shots, which translates into approximately 2 years of continuous operation. Together, the lasers are designed to meet the 3-year minimum lifetime requirement with a 5-year goal. However, after more than a month of operation, failure of the first laser forced the mission to handle a revised measurement campaign strategy. According to the Independent GLAS Anomaly Review Board Executive Summary [30], an unexpected failure mechanism in the laser's pump diode caused increased degradation of laser 1 and subsequent failure after only 37 days of operation. The design fault was likely to exist in the other two lasers as well, so a decision had to be made whether to continue taking continuous laser measurements and consequently face a premature ending of the mission, or to carry out discrete measurement campaigns and consequently be able to measure for a larger time frame. Since taking measurements over a long period of time is of higher importance for the mission, the choice was made to measure for discrete periods. As a result, the orbit of ICESat was changed to a 91-cycle

orbit with a 33 day sub-cycle. From that defining moment on, the instrument has been operational for periods of approximately one-month, with campaigns conducted every three to six months. At the time of writing, ICESat already has been measuring for more than six years, recording billions of measurements to date.

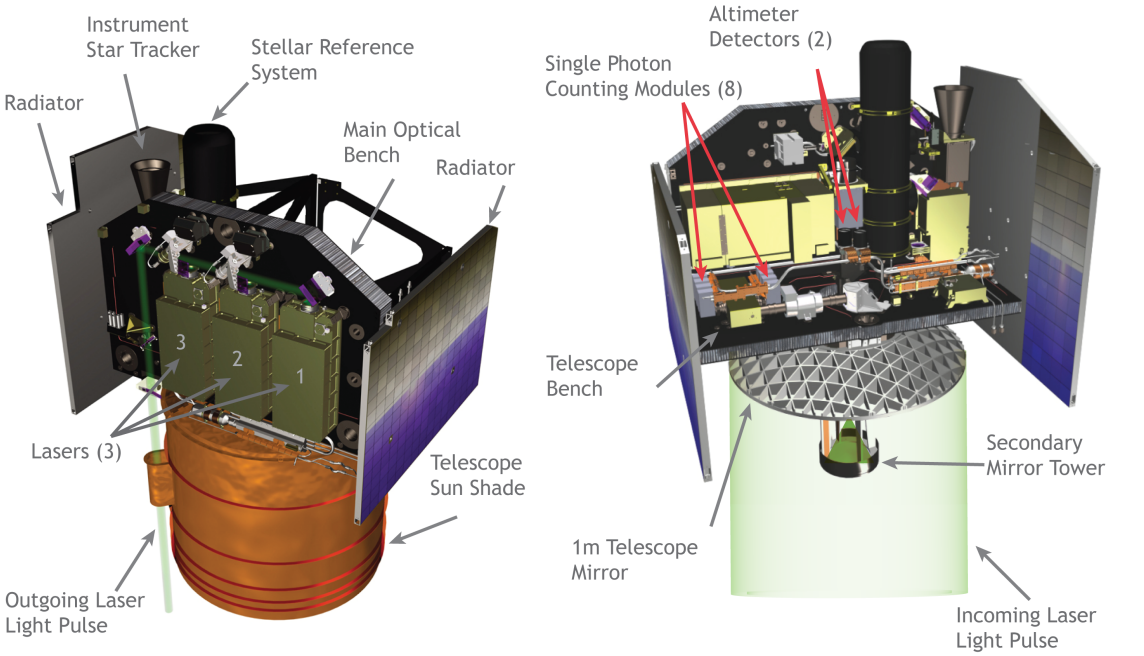

Figure 2.1: The GLAS instrument with annotations. The red arrows indicate the main signal measuring sub-instruments. Illustrations taken and adapted from [24, 13].

The GLAS instrument also consists of a telescope of about $1m$ in diameter collecting the reflected photons of the laser pulse. A mirror reflects the incoming beam onto the secondary mirror tower, which reflects the beam to the detectors. The incoming reflected laser light is then split into the $1064nm$ and $532nm$ signals. The $1064nm$ signal is directed to an altimeter detector and a dense cloud and aerosol detector, both subject to a bandpass filter of $0.8nm$. The first one uses an Avalanche Photo Diode detector which digitizes both the transmitted and surface backscatter signals at $1GHz$. This results in a time resolution of $1ns$, equivalent to a range resolution of $15cm$ taking into account the two-way travel. This is the only detector system used for altimetry measurements, recording the full waveforms that are used in this research study. The second one, i.e. the aerosol detector, separately filters and digitizes the signal at $1.953MHz$, resulting in a time resolution of $512ns$ and a vertical resolution of $76.8m$ [14]. This detector is used for the detection of aerosols and dense clouds and is complementary to the $532nm$ signal for atmospheric measurements. The $532nm$ laser pulses are measured with the more sensitive photon counting channel, consisting of eight separate photon counting detectors for improved dynamic range and redundancy. The vertical resolution is also $76.8m$. This channel uses a bandpass filter of about $0.03nm$, which is considerably lower than the altimeter detector, and consequently

provides the highest quality data for measuring both aerosol and cloud returns. The reason that the $532nm$ channel is used for atmospheric measurements is because the backscatter at this frequency is much more sensitive to weak scattering layers in the atmosphere [33].

Next to the lasers and detectors, GLAS includes an instrument star-tracker system with a stellar reference system for orbit and attitude determination. The ICESat bus carries a GPS receiver, coarse sun sensors, gyro system, and a laser retroreflector for additional orbit and attitude determination. Post-processing on-ground then determines the precise orbit and precise attitude determination, which is of particular importance for accurately estimating surface elevations and carrying out atmospheric profiling measurements. Heat pipes and two radiators supplemented by heaters take care of the thermal control. The optical and telescope benches serve as supporting platform for the sub-instruments.

The footprint that is illuminated by the laser was supposed to be about $70m$ in diameter, with spatial intervals of $170m$. However, due to the inconsistent performance of the laser systems, the footprint geometry changed significantly over the laser operating time. Also, the degradation caused the laser energy to decrease over time. For example, the footprint's major axis was roughly $160m$ for laser 1 measurements and roughly $60m$ for late laser 3 measurements. The transmit power degraded, for example, from $67mJ$ (early laser 3 measurements) to $32mJ$ (late laser 3 measurements) [45]. Another unfortunate effect of the laser degradation is that the atmospheric measurements became less reliable in the later laser campaigns [46].

As mentioned before, the GLAS instrument is equipped with three identical lasers. Each of these lasers fire laser pulses during several campaigns. Table 2.1 shows the several laser campaigns that have been conducted to date. The name of the laser campaigns are composed as follows, with L3D taken as example: the 'L' stands for laser, number '3' for laser number (i.e. 1 to 3), and the letter 'D' for the campaign number. The table shows that the first laser operated for somewhat more than a month, the second for half a year and the third for five years. At the time of writing, the third laser failed during camping L3K and the second laser was used to complete the initially planned ground tracks.

Table 2.2 shows some important specifications of the GLAS instrument, as expected before the launch of ICESat. In reality, the pulse transmit energy is not constant and fluctuates during the laser campaigns. The altimeter detector collects about 4.5 million $1ns$ samples for each transmitted pulse. However, only 544 samples over ice sheet and land surface, and 200 samples over sea and oceans are sent to the ground stations to reduce the data stream [3]. The pointing of the laser beam is carried out by rolling and/or pitching the spacecraft and can be controlled to within $\pm 35m$ on the surface [14]. In nominal mode, ICESat is near-nadir pointing with a beam co-elevation angle of about $0.3°$, pitched upward in the direction of the spacecraft motion. The reason for this off-nadir pointing is to protect the receiver from specular reflection of the laser pulse from flat, highly-reflected surfaces [45]. The pointing angle can be increased to about $5°$ for measuring targets of opportunity, such as validation and calibration sites [14].

Campaign number	Start date of operation	End date of operation	Days of operation
L1A	20/02/2003	29/03/2003	38 days
L2A	25/09/2003	19/11/2003	55 days
L2B	17/02/2004	21/03/2004	34 days
L2C	18/05/2004	21/06/2004	35 days
L3A	03/10/2004	08/11/2004	37 days
L3B	17/02/2005	24/03/2005	36 days
L3C	20/05/2005	23/06/2005	35 days
L3D	21/10/2005	24/11/2005	35 days
L3E	22/02/2006	28/03/2006	34 days
L3F	24/05/2006	26/06/2006	33 days
L3G	25/10/2006	27/11/2006	34 days
L3H	12/03/2007	14/04/2007	34 days
L3I	02/10/2007	05/11/2007	37 days
L3J	17/02/2008	21/03/2008	34 days
L3K	04/10/2008	19/10/2008	15 days
L2D	25/11/2008	17/12/2008	19 days

Table 2.1: The laser campaigns until the time of writing, with the start date, end date of operation and days of operation.

Specification	Surface	Atmosphere	Atmosphere
Wavelength	$1064nm$	$1064nm$	$532nm$
Laser pulse energy	$75mJ$	$75mJ$	$35mJ$
Laser pulse rate	$40Hz$	$40Hz$	$40Hz$
Laser pulse width	$5ns$	$5ns$	$5ns$
Telescope diameter	$1.0m$	$1.0m$	$1.0m$
Receiver FOV	$0.5mrad$	$0.5mrad$	$0.17mrad$
Vertical sampling resolution	$0.15m$	$76.8m$	$76.8m$

Table 2.2: The intended specifications of the GLAS instrument, taken and adapted from [33] and [14].

Since there exists a spatial gap between the footprints centers of about $170m$, the ground track is not illuminated by the laser continuously. Also, since the orbit is inclined by 94°, the ground tracks of the satellite are more densely distributed in the polar regions. The polar region is not covered completely due to the 4° polar gap, as illustrated in Figure 2.2.

The data and information support segment basically provides the infrastructure to distribute data to and from the satellite. Not only scientific data is accounted for, but also satellite housekeeping data and commands (e.g. for orbital corrections and activation of the GLAS instrument). It also provides the data portal to the end user of the data products. This segment is mainly operated by the Goddard Space Flight Center (GSFC). The National Snow and Ice Data Center (NSIDC) archives and distributes the data products to the users [38].

The mission control segment is a general term for all ground activities related to the functioning of the satellite itself. The launch vehicle was provided by the NASA Kennedy Space Center, the GLAS instrument mainly by the GSFC and the spacecraft by Ball

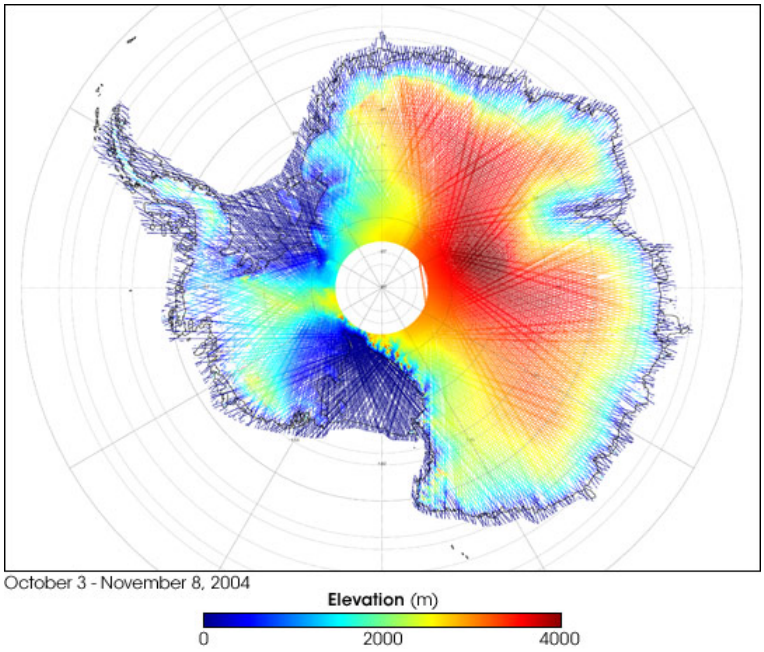

Figure 2.2: ICESat's ground tracks covering Antarctica to illustrate the denser coverage at the polar regions and the polar gap. The variable colors indicate elevation values, taken from measurements from October 3 through November 8, 2004 [20]. Image taken from the NASA Earth Observatory [26].

Aerospace. In orbit, the satellite is operated by Ball Aerospace. The housekeeping data of the satellite is monitored to check the satellite's status. In addition, here the orbital maneuvers are planned and prepared, and flight dynamics are monitored (e.g. orbit and attitude determination). Basically, Ball Aerospace monitors the satellite and compiles the commands to be sent to ICESat. The actual sending is then carried out via the ground stations that are part of the data and information support segment. The mission control segment works closely together with the science segment to discuss the status and plan schedules and scenarios for the satellite.

The science segment provides the in-depth research of the data products and basically forms the last step before providing the results to the users. Research is conducted on the data itself, the scientific methods used to process the data, potential optimizations, and data related results, such as estimating the actual ice mass balance changes. Reports, documents, publications and papers are produced by this segment to maximally exploit the available data. Also, this segment offers an advising role to the mission control segment for decision making regarding the functioning of the satellite. The key player in the science segment is the science team of ICESat, which comprises of scientists located at various research institutes like NASA departments, several universities and research centers [23]. The Center for Space Research (CSR) of the University of Texas at Austin and the GSFC act as the science data system (SDS) centers that primarily handle data processing. Here,

the raw data is processed and data products are created. Finally, the data products are provided to the data information and support segment, where the NSIDC archives the data products and distributes them to the user.

2.3 Data products

In total 15 data products are produced by the ICESat mission, each of them offers a unique set of parameters. Essentially all the relevant information of ICESat is incorporated in the data products, from raw satellite parameters to computed elevation estimations and parameter corrections. For this research, three data products are of special interest, namely the altimetry data, the global elevation data, and the global backscatter data.

All data products are available through a data portal of NSIDC; however, not all data sets are relevant for this thesis project. Three of them are, with the global altimetry data GLA01 being the first, since it includes the transmitted and received full waveforms. The second important data product is the global elevation data, GLA06, which incorporates, amongst others, accurate geolocation parameters of the footprints, the estimated elevations, correction parameters and waveform statistics. Finally, GLA07 contains atmospheric backscatter profiles that are used for classification analysis. Table 2.3 gives an overview of all the available data products.

Data product ID	Processing level level	Description
GLA01	1A	Global altimetry data
GLA02	1A	Global atmosphere data
GLA03	1A	Global engineering data
GLA04	1A	Global laser pointing data
GLA05	1B	Global waveform-based range corrections data
GLA06	1B	Global elevation data
GLA07	1B	Global backscatter data
GLA08	2	Global planetary boundary layer and elevated aerosol layer heights
GLA09	2	Global cloud heights for multi-layer clouds
GLA10	2	Global aerosol vertical structure data
GLA11	2	Global thin cloud/aerosol optical depths data
GLA12	2	Antarctic and Greenland ice sheet altimetry data
GLA13	2	Sea ice altimetry data
GLA14	2	Global land surface altimetry data
GLA15	2	Ocean altimetry data

Table 2.3: The data products available of ICESat, together with a short description. GLA01, GLA06 and GLA07 are the most important data products used for this research.

The level-1A data products contain raw measurement data, such as satellite housekeeping data and waveforms, and some precomputed parameters, such as estimated footprint dimensions. The GLA05 level-1B products contain corrections data including parameters required to calculate surface slope and roughness. It is used in conjunction with GLA01

to create GLA06 and GLA12-GLA15 [34]. GLA06 is a global product, containing all elevations that are computed using the standard parameterization, the ice sheet algorithm. GLA12-GLA14 are regional elevation data products, which only cover the specified surface type [34]. There are two algorithms for estimating the elevations, the standard one is used for ice sheets, oceans, and sea ice. This algorithm fits no more than two Gaussian functions in the received waveform and uses the maximum peak location for elevation estimation. The so-called alternative algorithm is used for land and allows up to six Gaussian functions, which is especially useful for vegetated areas. It uses the centroid of the sum of the fitted Gaussian functions for elevation estimation. Both correct for atmospheric delays [35]. The returned waveforms are subject to these two algorithms for estimating the elevations. Which algorithm is used for data product computation is dependent on a predefined surface mask. For this research the GLA06 product is used, because it uses the right (standard) algorithm for polar regions and is available globally. The GLA14 product is more relevant when looking at the elevations over more complex regions, like vegetated and urban areas. GLA02 and GLA08-GLA11 represent global atmospheric data products and are not relevant for this research. Besides the extensive use of GLA01 and GLA06, also the Global backscatter data product GLA07 is used for analysis of the atmospheric signals. Please refer to NSIDC [35] for more information on these products and algorithms.

All the data products are provided as binary formatted files to reduce the file size. Some of the parameters are recorded per individual shot (at $40Hz$), while others are only available per 40 shots (at $1Hz$). The data products used for this research are linked together with an unique record number and shot number. Each record number contains 40 shot numbers (i.e. corresponding to the 40 shots per second). For this research the record numbers have been combined with the shot numbers, such that every individual shot has its own unique shotID. With the shotIDs, the corresponding parameters in different data products can be found. For this research, all laser campaigns up through L3J are used for analysis. Only the data products of data Release 28 are considered for this research. Each data release incorporates its own corrections and recomputing techniques for improved data usage.

For full waveform analysis of the ICESat laser data, the Global altimetry data (GLA01) is the main source since this data set contains the transmitted and received waveforms from the GLAS instrument. Both the received echo and the corresponding transmit pulse are available in the GLA01 data set. Since the returned waveform is dependent on the laser output waveform, see Chapter 5.3, both waveforms are taken into account. Special care is paid to filtering criteria and corrections that are incorporated and applied to the data. For this research two sets of attributes are used: data product derived attributes and waveform derived attributes. Both sets are used for filtering and correction techniques as well as for surface topography and classification analysis later in this research. In Chapter 4, the data product derived attributes will be discussed and in Chapter 5, the waveform derived attributes will be discussed. The main difference between the two sets of attributes is the following.

The data product derived attributes includes all attributes that are extracted from data products but are not directly related to the received or transmitted waveform. For example, attributes related to geolocation, footprint dimensions and atmospheric signals are included. As a consequence, only GLA06 and GLA07 attributes are used in this discussion. The waveform derived attributes includes data product derived attributes as well, but only the ones that are directly related to waveforms. Since the waveforms are extracted from GLA01, this data product is the primary data product in this discussion. However, there are several GLA06 derived attributes that relate to the waveform as well, and as a consequence both GLA01 and GLA06 attributes are used in this discussion. Additionally, several self-computed (waveform) attributes are included as well.

As indicated by the report outline and research outline (Figure 1.1 and Figure 1.2, respectively), the chapter on validation data is essential for continuing with the research process. As a results, this chapter is discussed next. After the validation data chapter, the attributes that are used for the research will be discussed. Since these two chapters, i.e. 'Data product derived attributes' and 'Waveform derived attributes', are directly related to this chapter (see the report outline) both will be introduced briefly with Figure 2.3. This flow diagram shows the steps to compute the the 'Pre-filtered and corrected ICESat footprints and waveforms', that is included in the research outline.

The importance of presenting this flow diagram now is because of the contribution of the current chapter to later discussion on the ICESat data. As the diagram indicates, first the binary files of the data products are requested using the NSIDC data subsetter [38]. With the subsetter, spatial and temporal preferences can be specified. This is preferred above the other data portals, because only for the data subsetter the spatial subsetting of the GLA01 data product is enabled. After the data sets have been subsetted, the sets can be retrieved from FTP. All operations after the data sets have been downloaded are performed with IDL, ENVI or MATLAB. The reason that the box 'spatial and temporal preferences' is included in the second segment in the flow diagram as well is because the data sets from NSIDC should be customized with IDL or MATLAB for specific purposes based on finer preferences. The rest of the flow diagram is explained in more detail in Chapter 4 and Chapter 5. The waveform correction will be explained in Chapter 6.

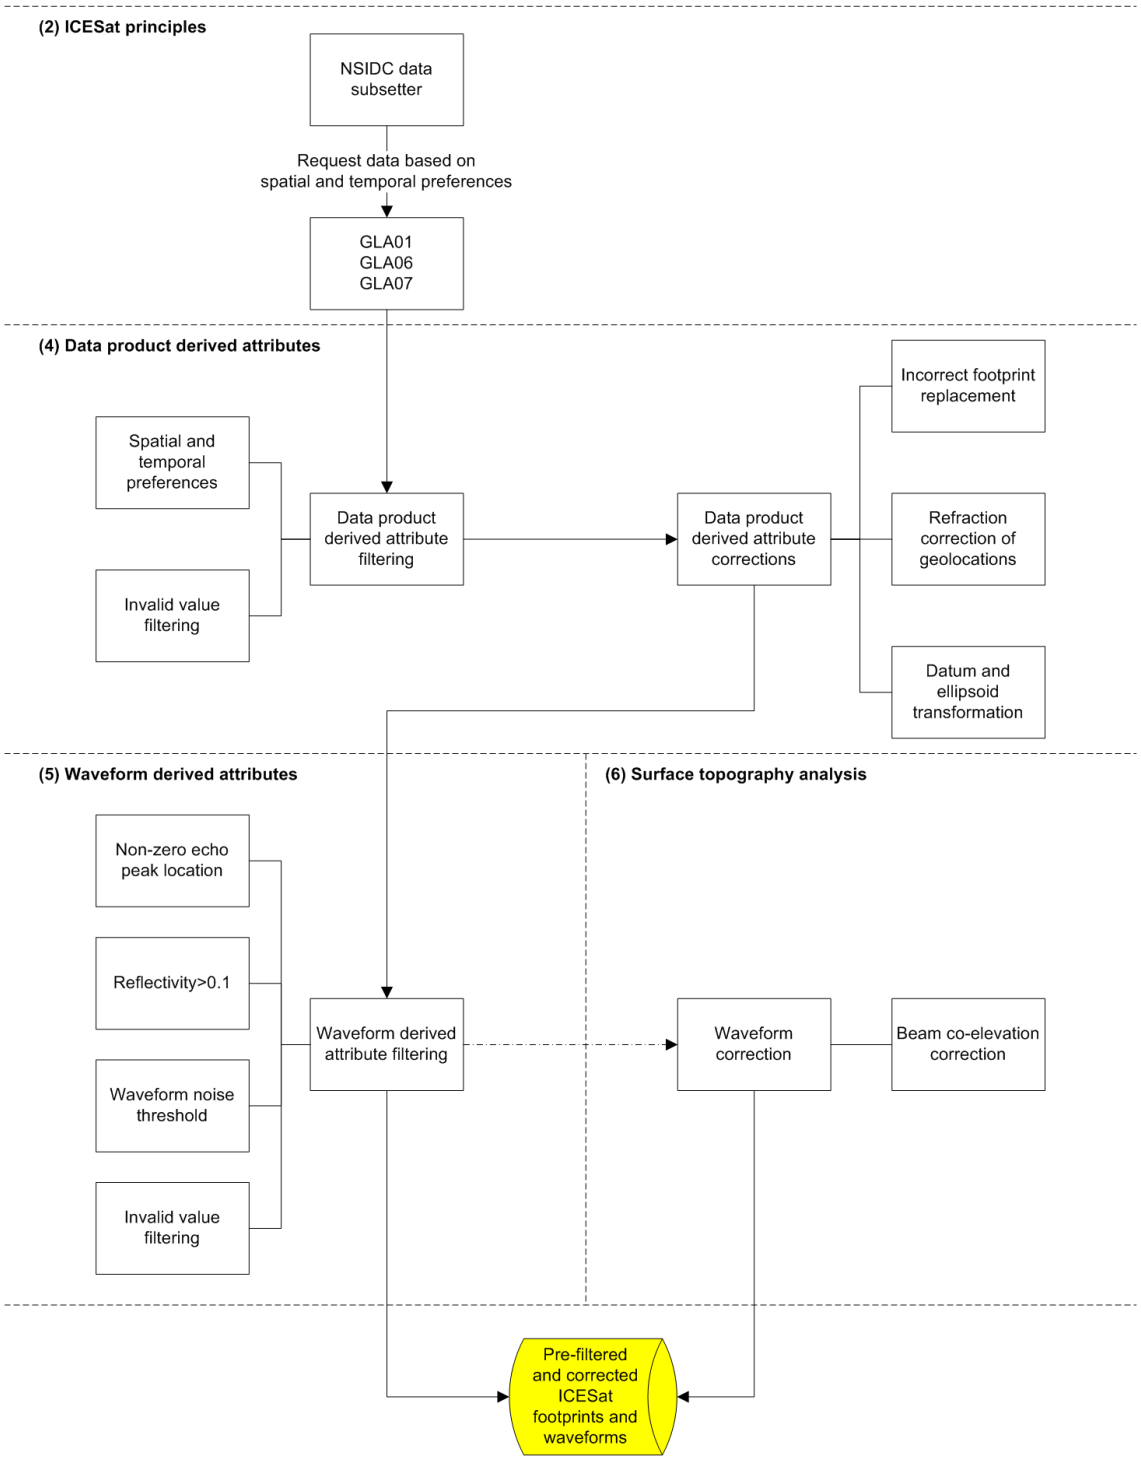

Figure 2.3: The flow diagram combining the three chapters that are responsible for GLAS data retrieval and operations, and one chapter that explains the waveform correction technique. The corresponding chapter numbers are indicated between parentheses.

Chapter 3

Validation data

In this chapter the validation areas and related data sets will be discussed. The detailed analysis that is carried out on the data is not described here, but in Chapter 6 and Chapter 7. Validation forms an essential part of the research, because it offers information about what ICESat's returned waveforms actually describe. Land cover type classification and surface topography analysis based on full-waveform analysis play a central role in this research, and require accurate and suitable validation data sets. In Section 3.1 the Dry Valleys in Antarctica will be discussed that serve as the primary validation site. In addition, some basic information is given about Lake Victoria that serves as the validation area for water. The classification analysis is carried out on both areas; the surface topography analysis is only carried out on the Dry Valleys. In Section 3.2 the validation data sets that will be used for classification purposes of both validation areas are discussed. This section is subdivided into two paragraphs to emphasize the importance of the different data sets. Finally, in Section 3.3 the validation data sets are described that will be used for surface topography analysis.

3.1 Validation areas

The Dry Valleys in Antarctica and Lake Victoria in Africa form the main validation areas of this research. For the validation of rock, glacier ice, lake ice and snow covered regions the Dry Valleys are used, for the validation of water Lake Victoria is used.

For classification and surface topography analysis, validation data sets are used that cover the Dry Valleys region in Antarctica. The region is located near the coast of the Ross Sea, see Figure 3.1. It consists of several valleys, mountain ridges and glacier formations. Figure 3.2 shows the part of the Dry Valleys that is used for this research. The valleys are 5-10km wide and 15-50km long, are mostly flat in the middle and sloped towards the ridges. Height differences between the valleys and neighboring mountains range up to several kilometers. The Dry Valleys offer some unique features that make the region suitable for various cryospheric research studies. The valleys get virtually no snowfall or moisture

and, as a consequence, the valleys form the largest relatively ice-free area of Antarctica [17]. The area is also known as the coldest and driest place on Earth [6]. In the valleys, several near-perennial ice covered lakes are located, which occasionally melt during the Antarctic mid-summer. Generally speaking, the small glaciers do not invade into the valleys because of regional and local weather patterns. Small inlet glaciers on the heads of the western part of Wright Valley, and the Mckelvey and Barwick Valleys, connect the Dry Valleys with the snow covered polar plateau and inland domes. These interesting features make the region attractive for various research studies and, as a result, the region has been well documented by various measurement missions. It serves as a validation site for several remote sensing missions as well, such as the ICESat mission. For this research, the region is used as a validation site for four reasons in particular: the meteorological stability, the presence of the land cover types that are required for this research study, the high density of ICESat measurements and overlapping tracks, and the presence of weather stations (see Section 3.2). The land cover types lake ice, glacier ice, snow and bare rock can all be found in the Dry Valleys with minimal change over time, due to little precipitation and the rapid sublimation of snow. The soil type in this part of the Dry Valleys is furthermore almost homogeneously covered by a material named Typic Anhyorthel [16]. This soil type is described as stony, gravely sands that is composed of basalt with ice-cemented permafrost below about $30cm$ [9]. No detailed study was found on which soil types cover the cryospheric regions, and the Typic Anhyorthel is assumed to be representative for polar regions other than the Dry Valleys as well. This assumption is based on the fact that Typic Anhyorthel is the main soil type of the Dry Valleys and is based on the similarity between LANDSAT images of bare rock covering the Dry Valleys and LANDSAT images of bare rock covering Greenland. In addition, the main objective of this research study is to find a methodology to discriminate between ice, rock, snow and water, rather than between various soil types.

Figure 3.1 shows the position of Antarctica and the neighboring continents. The location of the Dry Valleys, encircled with red, is located less than $100km$ from the Ross Sea, which partly freezes over winter time. The plain inland area towards the south pole is called the polar plateau and consists mainly of perennial snow covered inland domes. The region of the Dry Valleys that is used for this research is shown in Figure 3.2, which is the main part of the Dry Valleys.

As Figure 3.2 shows, the Dry Valleys offer a validation area for the land cover types glacier ice, lake ice, bare rock and snow. The fifth land cover type that is analyzed for this research is water. The Dry Valleys are located close to the Ross Sea, which partly freezes during the winter season. Between winter and the summer time, sea ice covers the seas, which makes the Antarctic region not optimal for analysis of waveform covering water. Therefore, another validation site is selected for water waveforms. To mitigate the probability of sea ice, Lake Victoria was selected. Figure 3.3 shows the location of Lake Victoria and the spatial boundary box used to subset ICESat data.

The bounding box has an upper latitude of -1.60°, lower latitude of -2.10°, western

Figure 3.1: The Antarctic continent with the location of the Dry Valleys region highlighted with the red circle. Illustration taken and adapted from [52].

longitude of 32.00° and eastern longitude of 32.60°. These spatial limits are taken to avoid land to be included in the image. Although there is the possibility that man-made structures and unnoticed natural features are present in the selected region, it is assumed that the bounding box encloses nothing but a surface of water. The lake is the world's second widest freshwater body and has a mean depth of $40m$. The average temperature ranges from 21°C to 24°C [15]. Since the lake has different characteristics compared to other lakes and seas around the world, it cannot perfectly represent all water bodies. For example, it may differ in salinity, wave height and temperature, but for this research these influences have not been analyzed. It is assumed that the waveforms over this region form a representation for all water bodies, from (ice-free) seas close to land to (ice-free) lakes inland. Ice covered lakes and seas are represented by the frozen lakes in the Dry Valleys.

3.2 Classification validation data

For the classification part of this research, the influences of land cover types on waveforms shapes and characteristics are analyzed. In order to know what land cover type each waveform describes, validation data sets are required.

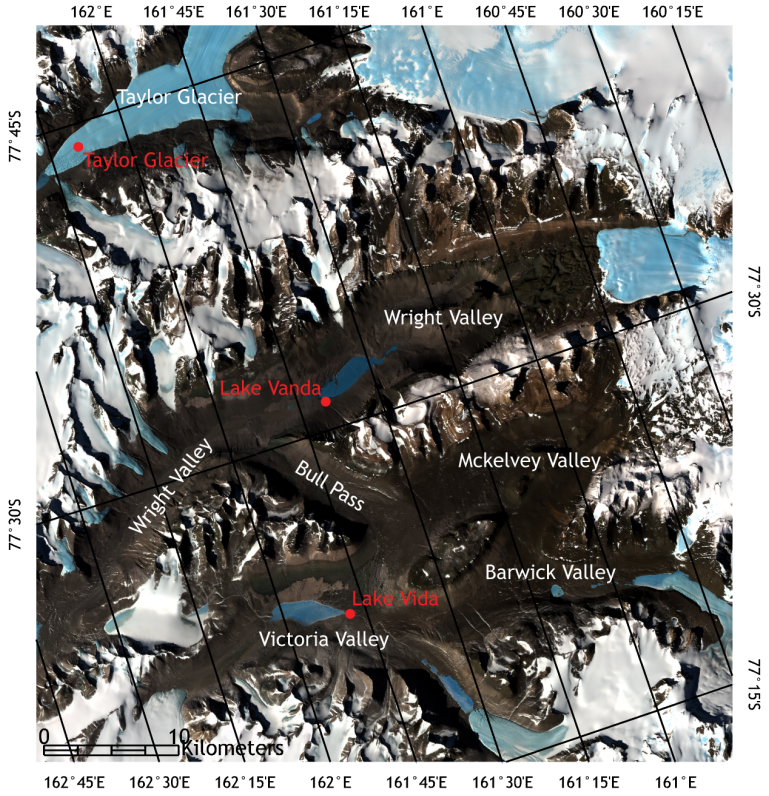

Figure 3.2: The Dry Valleys area in Antarctica that is used for the research. The black lines represent geographic grid lines, the red dots and names refer to the meteorological LTER stations, see Paragraph 3.2.2. The naming of the valleys and glacier is consistent with the naming used by US Geological Survey (USGS), [51]. Data taken and adapted from LANDSAT Image Mosaic of Antarctica (LIMA) [52].

The Dry Valleys form an important region for various research studies and there are several relevant data sets available for this area, including satellite images and on-site weather station data. One of the most important criteria for the validation data is that is should have sufficient resolution. With ICESat's laser footprint areas on the Earth's surface ranging from 2,800m^2 to 20,000m^2 (which is equivalent to a third of and two and half times a full soccer field respectively) and a distance between footprint centers of only 175m, the resolution of the validation data should be such that the land cover class can be identified for every individual footprint. LANDSAT images offer the resolution to achieve this and, in combination with the availability of the spectral information in various bands, these form a suitable validation data set. In Paragraph 3.2.1, this validation data set is discussed. A meteorological data set obtained by weather stations in the Dry Valleys form a complementary validation data set and is discussed in Paragraph 3.2.2. For Lake Victoria, only LANDSAT images are used.

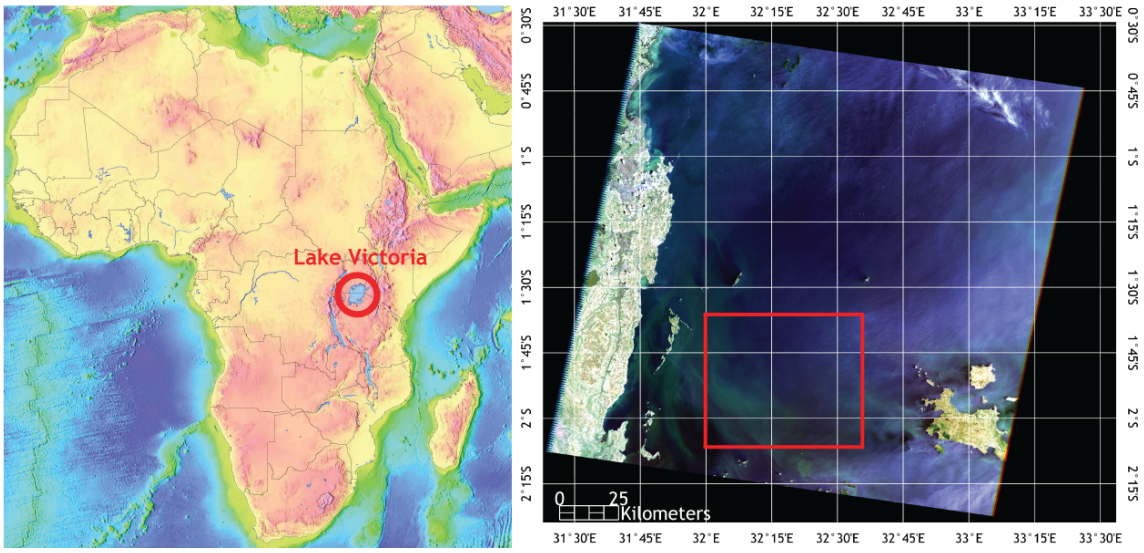

Figure 3.3: The location of Lake Victoria (left) and a LANDSAT image of part of the lake (right). The red bounding box in the right image indicates the area for subsetting ICESat data. The white lines in the LANDSAT image represent geographic grid lines. Left image taken and adapted from USGS Coastal and Marine Geology InfoBank [48], data of right image taken and adapted from the USGS Earth Resources Observation and Science (EROS) data portal [49].

3.2.1 LANDSAT images

In total, seven LANDSAT satellites have been launched over the last 36 years. Only images of the most recent one, LANDSAT 7, are used for validation purposes for this research. LANDSAT 7, which is still operational, was launched in 1999 and is a joint initiative of the USGS and NASA. The sensors collect data in eight spectral bands and images at resolutions in the range of $15m$ to $60m$. The satellite acquires the data over a $185km$ swath, in visible and infrared bands at *30m* spatial resolution, in the pan-chromatic band at $15m$ and in the thermal band at $60m$ [53].

Table 3.1 shows the band designations of LANDSAT with the corresponding spatial resolution values.

In support of the International Polar Year 2007-2008, a joint effort of USGS, British Antarctic Survey and NASA created the LANDSAT Image Mosaic of Antarctica (LIMA). Over 1000 LANDSAT 7 images have been combined into LIMA, which is a seamless, virtually cloud-free mosaic that offers high resolution images of Antarctica, except the center of the continent [52]. LIMA offers various pre-processed versions of the images, each highlighting different features of the landscape. For this research, the most interesting version is the false-color, pan-sharpened product, which incorporates infrared data. The bands in the green (band 2), the red (band 3) and near-infrared (band 4) at $30m$ resolution (i.e.

Band	Wavelength [nm]	Resolution [m]
1	450-520	30
2	520-600	30
3	630-690	30
4	770-900	30
5	1550-1750	30
6	10400-12500	60
7	2090-2350	30
8	520-900	15

Table 3.1: The band designations and resolution of LANDSAT 7.

false-color composite) have been merged with the grayscale panchromatic band (band 8) at $15m$ (i.e. pan-sharpened). As the LIMA website [52] points out, the pan-sharpening maintains the false-color information while keeping the high resolution of $15m$. Through resampling techniques of the pixels, a $15m$ spatial resolution image is produced. The infrared component offers a better discrimination between ice and snow [52]. Figure 3.4 shows a single LANDSAT image displayed as a false-color composite (left) and as a normal color composite (right). The difference in snow and ice is improved by incorporating the near-infrared band, which results in the icy part being displayed as more blue-colored.

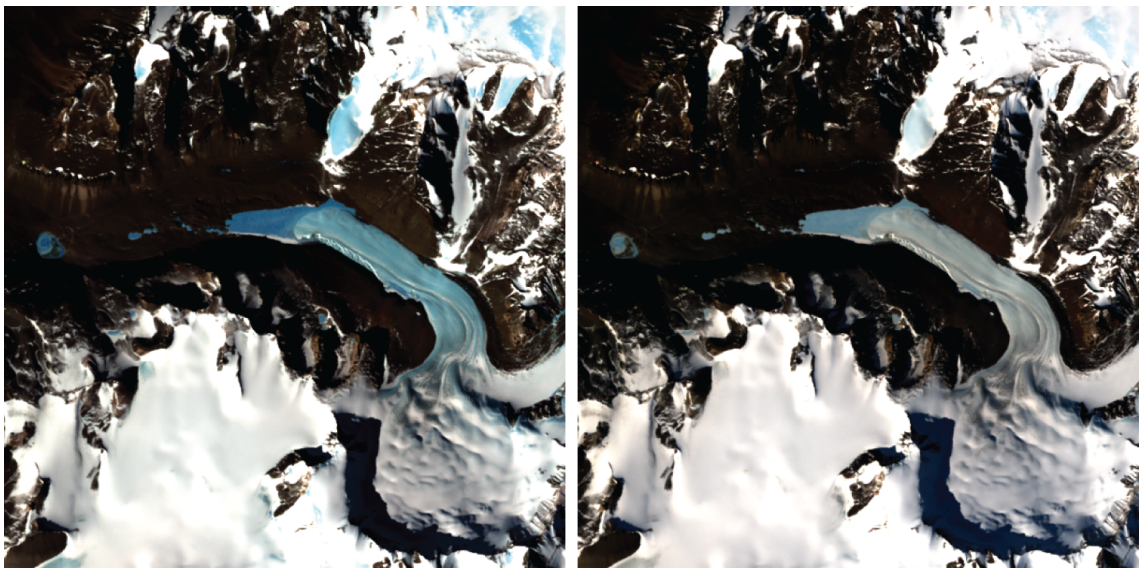

Figure 3.4: The same scene acquired by LANDSAT, using the false-color composite in the left image and the normal color composite in the right image. The images demonstrate the improved discrimination between snow (white) and ice (blue) using false-color composites. The scene shows an inlet glacier invading Barwick Valley with lake ice at the tip of the glacier and snow in the higher elevated areas, taken on the 13th of November 2008. Data taken and adapted from USGS EROS [49].

The entire spatial subset of the false-color, pan-sharpened LIMA image of the Dry Valleys that is used throughout this research for classification analysis, is in fact shown in

Figure 3.2. It shows that ice parts like Taylor Glacier and the lakes are easily identifiable. This image, without the annotations, serves as the primary validation image for classification purposes of the Dry Valleys. Its acquisition date is the 12th of November 2001, which is well before the launch of ICESat. The image contains no cloud cover and offers, as mentioned earlier, $15m$ resolution with improved discrimination between ice and snow through the incorporation of the near-infrared band. As will be explained in Chapter 7, this LIMA image is used for generating a classification map for validation purposes. The term 'LIMA' is used throughout this report to refer to the Dry Valleys LANDSAT image taken on the 21st of November, 2001. This reference makes it easier to avoid confusion with LANDSAT images taken at other dates.

Besides the single LIMA image, LANDSAT images acquired at later dates are available as well, corresponding better to ICESat's laser campaign dates. However, after May 2003, all images are heavily affected by the faulty Scan Line Corrector (SLC). The SLC is an electromechanical device correcting for the forward motion of the satellite [50]. The defect results in data gaps appearing as black lines running though the images. Figure 3.5 shows an example LANDSAT image taken of the Dry Valleys after the 2003 incident. It is clearly visible that the black stripes form an obstruction for automatic classification. However, the affected images are still suitable for validation purposes, because a significant part of the land cover is still visible. It is especially relevant for visually validating the surface state on different locations. For many occasions, the acquisition date of the LANDSAT images does not coincide with the acquisition dates of the specific ICESat shots, so assumptions have been made with respect to the land cover type. However, for some occasions, LANDSAT images coincide almost with ICESat acquisition dates and can be used for complementary validation.

The reason that later LANDSAT images are not used for the construction of multiple supervised classification maps is because of the date differences and the black stripes. The images serve only for post-classification analysis. The LIMA and LANDSAT images are displayed in polar stereographic projection, which is a conformal projection that assumes the projection plane is tangent to the Earth at the pole. The coordinates are defined in the WGS84 datum and WGS84 ellipsoid.

The individual LANDSAT images consists of the eight individual bands, in contrast to the LIMA image which is a pre-computed false-color, pan-sharpened image. For the post-classification analysis using LANDSAT images, the second (green), third (red) and fourth (near-infrared) bands are used. No pan-sharpening is carried out, because the images serve for manual post-classification purposes only, instead of being used for a high-resolution supervised classification map.

To summarize the difference between the LIMA and LANDSAT images: the LIMA image serves as the main validation data set for classification and the LANDSAT images serve as a complementary validation data set in time. Therefore, the supervised classification map that is estimated from the LIMA image only represents the land cover situation

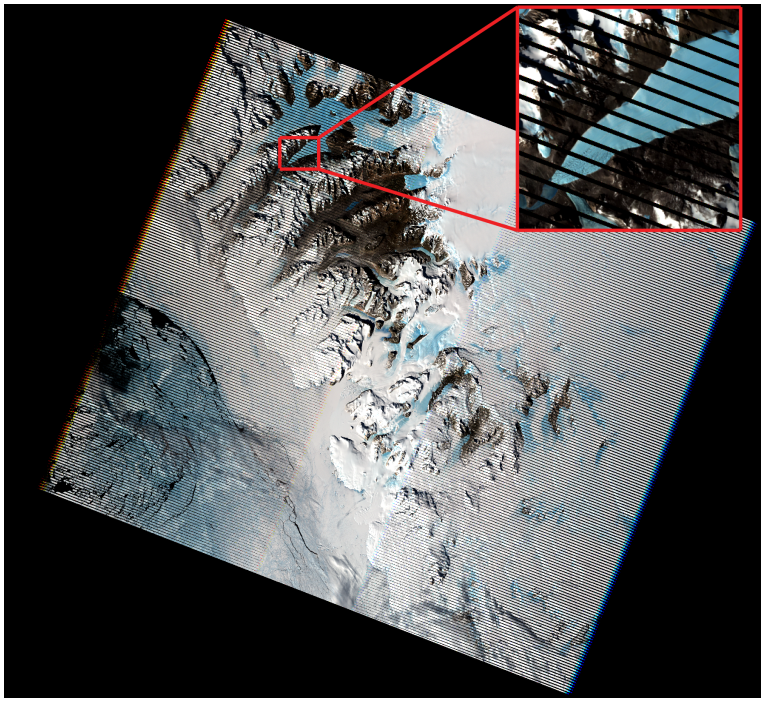

Figure 3.5: A LANDSAT image of the Dry Valleys and surrounding area. The width the black lines resulting from the SLC defect increase from center of the image towards the edge. The inset shows a close-up of the Taylor Glacier, where the effects of the black lines are clearly visible. Data taken and adapted from USGS EROS [49].

at one instance in time. To minimize the temporal effect as much a possible, only those ICESat campaigns are selected for classification analysis that were carried out in approximately the same calendar month as the LIMA acquisition month, November. Since the Dry Valleys form a stable region where little change in land cover type takes place, see Section 3.1, it is assumed that the single LIMA image can be used as a primary validation data set. In addition, regions of the Dry Valleys that are assumed to be less stable are left out of the analysis. Chapter 7 explains the methods to deal with temporal inconsistencies in more detail.

3.2.2 LTER

In addition to the LANDSAT images, another complementary validation data set will be used as well, which is offered by the Long Term Ecological Research (LTER) project. As explained before, post-classification analysis of the ICESat waveforms sometimes requires validation at specific dates. For post-classification validation, LANDSAT images are used in combination with LTER data of the Dry Valleys. The LTER data set consists of meteorological data acquired by weather stations in the Dry Valleys region. Figure 3.2 shows the locations of the three weather stations, highlighted in red, that are used for this research.

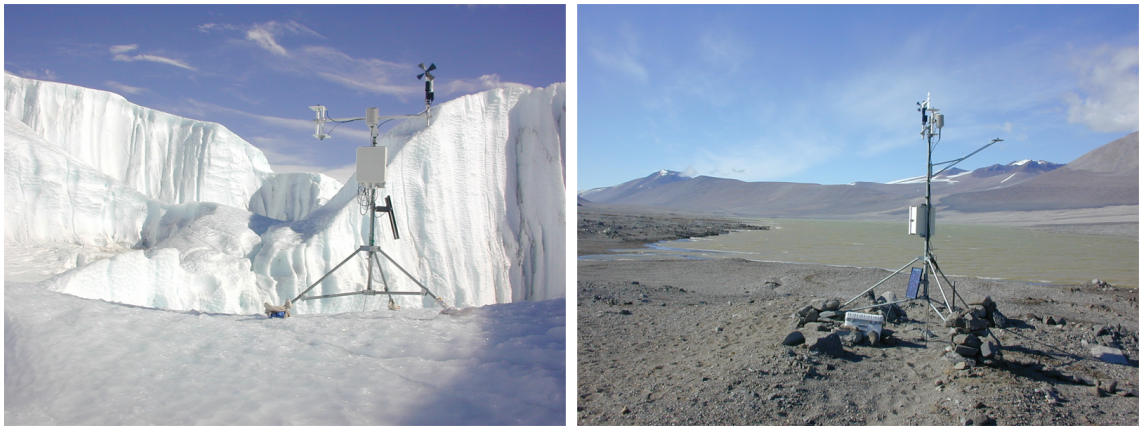

Figure 3.6: Two LTER meteorological stations in the Dry Valleys. The left picture shows a station on Taylor Glacier, the right picture on Lake Vida. Pictures taken and adapted from LTER [17]. ©2006 McMurdo Dry Valleys LTER, both pictures courtesy Thomas Nylen.

These stations are located at two lakes and one glacier. Figure 3.6 shows two examples of meteorological stations used in the Dry Valleys.

The meteorological stations acquire a wide range of measurements every 15 minutes. At the two lakes (i.e. Lake Vanda and Lake Vida), the measurements that are relevant for this research are air temperature, soil temperature and snow height/soil surface change. The measurements that are used at the Taylor Glacier station are air temperature, ice surface temperature and ice surface change. For measuring the air temperatures a scientific temperature probe is used, at $3m$ height for the lake stations and at $1m$ height for the glacier station. The ice temperatures are measured with an infrared thermometer system. For the lake stations the soil temperatures are measured at the soil surface and the glacier station measures the temperatures on the ice surface.

For the lake stations, ultrasonic distance rangers are used for snow height/surface change measurements. From correspondence with the Department of Geology at Portland State University [11], the exact method of measuring these changes can be explained as follows. The distance sensors are mounted directly on the weather station tripods at a fixed height above the surface. The sensors measure the distance to the surface in meters which is then subtracted from the original measured height. As a result, the height is relative to the original measured height. The sensors have an accuracy to $1cm$. Since the ground surface changes over time because of blowing sediment, the values could become slightly negative. In case of the weather station on glacier, the tripods are mounted on wooden pods which are drilled into the ice and are consequently fixed to the ice. Each five years these station pods are reset. Even though the measurements are not available for all years of the ICESat campaigns, the data set offers a valuable source of information for post-classification analysis. A major drawback however is that the information is very local.

3.3 Surface topography validation data

For surface topography analysis of ICESat waveforms, validation data of surface elevation is required. In December 2001, an extensive survey by NASA's Airborne Topographic Mapper (ATM) of the Dry Valleys resulted in several digital elevation models (DEM). The surface elevations were estimated using an airborne laser altimetry system, after which a DEM grid was interpolated through the laser points.

In total 18 ATM DEMs were produced, the majority having $2m$ and others having $4m$ spatial resolution. For this research, the DEMs of Barwick Valley, Victoria Valley, Mckelvey Valley, Bull Pass, Wright Valley and Taylor Glacier are used as validation data sets. All six DEMs offer elevations estimates at a spatial resolution of $2m$, measured relative to the WGS84 reference ellipsoid [42]. The accuracy of the DEMs was verified by using independent GPS data obtained by on-ground stations. Interpolated DEM elevations at the position of the GPS points were compared with the GPS elevations for the accuracy analysis [42]. Approximate elevation differences at the GPS points are on the order of centimeters to decimeters, which is acceptable considering the vertical accuracy of ICESat of $15cm$ at best and the validation purpose of the DEMs for this research.

Figure 3.7 shows the six DEMs that are used for this research. The names in the figure correspond to the names in Figure 3.2. As the figure shows, the DEMs do not cover the areas homogeneously. During the measurement campaigns, only elevation measurements were taken of those areas that were not covered by snow. The most plausible reason for this is that elevations of the snowy parts are likely to change. The areas that are measured cover the more stable lower elevated rock and ice parts of the Dry Valleys. However, even in these parts, gaps in the DEMs are present, most likely due to non-overlapping ground tracks and outliers that were already removed by the provider of the data [42].

The DEM observations are defined in the WGS84 datum and WGS84 ellipsoid and are map projected using the Lambert Conformal Conic projection. As the projection name indicates, the projection is conformal and consequently preserves angles. Although this projection is different from the projection of LIMA and LANDSAT, the observations of both data sets are defined in the same datum and reference ellipsoid. As a consequence, the data sets can be compared in geographical coordinates without any coordinate transformation. However, when comparing observations in map projected coordinates one should take care of the projection differences. The classification analysis and topography analysis are carried out separately, meaning ICESat footprints are projected according to the projection of the data set that is analyzed. Thus, when analyzing classification, the footprints are projected in the polar steographic projection; when analyzing topography, the footprints are projected in the Lambert Conformal Conic projection.

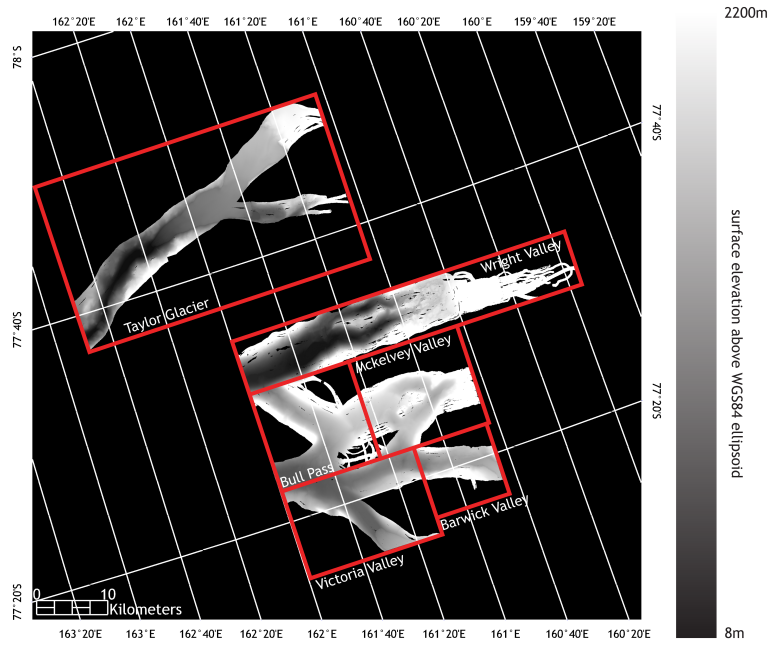

Figure 3.7: The six DEMs that are used for elevation validation, reprojected into polar stereographic projection for direct matching with the LIMA image. The red lines mark the boundaries of each DEM. Data taken and adapted from USGS Atlas of Antarctic Research - United States Antarctic Resource Center [51].

Chapter 4

Data product derived attributes

For this research two sets of attributes are used: data product derived attributes and waveform derived attributes. Both sets are used for filtering and correction techniques, and subsequently used for surface topography and classification analysis. In this chapter, the data product derived attributes will be discussed; in Chapter 5, the waveform derived attributes will be discussed. Data product derived attributes are all attributes that do not directly relate to waveforms. The data products contain many attributes, including raw satellite housekeeping values, pre-computed elevation estimates, quality flags and attribute corrections. Before relevant data attributes can be used for this research, various filters and corrections are applied to get reliable data. This is especially important considering the sensitivity of unwanted influences on the waveforms and the importance of characteristic parameters.

This chapter is divided into three parts: an overview of the data product derived attributes used for this research (Section 4.1); the filtering techniques (Section 4.2) and the correction techniques (Section 4.3); and finally the transmitted energy analysis (Section 4.4) and atmospheric channel analysis (Section 4.5). The three chapters 'ICESat principles', 'Data product derived attributes' and 'Waveform derived attributes' (Chapter 2, Chapter 4 and Chapter 5, respectively) contribute to the yellow colored box 'Pre-filtered and corrected ICESat footprints and waveforms' of the flow diagram of Figure 1.2. Figure 2.3 can be used as a guide for these chapters. This flow diagram is especially focused on the retrieval, filtering and correction techniques on the data.

4.1 Overview of data product derived attributes

The data product derived attributes that are used for this research consist of descriptive values and corrections. They are selected in agreement with research literature and in correspondence with experts on ICESat data analysis.

The descriptive values and corrections used for this research come from the higher-

level processed data product Global elevation data (GLA06) and Global backscatter data (GLA07). The higher-level processed attributes include the precise geolocations of the footprints and beam geometry attributes. The initial data filtering is based on GLA06, because this product offers the most relevant attributes for spatial subsetting and indications of bad data. As depicted in Section 2.3, this data product is commonly used for ice sheet elevation analysis. After filtering and applying the corrections, the selected unique record index numbers and shot numbers (combined called the *ShotID*) are used to extract the actual waveforms from GLA01. Only a fraction of the variables available in the GLA06 product files have been used for data filtering, because not all are relevant for the research purposes discussed here.

To give a complete overview of the data product derived attributes, Table 4.1 not only provides the attribute name and official ID, but also the corresponding data product and the frequency rate in the data product (i.e. $1Hz$ or $40Hz$). The official ID of the original attribute is a reference to the data product description of the NSIDC website [35]. The relevant attributes of GLA07 are also included in the table; these are discussed in Section 4.5. All data product derived attributes are extracted from the data product binary files using IDL code and software (see IDL website [40] for more information on IDL). The original code has been taken from the NSIDC website [34] and modified to serve the purpose of this research.

Derived attribute ID	**Data product**	**Frequency**	**Attribute name**
shotID	GLA01/GLA06/GLA07	40Hz	record index + shot number
Original attribute ID	**Data product**	**Frequency**	**Attribute name**
i_lat	GLA06	40Hz	latitude of footprint center [degrees north]
i_lon	GLA06	40Hz	longitude of footprint center [degrees east]
i_elev	GLA06	40Hz	surface elevation [meter]
i_tpazimuth_avg	GLA06	1Hz	transmitted pulse azimuth - frame average [degrees]
i_tpeccentricity_avg	GLA06	1Hz	transmitted pulse eccentricity - frame average [-]
i_tpmajoraxis_avg	GLA06	1Hz	tramsitted pulse major axis - frame avg [meter]
i_beam_coelev	GLA06	40Hz	beam co-elevation [degrees]
i_beam_azimuth	GLA06	40Hz	beam azimuth [degrees]
i_gdHt	GLA06	2Hz	height of the geoid above the ellipsoid [m]
i_DEM_elv	GLA06	40Hz	SRTM30 DEM height above the ellipsoid [m]
i40_g_bcsc	GLA07	40Hz	532nm merged attenuated backscatter profile from 10.3km to -1.0km [1/(m·sr)]
i40_ir_bscs	GLA07	40Hz	1064nm attenuated backscatter profile from 10.3km to -1.0km [1/(m·sr)]
i40_g_TxNrg_EU	GLA07	40Hz	532nm laser transmitted energy [1e-5 J]
i40_ir_TxNrgEU	GLA07	40Hz	1064nm laser transmitted energy [1e-5 J]

Table 4.1: The data product derived attributes that are used for the research. All distances are relative to the geoid. The frequency represents the rate per second of attribute occurrence in the specified data product. Units of the parameters are indicated between brackets.

The detailed descriptions of the attributes of Table 4.1 will be given throughout the report; please refer to the index at the end of this report for the specific page numbers. The table makes a distinction between derived and original attributes. The derived attribute,

shotID, is not extracted directly from any data product, but is a combination of the original attributes unique record number and shot number. As a result, every individual shot has its unique number that can be used for referencing shots between different data products. The original attributes are directly extracted from GLA06 and GLA07 data products.

4.2 Filtering techniques

The filtering techniques that are applied on the data product derived attributes consist of data subsetting in the temporal and spatial domain, and of invalid value detection.

The geolocation parameters, *i_lat* and *i_lon*, describe the center of the laser footprint in geodetic latitude and geodetic longitude, respectively. The surface elevation, *i_elev*, is the surface elevation with respect to the Topex/Poseidon ellipsoid. The three attributes have been computed after precise orbit and attitude determination, and after instrument corrections, atmospheric delays and tides have been applied. Furthermore, the coordinates are computed with the the standard, i.e. ice-sheet, algorithm. See Section 2.3 for more information about this algorithm. Before using these coordinates, an atmospheric correction, and a datum and ellipsoid transformation are applied for correct assimilation with the validation data sets. Section 4.3 explains these correction techniques in more detail. The corrected footprint geolocations are used for data subsetting in the spatial domain. The derived attribute *ShotID* is used for temporal subsetting.

The validity of important attributes is, next to the spatial and temporal filtering, the only filtering that is discussed in this chapter. The reason that measurements are filtered out if an important attribute has an invalid value is simply because it cannot be used for the waveform analysis. Particularly geolocation of the footprint and beam geometry attributes are essential for analysis of the data and consequently all shots are filtered out that do not hold valid values for these attributes. This means that *i_lat*, *i_lon*, *i_elev*, *i_beam_coelev*, *i_beam_azimuth*, all should represent valid values. Generally speaking, these attributes are invalid when the elevation estimation was not carried out successfully. The reason that the other attributes like footprint geometry and atmospheric backscatter attributes are not filtered out when invalid, is because these are not of highest importance for the analysis: the footprint geometry can usually be derived from neighboring shots; the atmospheric channels attributes are used for classification analysis only and consequently their validity is analyzed separately.

As the flow diagram in Figure 2.3 shows, more filtering techniques will be described in Section 5.4. These are not discussed here, because they relate to waveform derived attributes.

4.3 Correction techniques

To make the data suitable for analysis, three corrections are applied. All corrections relate to the characteristics of the laser footprints. The first correction is the replacement of bad footprint geometry values for a particular laser campaign, the second is the correction of the footprint geolocation due to atmospheric refraction correction, and the third is the datum and ellipsoid transformation of the footprint geolocations.

4.3.1 Footprint geometry correction

The attributes *i_tpazimuth_avg*, *i_tpeccentricity_avg* and *i_tpmajoraxis_avg* (transmitted pulse azimuth angle, eccentricity and major axis, respectively) all relate to the footprint orientation and geometry and are embedded in the GLA06 data product. Figure 4.1 illustrates how these attributes define the footprint geometry. The values of the attributes are based on so-called laser profiling array (LPA) images. The figure also shows some examples of these images (up to laser campaign L3B).

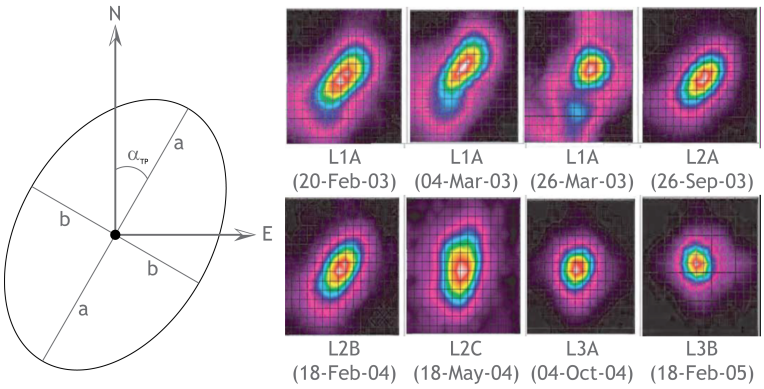

Figure 4.1: The laser footprint geometry is described by the GLA06 attributes transmitted pulse azimuth (α_{TP}), the major axis (a) and eccentricity, see the left part. The arrows pointing in north and east directions are included for orientation purposes. The footprint geometry is determined based on LPA images. Some example LPA images up to laser campaign L3B are also shown (right), courtesy NASA/ICESat.

In Figure 4.1, a is the semi-major axis, b is the semi-minor axis and α_{TP} is the transmitted pulse azimuth angle. Since the major axis and the eccentricity, e, are provided in the GLA06 product, the semi-minor axis can be computed with the following equation.

$$b = a\sqrt{1 - e^2} \tag{4.1}$$

The transmitted pulse azimuth angle is defined as the angle clockwise from north of the major axis of the transmitted pulse as seen by the LPA. LPA images are recorded at $40\,Hz$

by the GLAS instrument when transmitting the laser pulse. It is a far-field projection of the laser spot that is used for footprint geometry and orientation derivation. Pixels in the array with intensity more than $\frac{1}{e^2}$ ($\sim$13.5%) of the maximum pixel intensity are used for the footprint size computation [31], where e is Euler's number. The LPA images in Figure 4.1 illustrates the non-uniform power distribution within the footprint and is, for this research, approximated by a two dimensional Gaussian distribution. From correspondence with ICESat data experts at the CSR [44], it turns out that the GLA06 data product of laser campaign L3A (Release 28) holds incorrect values for the footprint geometry parameters. As a consequence, these values are replaced by corrected values provided by the CSR.

4.3.2 Geolocation correction

Another correction is the footprint geolocation correction due to atmospheric refraction. This is of particular concern for large off-nadir pointing angles of the laser beam, due to the bending of the ray in the atmosphere [43]. Figure 4.2 illustrates the geometry of the laser ray path, both undisturbed by refraction (straight line from satellite to *P1*) and disturbed by refraction (dashed line from satellite to *P2*). The reason that the path bends more at lower elevations is because of an increase in density of the atmosphere near the Earth's surface.

In Figure 4.2, *P1* is the expected ground location (not taking refraction into account) and *P2* the real ground location (taking the refraction into account). The figure shows that the footprint is shifted by a distance towards the sub-satellite point (i.e. the point at nadir, where the dotted line would intersect the Earth's surface) in the horizontal plane. This distance is computed with Equation 4.2.

$$d = h \cdot (\tan(\beta) - \tan(\beta - \delta\beta)) \tag{4.2}$$

Here, d is the distance between the expected (*P1*) and the real (*P2*) ground location of the footprint, h is the satellite's altitude above ground, β is the beam co-elevation angle in degrees and $\delta\beta$ is difference between the expected beam co-elevation angle in degrees (related to *P1*) and the real beam co-elevation angle in degrees (related to *P2*). This difference is approximated by Equation 4.3 assuming standard surface conditions [43].

$$\delta\beta = 0.016^\circ \cdot tan(\beta) \tag{4.3}$$

As a result, when taking $600km$ as the satellite's altitude and taking the nominal beam co-elevation angle of 0.3°, the horizontal shift is about $0.9m$ which is insignificant considering the geolocation accuracy of about $4.5m$ [14] and the much larger footprint dimensions. However when taking the maximum off-nadir beam co-elevation angle of 5°, the shift increases to about $14.7m$, which is significant. Even though the shift at the nominal angle is well within the geolocation accuracy, the refraction correction is taken into account for all footprints to make sure that the locations of the footprints are as accurate as possible.

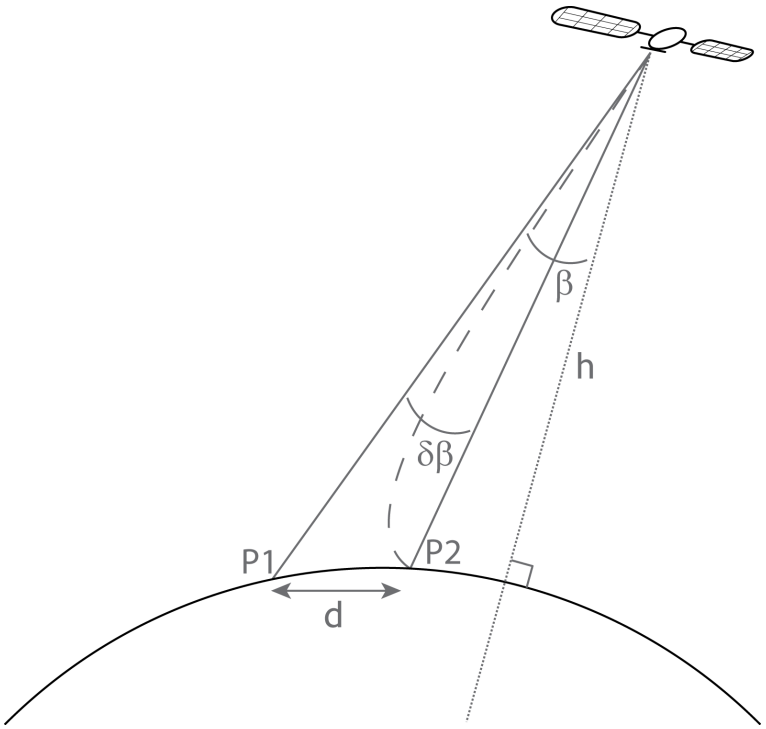

Figure 4.2: The geolocation difference of the footprint due to atmospheric refraction of the laser ray path. The intersection of the dotted line with the Earth's surface represents the sub-satellite point. The dashed line represents the real path, the solid line from satellite to *P1* represents the undisturbed path. Adapted from GLAS Algorithm Theoretical Basis Document (ATBD) about atmosphere delay correction [43].

4.3.3 Datum and ellipsoid transformation

The third and final correction relates to the datum and ellipsoid transformation of the footprint's geolocations. ICESat's laser footprints are expressed in geodetic latitude, longitude and elevation. These coordinates are referenced to the Topex/Poseidon ellipsoid. The validation data sets (see Chapter 3) that are used for this research are usually described in the WGS84 datum with the WGS84 reference ellipsoid. When merging the ICESat data with the validation data for analysis, both should be described in the same reference frame to avoid position inconsistencies. Since WGS84 is a more commonly used reference frame, the ICESat geolocated footprints are transformed into this reference frame.

The measurements that ICESat carries out are defined in the International Terrestrial Reference Frame (ITRF) with the Topex/Poseidon reference ellipsoid [5]. The reason that the ICESat footprints are described in this particular reference ellipsoid is to enable easy comparison with the Topex/Poseidon and Jason-1 altimetry missions [37], which both work with this rather specialized reference ellipsoid. Table 4.2 compares the characteristic values of the Topex/Poseidon and WGS84 ellipsoids. For the transformation, it is assumed that the difference between the center of the Earth defined by the ITRF used for the WGS84 and

Topex/Poseidon ellipsoids is negligible, taking into account the accuracy of the geolocation.

Parameter	Topex/Poseidon	WGS84
Equitoral radius	6378136.300000	6378137.000000
Polar radius	6356751.600563	6356752.314245
Reciprocal flattening	298.25700000	298.25722356
Eccentricity	0.081819221456	0.081819190843

Table 4.2: The characteristic values of the Topex/Poseidon and WGS84 reference ellipsoids, taken from NSIDC [37].

The Topex/Poseidon ellipsoid is $70cm$ smaller than the WGS84 ellipsoid. As a consequence, the same point on the Earth's surface will have different geodetic coordinates when using the different ellipsoids. As the NSIDC [37] points out, the horizontal shift is only a few centimeters, which is neglectable considering that the GLAS geolocation accuracy in the horizontal direction is on the order of a couple of meters. However, the dominant difference is in geodetic height (elevation) and should be accounted for when comparing with digital elevation models defined in the WGS84 datum. To make sure ICESat data can be properly validated against other data sets, all sets are defined in the WGS84 datum. For ICESat's footprint geolocations and some validation data sets, ellipsoid transformations are carried out.

To transform the geolocations in WGS84, a standard script provided by the CSR is used. Part of the script uses the theory described by Tscherning et al., [7]. First, the script converts the geodetic coordinates to Earth fixed coordinates, expressed in Cartesian coordinates (x,y,z), based on the Topex/Poseidon ellipsoid semi-major axis and flattening parameters. Subsequently, the Cartesian coordinates are transformed to geodetic coordinates using a first order Taylor series expansion with the WGS84 ellipsoid parameters. The process to estimate the geodetic coordinates is iterated as long as the difference between the estimations is above certain thresholds. Finally, the geodetic coordinates in latitude, longitude and height referenced to the WGS84 datum can be used in combination with the validation data sets.

4.4 Transmitted energy analysis

The satellite was designed to transmit green and infrared signals at a constant energy level over ICESat's lifetime. For the green signals, the instrument was designed to transmit pulses at $35mJ$; for the infrared signals, at $75mJ$. However, as was explained in Chapter 2, an unexpected defect in the laser's pump diode caused the energy to decrease over time. This design fault even ended in premature failure of the lasers. In the classification analysis, this varying transmitted energy can have considerable effects on attributes.

Figure 4.3 shows the transmitted energy values of both the green and the infrared laser over the Dry Valleys. These values are extracted from the GLA07 data product that holds the transmitted energy at $40Hz$ for both wavelengths, i.e. the attributes *i40_g_TxNrg_EU*

and *i40_ir_TxNrgEU* of Table 4.1. It illustrates the high variability of the energy between campaigns and during campaigns. It also shows the relatively small operational period of the first laser compared to the other two. As for the green laser signals, the energy decreased from about $35mJ$ to less than $5mJ$ in just two campaigns using the second laser. Apparently, for the third laser system it was decided to begin the measurements with a lower energy level to extend its lifetime. The right plot of the figure shows the dramatic decrease in transmitted energy for the infrared lasers, especially for the L2C campaign. Even though the energy levels for L2C approach 10% of the initial level, they remain substantially higher than the energy of the green laser. For some instances, the plots show large short-period fluctuations, each of these concerning individual shots and represent a very small minority of the total values.

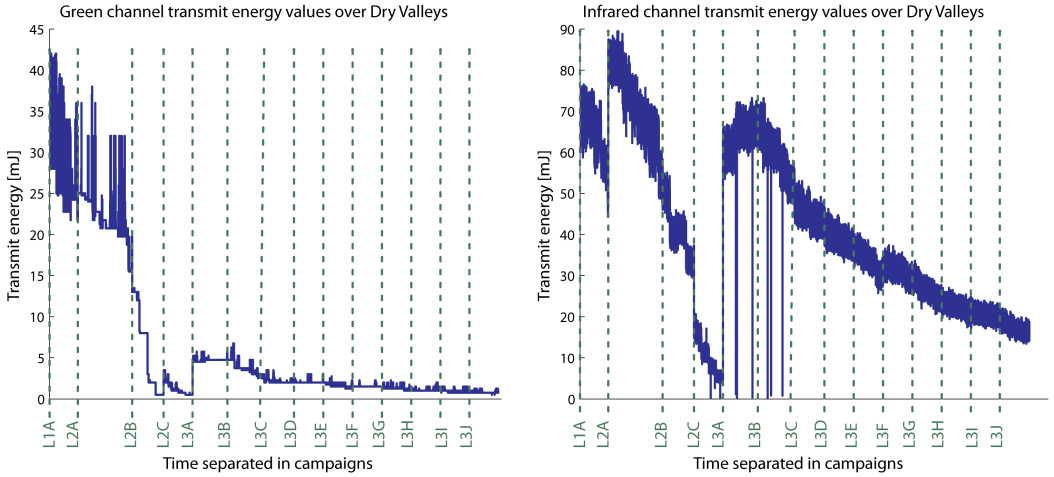

Figure 4.3: The transmitted energy values of the green and infrared signals over the Dry Valleys, separated by campaign.

4.5 Atmospheric channel analysis

The GLAS instrument transmits laser pulses in the green ($532nm$) for atmospheric measurements and in the infrared ($1064nm$) for altimetry measurements as well as for atmospheric measurements. For analyzing the classification potential of ICESat's attributes, not only is the infrared returned signal that is recorded by the altimetry channel (i.e. the full waveform) taken into account, but also the ground return signals of the atmospheric channels. In this section, the relevance of the atmospheric signals will be discussed first. In addition, the meaning of the atmospheric signals will be explained by discussing how the signals are computed. Finally, the retrieval of the atmospheric signals at ground level will be addressed. These ground return signals are used for classification analysis later in this research.

4.5.1 Relevance of atmospheric signals

Incorporating the infrared band in the LIMA image significantly improves the discrimination between snow and ice, as was discussed in Section 3.2. In Figure 3.4 it is clearly visible that incorporating the infrared component makes the the icy parts more blue-colored than the snowy parts, while the normal color composite displays the icy parts more white and more similar to snow. Since GLAS transmits and receives laser pulses both in the infrared and in the green, the two bands could be used to construct a spectral signature of each land cover type. The relevance of incorporating the green channel for classification analysis is demonstrated with Figure 4.4.

In the infrared band (left image), the icy parts and bare rock are darker (lower reflectance) and are easy to distinguish from the snowy parts that are brighter (higher reflectance). However, the ice and bare rock areas are more difficult to distinguish from each other due to their low reflectance. Moreover, the small icy lake left of the glacier is nearly invisible from the surrounding bare soil. Looking at the same area in the green (right image), the difference between ice and bare rock is more clear. For example, the lake ice is clearly distinguishable from the bare rock surroundings. In addition, the glacier that invades the valley shows high reflectance compared to rock as well. On the other hand, the difference in reflectance in the green between ice and snow is less clear and are consequently harder to distinguish from each other.

From the figure it is clear that ice shows low reflectance in the infrared and high reflectance in the green. Thus, combining the reflectance of a laser pulse in the infrared with the reflectance of a laser pulse in the green might result in an improved classification algorithm.

4.5.2 Explanation of atmospheric signals

For reflectance of the infrared pulse, the parameter *i_reflctUncorr*, 'reflectivity not corrected for atmospheric effects', is used for classification analysis. See Section 5.1 for more information on this attribute. However, a similar attribute in the data products for the green signal does not exist. To still incorporate the green signal in the classification analysis, the green atmospheric backscatter profiles are used for classification analysis. In addition, since the method of extracting the backscatter profiles is similar for the infrared signals, the infrared atmospheric backscatter profiles will be incorporated as well. Both profiles, *i40_g_bscs* and *i40_ir_bscs* for the green and infrared signals respectively, are available in the Level-1B atmospheric data product GLA07, see Table 4.1.

Both frequencies are measured with a vertical resolution of $76.8m$, but differ in the way they are measured. The $532nm$ signals are measured with photon counting modules; the

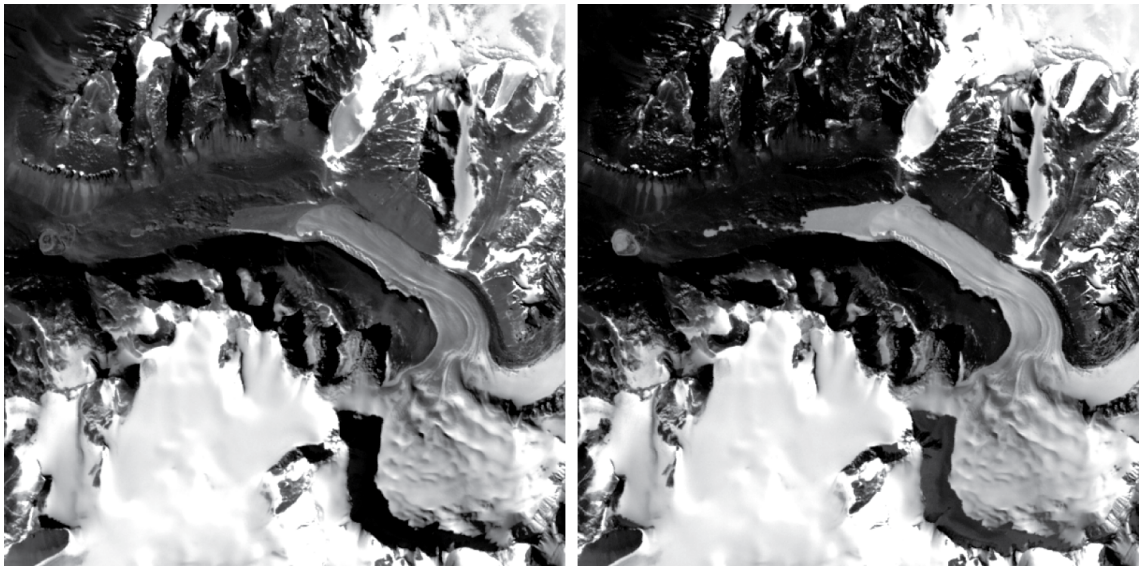

Figure 4.4: The same scene acquired by LANDSAT, using the infrared band in the left image and using the green band in the right image. The images demonstrate the spectral difference between the different land cover types. It is the same scene as Figure 3.4. Data taken and adapted from USGS/EROS, [49].

1064nm signals are measured with an analog-to-digital converter. The photon counting module measures the green signal in photons; the analog-to-digital converter measures the signal in digital counts, which are converted back to voltages using a pre-defined look-up table [33].

Table 4.3, together with the attenuated backscatter profiles of Table 4.1, shows the most important attributes that are used for the detailed discussion on the atmospheric signals later in this section. These attributes of Table 4.3 are not extracted from the data products, but only serve for discussion purposes.

Original parameter ID	**Data product**	**Frequency**	**Parameter name**
i40_g_lid	GLA02	40Hz	532nm (normalized) lidar data from 10.3km to -1.0km [(photons/bin)· km^2/J]
i40_ir_lid	GLA02	40Hz	1064nm (normalized) lidar data from 10.3km to -1.0km [(W· km^2)/J]
i_g_cal_cof	GLA07	1Hz	532 nm backscatter calibration coefficient [(photons/bin)·(km^3/J)·sr]
i_ir_cal_cof	GLA07	1Hz	1064 nm backscatter calibration coefficient [W· km^3/J·sr]

Table 4.3: Attributes relevant for the discussion on the green and infrared backscatter at ground level. These are not extracted from the data products. All distances are relative to the geoid. The frequency represents the rate per second of attribute occurrence in the specified data product. Units of the attributes are indicated between brackets.

The units of the backscatter profiles are $[m^{-1}sr^{-1}]$ where sr stands for steradian, the International System of Units (SI) unit of solid angle measure [36], see Table 4.1. The profiles basically represent the backscatter intensity per spatial column unit. The solid angle (in steradian) is the ratio between the fractional area subtended by a spherical surface of the beam at range r (i.e. A_{fr}, which is a spherical cap) and the total area of the sphere at r (i.e. A_t), multiplied by 4π. In Figure 4.5, the range and the areas are illustrated. In this explanation, r is the radius of the sphere, i.e. the distance from the point where the beam originates to the fractional area. The ratio is multiplied by 4π, such that the solid angle of a sphere measured from a point in the sphere is $4\pi[sr]$: $[4\pi r^2/(4\pi r^2)] \cdot 4\pi = 4\pi[sr]$. It is a measure of the narrowness of the beam. For example, consider a circular footprint on the ground surface with radius $50m$ and assume its flat surface area is approximately the same as the area of sphere that it occupies (i.e the cross-section is approximately equal to the spherical cap). Taking a distance from satellite to ground surface of $600km$ results in: $[A_{fr}/A_t] \cdot 4\pi = [\pi \cdot 50^2/(4\pi \cdot (600 \cdot 10^3)^2)] \cdot 4\pi = 2.18 \cdot 10^{-8} sr$. The solid angle of ICESat's transmit pulse is larger in reality due to the fact that the laser pulse actually illuminates a larger area. The reason that the surface area is larger is because the footprint diameter of $50m$ is based on an intensity threshold of a LPA image, see Section 4.3.

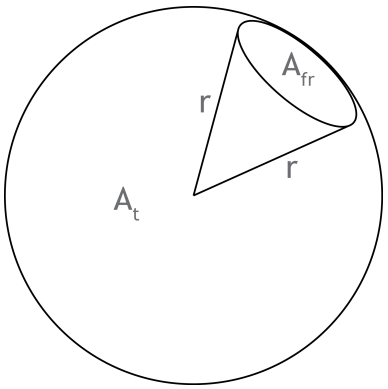

Figure 4.5: An illustration of the computation of the solid angle. The fractional area, A_{fr}, is a spherical cap and the total area, A_t, is the area of the sphere.

In the following paragraphs an explanation will be given of how the attenuated backscatter profiles, *i40_g_bscs* and *i40_ir_bscs* of Table 4.1, are computed using, together with other variables, the attributes of Table 4.3. Please note that these attenuated backscatter profiles are readily available from the GLA07 and consequently the following discussion is not strictly required for understanding how the profiles have been acquired. The discussion is only meant for the reader to understand how the profiles have been computed by ICESat's data product algorithms. By providing this background information, the meaning of the profiles can be understood better. The information on the attenuated backscatter profiles profiles has been taken and adapted from the GLAS ATBD of GLAS Atmospheric Data Products [33].

Before computing the attenuated backscatter profiles, the meaning of the normalized signals, *i40_g_lid* and *i40_ir_lid*, will first of all be explained. The computation principle of these signals is the first step towards the computation of the attenuated backscatter profiles. For both $532nm$ and $1064nm$ signals the normalization of the raw data consists of applying fundamental corrections and range and laser energy normalization. Although the raw signals of both frequencies are subject to essentially the same normalization algorithm, both require a slightly different approach. The normalization of $532nm$ channel will be discussed first. The signal is normalized for range and laser energy and is corrected for the ambient background signal and dead time, see Equation 4.4. Simply put, the dead time correction is the correction for non-continuous recording of the signal. The second part of the equation is important for understanding the normalization procedure. The third part is important for understanding the meaning of the green backscatter profile, *i40_g_bcsc* of Table 4.1.

$$P'_{532}(z) = \frac{(\mathrm{DC}[S_{532}(z)] - \mathrm{DC}[B_{532}])R^2}{\alpha E_{532}} = C_{532}\beta_{532}(z)T^2_{532} \tag{4.4}$$

Here, P'_{532} is the green normalized signal at an altitude z (i.e. *i40_g_lid*). The second part consists of the dead time correction (DC) of the measured signal $S_{532}(z)$ and the background signal, $B_{532}(z)$, the range R from the satellite to altitude z, the ratio α of the input and output energy passing through the same bandpass filter (i.e. a measure for filter transmission coefficient), and the transmitted energy before the bandpass filter E_{532}. In other words, the normalization consists of subtracting the background from the atmospheric return signal, then multiplying by the square of the range and finally dividing by the effective transmitted energy. This second part shows which measurements are required for computing the normalized signal. The third part of the equation basically combines the second part of the equation in other variables. Here, C_{532} is the lidar calibration constant, β_{532} is the total atmospheric backscatter cross section at altitude z and T^2_{532} is the two-way transmission of the top of the atmosphere to altitude z.

The algorithm for the $1064nm$ channel is essentially the same, using slightly different variables. Equation 4.5 shows the algorithm for computing the infrared normalized signal.

$$P'_{1064}(z) = \frac{D_r(V_s(z) - V_b)R^2}{E_{1064}} = C_{1064}\beta_{1064}(z)T^2_{1064} \tag{4.5}$$

Here, P'_{1064} is the infrared normalized signal at an altitude z (i.e. *i40_ir_lid*), D_r is the detector responsivity factor (which is a constant that converts Volts to Watts), $V_s(z)$ is the return signal at altitude z, V_b is the background signal (both converted from digital counts to Volts), and E_{1064} is the laser transmitted energy. The third part of the equation follows the same principle as before, where C_{1064} is the lidar calibration constant, β_{1064} is the total atmospheric backscatter cross section at altitude z and T^2_{1064} is the two-way transmission of the top of the atmosphere to altitude z.

As depicted in Table 4.3, the units are $[(photons/bin) \cdot km^2/J]$ for the green normalized signal and $[(W \cdot km^2)/J]$ for the infrared normalized signal. Both signals are available in the GLA02 data product. However, the retrieval of these parameters from the GLA02 data products is not as straightforward as the retrieval of parameters from other products. GLA02 is a near-real-time atmospheric data product and the files become very large. Also, the data subsetter of NSIDC does not support GLA02 data subsetting based on fine spatial and temporal preferences. The main reason for this is that the product only contains the predicted latitude and longitude of the footprints assuming nadir-pointing only, which can differ tens of kilometers from the actual locations in case of off-nadir pointing. As a consequence, for this research it is decided that the normalized signals are not used for analysis, but instead the attenuated backscatter profiles. Basically, these profiles only differ from the normalized signals in a scaling factor (the calibration constant) and the correction for transmission losses. Another advantage of this approach is that both the green and infrared backscatter profiles can be directly compared due to the same units, i.e. $[m^{-1}sr^{-1}]$, see Table 4.1.

The two equations described above basically show how the attenuated backscatter profiles can be computed. However, first the lidar calibration constants should be computed.

The laser light is backscattered from atmospheric clouds, aerosols and molecules [33]. The lidar calibration constants, C_{532} and C_{1064}, are time-dependent and require, next to the normalized lidar signals, the average molecular backscatter cross-section and the two-way transmission of the atmosphere. This means that, for the computation of the lidar calibration constants, β_{532} and β_{1064} in Equation 4.4 and Equation 4.5 become the molecular profiles, and P'_{532} and P'_{1064} become the vertically integrated normalized signals. These molecular profiles and normalized signals are only required for a specific calibration layer of $2km$. The molecular backscatter profiles, that are dependent on atmospheric temperature, pressure and humidity, are determined from meteorological data. The transmission term is the two-way transmission from the top of atmosphere to the calibration height and is composed of a Rayleigh (molecular scattering [54]) component and an ozone component. When computing this two-way molecular transmission term, it is assumed that there is no absorption due to aerosols. The molecular transmission is assumed to be equal to one for the infrared laser and is close to one for the green channel. For the infrared channel, the ozone absorption is negligible and consequently the total transmission term for the infrared channel is equal to one. However for the green channel, the ozone absorption should be accounted for and is calculated using ozone mixing ratios. These ratios are dependent on latitude and month of year and are found in pre-defined look-up tables. The molecular transmission term for the green channel is determined with meteorological data as well. More detailed information on the constant and transmission calculations can be found in the ATBD [33].

Thus, briefly summarized, the lidar calibration constant is dependent on the normalized lidar signal that is measured and the molecular backscatter profile that is found in look-up tables. The signal and backscatter profile are retrieved only for a specific calibration layer

of $2km$. For the infrared channel, the signal and profile are sufficient, for the green channel, also the two-way transmission term from top of the atmosphere to the calibration height should be determined with look-up tables. After measuring the signal and determining the profile and transmission term, the lidar calibration constants can be computed with Equation 4.4 and Equation 4.5.

Once the lidar calibration constants are known, the attenuated backscatter profiles can be computed using the normalized lidar signals from GLA02 and the two-way transmission term over the entire profile range (instead of only the calibration layer), again using Equation 4.4 and Equation 4.5. These are rewritten and shown in Equation 4.6. Here too, it holds that the transmission term is equal to one in case of the infrared channel and dependent on altitude for the green channel. The reason that the backscatter profiles are renamed as attenuated backscatter profiles is because of the attenuation of the laser pulses due to molecular and aerosol/cloud scattering effects along the laser's path line-of-sight.

$$\begin{aligned} \beta'_{532}(z) &= \frac{P'_{532}(z)}{C_{532}T^2_{o,532}(z)} \\ \beta'_{1064}(z) &= \frac{P'_{1064}(z)}{C_{1064}} \end{aligned} \tag{4.6}$$

Here, β'_{532} and β'_{1064} are the attenuated backscatter profiles, P'_{532} and P'_{1064} are the normalized signals at an altitude z, C_{532} and C_{1064} are the lidar calibration constants and $T^2_{o,532}(z)$ is the two-way ozone transmission from the top of the atmosphere to altitude z. Even though the GLA07 product documentation [36] specifies that the green signals are replaced by infrared signals when the green channel becomes saturated (i.e. in case of thick clouds or ground), correspondence with the GSFC [41] brought forth that this was never implemented and consequently is not relevant.

As Table 4.1 shows, the green and infrared backscatter profiles are computed between $10.3km$ and $-1.0km$ above the geoid at $40Hz$. Since their vertical resolution is $76.8m$, each of the $40Hz$ backscatter profiles contain 148 bins. The green laser footprint is identical to the infrared footprint, with a major axis ranging from $50m$ to $120m$. Since the repetition rate is $40Hz$, the horizontal resolution is the same as the horizontal resolution of the altimeter signals, i.e. about $175m$.

4.5.3 Retrieval of atmospheric signals at ground level

Now that the backscatter profiles from $10km$ to $-1km$ above the geoid are known, the ground return values can be extracted. For this, the height of the geoid, *i_gdHt*, and the height of the surface above the reference ellipsoid are required. To get the most reliable results, the altimeter detector estimated elevation, *i_elev*, is used as surface height. See Table 4.1 for more information on these attributes.

In case the surface elevation, *i_elev*, is invalid (i.e. in most cases due to failed elevation estimation), the SRTM30 DEM, *i_DEM_elv*, is used. This is a near-global DEM that combines the dataset of the Shuttle Radar Topography Mission (SRTM), flown in February 2000, and the dataset of USGS's GTOPO30 [22]. The reference frame of this DEM has already been changed to the Topex/Poseidon ellipsoid to make it consistent with the ICESat derived elevations. Now that the geoid height and surface elevation above the ellipsoid are known, the surface height above the geoid can be determined by simple subtraction. Next, the backscatter values are extracted from the bin numbers corresponding to elevations ranging from $300m$ to -$300m$ with respect to the ground surface. In the rare case that both the ICESat estimated elevation and the DEM elevation are invalid, a maximum height of $4km$ and a minimum height of -$1km$ above the geoid is chosen, and consequently more bin values are extracted. The maximum value of the backscatter profiles that occurs in the extracted bin values is assumed to be the ground return value. Figure 4.6 shows typical examples of the green and infrared backscatter return values from $10km$ to -$1km$ above the geoid, also indicating $300m$ and -$300m$ above the ground surface.

As the plots in Figure 4.6 show, the backscatter values of the green channel are roughly an order of magnitude lower than those of the infrared channel, most likely due to the lower transmitted energy, which decreases the signal-to-noise ratio (SNR), and higher atmospheric absorption. For the infrared channel, the ground returns are clearly identifiable with respect to the rest of the backscatter profiles (i.e. atmospheric values). However the ground returns of the green channel show more fluctuations. Some green return values are clearly higher than the rest at ground level, while others are closer to the backscatter values of the atmosphere, see the top two plots of Figure 4.6.

Over the Dry Valleys, in total 67,318 valid green and 73,160 valid infrared backscatter returns at ground level are available, not taking into account any filtering criteria other than the return signals being a valid number. The difference in number of returns at ground level is accountable to the high number of invalid green backscatter profiles for laser campaign 1A, most probably due to erroneous backscatter profile computations.

The attenuated backscatter values at ground level of the green and infrared signals will be used for classification analysis. The combination of the green ground return signals with the infrared signal attributes might improve the classification algorithm, because ice in particular shows different reflectance for both wavelengths. Even though theoretically the classification algorithm might be improved by taking into account the green signals, in practice however the dramatic decrease in transmit energy might pose a problem. Whether the green and infrared ground signals will actually improve the classification algorithm will be discussed in Chapter 8.

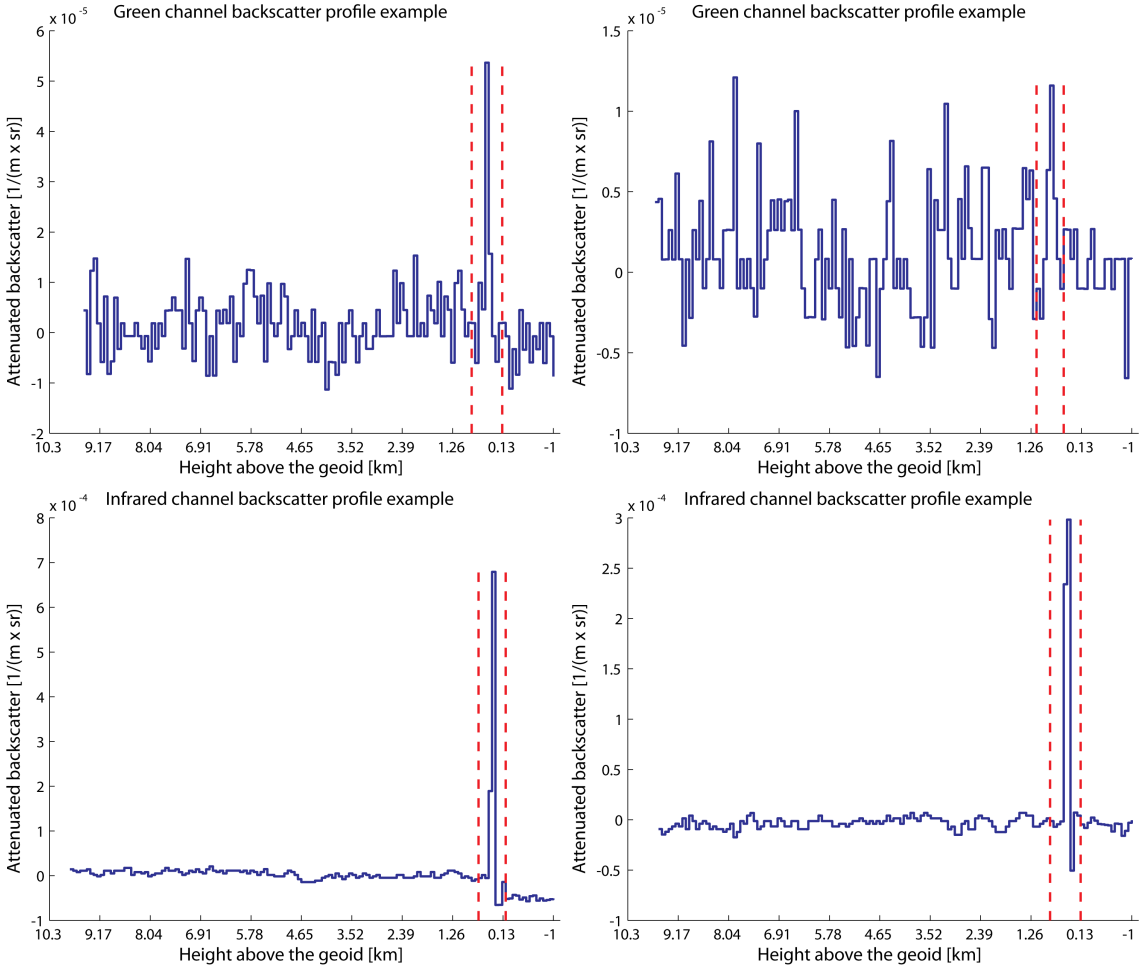

Figure 4.6: Four typical backscatter profiles of the green and infrared channels over the Dry Valleys. The red lines indicate $300m$ above and below the geoid.

Chapter 5

Waveform derived attributes

In addition to the data product derived attributes, waveform derived attributes are used for the surface topography and classification analysis as well. With the filtering techniques that are described in this chapter, together with the filtering and correction techniques described in Chapter 4, the subset of shots to be used for further analysis is determined. The filtering techniques in this chapter are included in Figure 2.3. The beam co-elevation angle correction technique will be explained in Chapter 6, because this correction can be better explained in the context of surface topography analysis.

All attributes in this chapter are related to the received and transmitted waveforms. In addition to how these attributes are used for filtering techniques, this chapter will discuss how surface topography and land cover types might influence ICESat's laser pulses as well. First of all, in Section 5.1 an overview of the waveform derived attributes will be given. Brief descriptions of several of these attributes are provided in this section as well. Next, to make the reader familiar with the waveform principles and its composition, Section 5.2 addresses some background information about the returned waveform. With the knowledge of the returned waveform and the transmitted waveform, the so-called deconvolution technique can be applied. In Section 5.3 this technique will be explained. The filtering technique of the waveforms will be discussed in Section 5.4. Finally, the waveform is subject to statistical analysis, which will be explained in detail in Section 5.5.

5.1 Overview of waveform derived attributes

The waveform derived attributes include data product derived attributes and self-computed attributes. The data product derived attributes that are explained here are all related to transmitted and received waveforms.

The transmitted waveforms, *i_tx_wf*, and received waveforms, *i_rng_wf*, are extracted from the Global altimetry data product, GLA01. Both waveforms are embedded in the data product and are expressed as digitizer counts. With IDL code and software these

waveforms are extracted and counts are transformed into energy units, Volts, using a simple look-up table. This table holds for each digitizer count the corresponding value in Volts. Detailed information about GLA01 data product usage guidance can be found on the NSIDC website [34].

Derived attribute ID	**Data product**	**Frequency**	**Attribute name**
shotID	GLA01/GLA06/GLA07	40Hz	record index + shot number
Original attribute ID	**Data product**	**Frequency**	**Attribute name**
i_rng_wf	GLA01	40Hz	1064nm range waveform [counts]
i_tx_wf	GLA01	40Hz	sampled transmitted pulse waveform [counts]
i_RecNrgLast_EU	GLA01	40Hz	1064 nm laser received energy [atto J]
i_EchoPeakLoc	GLA01	40Hz	Echo peak location [ns]
i_reflctUncorr	GLA06	40Hz	reflectivity not corrected for atmospheric effects [-]
i_gval_rcv	GLA06	40Hz	gain value used for received pulse [counts]
i_satNrgCorr	GLA06	40Hz	saturation energy correction [0.1fJ]
i_kurt2	GLA06	40Hz	kurtosis of the received echo (standard) [-]
i_skew2	GLA06	40Hz	skewness of the received echo (standard) [-]
Self-computed attribute ID	**Related data product**	**Frequency**	**Attribute name**
Width	GLA01	40Hz	Waveform width [ns]
FWHM	GLA01	40Hz	Waveform Full Width at Half Maximum [ns]
Maximum	GLA01	40Hz	Waveform maximum [V]
Summation	GLA01	40Hz	Waveform summation [V]
Risetime	GLA01	40Hz	Waveform risetime [ns]
Number of peaks	GLA01	40Hz	Waveform number of peaks [-]
Mean	GLA01	40Hz	Waveform mean [V]

Table 5.1: The waveform derived attributes that are used for the research, all relate to waveforms. The frequency represents the rate per second of parameter occurrence in the specified data product. Units of the parameters are indicated between brackets.

Table 5.1 shows the waveform derived attributes that are used for this research. It differentiates between one derived attribute, original (data product derived) attributes and self-computed attributes. All original and self-computed attributes relate to ICESat's waveforms. The first attribute, *ShotID*, is used for referencing the shots between the data products. The original attributes are explained next. However, waveform kurtosis, *i_kurt2*, waveform skewness, *i_skew2*, and all self-computed attributes will be explained in Section 5.5.

The echo peak location, *i_EchoPeakLoc*, is used for this research as a measure for the presence of a valid return peak in the received waveform. It is used for initial filtering of the waveforms, see Section 5.4. The reflectivity not corrected for atmospheric effects, *i_reflctUncorr* or hereafter referred to as reflectivity, is calculated as the ratio of the received energy and the transmitted energy and it has been scaled for range. The received energy is calculated as the area under the waveform after noise subtraction and scaled by the receiver gain. More information on the receiver gain, later in this section. The transmitted energy is calculated as the area under the transmitted waveform and scaled by the transmitted gain. Both are also scaled by the corresponding optical to detector Volt efficiency [3]. Equation 5.1 shows the equation for computing ρ_{surf_uncor}, the reflectivity uncorrected for atmospheric effects.

$$\rho_{surf\_uncor} = \frac{\pi E_{rec} R^2}{E_{trans} A_{telescope} \tau_{opt}} \tag{5.1}$$

Here, E_{rec} is the received energy, R is the range in meters, E_{trans} is the transmitted energy, $A_{telescope}$ is the telescope area ($0.709m^2$), and τ_{opt} is the optics transmission (67%). GLA06 also offers the attribute *i_reflCorr_atm*, reflectivity corrected for atmospheric effects; however, this parameter is not used for this research, because the atmospheric correction is not reliable throughout the laser campaigns due to the degradation of the green laser [46]. Since *i_reflctUncorr* is not corrected for the atmospheric effects, this value should be interpreted with caution and only used when the cloud presence is minimal. The test for cloud presence is addressed later in this chapter. The received energy is the energy of the received laser pulse between the thresholding crossing before and after the maximum amplitude in energy units [J]. It used in Chapter 8 for the discussion on correcting the reflectivity attribute, *i_reflctUncorr*, for saturation. More information on saturation later in this section.

For each laser pulse the GLAS instrument transmits, the received energy is mainly dependent on atmospheric and ground reflecting effects. Based on the amount of received energy, the voltage gain of the detector pre-amplifier, *i_gval_rcv*, is adjusted. This is carried out automatically to amplify the backscattered signal such that a proper waveform is recorded and Gaussian distributions can be fit in order to estimate the elevation. To get this proper received waveform shape, the gain is adjusted depending on the maximum received amplitude [3]. If the received energy is low, the gain will be increased and if the received energy is high, the gain will be decreased. The minimum and maximum values are 13 counts and 250 counts, respectively (although on the NSIDC data description page [35] a maximum of 200 [counts] is posted). However, since the response time of the gain is limited, in some cases the recorder becomes saturated. As will be demonstrated with case studies in Section 8.1, the response time of the receiver gain is about 10 shots (i.e. 0.25 seconds) from high gain to low gain and about 5 shots (i.e. 0.125 seconds) for low gain to high gain. Basically, saturation means that the GLAS instrument may not properly record the received energy values because gain over-amplifies the received pulse. This occurs generally for the following case. The gain has been set at a high value (i.e. large amplification) when the returned energy is low. In case the received energy suddenly increases, due to changing land cover type for example, the instrument should adjust to lower gain (i.e. low amplification) in order to keep proper output signals for the digital waveform recorder. However, since the response time of the gain is non-optimal (too slow), the gain does not adjust properly and consequently the signal output is over-amplified. The waveform recorder has a limited recording range and the over-amplification results in the recorder to become saturated. Consequently, the samples that are recorded may not represent the actual waveform values. This over-saturation of the detector can cause the waveform peaks to be clipped [45] and the upper part of the received waveform is omitted. When the gain eventually adjusts, the waveforms are recorded normally again. Conversely, in case the received energy jumps from high returned energy to low returned energy and

the gain cannot increase properly, the recorder does not become saturated but instead a weak signal may be recorded. In extreme cases, no echo peak could be found and the echo peak location, *i_EchoPeakLoc*, becomes equal to zero.

Figure 5.1 shows two waveforms, the left is non-saturated and the right is saturated. When the returned pulse causes the receiver to be saturated, the received waveform could be distorted and does not represent the actual returned voltages. To correct for this effect the saturation energy correction, *i_satNrgCorr*, can be applied. This attribute indicates the energy that should be added to the recorded received energy in case of saturation. The detailed discussed with real examples of the variability of the gain and saturation energy correction with varying transmitted energy and land cover type, is provided in Chapter 8.

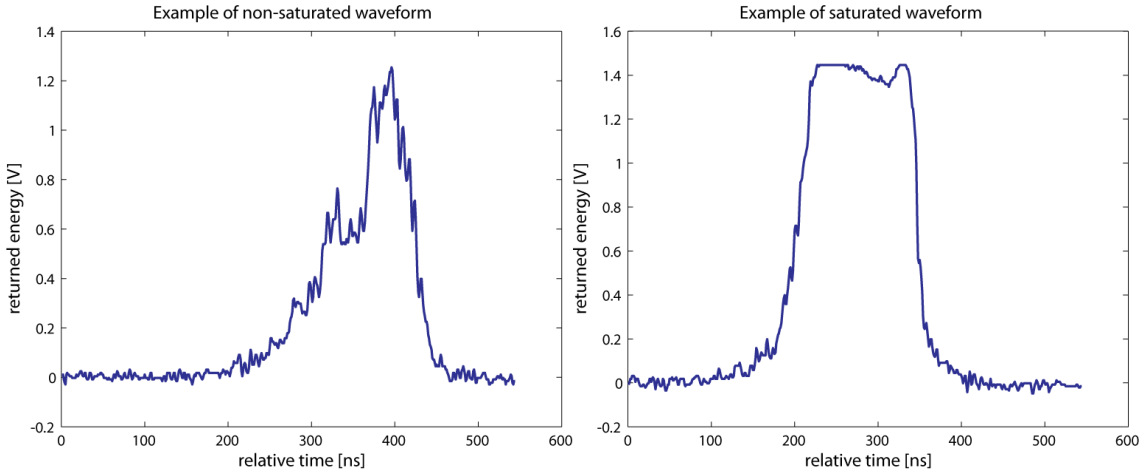

Figure 5.1: Example of a non-saturated waveform (left plot) and a saturated waveform (right plot). Notice the clipped peak when saturated.

5.2 Returned waveform composition

Between the time the GLAS instrument transmits a laser pulse and receives the returned pulse, the pulse is subject to various attenuation and distortion factors. Not only the surface that is illuminated by the laser beam contributes to pulse distortion, but also atmospheric disturbances. Besides these external factors, two internal factors are taken into account as well, namely the transmitted waveform shape and the receiver impulse response effects. As a result, the returned waveform is a function of the scattering elements along the traveled path of the laser and internal GLAS instrument related factors [3].

After the waveform is extracted from GLA01 with IDL software, the waveform basically represents the energy of the laser pulse as a function of time. From the moment the laser

pulse is generated at the satellite, it is subject to many complex influences. A simplified overview of the processes the laser pulse is subjected to from transmittance until receptance is illustrated in Figure 5.2. The waveforms show an example sequence of what the pulse could experience. The landscape in the figure is actually a part of the Dry Valleys in Antarctica validation site and is used here only for visualization purposes. See Section 3.1 for more information on the Dry Valleys.

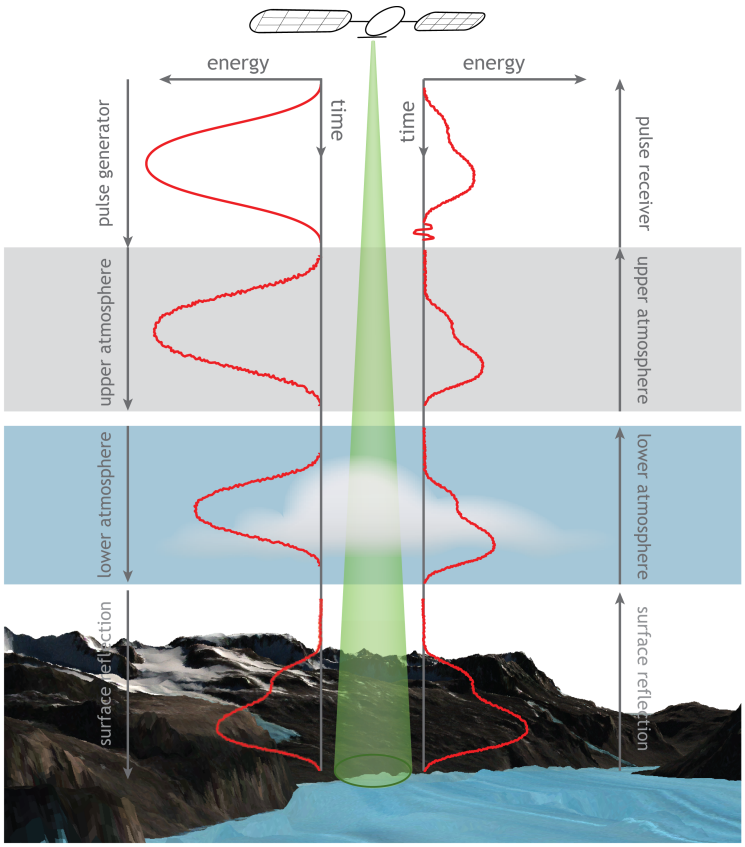

Figure 5.2: A simplified overview of waveform propagation from transmittance till receptance. The changes between the waveforms are not to scale and only meant for visualization purposes. The upper atmosphere includes the ionosphere and mesosphere, the lower atmosphere includes the stratosphere and troposphere.

Figure 5.2 should be interpreted as follows. Eight waveforms are indicated and travel in the direction of the outer arrows, so in the left part the pulse travels from satellite to ground and in the right part the pulse travels from ground to satellite. Each waveform is expressed as energy versus time, and for each block (i.e. pulse generator, upper atmosphere, etc.) the axes of the waveform correspond to the axes shown on top of that part. Thus, for the left and right waveforms the time axis is orientated downwards. This is because the first sample that is recorded after transmission in the left part is also the first

sample that is recorded in the right part. The first internal factor that contributes to the shape of the received waveform is the shape of the transmitted waveform. Generally speaking, the transmitted waveforms are similar, each describing a near-perfect single Gaussian pulse with a standard deviation of about $3ns$ [14]. Although not visible in the figure, the transmitted energy is in reality an order of magnitude twelve times higher than the received energy, mainly due to attenuation and scattering. The first external disturbance (the free space loss ignored) is the upper atmosphere, roughly defined as a combination of the ionosphere and mesosphere. The upper atmospheric attenuation is very small and delay corrections can be neglected [45]. Since the entire pulse is short ($5ns$) and travels with the speed of light, it penetrates the atmosphere extremely quickly and it is assumed that the entire pulse is affected homogeneously over the entire waveform. The only influences are a slight decrease in signal energy and the addition of small scale noise over the entire waveform. This small scale noise is assumed to come from background radiation having the same frequency as the laser pulse. In Figure 5.2 the lower atmosphere, roughly defined as a combination of the stratosphere and troposphere, is mentioned separately, because larger particles forming clouds and aerosols exist here. These particles can distort the waveform significantly by severe scattering. Heavy clouds completely block the laser signal and as a result no surface elevation can be estimated but only cloud information is measured [45]. Thinner clouds and aerosols cause forward scattering, which distorts the waveform and shifts the centroid of the return waveform to later times [3]. Especially with the low laser energy at later ICESat campaigns, cloud cover can prohibit the signal to touch the ground at all. Since it is important for this research to work with 'clean' waveforms (affected by unwanted effects as little as possible) a cloud filtering threshold is proposed, see Section 5.4 and Section 8.2.

After the pulse has penetrated the atmosphere the most important interaction takes place, namely the interaction with ground features. Higher elevated objects are illuminated first and, as a consequence, these scatter back (part of) the energy first. Lower elevated objects scatter back (part of) the energy at a later stage and this results in a complex returned waveform. Important to know in this discussion is that the ground as a three dimensional surface affects the waveform in a two dimensional way. Since the receiver on the satellite only records counts (indirectly energy levels) versus time, rather than an (spatial) energy image over time, the laser beam becomes rotationally independent around the beam center line. As a result, in the transformation from surface to waveform, the information looses one dimension. Figure 5.3 illustrates such an example where it does not matter for the returned waveform where the object is located, as long as the same area is illuminated and the distance between the object and beam center line remains the same. In other words, if the object would have been exactly on the opposite side, the received waveform would be exactly the same. The reason that the distance from the object to the beam center line should remain the same for the waveform to be identical is because of the non-uniform power distribution (assumed to be Gaussian) of the beam with respect to the beam cross-section (i.e. the plane perpendicular to the beam center line). If the same illuminated area of the object is located closer to the beam center, the object would cause a higher local waveform peak because more energy is returned at that particular elevation.

In case the object would be sufficiently lower (i.e. closer to the ground), the peaks in the waveform would merge. The maximum separability of features in the waveform is $15cm$, see Section 2.2. Therefore, an object that is elevated $15cm$ or more above the ground or other objects, should theoretically be distinguishable in the waveform. However, receiver noise makes this ideal resolution worse in reality. Also, more objects illuminated by the beam result in a more complex waveform.

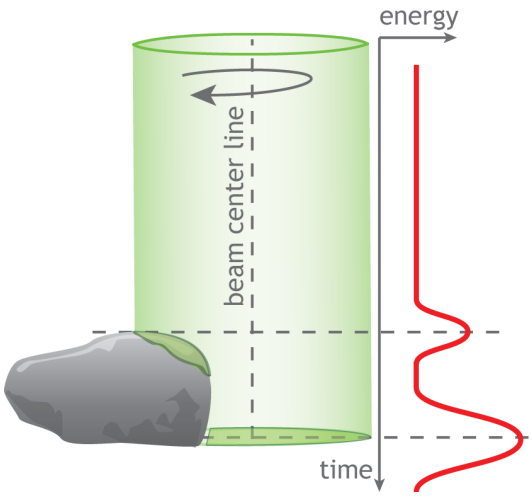

Figure 5.3: The effect of elevated objects on the received waveform. In this case only an object and the ground are considered. The beam center line represents the axis of beam rotational independence, indicated by the rotating arrow.

After the laser pulse is reflected off objects and (possibly) ground, the pulse is scattered back into the atmosphere. Only a fraction will then be scattered in such a way that it is within the field of view of the GLAS telescope. It will again travel through the difference atmospheric layers, resulting in even a lower returned energy level. Finally, the pulse is recorded by the GLAS receiver, as indicated in Figure 5.2. In the figure, a potential unwanted effect of the receiver is incorporated as well, namely the ringing effect after the actual waveform. As the GLAS ATBD on range derivation [3] mentions, this always happens when detector saturation occurs. Since no filtering based on saturation energy correction is carried out, the possibility of ringing effect is taken into account (see Section 5.5). The exact cause behind the ringing effect is unknown, other than that it is an unwanted effect caused by the detector. In all cases of saturation, the ringing does not represent the surface [35]. Especially for waveform analysis this effect is important, because it can influence the derivation of waveform statistics.

5.3 Waveform deconvolution

The returned waveform is mainly dependent on the transmitted waveform, the influences along the laser's path line-of-sight and the receiver impulse response. The surface response function is of special interest to this research since it holds information about the illuminated features. A technique called deconvolution can estimate this function using the received waveform and the transmitted waveform. In this section, this technique is explained and its relevance is discussed. Even though deconvolution is a very promising technique for vegetated areas, it is concluded that for ice sheet applications the technique is not of significant importance.

The transmitted waveform experiences many interactions along the laser's path line-of-sight, basically consisting of the atmospheric and surface response functions, and the receiver impulse response. Since the atmospheric effects are hard to model, especially with the degradation of the green laser pulses (see Section 2.2), preferably these effects should be minimized. In Section 5.4 and Section 8.2 a method is proposed to filter out data that most likely is affected by significant cloud coverage. As a consequence, the returned waveform, assuming it is not biased by significant atmospheric disturbances, represents a convolution of the transmitted signal, the surface response function and the receiver response function. Although it does not mean that the waveforms are actually 'free' from atmospheric effects, it does imply that the waveforms are affected by atmospheric effects in a more or less similar way. When ignoring any receiver instrument effects, the simple convolution is expressed by Equation 5.2 [4].

$$p(t) = s_m(t) * h(t) = h(t) * s_m(t) \tag{5.2}$$

Here, the '$*$' symbol represents the actual convolution operation, $p(t)$ is the returned waveform as function of time, $s_m(t)$ is the (measured) transmitted waveform a as function of time and $h(t)$ is the surface response function as a function of time. To compute an estimate of the surface response function, the deconvolution technique can be applied. Since Equation 5.2 describes the convolution in the time domain, the deconvolution operation has no direct mathematical definition other than being the inverse [32]. In order to transform the time domain convolution into a multiplication, the variables should be transformed into the frequency domain. So, the known functions $p(t)$ and $s_m(t)$ should be fast-Fourier transformed (*fft*) into $P(f)$ and $S_m(f)$, respectively, where the variable f denotes the variable in the frequency domain. Equation 5.3 shows this simple transformation from the time domain to the frequency domain. The fast-Fourier transforms have been normalized by the length of the signal and cut off at half the fft, because the frequency spectrum of the *fft* is symmetric.

$$\begin{aligned} P(f) &= \text{fft}(p(t)) \\ S_m(f) &= \text{fft}(s_m(t)) \end{aligned} \tag{5.3}$$

With the transmitted and received waveforms expressed in the frequency domain, the actual deconvolution equation can be stated, see Equation 5.4.

$$H(f) = \frac{P(f)}{S_m(f)} \tag{5.4}$$

$H(f)$ is the spectral surface function, $P(f)$ is the spectral received function and $S_m(f)$ is the spectral transmitted function. This equation exactly describes what is needed for deconvolution in the ideal case. However, in practice, the received and transmitted waveform are biased with a receiver noise term. Without taking this noise into account, large numerical errors could result [4]. Equation 5.5 shows the additional noise ($n(t)$), which is simply added to the signal in the time domain.

$$\begin{aligned} s_{m,n}(t) &= s_m(t) + n(t) \\ p_n(t) &= p(t) + n(t) \end{aligned} \tag{5.5}$$

The $s_{m,n}(t)$ and $p_n(t)$ of Equation 5.5 represents the raw transmitted and received waveform that are recorded by the GLAS instrument and are available in the GLA01 data product. To account for the noise term, two filtering techniques are applied, one in the time domain and one in the frequency domain. In the frequency domain the Wiener filter is used, which is a statistically optimal filter [28] that minimizes the error between the uncorrupted surface response function $H(f)$ and the estimated surface response function $\hat{H}(f)$ [4]. Equation 5.6 shows the Wiener filter in the frequency domain.

$$W(f) = \frac{|P(f)|^2}{|P(f)|^2 + |N(f)|^2} \tag{5.6}$$

The Wiener filter takes the absolute of both frequency functions to describe a real function. $N(f)$ is the noise term in the frequency domain and $P(f)$ is again the uncorrupted received waveform in the frequency domain. The noise frequency spectrum is computed by fast-Fourier transforming the first 150 samples of the waveform in the time domain into the frequency domain. These 150 samples usually represent the noise component of waveforms having 544 samples (i.e. over ice sheet and land) [12]. In case the entire waveform is 200 samples long, i.e. over ocean and sea ice, the first 75 samples are used for noise estimation. Since only the noise corrupted received signal is available, the uncorrupted signal should be estimated. For this estimation a linear binominal filter is used as a low pass filter to suppress the high frequency noise [4]. The filter uses one row of Pascal's triangle and is normalized over the sum of the row. The matrix below shows the the first six rows of Pascal's triangle.

$$\begin{array}{ccccccccccc} & & & & & 1 & & & & & \\ & & & & 1 & & 1 & & & & \\ & & & 1 & & 2 & & 1 & & & \\ & & 1 & & 3 & & 3 & & 1 & & \\ & 1 & & 4 & & 6 & & 4 & & 1 & \\ 1 & & 5 & & 10 & & 10 & & 5 & & 1 \end{array}$$

For every received waveform, the row number is determined based on the width of the returned pulse. How this width is determined, see Section 5.5. For example, if the sixth row is used for filtering, the row $[1, 5, 10, 10, 5, 1]$ would first be normalized, resulting in $[1, 5, 10, 10, 5, 1]/32 = [0.313, 0.1563, 0.3125, 0.3125, 0.1563, 0.0313]$. The values form an approximation to the Gaussian [18]. By convolving the filter with the raw received waveform, the noise is suppressed. Equation 5.7 shows this simple convolution, where *binom_filter* represents the normalized binominal filter. It basically acts as a Gaussian smoothing filter. The effect of the binominal filter on an example raw waveform is shown in Figure 5.4.

$$\begin{aligned} p(t) &= p_n(t) * \mathit{binom\_filter} \\ P(f) &= \mathrm{fft}(p(t)) \end{aligned} \tag{5.7}$$

With the noise-suppressed received waveform in the frequency domain, the Wiener filter is multiplied with Equation 5.4, resulting in Equation 5.8.

$$\begin{aligned} \hat{H}(f) &= \frac{P(f)}{S_m(f)} \cdot W(f) \\ \hat{h}(t) &= \mathrm{ifft}(\hat{H}(f)) \end{aligned} \tag{5.8}$$

Here, $\hat{H}(f)$ again represents the estimated surface response function in the frequency domain. After applying an inverse fast-Fourier transformation (*ifft*), the estimated surface response function is transformed into time domain ($\hat{h}(t)$). As Neuenschwander [2] points out, some of the advantages of deconvolution over Gaussian waveform decomposition are that deconvolution identifies more peaks compared to the Gaussian decomposition since no *a priori* information on the number of peaks is required. Based on this information, the deconvolution technique opens a very interesting way for modeling the waveforms and consequently for classification, waveform parameter analysis and slope correction purposes. The right plot of Figure 5.4 demonstrates the deconvolution result of the raw waveform illustrated in the left plot. The deconvoluted waveform is normalized with respect to the maximum peak. The middle plot shows the effect of the binominal filter on the raw received waveform.

With the estimated surface response function, the separability between the targets along the laser path is increased. The function, also known as reflectivity function, is again dependent on various complex relationships along the laser path's line-of-sight, such as surface topography, surface reflection, forward scattering, laser pointing, laser energy, footprint size, shape and orientation, and object height and position within the footprint [2]. As the right plot of Figure 5.4 demonstrates, the peaks are better identifiable. For this research, particularly the convolution of the surface topography and surface reflection encapsulated in the deconvolved waveform could be useful for defining rock, snow, ice and water terrain classification and for improved slope correction technique. To separate these factors and, more importantly, to deduce the classification and correction technique, ideally all of these influences are deconvoluted and decomposed into individual and identifiable

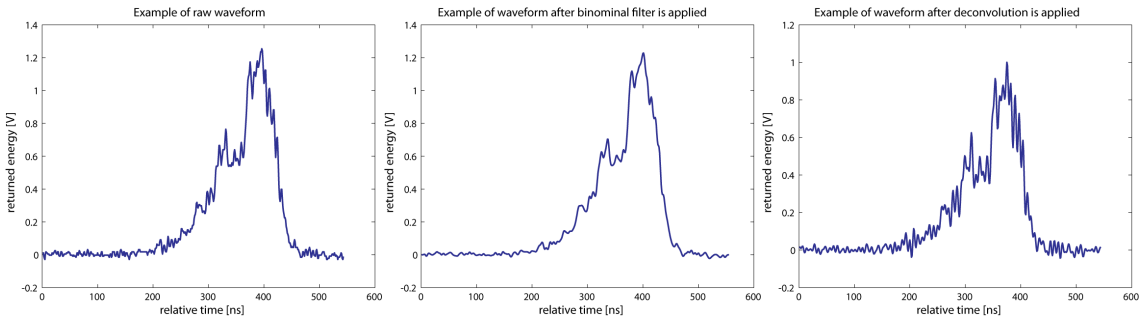

Figure 5.4: Example of a raw waveform (left), the same waveform after the binominal filter is applied (middle) and finally the normalized deconvolved waveform (right).

functions.

Especially for applications like vegetation and urban areas, this decomposition is important for accurately measuring the difference between the top and the bottom of the object. For example, detection of tree heights could be improved by carrying out the deconvolution technique on received waveforms. With this technique the separability between tree canopy and ground surface could be improved. However for this research, it turned out that the improved separability between features in the signal does not compensate for the extra amount of computation time needed to deconvolve each waveform. Since this research focuses on ice sheet and water waveforms only, it was concluded after comparison with raw received waveforms, that the deconvoluted received waveforms do not change significantly. Moreover, the transmitted waveform over time is consistent enough for this research (despite the problems with the transmitted energy) and it is assumed that the minor changes in shape of the transmit waveform does not affect the shape of the received waveform in any way that it compromises the waveform analysis of this research.

Figure 5.5 shows the variability of all transmitted waveforms over the Dry Valleys validation area. The waveforms are centered along the time axis such that the maximum peak of all waveforms are located on one vertical line. As indicated by the figure, the shape remains more or less the same over time. After comparison of this observation, together with the lack of need for high separability of peaks of ice sheet and water waveforms, with the extra required computation time, it was decided that the deconvolution is not carried out for this research.

5.4 Waveform filtering techniques

After the waveforms are extracted, they are tested on several criteria to filter out unreliable waveforms. This is necessary because in case waveforms have low SNR, the statistics that are derived from the waveforms, see Section 5.5, may give bad results. In addition to filtering based on data product derived attributes related to the waveform, a waveform filtering criterion based com-

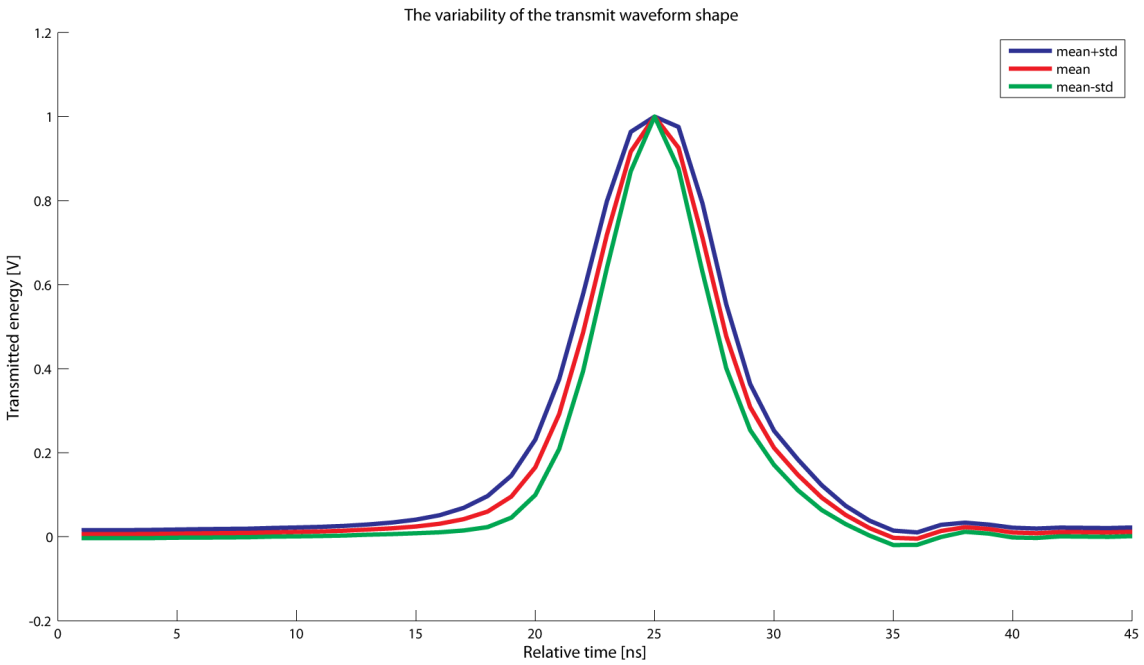

Figure 5.5: The variability of the transmitted waveform over the Dry Valleys, normalized over the maximum peak. The mean and the standard deviation (std) values are indicated.

parison between the returned waveform with background noise is applied as well. The filtering steps are included in the flow diagram of Figure 2.3.

The first attribute used for filtering out waveforms is the already discussed *i_EchoPeakLoc*, the echo peak location. If the GLAS algorithm does not detect a noise threshold crossing of the waveform, the echo peak location becomes zero. Consequently, all waveforms are filtered out that have an echo peak location value equal to zero. The second attribute that could be used for filtering is *i_gval_rcv*, gain value used for received pulse. According to several references, the gain value is an important measure for the cloud presence during the laser shot [55, 21, 46]. It is an alternative for using the cloud flag (*i_FRir_qaFlag*), which turned out to be unreliable during the mission lifetime (due to the green laser degradation). As outlined by Donghui et al. [55], a gain value of higher than 50 normally indicates clouded sky. However, as is discussed in Chapter 8 with detailed examples, this threshold setting may also exclude valuable waveforms that are most probably not affected by cloud cover. From the analysis in Chapter 8 a cloud filtering technique is used to filter out all shots that have a reflectivity (*i_reflctUncorr*) value of 0.1 or lower; no gain filtering is applied.

The third attribute that could be used for filtering out inaccurately recorded waveforms is the already discussed *i_satNrgCorr*, the saturation energy correction. It is a measure of the amount of saturation affecting the returned waveform. One possible and simple solution is that if the value is not equal to zero, and consequently saturation has occurred (resulting in possible waveform distortion), the shot is removed. However, as will be dis-

cussed in Chapter 8 as well, using this parameter to filter out shots results in the loss of valuable shots over particular land cover types. So no filtering is carried out based on the saturation energy correction.

Beside these initial filtering steps (i.e. based on reflectivity and echo peak location), another filtering technique is applied as well. This technique is based on the waveform noise level with respect to the waveform peaks. As mentioned in Section 5.3, the noise is taken from the first 150 samples of the waveform. Over ice sheet and land surface the waveform holds 544 bins of $1ns$ samples and from visualizing the waveforms it is clear that the first 150 samples represent the background noise [12]. In case of 200 bins long waveforms over sea ice and ocean, the first 75 samples are used to represent the noise. Equation 5.9 uses this noise to compute a noise threshold value, which is the sum of the mean (noise level) and three times the standard deviation [10].

$$threshold = mean(noise) + 3 \cdot std(noise) \tag{5.9}$$

If the waveform has no values larger than this threshold, the waveform is filtered out. See Figure 5.7 for an example of the noise threshold compared to the waveform. Also in case one of the first 150 or 75 samples is higher than the rest of the waveform, the waveform is filtered out. Finally, all waveforms are filtered out that have invalid values for the attributes given in the overview of Table 5.1 (except *i_satNrgCorr*).

5.5 Waveform statistics

An essential part of full waveform analysis is the computation of waveform statistics. These statistics refer to all attributes that are directly related to the returned waveform's shape. Each waveform statistic describes a characteristic of the waveform and can be used for discrimination purposes between land cover types and to analyze the surface topography within the footprint. Two kinds of waveform statistics are used for this research: pre-computed statistics that are readily available in the GLAS data products and self-computed statistics that are computed for this research. All attributes discussed in this Section are included in the overview of Table 5.1 as well.

The pre-computed waveform statistics used for this research are kurtosis, *i_kurt*2, and skewness, *i_skew*2. The kurtosis is a measure for the peakedness and flatness of the distribution and skewness is a measure for the asymmetry of the distribution around the mean [3]. Both are available from GLA06 and are computed using the standard parameterization, which fits up to two Gaussian distributions in the returned waveform. Hence the number *2* in the ID of the attributes. Based on the location(s) of these Gaussian(s) and the raw waveform, the two attributes are calculated. As the GLAS ATBD on the range derivation [3] describes, kurtosis and skewness are based on equations that are dependent

on the mean and standard deviation of the waveform. Equation 5.10 describes the equations behind these pre-computed attributes.

$$\begin{aligned} kurtosis &= \frac{1}{\sigma^4} \frac{\sum_{i=start}^{i=end} (i - mean)^4 w[i]}{\sum_{i=start}^{i=end} w[i]} - 3 \\ skewness &= \frac{1}{\sigma^3} \frac{\sum_{i=start}^{i=end} (i - mean)^3 w[i]}{\sum_{i=start}^{i=end} w[i]} \end{aligned} \tag{5.10}$$

In this equation, $w[i]$ is the power at gate i of the waveform, *mean* is the mean and σ is the standard deviation of the waveform. The kurtosis and skewness are the 4th and the 3rd standardized moments, respectively. The 'minus three' in the kurtosis equation is a correction to make the kurtosis of the normal distribution equal to zero. Equation 5.11 provides the computation for the mean and standard deviation.

$$\begin{aligned} mean &= \frac{\sum_{i=start}^{i=end} iw[i]}{\sum_{i=start}^{i=end} w[i]} \\ \sigma^2 &= \frac{\sum_{i=start}^{i=end} (i - mean)^2 w[i]}{\sum_{i=start}^{i=end} w[i]} \end{aligned} \tag{5.11}$$

In case the returned waveform describes a perfect normal distribution, both the kurtosis and skewness are zero. If not, the ATBD on the range derivation [3] provides the methods that are used to determine the start and end values used for the computations. Figure 5.6 illustrates the relationship of kurtosis and skewness on probability density functions.

The subfigures in Figure 5.6 show probability density functions with varying kurtosis and skewness values. The distributions in each subfigure have the same standard deviation and are normalized with respect to the area under the distributions to make proper comparisons. The distributions illustrating the different skewness values are shifted horizontally to show the distributions as clearly as possible. As the figure shows, the larger the

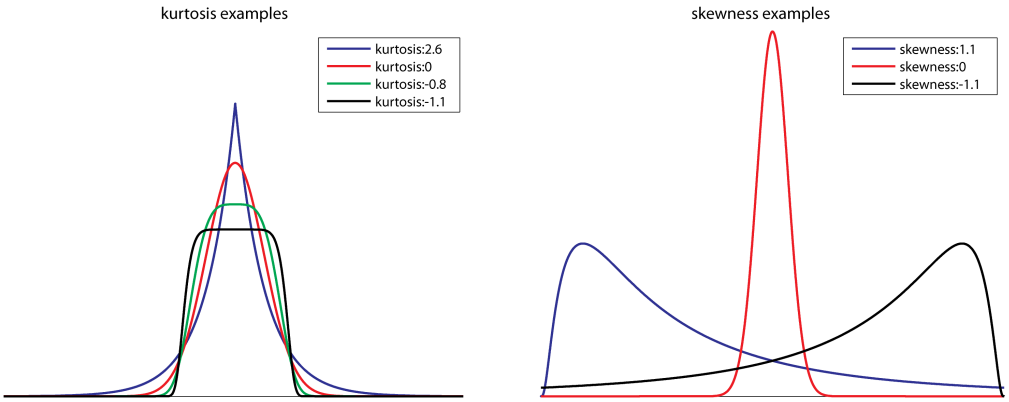

Figure 5.6: Example probability density functions with varying kurtosis and skewness values. In both subfigures, the red line represents the normal distribution having zero kurtosis and zero skewness. All distributions in each subfigure have the same standard deviation and are normalized with respect to the area under the distribution. The figure is only meant for visualization purposes.

kurtosis, the more peaked the distribution is. The distributions become skewed to the left in case of positive skewness values and vice versa. The physical explanation behind varying kurtosis and skewness can be explained as follows. A laser pulse that illuminates a sloped or rough surface will result in a broad peak. Consequently, the kurtosis will decrease. Conversely, the received waveform will be sharply peaked over flat surfaces, resulting in more positive kurtosis values. As for the skewness, the physical interpretation could be found in terms of non-uniform illuminated surfaces. For example, a tree with increased canopy width towards ground surface will result in increased waveform energy values with the main lobe at the end of the waveform. Consequently, the received waveform skewness will have negative skewness values. Positive skewness values will result from opposite surface geometries, i.e. a large cross-sectional area elevated above the ground (such as a blunt hill).

The kurtosis and skewness are the only pre-computed waveform attributes (i.e. available from the data products). The self-computed waveform attributes are not extracted from any data product, but are computed for this research. These waveform statistics form an essential part in full waveform analysis because they indicate valuable characteristics of the waveform. Eventually, these characteristics are used to get correlations between surface types and waveform statistics, in order to come up with an algorithm to discriminate between surface types based on full waveform statistics. It also plays an important role for surface topography analysis. Below, each of these statistics are explained individually using Figure 5.7 as an illustrative example. The figure shows the waveform as returned energy versus relative distance and returned energy versus relative time. Originally in the GLA01 data product, the waveform is described as a function of relative time in $[ns]$. By multiplying this time by half the speed of light (due to the two way ranging), the waveform can be described by relative distance.

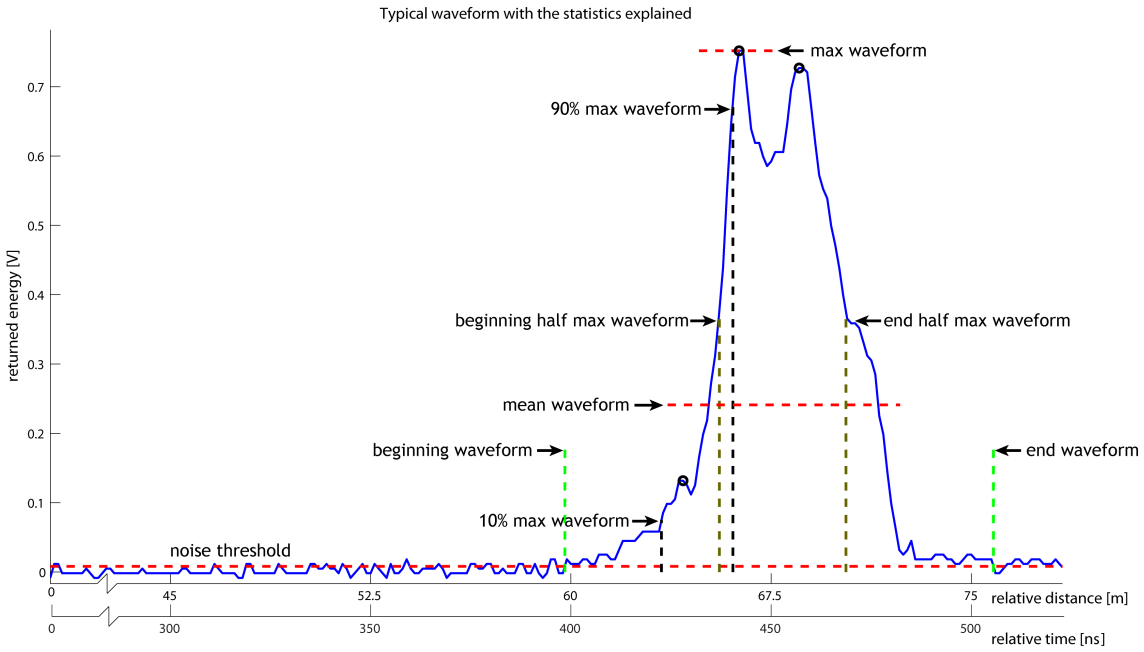

Figure 5.7: Example waveform with the indications of waveform statistics. The black rounds indicate peaks within the waveform. The waveform is expressed in returned energy versus relative distance and returned energy versus relative time.

The waveform threshold is meant to distinguish between the initial noise signal and the returned waveform. Identical to the explanation in Section 5.4, this value is computed by taking the sum of the noise mean and three times the noise standard deviation of the first 150 or 75 samples. In Figure 5.7 the threshold is indicated with the lowest horizontal red dashed line. Below this line, no statistics are computed and it serves as the division line between the background noise and the returned waveform.

In the following paragraphs, the computation of the self-computed attributes are explained. In the explanations, sometimes 'interpolation' between specific samples is mentioned, with which the following is meant. Since the samples have interval values of $1\,ns$ and the waveforms can be narrow (width as small as $5ns$) the waveform statistics are computed with sub-nanosecond values. Consequently, interpolation between samples before and after the specific value is carried out. For example, if the noise threshold is $0.01\,V$ and the sample (at $100ns$) before the threshold crossing has a value of $0.008\,V$ and the subsequent sample (at $101ns$) has a value of 0.02, the interpolated sample is $(0.01 - 0.008)/(0.02 - 0.008) + 100 = 101.17ns$.

The *width* of the waveform is the time difference between the beginning and the end of the returned waveform. The beginning is defined as the interpolation between the last sample under threshold before the maximum of the waveform, and the subsequent sample.

Similarly, the end is defined as the interpolation between the first sample under threshold after the maximum of the waveform, and the previous sample. Consequently, the beginning is the first sample of consecutive samples above the threshold, the end is the last sample of same consecutive samples above the threshold. In Figure 5.7 the two vertical green dashed lines indicate the start and end point of the waveform. If a sample rises above the threshold and there exists no other sample that is below the threshold between this sample and the maximum, this sample (interpolated with the previous sample) defines the beginning of the waveform. The same principle holds for the end of the waveform, but reversed. This statistic can indicate the elevation difference between the illuminated objects within the beam footprint. For the example of Figure 5.7, the *width* is $16m$ or $107ns$.

The full width at half maximum ($FWHM$) is defined as the time difference between the beginning at half of the maximum and the end at half of the maximum of the waveform. In Figure 5.7 the two vertical brown dashed lines indicate the start and end point of the $FWHM$ of the waveform. The beginning and end samples are defined with the same principle as before: the beginning is the interpolation between the last sample under the half maximum value before the maximum, and the subsequent sample. The end is the interpolation between the first sample under the half maximum value after the maximum, and the previous sample. This statistic is also a measure for the elevation difference within the footprint, however ignoring the lower returned energies. Higher energies values may be related to the object close to the beam center due to the Gaussian power distribution of the beam and/or from more scattering elements within the footprint. The higher the maximum of the waveform, the less energy is distributed over time and the more energy is concentrated near the peak. This will then result in a smaller $FWHM$. In Figure 5.7 the $FWHM$ is $4.7m$ or $32ns$, which could indicate the elevation difference of that part within the footprint contributing to the maximum peak.

The *maximum* of the waveform is simply the value of the maximum peak of the waveform expressed in energy units. In this case, the figure indicates a returned maximum recorded energy value of $0.75\,V$. This value should be interpreted with caution, because the receiver gain determines the height of the waveform recording. In case the received energy is very low and consequently the gain is large, the recorded maximum value of the waveform could be the same in case the received energy is high and consequently the gain is low. Especially when saturation occurs and the recorded waveform gets clipped, the maximum value is not very reliable. Although this statistic is discussed here, it is not used for classification or slope analysis. However, it is important for risetime computations that are discussed later on.

The *summation* of the waveform, not indicated in Figure 5.7, is simply the summation of the waveform values that are above the threshold. No integration takes place, because the values are not necessarily continuously ordered in the waveform (i.e. some peaks are identified before the waveform beginning). In addition, all waveforms are sampled at $1ns$ ($0.15m$) intervals, resulting in the same summation computation. The statistic is a measure for the returned energy above threshold and can be used for a SNR estimate by taking

the ratio between the sum of the waveform and the noise threshold value. For example, different surface types could result in different SNR values. In case of Figure 5.7, the *summation* is $76\,V$ compared to a noise threshold of $0.008\,V$. The *summation* should be taken with caution as well, because it does not represent the actual received energy, for the same reasons as explained above.

The *risetime* of the waveform is defined as the time difference between 10% and 90% of the maximum peak value. Figure 5.7 illustrates these points as '10% max waveform' and '90% max waveform' (the two vertical black dashed lines). The '10% max waveform' is the interpolation between the last sample under the 10% value of the maximum before the maximum value, and the subsequent sample. Similarly, the 90% max waveform is the interpolation between the last sample under the 90% value of the maximum value before the maximum, and the subsequent sample. This statistic is a measure for the slope of the maximum peak. It could indicate the signal penetration of objects or the slopeness of the illuminated surface within the footprint. For example, a tree canopy that allows the signal to gradually be returned will have a lower risetime than a flat solid surface type. A flat surface will result in a distict peaked waveform return, whereas a sloped surface will result in a more broadened waveform return. The example of Figure 5.7 has a *risetime* of $2.7m$ or $18ns$.

The *number of peaks* of a waveform is dependent on the step width of the detection algorithm, the noise threshold and the location in the waveform. First of all, the step width is taken to be $2ns$, which places constraints on the two samples before and the two samples after the peak. The first sample before and the first sample after the peak should both have values smaller than the peak value. In addition, the second sample before and the second sample after the peak should both have values smaller than the first sample before the peak and the first sample after the peak, respectively. In other words, the five values (where the middle is the peak) should describe a '$\wedge$'. This step width is taken into account to deal with the additive noise term in the waveform (i.e. very small local spikes are not indicated as peaks). Only the peaks that are above the 10% of the maximum value are considered, because peaks in the noise sections before and after the waveform are not relevant. Finally, the location of the peak in the waveform is used to deal with the ringing effect, which was already mentioned in Section 5.2. Detector saturation could result in oscillations after the major waveform that rise above the threshold. As a consequence this could negatively influence the *number of peaks*. To exclude the ringing peaks, all values after the waveform end are left out in the peak identification. The *number of peaks* could be a measure for the roughness of the terrain and different elevated objects within the illuminated surface of the footprint. The example of Figure 5.7 has three peaks.

Finally, the *mean* of the waveform is the mean value between the beginning and the end of the waveform. This statistic is also a measure of the spreading of the energy over time when comparing to the maximum value. The smaller the mean over the maximum, the more energy is spread over a wide waveform time frame. The higher the mean over the maximum, the more the energy is concentrated in a narrow waveform time frame (for

example the waveform describes one single peak). Figure 5.7 has a *mean* of $0.24\,V$.

Most of the the described statistics are commonly used for signal waveform analysis. They provide the basis for this research in order to analyze the effects of various surface types and slopes on the returned waveform. In Chapter 6 and Chapter 8, these statistics will be used.

Chapter 6

Surface topography analysis

Slope and roughness of the illuminated surface form two major influences on the returned waveform shape. In addition, the waveform's shape is affected by the angle under which the laser beam illuminates the surface as well. This angle is referred to as the beam co-elevation angle. The combination of the slope and beam co-elevation angle together form the so-called incidence angle. In this chapter the effects of slope, roughness and beam co-elevation angle on the waveform width will be analyzed. This is important because these effects should be understood properly before using waveform analysis for land cover type classification.

In Section 6.1, a method of estimating slope and roughness will be explained. The method of combining the beam co-elevation and slope to get the incidence angle is explained in Section 6.2. In Section 6.4 the influence of incidence angle and roughness on waveform width are analyzed. To correct for the effect of beam co-elevation angle on waveform derived attributes like width and risetime, Section 6.3 explains how this angle can be used to compress the waveform. For this research, the surface slope is defined as the inclination angle of a plane with respect to the horizontal plane in the projected map coordinate system. The incidence angle is defined as the angle the laser beam makes with the line perpendicular to the surface footprint plane. The roughness is defined as the root mean square (rms) of the surface deviations from a fitted plane through the elevations. The angles and roughness will be explained in more detail in this chapter.

6.1 Slope and roughness estimation

From the six DEMs over the Dry Valleys (see Section 3.3) together with the ICESat footprint dimensions and geolocations, the slope and roughness parameters can be computed. The information of these parameters is essential for analyzing their effect on the waveform's shape.

To compute the essential surface topography parameters (slope and roughness) the

ATM validation data set is used over the Dry Valleys, see Section 3.3. The set consists of six individual DEMs, all having a spatial resolution of $2m$. The models primarily cover the valleys and Taylor Glacier, where bare rock, lake ice and glacier ice are the main land cover classes. No elevations of regions covered by snow at the time of the DEM acquisition dates are available. Apart from gaps between the DEMs, also data gaps in the DEMs themselves occur. Nevertheless, the data set forms a valuable and relevant validation information source. Since the Dry Valleys is a meteorologically stable region, this region is used for topography analysis for several ICESat laser campaigns.

For the estimation of the surface topography parameters, the DEM elevations that are analyzed are located within the laser footprint's border. The geographical coordinates (latitude, longitude) of this border are computed with the major axis and eccentricity and are then projected into Cartesian coordinates using the same map coordinate system of the DEMs (Lambert Conformal Conic projection). To compute the slope and roughness parameters, a plane is estimated through the DEM elevations located within the footprint. The estimation is carried out with the least-squares method. Equation 6.1 shows the equation for the linear system.

$$z \approx Ax + \epsilon \tag{6.1}$$

Here, z is the measurement vector, A is the matrix of coefficients, x is the parameter vector, ϵ is the measurement error vector [29]. The $N \times 1$ measurement vector consists of N DEM observations located in the footprint. The $N \times 3$ matrix of coefficients consists of the two $N \times 1$ location vectors of the observations and a $N \times 1$ vector of ones. The $N \times 1$ parameter vector consists of the unknown parameters to be estimated. The observation model looks as follows:

$$\begin{bmatrix} z_1 \\ \vdots \\ z_N \end{bmatrix} = \begin{bmatrix} xpos_1 & ypos_1 & 1 \\ \vdots & \vdots & \vdots \\ xpos_N & ypos_N & 1 \end{bmatrix} \begin{bmatrix} xslope \\ yslope \\ offset \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \vdots \\ \epsilon_N \end{bmatrix}$$

The least-squares estimation of the parameter vector is denoted in Equation 6.2.

$$\hat{x} = (A^T A)^{-1} A^T z \tag{6.2}$$

One problem with the measurement vector is that the values (elevations) are quite large. As a consequence, the 3×3 matrix that follows from $A^T A$ in Equation 6.2 could come close to singular and could violate the working precision of MATLAB. To deal with this, a local reference frame is defined, which has its origin at a fixed location of one of the analyzed observations and it has the same axes directions. In this way, the new measurement vector and the location vectors contains only the differences with the original vectors resulting in much smaller values and the working precision is not pushed to its limits.

The parameter vector consists of the two gradients expressed in $[m/m]$ and the offset in $[m]$ from the origin. The two gradients essentially give the slope of the fitted plane in two perpendicular directions, i.e. in positive x and in positive y directions. To link the direction of the gradients to positive latitude and longitude (required for later analysis), the map coordinate system of the DEMs is used. In this case (Dry Valleys DEMs), the map axes are defined as being positive in local north and local east directions, corresponding to positive x and positive y map directions. Since the DEMs are map projected using a conformal projection and cover a rather small area, these directions are preserved [19]. Therefore the gradients can be directly applied to compute the slopes in north and east directions. Equation 6.3 shows the relationship between the gradients and the slopes.

$$\begin{aligned} \sigma_N &= \arctan \hat{x}_N \\ \sigma_E &= \arctan \hat{x}_E \end{aligned} \tag{6.3}$$

With the estimated parameter vector, the least-squares solution of the observations can be computed, see the first line of Equation 6.4. The least-squares residuals vector yields the discrepancy between the adjusted elevations and the actual elevations, shown in the second part of the equation.

$$\begin{aligned} \hat{y} &= A\hat{x} \\ \hat{\epsilon} &= y - \hat{y} \end{aligned} \tag{6.4}$$

Here, $\hat{y}$ is the least-squares adjusted elevations at the same horizontal location as the original DEM locations and $\hat{\epsilon}$ is the least-squares residual vector. As mentioned before, the slope of the plane is one of the three major influences on the shape of ICESat waveforms. The other topography parameter is the roughness of the illuminated surface. For this research the roughness is defined as the rms value of the residual vector between the actual DEM elevations and the estimated plane elevations within the footprint, i.e. the rms of $\hat{e}$, see Equation 6.5.

$$roughness = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \tag{6.5}$$

Similar as before, N is the number of DEM elevations that are present in the footprint, y_i is the ith DEM elevation in the footprint and $\hat{y}_i$ is the ith plane adjusted elevation at the same horizontal location. This roughness parameter gives the quadratic mean of the DEM deviations from the estimated plane.

Figure 6.1 illustrates the DEM observations and least squares fitted plane in one of the footprints over Victoria Valley. For visualization purposes, Delaunay triangulation has been carried out on the adjusted elevations to make the fitted plane more distinct from the individual DEM observations. The left plot shows the front face of the fitted plane

and part of the DEM observations. In the right plot the figure is rotated to illustrate the deviations of the DEM observations with the plane. Both plots are defined in a local horizontal reference frame for visualization purposes; the vertical axis expresses the elevations with respect to the ellipsoid. For the slope computations, the fitted plane is compared with the local horizontal plane (not shown in the plots); the roughness is dependent on the difference between the DEM observations and the adjusted elevations.

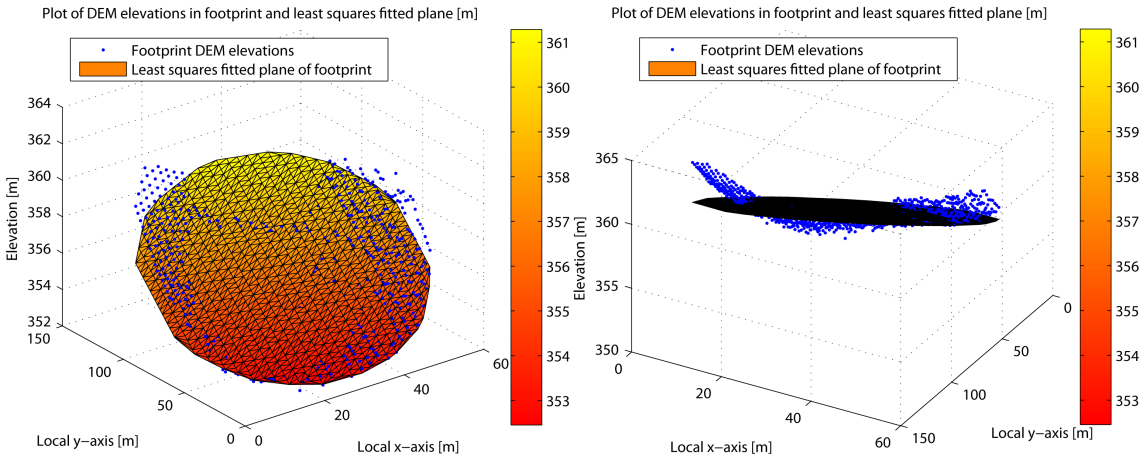

Figure 6.1: Illustration of slope and roughness analysis, demonstrating the DEM observations within the footprint and the least squares fitted plane based on the adjusted elevations in the footprint. The left plot shows primarily the face of the plane, the right plot shows primarily the deviations of the DEM observations with the plane.

The third major influence on the waveform shape is the beam orientation relative to the surface footprint plane. The combination of the slope and beam orientation then yields the beam incidence angle. The next section explains the geometry and equations involved in the incidence angle computation.

6.2 Incidence angle estimation

If the laser beam would hit the surface parallel to the local vertical, i.e. exactly nadir pointing, the slope and roughness parameters could directly be used for analysis on the waveform's shape. However, this is not the case for ICESat laser pointing and consequently the beam orientation also affects the shape. The combination of the beam orientation and slope angle of the footprint plane results in the incidence angle.

Figure 6.2 and Figure 6.3 illustrate the geometry that is involved in the incidence angle estimation. The left part of Figure 6.2 illustrates the beam co-elevation angle and the azimuth angle of the laser beam from the satellite to the Earth's surface. Here, α_B is

the beam azimuth angle minus 180° and β is the beam co-elevation angle. Both angles are extracted from the GLA06 data product, see Section 4.1. The beam azimuth angle, *i_beam_azimuth*, is defined as follows: *"the direction clockwise from north of the laser beam vector as seen by an observer at the laser ground spot viewing toward the spacecraft (i.e. the vector from the ground to the spacecraft)"* [35]. The north (increasing latitude) direction and east (increasing longitude) direction are also indicated in both figures. To get α_B, 180° is subtracted from the beam azimuth angle, because as a consequence, α_B can directly be used for decomposition of the beam in north and east components, see the next paragraph. The beam co-elevation angle (β), *i_beam_coelev*, is defined as follows: *"the direction from vertical of the laser beam as seen by an observer located at the laser ground spot"* [35].

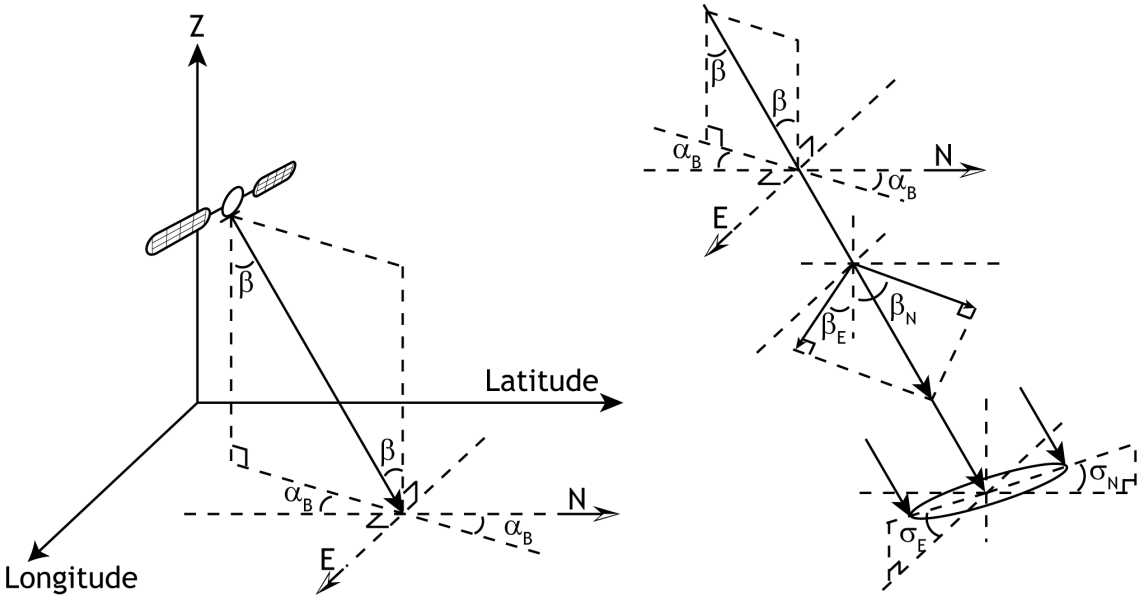

Figure 6.2: Angles of the laser beam with respect to a local terrestrial reference system (left), and beam and slope angles decomposed into north and east components (right).

In the right part of Figure 6.2, both angles α_B and β are indicated again. The geometry of the left part of the figure (where the beam hits the ground plane) is shown again. Just below this, the beam co-elevation angle is decomposed into a north and east component, β_N and β_E respectively. Equation 6.6 shows how these components are computed.

$$\begin{aligned} \beta_N &= \beta \cdot \cos \alpha_B \\ \beta_E &= \beta \cdot \sin \alpha_B \end{aligned} \tag{6.6}$$

The bottom right part of the figure shows the slope angle decomposed in north and east components of the adjusted plane within the footprint. This plane is illustrated as an ellipse with the north slope component σ_N and the east slope component σ_E. As mentioned in Section 6.1, the plane is estimated with least-squares analysis of the DEM elevations

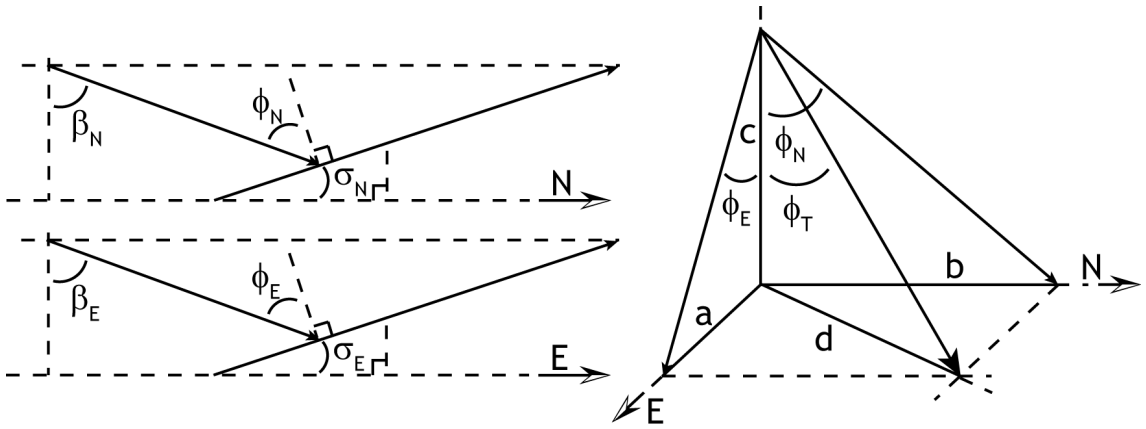

Figure 6.3: The north and east components of the beam co-elevation angle and the slope angle result in the north and east components of the incidence angle (left), and the combination of the incidence angle components resulting in the total incidence angle (right).

within the footprint. With the north and east components of the beam co-elevation angle and the slope angles, the incidence angle can be computed. Figure 6.3 illustrates the combination of these angles with the sloped plane.

As indicated in the figure, the incidence angle is defined as the angle the beam makes with the line perpendicular to the surface footprint plane. Or, in other words, the off-nadir angle with respect to the surface footprint plane. Equation 6.7 shows how the north and east components of the beam co-elevation (β_N and β_E) and the north and east components of the slope of the plane (σ_N and σ_E) combine to get the components of the incidence angle, ϕ_N and ϕ_E.

$$\begin{aligned} \phi_N &= |(90° + \beta_N - \sigma_N) - 90°| &= |\beta_N - \sigma_N| \\ \phi_E &= |(90° + \beta_E - \sigma_E) - 90°| &= |\beta_E - \sigma_E| \end{aligned} \tag{6.7}$$

From the definition of the incidence angle and Equation 6.7, the incidence angle will be zero if the beam is exactly nadir pointing from the spacecraft with respect to the horizontal map plane ($\beta = 0$) and the footprint plane is not tilted with respect to the horizontal map plane ($\sigma = 0$). On the other hand, the beam is also nadir pointing with respect to the ground plane if the north and east components of the beam co-elevation angle are equal to the north and east components of the slope, respectively. In the latter case, the beam orientation cancels out the slope orientation and consequently the beam is nadir pointing with respect to the footprint plane.

Having computed the north and east components of the incidence angle, the total incidence angle can be derived. The right part of Figure 6.3 illustrates the construction of the total incidence angle. The sides a, b, c and d help visualize the construction of this

angle. Equation 6.8 shows the simple dependency between the sides and the incidence angle components.

$$\begin{aligned} a &= c \cdot \tan \phi_E \\ b &= c \cdot \tan \phi_N \\ d &= c \cdot \tan \phi_T \quad = \sqrt{a^2 + b^2} \end{aligned} \tag{6.8}$$

From Equation 6.9 the simple expression for the incidence angle can be derived, see Equation 6.9.

$$\phi_T = \arctan(\sqrt{\tan^2 \phi_E + \tan^2 \phi_N}) \tag{6.9}$$

6.3 Beam co-elevation angle correction

The GLAS instrument of ICESat mostly operates in nominal mode, i.e. near-nadir pointing of the laser with a beam co-elevation angle of about $0.3°$, pitched upward in the direction of the spacecraft motion. However, this angle can be increased up to about $5°$ (in any direction) for illuminating targets of opportunities, like calibration and validation sites. The Dry Valleys area in Antarctica also serves as a validation site and ICESat deliberately carries out off-nadir pointing to make the tracks overlap as much as possible. Before some of the waveform derived attributes are computed, beam co-elevation correction is carried out.

The shape, and in particular the width, of the returned pulse is mainly affected by the incidence angle and to a lesser extent the roughness. The exact relationship between these surface topography parameter and the width is discussed in Section 6.4. Since incidence angle is dependent on the beam co-elevation angle and the slope angle, keeping the slope angle constant and increasing the beam co-elevation angle will change the width of the returned waveform. The amount of this increase is important, because it could significantly change the attribute values that are analyzed per land cover type. For example, a beam co-elevation of $0.3°$ over a flat lake will result in a narrow return peak, while a beam co-elevation angle of $4°$ illuminating the same surface will result in a wider return peak. Consequently, when looking at the waveform characteristics over lake surfaces, the width values indicate that the waveform width over the lake surface changes. Correcting for the beam co-elevation angle would remove this effect and consequently only the slope and roughness of the footprint influences the waveform width.

The effect can be modeled by looking at the direct effect of a tilted beam hitting the surface. For the explanation below, a flat horizontal surface is assumed, while varying the beam co-elevation angle, see Figure 6.4. The footprint plane of the beam is perpendicular to the beam center line, which means that the time difference between the first and the

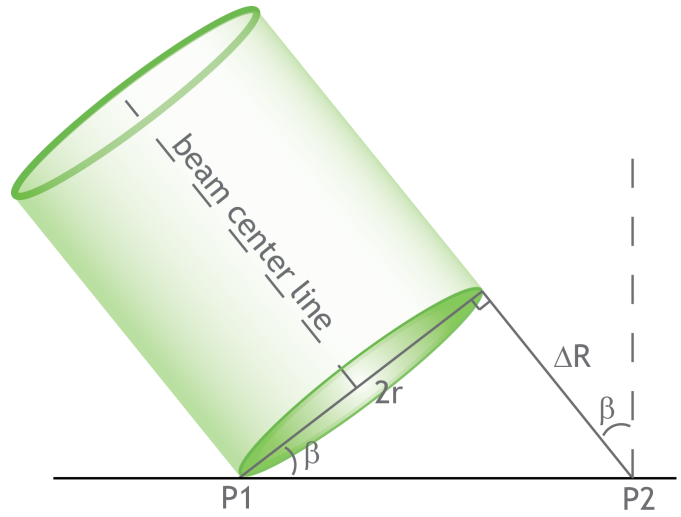

Figure 6.4: The incoming laser beam hitting the surface, the dark green plane is the footprint of the beam. The extra range of part of the incoming laser beam is induced by a beam co-elevation angle.

last beam ray hitting the ground is directly dependent on the range difference between the satellite receiver and point P1, and the satellite receiver and point P2.

The amount of range difference, ΔR, is dependent on side *2r* and beam co-elevation angle β, as illustrated in the figure. The size of side *2r* is dependent on the footprint dimensions and its orientation with respect to the laser beam's azimuth angle. Figure 6.5 shows the geometry involved in the computation of side *2r*.

The left part of the figure shows the footprint in a horizontal plane with the north, east and z-directions for reference purposes. As used before, α_B is the beam azimuth angle minus 180° and β is the beam co-elevation angle. The transmit pulse azimuth angle is α_{TP} and is defined as the angle the major axis makes with north, measured clockwise. In Figure 6.5 a and b are the semi-major and semi-minor axes, respectively. The polar angle θ in the right part of the figure is defined as the angle the radius r of the ellipse makes with respect to the north direction, also in a clockwise (east) direction. This is in agreement with the definition of the beam azimuth angle and the transmit pulse azimuth angle, both measured clockwise from the north direction. The angle θ is simply the difference between α_B and α_{TP}. When expressing the equation of the ellipse, first part in Equation 6.10, in polar coordinates, second part, the radius r can be computed.

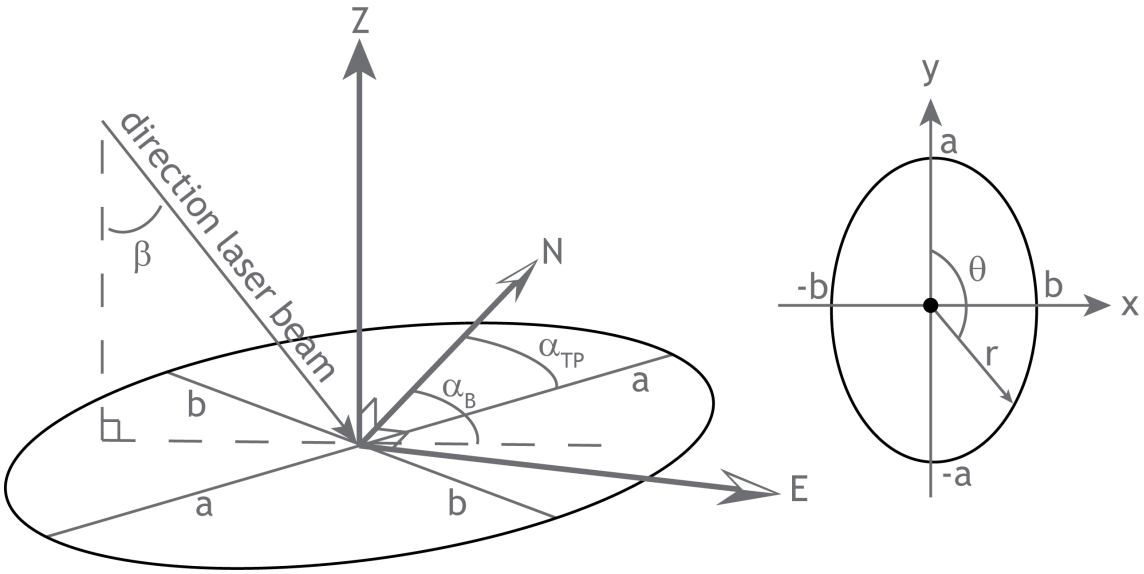

Figure 6.5: The geometry of the incoming laser beam combined with the orientation of the footprint axes, shown in the left part. The right part shows the geometry involved in the computation of side *2r*.

$$\begin{aligned} \frac{y^2}{a^2} + \frac{x^2}{b^2} &= 1 \\ \frac{(r\cos\theta)^2}{a^2} + \frac{(r\cos\theta)^2}{b^2} &= 1 \\ r &= \sqrt{\frac{1}{\frac{\cos^2\theta}{a^2} + \frac{\sin^2\theta}{b^2}}} \end{aligned} \tag{6.10}$$

The variables a and b are the footprint semi-major and semi-minor axes, respectively. The range and time difference, ΔR and Δt respectively, can consequently easily be calculated with Equation 6.11.

$$\begin{aligned} \Delta R &= 2r \cdot \tan\beta \\ \Delta t &= 2 \cdot \frac{\Delta R}{c} \end{aligned} \tag{6.11}$$

The time difference computation in Equation 6.11 incorporates the two-way travel time of the beam and the speed of light, c. Since the samples of the returned waveform represent the returned energy per nanosecond, the time difference should be expressed as nanoseconds as well for direct correction of the returned pulse.

Not the entire waveform as available in the GLA01 product is relevant for width correction, since it also includes noise sections. Therefore, only that part of the waveform

which is between the beginning and end sample of the width is corrected. See Section 5.5 on how these two samples are determined. To make proper reference to this waveform part in this discussion, the term 'returned pulse' is used and basically represents the continuous part of the entire waveform that is between the beginning and end sample of the width. It is assumed that this is the part that results from the reflecting of the transmit pulse off objects along the lasers path line-of-sight and does not include the near-flat noise part before and after. Figure 6.6 shows the flow diagram of the correction algorithm to illustrate each step in the process with numbered references. This process is explained step by step in the next paragraphs. The identification of the returned pulse is referred to as **step 1**.

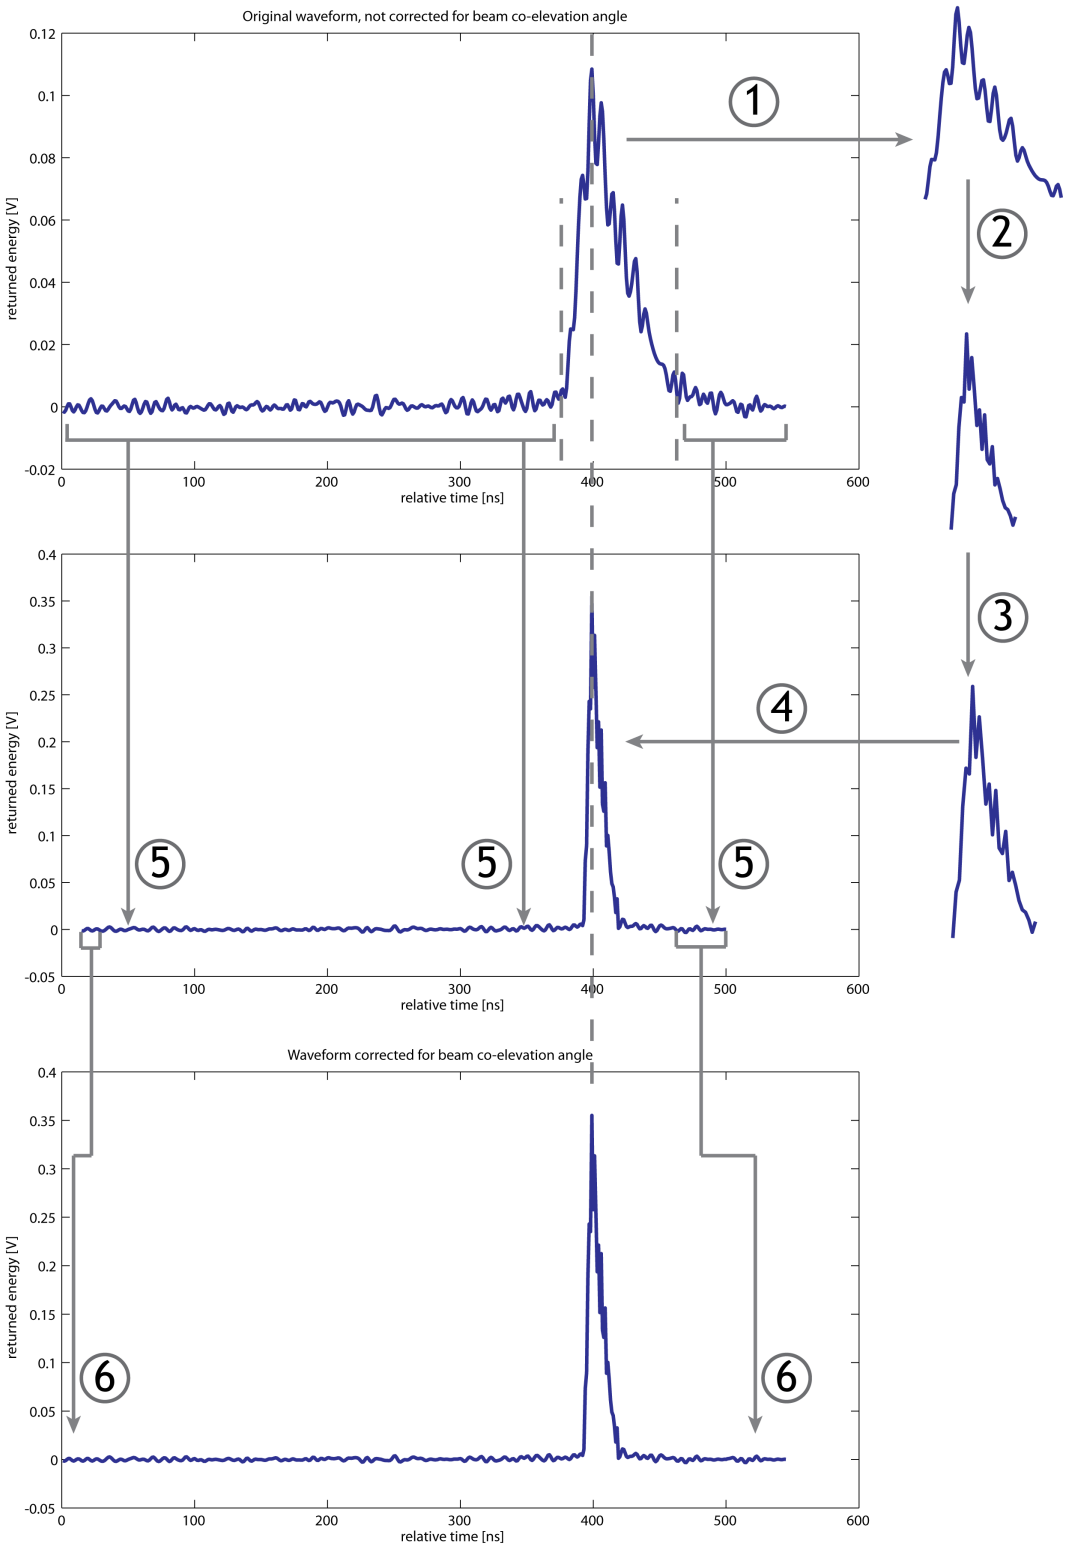

Figure 6.6: The flow diagram illustrating the steps that are taken to correct waveforms for beam co-elevation angles.

The returned pulse is then compressed to correct for the widening effects caused by the beam co-elevation angle. The compression ratio is defined as the width of the original returned waveform over the corrected width. Equation 6.12 shows the computation of the ratio, with Δt as determined with Equation 6.11

$$compression\ ratio = \frac{width_{returned\ pulse}}{|width_{returned\ pulse} - \Delta t|} \tag{6.12}$$

The original returned pulse is then compressed by linear interpolation of the original returned pulse at intervals equal to the compression ratio. In the flow diagram this step is referred to as **step 2**. When compressed, the returned pulse will have approximately the same return energy values but over a smaller time frame. As a consequence, the sum of the compressed returned pulse energy values is not the same as the sum of the original returned pulse energy values. To correct for this, a simple factorization is applied, see Equation 6.13. Please note that the returned energy values are based on the recorded waveform values and are not the same as the actual received energy because the waveform values have not been scaled to receiver gain.

$$pulse_{corrected} = pulse_{compressed} \cdot \frac{\sum_{i=1}^{N} pulse_{i,original}}{\sum_{i=1}^{M} pulse_{i,compressed}} \tag{6.13}$$

Here, the $pulse_{corrected}$ represents the energy values of the compressed returned pulse corrected for the sum of returned energy, $pulse_{i,compressed}$ is the ith energy value of the compressed pulse uncorrected for energy, $pulse_{i,original}$ is the ith energy value of the pulse from the original waveform, and N and M are the number of energy values of the two corresponding pulses. This step is referred to as **step 3** in the flow diagram.

To simulate the corrected entire waveform as accurately as possible, the waveform should be constructed similarly to the original entire waveform. As such, the maximum peak of the compressed returned pulse is aligned with the maximum peak of the original returned pulse. This is **step 4** in the flow diagram. Next, the part of the original waveform in front of and after the returned pulse (i.e. the noise parts) is pasted in front of and after the corrected pulse, **step 5** in the flow diagram. However, these added parts will not cover the entire waveform, because of the compression of the returned pulse results. So the beginning and end samples that are left to be filled are filled with a copy of begin and end samples of the original waveform, see **step 6**. These samples are not relevant for waveform analysis, they only contribute slightly to the computed threshold value used for statistical analysis.

Finally, the result is a waveform that is corrected for beam co-elevation angle while keeping the summation of the returned pulse values and preserving the length of the corrected waveform with respect to the original waveform. With this algorithm, the extra width induced by the co-elevation angle is corrected for, such that the total width of the

corrected returned pulse is only dependent on the shape of the transmitted pulse, the slope and roughness of the laser illuminated surface, and possible objects along the laser's path line-of-sight.

Please keep in mind that this beam co-elevation angle only affects the returned waveform signal. It does not have any influence on data product extracted values like the reflectivity, receiver gain, kurtosis, skewness and returned energy. These original values are still dependent on the returned pulse affected by the total incidence angle and consequently the beam co-elevation angle. The beam co-elevation correction is primarily used to correct for the extra width of the waveform induced by the beam co-elevation angle, because otherwise the same location on Earth would return different waveform widths and risetimes for different beam co-elevation angles while in reality the location itself does not change. Each time reference is made to waveform derived attributes like width and risetime, it will be explicitly indicated whether the corrected or raw waveform is used for computations.

To give the reader an idea about the difference in waveform width between a raw and a compressed waveform, the following examples are given. For nominal beam co-elevation angles, the beam angle correction is small (in the order of 0.3° to 0.4°) and the corrected width is similar to the width derived from the original waveform. Only in case of large beam co-elevation angles, the waveform will be significantly compressed. For 90% of all shots covering the Dry Valleys, the beam co-elevation angle is below 0.5° and consequently in most cases the beam co-elevation corrected width is similar to the width computed from the raw waveforms. Only in the cases the beam co-elevation angles are significantly larger (i.e. larger than 1°) the correction becomes increasingly effective. For example, a beam co-elevation angle of 0.3° and a circular footprint with radius $50m$ illuminating a flat surface results in a ΔR of $0.52m$ and a Δt (i.e. width compression) of $3.5ns$. A beam co-elevation angle of 4° with the same footprint and illuminating the same surface results in a ΔR of $7.0m$ and a Δt (i.e. width compression) of $47ns$. Based on the Δt, the amount of waveform compression is computed. Alternatively, when only interested in the waveform width, the original waveform width can be directly subtracted by Δt.

6.4 Incidence angle and roughness results

The previous sections described the computations involved in the incidence angle, roughness and beam co-elevation correction. This section discusses some results of the ATM DEM derived incidence angles over the Dry Valleys. The effect of the incidence angle on the returned width is analyzed. This section is of particular importance for the waveform derived attribute: waveform width. Since this attribute is eventually used for classification purposes, its dependence on surface topography should be understood properly.

This section is subdivided into three paragraphs. In the first paragraph the histograms of incidence angle and roughness are discussed. In the second paragraph the incidence angle and roughness results over the Dry Valleys are demonstrated. Finally, the third paragraph shows the effect of incidence angle and roughness on the non-corrected width.

6.4.1 Incidence angle and roughness histograms

For the incidence angle and roughness computations the ATM DEMs of the Dry Valleys are used. In total, 24408 ICESat footprints cover the Dry Valleys area, of which 6099 footprints cover the ATM DEMs. For each DEM overlapping footprint the incidence angle and roughness parameter is estimated.

The histograms of the incidence angle and roughness parameter are given by Figure 6.7. The incidence angle histogram shows that the frequency is highest around 1.5°. Since mostly a beam co-elevation angle of around 0.3° is used for the footprints in the Dry Valleys, most footprints cover low-sloped areas. The near-zero incidence angles occur at the lakes and flat glacier parts, which cover a smaller part of the Dry Valleys than the low-sloped valleys. As a consequence, the incidence angles show a drop in frequency close to zero. The more sloped parts are located towards and on the ridges. The roughness histogram also shows a drop in frequency close to zero from the maximum frequency point of view. Again, this can be explained by the fact that the flat areas are lakes and glaciers are covered by a smaller number of footprints than the rock covered areas.

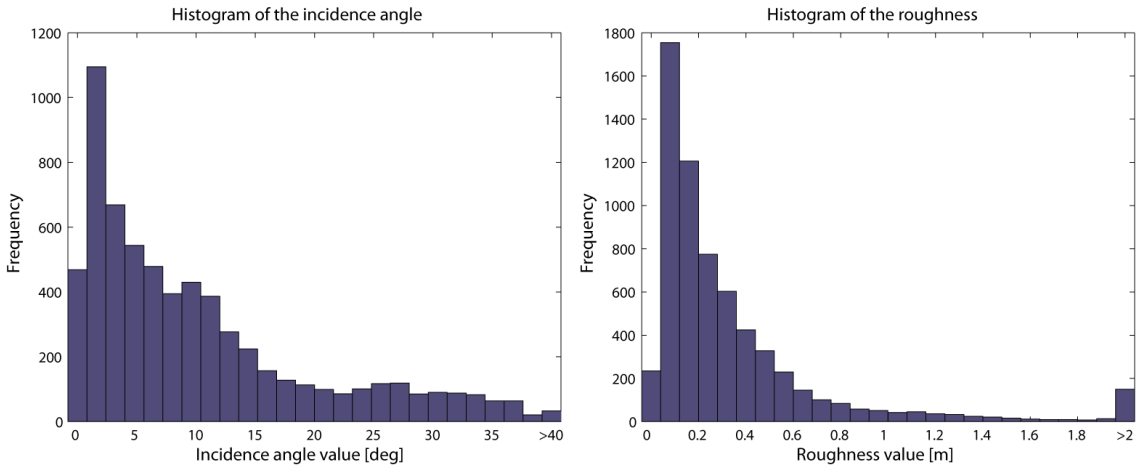

Figure 6.7: The histograms of the incidence angle and the roughness over the Dry Valleys. Both are truncated at 40° and $2m$ for visualization purposes, resulting in the exclusion of 0.4% of the incidence angle values and 2.2% of the roughness values.

6.4.2 Incidence angle and roughness examples

Figure 6.8 shows the Dry Valleys area with all footprints that cover the ATM DEMs. The colored points vary in the left plot with incidence angle and in the right plot with roughness. Especially on the lakes and glaciers, the incidence angle and roughness come close to zero. Also the valley's inner parts show small incidence angle values. Towards the ridges, the angle increases even to a maximum angle of 56°.

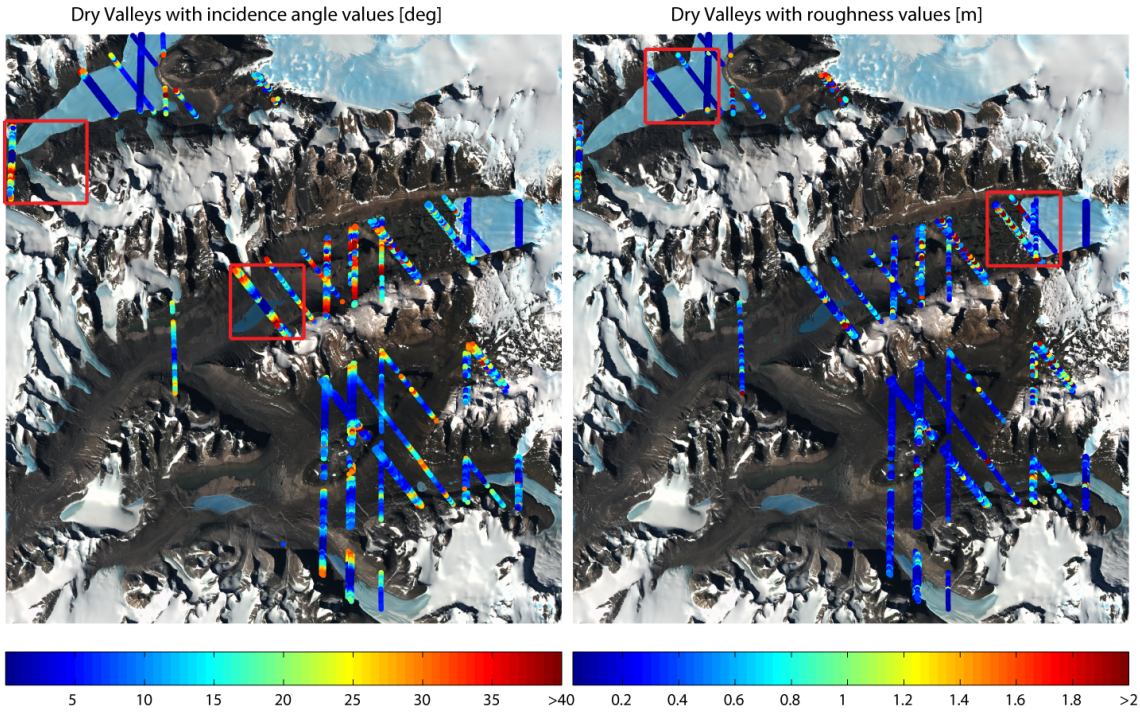

Figure 6.8: The Dry Valleys with the DEM covering footprints on top. The left plot shows varying incidence angles, the right plot shows varying roughness. Bounding boxes enclose the enlargements of Figure 6.9. The colorbars are truncated at 40° and $2m$ for visualization purposes.

To show examples of what effects land cover types can have on the incidence angle and roughness, some enlargements of interesting areas are displayed in Figure 6.9. The upper left plot shows the tip of Taylor Glacier and its adjacent lake; the upper right plot shows Lake Vanda. The increased incidence angles at the upper edge of Lake Vanda indicate that the ground surface deflects just in front of the lake. Towards the bottom of the left plot, the incidence angles become higher and near the elevated plateau even more than 35°. The same holds for the ridge in the lower part of the upper right plot. The lower plots demonstrate the effects of land cover on roughness. The lower left plot shows the glacier that invades Wright Valley. As the values demonstrate, the glacier ice generally causes lower roughness and the varying terrain left causes a mix between low and high roughness values. The lower right plot shows a part of Taylor Glacier. Interesting of this plot is the

increased roughness values right on the edges between the glacier ice and rock.

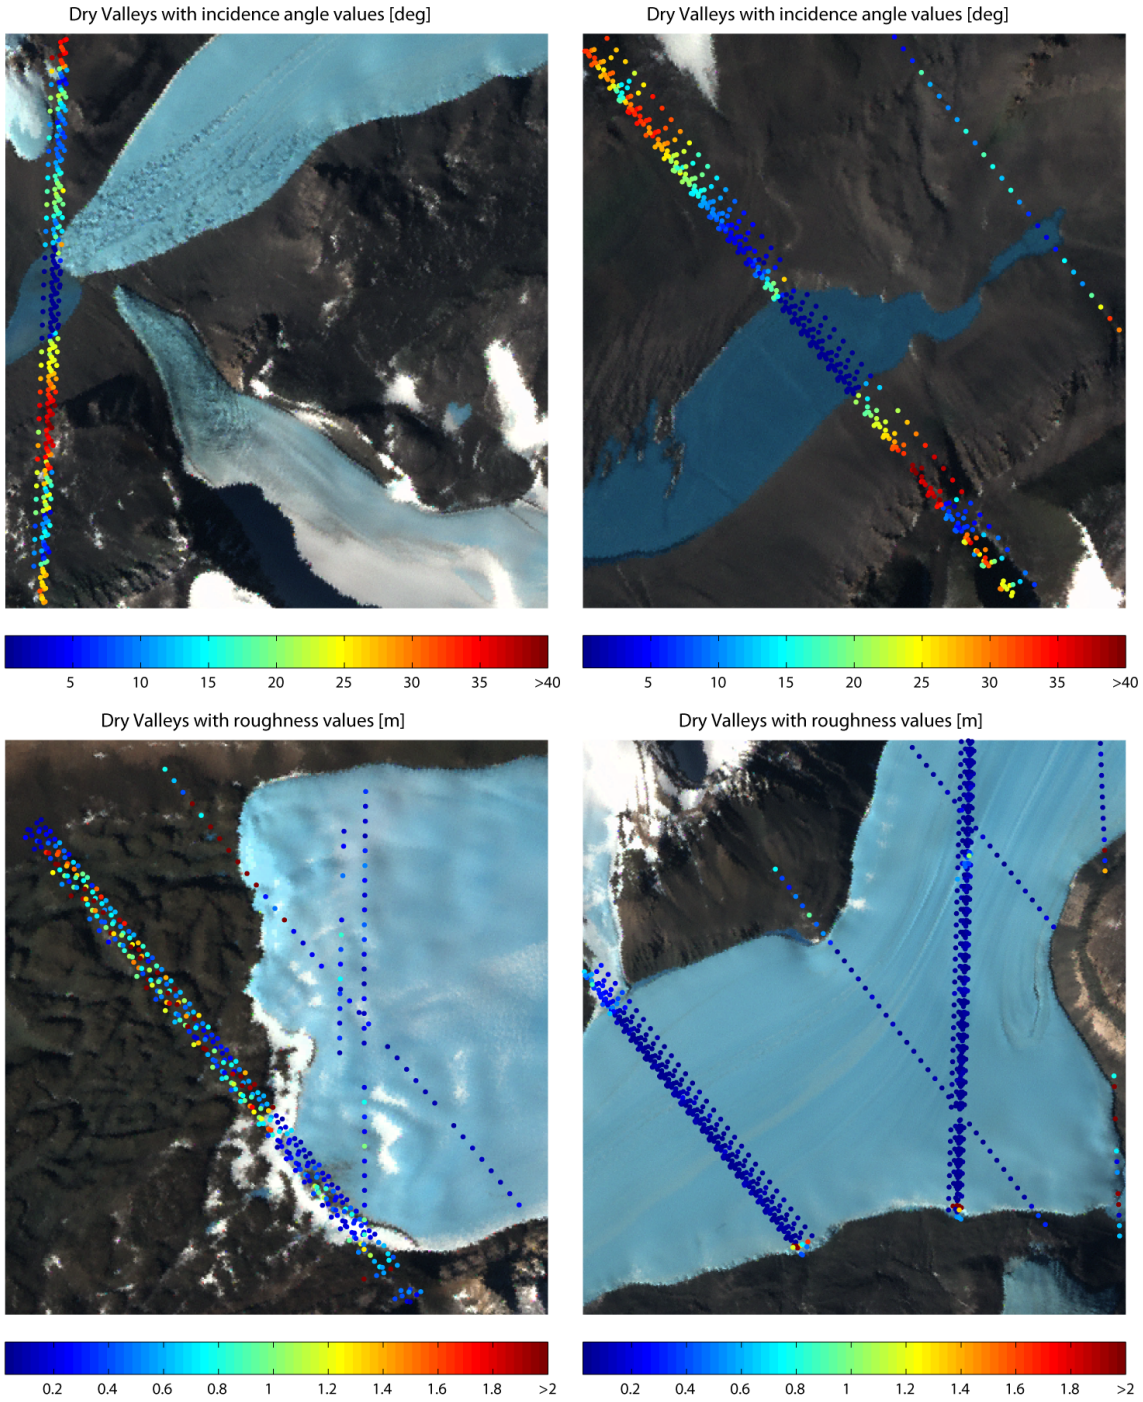

Figure 6.9: Enlargements of bounding boxes indicated in Figure 6.8. The upper two plots demonstrate the effects of land cover types on incidence angle, the lower two plot demonstrate the effects of land cover type on the roughness.

6.4.3 Effect of incidence angle and roughness on waveform width

The width of the waveform is an important indication of the incidence angle, as explained in Section 5.5. Figure 6.10 shows the incidence angles versus the width and the incidence angle induced range difference versus the width. No beam co-elevation correction has been applied to the width because this discussion and the figures concern the incidence angle. The incidence angle induced range that is used in the right plot is simply the extra range that arises from a non-zero incidence angle. It is essentially the same as the extra range induced by the beam co-elevation angle, see Figure 6.4, but instead it is induced by the incidence angle. Consequently, both beam co-elevation angle and slope angle are included in this computation. The result is the incidence angle induced range expressed in meters. This range also accounts for the two-way travel of the beam.

The induced range is computed with the mean of the footprint major axis and minor axis, and the incidence angle. For example, a footprint with a major axis of $136m$ and a minor axis of $21m$ (mean is $78.5m$) illuminating the surface under an incidence angle of $5°$, results in an incidence angle induced range of $78.5 \cdot \tan 5° \cdot 2 = 13.7m$. When the returned waveform width is $50ns$, this width in metric units is $50 \cdot 10^{-9} \cdot 3 \cdot 10^{8} = 15m$. The received waveform width (right plot) can be directly compared with the extra range that is induced by the incidence angle. One would expect that the incidence angle induced range would be about the same as the width in meters. For the first part of the figure the linear relationship indeed holds, however there is a factor two difference. The reason for this difference is not exactly known, but some possible explanations will be given later in this section. The mean value of both ellipse axes is used because the beam is rotationally independent, as explained in Section 5.2. Since the footprint slope is unknown in real-life, there is no way to tell whether a certain width is induced by the major axis covering a smaller slope angle or the minor axis covering a larger slope angle. Thus, the footprint is assumed circular for this incidence angle induced range because the orientation of the surface slope is unknown in reality. Possibly a global DEM could be used for complementary surface slope information, however this has not been considered for this research.

For both plots the mean and standard deviation are indicated for 50 bins, to show the increased variability of the values for higher incidence angle and incidence angle induced range. As can be seen in the left plot of Figure 6.7, about 60% of the incidence angles are below $10°$. In both plots of Figure 6.10, the standard deviation is relatively small below $10°$ and tends to diverge for higher angles. The mean of the binned values show that the width of the waveform does not follow a linear trend towards the higher angles. A two degree polynomial fit based on least-squares (red line) is chosen to approximate the overal trend of the incidence angle and incidence angle induced range versus the width. Since the estimation of the incidence angle from the waveform width is not the primary goal of this research, the exact reason for this quadratic trend is not analyzed. One possible reason is that the energy of the received waveform becomes low compared to the noise level for high incidence angles, such that the begining of the waveform stays below the noise threshold value for a longer time and therefore the computation of the waveform width results in a

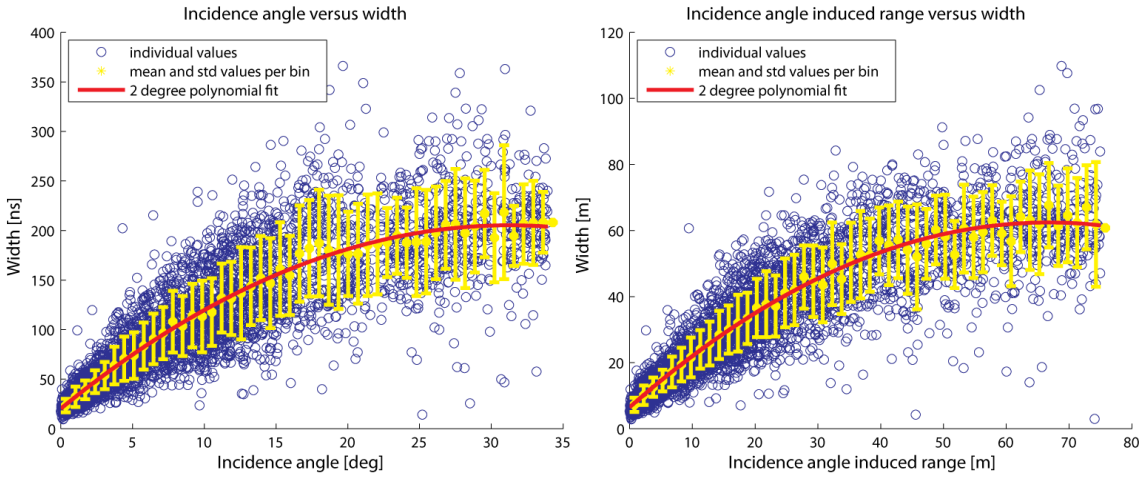

Figure 6.10: The incidence angle versus the waveform width (left plot) and the incidence angle induced range difference versus the waveform width (right plot). The yellow lines indicate the mean (dot), and mean plus and minus the standard deviation (horizontal bar). For visualization purposes, the plots are truncated at about 40° and 45m respectively, resulting in the exclusion of 1% of the total observations.

smaller width than actually is the case.

The advantage of the right plot is that the width of the waveform and the incidence angle induced range can be directly compared. As the plot shows, for small incidence angles the induced range agrees with the observed waveform width (i.e. the two variables have approximately the same values). However, for larger incidence angles the width tends to converge and even decrease. Again, this could very well arise from the conflict between decreased energy of the waveform and the threshold value. Also the standard deviation diverges with higher incidence angles. The red line of both plots can be used to indicate the relation between the waveform width and the incidence angle. What is important in this analysis is the strong correlation between the width and incidence angle. This is required for proper analysis of the effect of land cover types and its curvatures on waveform attributes like gain and reflectivity. The reason that the width and incidence angle induced range differ a factor two, is most probably due to the fact that the footprints are larger in reality. Since the footprint dimensions are based on an intensity threshold of the LPA images (see Section 4.3), the footprint in reality will be larger. As a consequence, a footprint that covers a sloped surface will result in an induced range that is larger in reality than the induced range that is computed based on the LPA derived footprint dimensions.

The roughness of the surface also affects the waveform width. Figure 6.11 shows the roughness versus the width, where the difference between the plots is only the units of the width. Since the roughness is based on the rms of the residuals, the larger the roughness, the more the surface deviates with respect to the fitted plane. As shown in the plots and in Figure 6.7, the roughness of the Dry Valleys stay close to zero. In fact, 74% of the observa-

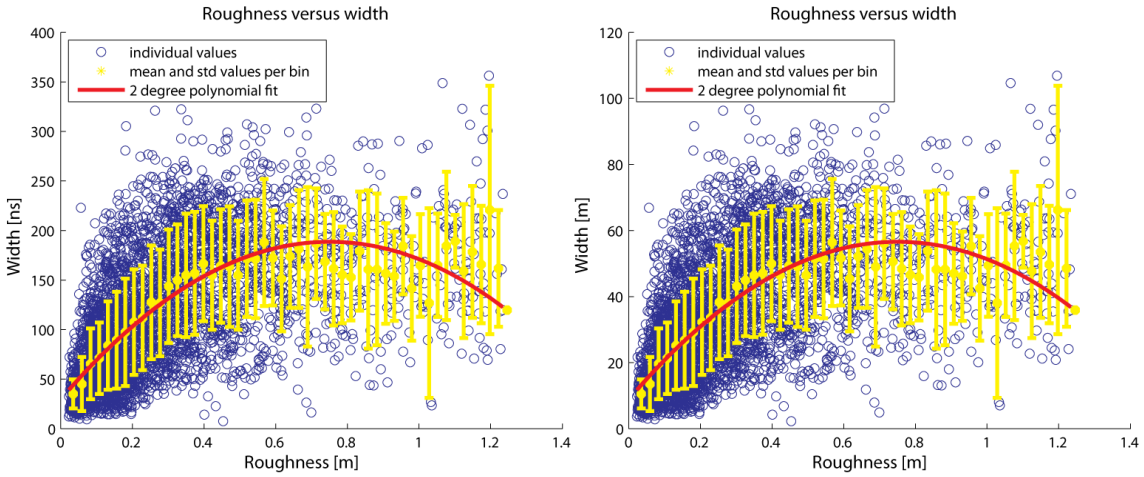

Figure 6.11: The roughness versus the width in $[ns]$ (left plot) and the roughness versus the waveform width in $[m]$ (right plot). The yellow lines indicate the mean (dot), and mean plus and minus the standard deviation (horizontal bar). For visualization purposes, the plots are truncated at about $1.3m$, resulting in the exclusion of 5% of the total observations.

tions describe roughness values below $40cm$ and, when taking into account the waveform sample recording of $1ns$ results in a maximum seperability of $15cm$, most roughness values will have negligable effect on the waveform width.

Similar to the incidence angle plots, the larger the roughness values the larger the standard deviation and the less the mean values agree with a linear trend. However, the roughness plots even show less correlation between the roughness and the waveform width. One explanation of this smaller correlation is because larger roughness occurs at the more sloped planes. The flat valleys of the Dry Valleys do not contain much deviating surfaces. And in case of large roughness, generally speaking also the surface plane is tilted. This results in the effect of roughness on the waveform width to be dominated by the effect of incidence angle on the waveform width. For example, when a footprint has an incidence angle of only 5°, the width of the waveform will already be increased with about $70ns$ or $21m$, based on Figure 6.10. A roughness value of $1m$ incorporated in the plane then contributes almost no additional width to the waveform, because the waveform is already broadened more by the incidence angle. Only in case of large surface deviations right at the point where the first or last ray of the beam hits the surface an additional broadening effect on the width results. Even though the width might not change that much with larger roughness, the waveform shape most probably will be affected especially in the increase of local peaks. However, this is beyond the scope of this research and will not be discussed further.

From the previous analysis and discussions, it is clear that the incidence angle of the laser beam with respect to the ground surface has significant effects on the waveform width. The roughness affects the waveform width as well, however in a less linear manner, due

to the more significant effect of the incidence angle in combination with the footprint dimensions. It should be noted that these observations only concern the polar regions (the Dry Valleys in particular) and will most probably not hold for vegetated areas. Since the beam co-elevation angle and surface slope angle directly contribute to the incidence angle, similar results can be expected for these angles. In the following chapters all waveform width values are compensated for beam co-elevation angle. From the results analyzed in this section, waveform width proves to be a useful waveform derived attribute for surface topography analysis. This is of particular importance for the classification analysis in the next chapters. Eventually, the width will be used in the classification algorithm to distinct shots between ice and rock for example. Thus, this chapter contributes directly to understanding the effects of surface topography on the waveform's width and this information can be directly used for classification analysis.

Chapter 7

Classification methodology

In this chapter the preparations are taken for the classification analysis. First of all, for each footprint that covers the Dry Valleys the land cover type is determined using a so-called supervised classification map. With this information, the effects of each land cover type on the waveform attributes can be analyzed. The understanding, purpose and how this map is created will be discussed in Section 7.1. The exact methodology to determine the land cover type for each footprint will be discussed in Section 7.2. After a land cover type is assigned to each footprint, the effects can be analyzed. One of the methods to analyze these is by computing so-called Kernel smoothed probability density estimates. The method of computing the distributions and how to interpret these will be explained in Section 7.3.

7.1 Supervised classification map

To analyze the influence of the land cover type on the returned waveform, validation data is required. For this research a so-called supervised classification map of the Dry Valleys is produced. With this map, each footprint is tested for the dominant land cover type it covers.

The LIMA image, acquired in November 2001, serves as the main validation image for classification purposes. The image is pan-sharpened resulting in a spatial resolution of $15m$ and incorporates near-infrared data for improved discrimination between ice and snow. See Section 3.2 for more information about the validation data sets. Furthermore, it offers a cloud-free, consistent and complete overview of the Dry Valleys region, which is used as the primary validation area for this research. Since the Dry Valleys do not offer a water basin that is ice-free year round, Lake Victoria serves as the area for validation of ICESat waveforms over water. In this section, the supervised classification map of the Dry Valleys is discussed. The footprint classification over water is not carried out, because all footprints covering the subset of Lake Victoria are assumed to cover only water.

A spatial subset of the LIMA image is used to create a classified image with the image

processing program ENVI. The classification is carried out by supervised maximum likelihood classification. The supervised part means that for each surface type one or more regions of interest (ROI) are defined manually. Each set of these ROIs represents a training class, group of pixels with spectral information, that belongs to one of the corresponding surface types, i.e. one ROI is defined that represents snow, one ROI is defined that represents rock, etc. Figure 7.1 indicates the defined ROIs superimposed on the LIMA image.

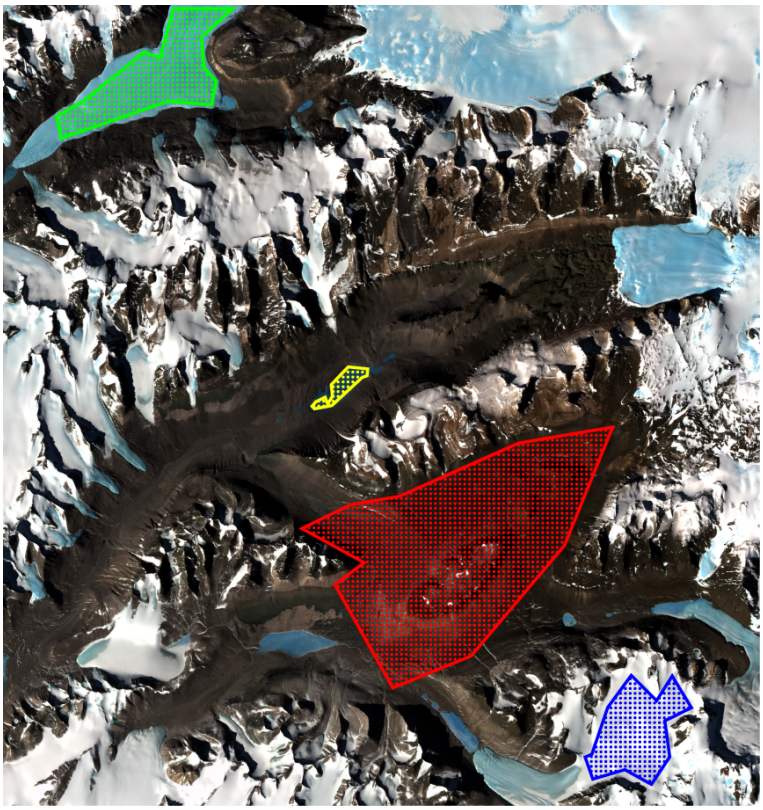

Figure 7.1: The LIMA image of the Dry Valleys, with the ROIs on top. The green polygon is the ROI used for glacier ice, red for bare rock, blue for snow, and yellow for lake ice.

Based on these ROIs, ENVI computes a classified map based on maximum likelihood estimation. The maximum likelihood algorithm estimates probability density functions based on the information of the training data that holds the spectral signature of each class. It assumes the probability is normally distributed and per pixel the probability is computed that it belongs to each class. The class with the highest probability is assigned to that particular pixel. After the maximum likelihood classification, post-classification analysis is carried out to make sure no essential parts of the LIMA image are misclassified. For this classification map only some misclassified glacier parts were assigned lake ice parts. Figure 7.2 shows the final classification map, where red represents bare rock, blue is snow, yellow is lake ice and green is glacier ice. The reason that glacier ice and lake ice are separately classified, is because these two land cover type are easily classified by the

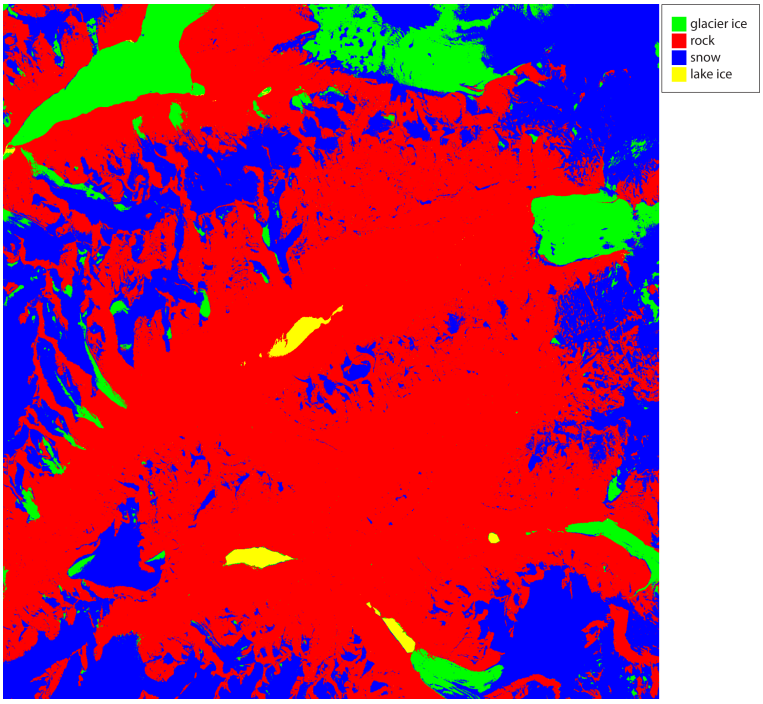

Figure 7.2: The classification map of the subset LIMA image after supervised maximum likelihood classification and post-classification.

supervised classification algorithm and post-classification. As a result, both ice types can be analyzed for the classification potential using ICESat's signals.

The accuracy of the classification map can be analyzed with a so-called confusion matrix, the overall accuracy and $\hat{\kappa}$-accuracy [12]. A confusion matrix basically represents an overview of the number of pixels that are classified correctly and the number pixels that are classified incorrectly. Table 7.1 shows the confusion matrix for the classification of the LIMA image.

		Ground truth					
Classification	**Class**	**Glacier ice**	**Rock**	**Snow**	**Lake ice**	**Total**	**Commission**
	Glacier ice	92073	4	2	67	92146	0.079%
	Rock	0	152888	24	103	153015	0.083%
	Snow	0	0	136741	0	136741	0.00%
	Lake ice	0	0	0	29820	29820	0.00%
	Total	92073	152892	136767	29990	411722	
	Omission	0.00%	0.00%	0.019%	0.057%		0.049%

Table 7.1: The confusion matrix for the classification map based on the LIMA image. The table is expressed in number of pixels and the corresponding error percentages.

For the computation of the confusion matrix, new ROIs are used to make sure that the results are not biased by using the same ROIs used for classification. In this way, the

confusion matrix offers an independent accuracy overview of the classification map. Thus, four independent sets of ROIs each representing a class are compared with the classification result. The columns represent the ground truth pixels in the independent ROIs, the rows represent resulting classification of these pixels. The table shows for example that 24 pixels are classified as rock, while they should have been classified as snow. A possible explanation for this is that the independent ROI of rock incorrectly included some snow pixels. The table also shows the commission and omission percentages. Here, the commission is defined as the percentage of the pixels that are classified as that particular class (row) but in fact do not belong to that class, i.e. the incorrect inclusion of pixels. To use rock as an example, 24 plus 103 pixels are classified as rock but should have been classified as snow and lake ice respectively. The omission is defined as the percentage of the pixels that are classified as another class but in fact belong to that particular class (column), i.e. the incorrect exclusion of pixels. For example for snow, 24 pixels are classified as rock and 2 pixels are classified as glacier ice but all should have been classified as snow. The overall percentage, 0.048% in this case, is the total number of pixels that are classified incorrectly compared to the total number of ground truth pixels. Equation 7.1 shows how the overall accuracy is computed.

$$overall\ accuracy = \frac{\sum_{i=1}^{r} x_{ii}}{N} \tag{7.1}$$

Here, r is the number of classes, x_{ii} is the value of the confusion matrix in row i and column i (on the diagonal), and N is the total number of observations (411722 in this case). The overall accuracy is equal to the sum of the four diagonal elements divided by the total number of ground truth pixels, which is equal to 100% minus the overall percentage. In this example, the overall accuracy is 99.951%. This percentage indicates the success of the classification. However this accuracy represents a biased indication of success, because it does not take into account to the chance agreement of a random classifier. The $\hat{\kappa}$-accuracy compensates for this effect and is basically the overall accuracy compensating for random classification. It is always smaller or equal to the overall accuracy and gives a more realitistic indication of the classification success. Equation 7.2 shows the conceptual equation as well as the mathematical equation of the $\hat{\kappa}$-accuracy [12]. Essentially, the $\hat{\kappa}$-accuracy indicates the success of the classification with respect to a random classifier.

$$\begin{aligned} \hat{\kappa}\text{-}accuracy &= \frac{overall\ accuracy - chance\ agreement}{1 - chance\ agreement} \\ \hat{\kappa}\text{-}accuracy &= \frac{N \cdot \sum_{i=1}^{r} x_{ii} - \sum_{i=1}^{r}(x_{i+} \cdot x_{+i})}{N^2 - \sum_{i=1}^{r}(x_{i+} \cdot x_{+i})} \end{aligned} \tag{7.2}$$

Similar to Equation 7.1, r is the number of classes, x_{ii} is the value of the confusion matrix in row i and column i (on the diagonal), and N is the total number of observations. In addition, x_{i+} is the total number of observations in row i and x_{+i} is the total number of observations in column i. In this case the $\hat{\kappa}$-accuracy is 99.93%, which is only a fraction smaller than the overall accuracy. This accuracy can be interpreted as that the supervised

classification performs 99.93% better than a classification from chance [12]. Such a high $\hat{\kappa}$-accuracy is acceptable for using the classification map for validation purposes. Even though this number indicates the success of this classification map, it does not in any way indicates the success of using the classification map for ICESat footprint validation at later dates.

Parallel to this process, a mask is created based on the same LIMA image. Since the LIMA image was acquired in November 2001 and no ICESat waveforms are available of that moment, the supervised classification map should be used with caution. To include only the stable regions, a mask is drawn over the LIMA image. This mask excludes the more instable parts, like the snowy mountain region that could be subject to changing snow cover. This unknown factor cannot be corrected for with the other validation data, because the LANDSAT images taken after 2003 do not coincide precisely with any of the acquisition dates of the tracks, and the LTER stations are too sparsely distributed in the Dry Valleys region. In addition, the later LANDSAT images are not suitable for automatic classification because of the black stripes. See Section 3.2 for more information about LTER and LANDSAT data. Therefore, the mask is required to exclude the instable regions. It is basically a binary layer that has value 1 over stable areas and 0 over instable areas. As a result, the combination of the mask and the classified image represents a sufficiently reliable source for classification validation. In addition, only ICESat footprints are selected that were acquired close to November, the acquisition month of the LIMA image. Figure 7.3 shows the mask. The parts that the mask excludes are made transparent red and laid over the LIMA image for visualization purposes. The clearly visible parts with the annotations indicate the areas that the mask includes for analysis.

In Figure 7.3 the numbered regions indicate the parts of the mask that are selected for analysis. Region number 1 mainly covers glacier ice on Taylor Glacier, region 2 covers bare rock and lake ice in the valleys, and region 3 covers snow which lies on a rather flat plateau and is surrounded by mountain ridges. From the figure it is also visible that the mountainous regions where snow and ice are mixed with rocky parts are excluded by the mask.

To summarize the means taken to make the classification analysis as reliable as possible:

- The Dry Valleys region is selected as a validation data set, because the area is one of the driest regions on Earth; very little precipitation occurs, snow that falls in the valleys sublimates quickly and temperatures are almost always below zero (except for the mid-summer). Consequently the region forms a stable validation area.
- Accuracy assessment of classification of the Dry Valleys indicates a $\hat{\kappa}$-accuracy of 99.93%, which is acceptable to serve as an initial validation map for ICESat data.
- A mask is created that excludes the more instable parts of the Dry Valleys, which are generally speaking the mountain ridges, where land cover types snow, ice and rock unpredictably change over time.

Figure 7.3: The mask superimposed on the LIMA image. The transparent red area covers the regions of the Dry Valleys that are left out for analysis. The numbers indicate the regions that are used for analysis.

- Only the ICESat waveforms and footprints are analyzed for Dy Valleys classification analysis that were taken in the same season as the acquisition month of the LIMA image, which is in the Antarctic spring. As a consequence, the selected campaigns are: L2A, L3A, L3D L3G and L3I.
- Post-classification analysis with LANDSAT images and LTER that were acquired close to ICESat campaign dates can be used as complementary information to the LIMA supervised classification map.

7.2 Footprint classification

After the estimation of the classification map and the construction of the mask, the resulting images are analyzed with MATLAB. In MATLAB only those footprints are selected that cover the valid parts of the mask and were acquired close to the LIMA acquisition month. To analyze the dependency of waveform characteristics on the land cover class, for each footprint the surface class is computed, based on the supervised classification map. This section explains

step by step the footprint classification procedure.

The supervised classified map is exported from ENVI, with which the supervised classification was carried out, as a GeoTIFF image, which is basically a georeferenced raster image. It consists of the image information, the indexed raster image and the colormap. The image information includes characteristics that describe the image properties and cartographic information. The most relevant fields incorporated in GeoTIFFs are the projection, geographical coordinate system, datum, reference ellipsoid, and the referencing matrix that is used to transform pixel row and column coordinates to map coordinates.

After the GeoTiff has been imported in MATLAB, first of all, the geographical coordinates of the footprint with respect to the WGS84 ellipsoid are computed according to the footprint major axis and eccentricity. These coordinates are then converted to Cartesian coordinates in the same map coordinate system of the classification map (polar stereographic projection). With these coordinates and the pixel size (which is $15m$), the bounding box around the footprint is defined. Using its coordinates and the pixel center locations, all pixels are selected that lie inside the bounding box. Figure 7.4 illustrates the how the land cover is determined of a footprint that covers multiple land cover classes. The bounding box around the footprint is enlarged with one pixel. The magenta dots show the outline of the footprint. Since the pixel center locations inside the bounding box are known, the pixel values can be retrieved as well. These values represent the class of the land cover at that location; they are colored in the figure according to the same color scheme used in the classification map of Figure 7.2.

The land cover classes, enclosed by the bounding box, with their locations are now available for further analysis. As Figure 7.4 illustrates, not all pixels are completely enclosed by the footprint and consequently are taken into account only partially. The percentage of the pixel overlap is computed by making bounding boxes around each individual pixel and determining the intersection of each pixel with the footprint outline. In Figure 7.4 this is illustrated by black lines enclosing pixels that do not lie inside the footprint and magenta pixels that do lie inside. Also partially intersected lines are colored likewise. For each pixel that is located partially or completely inside the footprint, the intersected polygon is defined. The area of these polygons, which is smaller or equal to the pixel area, is then divided by the total area of the footprint to compute the percentage of footprint overlap with the pixel. In other words, for each pixel that is partially or completely inside the footprint the percentage overlap is computed; the sum of the percentages is then 100%, i.e. filling the entire footprint. Since the percentages of overlap are now computed and the class each pixel belongs to is known, the sum of the percentages per class gives the footprint classification result, expressed in a percentage per class.

For further waveform classification analysis, only the footprints are selected that cover a single class for 95% or more. If this is the case, the prevailing class is assigned as a single class to the footprint, i.e. becomes 100%. In case of Figure 7.4, the footprint is classified as 24% rock, 53% glacier, 0% snow and 22% lake ice, and is consequently not selected for

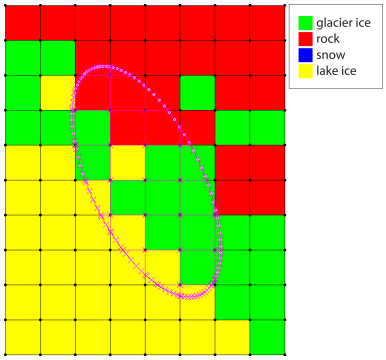

Figure 7.4: Footprint ellipse (colored magenta) over the classified pixels. The black boxes enclose the individual pixels; the magenta boxes enclose the (part of) pixels that are located in the footprint.

further analysis.

7.3 Kernel smoothed probability density estimate

To analyze the effect of land cover types on the ICESat's signals, several tools can be used to determine and visualize these effects. For example, when comparing the histograms of waveform attributes for each land cover type, the differences can be analyzed. For this research, not histograms are used for this kind of comparison, but Kernel probability density estimates using MATLAB's 'ksdensity' tool.

The 'ksdensity' tool is a function offered by MATLAB and is used for this research to visualize the distribution of important parameters per land cover type. The following information is interpreted from the MATLAB 'ksdensity' source code and its main reference [1]. It basically carries out histogram smoothing using a normal Kernel function and subsequently computes the probability density estimate. For optimal explanation of the function, the reflectivity parameter, *i_reflctUncorr* see Section 5.1 is used as an example. For this particular example, only the shots have been selected that cover snow and have a gain value lower than 50 counts.

The MATLAB tool requires two input vectors and offers the specification of many function control parameters, however only one control parameter has been used for this research. The first and most important function input is the vector to be analyzed, which is basically the attribute value vector of the specific land cover type. For the example, this vector is the reflectivity of snow: the vector contains 127 values with a mean value of 0.73 and a standard deviation of 0.099. Figure 7.5 shows two histograms of the vector, the first using 50 and the second using 10 bins between the minimum and maximum values. As the left plot shows, the small number of observations in combination with the narrow bin

width results in gaps appearing between some bins. One of the most probable causes is random variations in the data. One would expect, for example, that the gap at about 0.6 would be filled if more snow observations were available. On the other hand, when a larger bin width is taken, the detail is compromised resulting in loss of information, see the right plot. This trade-off between high level of detail and suppressing the influence of random variations is optimized by the 'ksdensity' tool. The tool essentially smooths the histogram of the observations and estimates its probability density function.

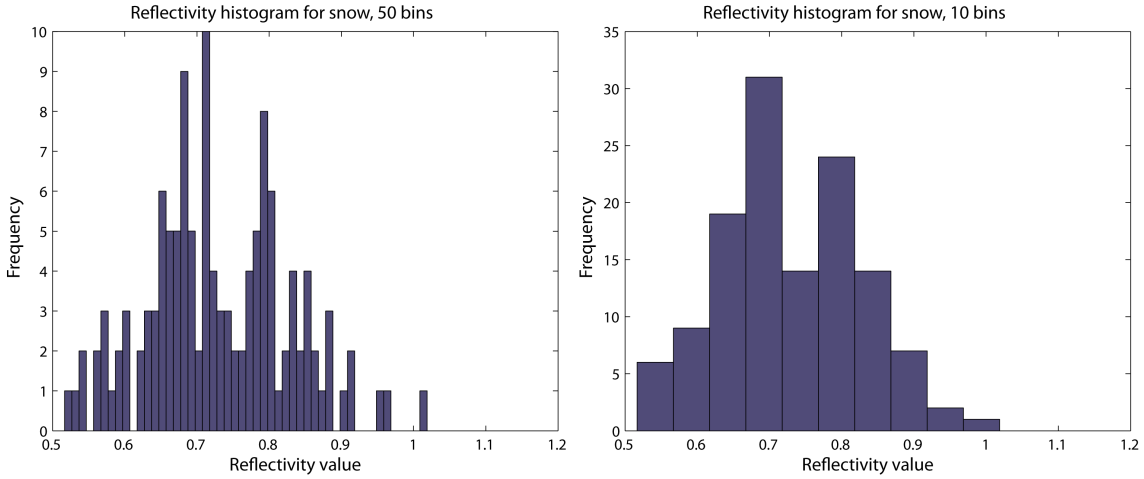

Figure 7.5: The histogram of the reflectivity of snow, using 50 bins in the left plot and 10 bins in the right plot.

The general form of the Kernel density estimator is given by Equation 7.3 [1].

$$\hat{f}(y) = \frac{1}{Nh} \sum_{i=1}^{N} K\left(\frac{y - y_i}{h}\right) \tag{7.3}$$

Here, $\hat{f}(y)$ is the Kernel density estimate at point y, h is the bandwidth also known as the smoothing parameter, N is the number of observations, K is the Kernel function whose variance is controlled by h, and $\{y_1, ..., y_N\}$ is the observation vector. The smoothing parameter, h, determines the amount of smoothing by the Kernel function. Since generally speaking the observations, for example the snow reflectivity, are assumed to be normally distributed, the smoothing parameter is used that is optimal for a normal distribution [1]. Equation 7.4 shows this smoothing window.

$$h = \sigma \cdot \left(\frac{4}{3N}\right)^{1/5} \tag{7.4}$$

Here, h is the smoothing parameter or bandwidth over which the smoothing is applied, σ is the standard deviation of the distribution and N is the length of the vector. For the

computation of h, a robust estimate of the standard deviation is preferred above the usual sample standard deviation, to account for long-tailed distributions and possible outliers. Therefore the so-called median absolute deviation estimator, $\tilde{\sigma}$ is used [1], see Equation 7.5.

$$\tilde{\sigma} = \text{median}\{|y_i - \tilde{\mu}|\}/0.6745 \tag{7.5}$$

Here, *median* simply means the median of the vector that is computed between the curly brackets, and $\tilde{\mu}$ is the median of the observation vector y_i. The constant 0.6745 is the value of the interquartile range of a normal distribution. It is used as a correction factor to make the estimate consistent with the standard deviation of a normal distribution [27]. In case the observation vector is perfectly normal distributed, $\tilde{\sigma}$ is equal to one. When the vector is one-sided long-tailed distributed (i.e. skewed to one side) or has outliers , $\tilde{\sigma}$ will be less affected than the usual standard deviation.

Now that the smoothing parameter is known, the range of observation values (in this case the range of reflectivity values) over which to evaluate the density is determined. The minimum value is equal to the minimum of the observation vector minus three times h; the maximum value is equal to the maximum of the observation vector plus three times h. The number of points over which the probability density estimate is evaluated is by default 100 points, but is changed to 1,000 points to make the estimate a smoother function. For this specification, one of the function control parameters is used. It does not affect the resulting density estimate, but only increases the number of intervals over which the density estimate is evaluated. As a result, the vector over which the density estimate is evaluated is a linearly spaced vector with minimum and maximum values dependent on the bandwidth together with the minimum and maximum values of the observation vector, and is filled with 1,000 linearly spaced points in between. In case of the example, the minimum and maximum values are 0.5168 and 1.0188, respectively. With h equal to 0.0423, the minimum value and maximum value over which the density estimate is evaluated becomes $0.5168 - 3 \cdot 0.0423 = 0.3899$ and $1.0188 + 3 \cdot 0.0423 = 1.1457$, respectively.

The most important part of the Kernel density estimator is the Kernel function. There are several Kernel functions possible, ranging from triangular and box to quadratic and higher order functions. For this research only the normal Kernel function is considered, which is defined by Equation 7.6.

$$K(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \tag{7.6}$$

Here, $K(z)$ represents the Kernel function that is applied on value z. It is identical to the the standard normal distribution with the mean equal to zero and standard deviation equal to one. Referring back to Equation 7.3, z is simply replaced by $\frac{y-y_i}{h}$. Equation 7.3 also includes the division over Nh, which normalizes the Kernel density estimator such that the summation of the values that are evaluated over points y multiplied by the interval width is equal to one. In other words, the integral over the density estimate is always one.

Briefly summarized, Equation 7.3 estimates the probability density over each point y. The range of these points is determined by the bandwidth h that is computed with Equation 7.4, and the minimum and maximum values of the observation vector. In between the begin and end values of the range, the density is estimated over 1,000 linearly spaced points. The most important part of the estimation is the Kernel function, which computes per observation y_i the Kernel density estimate at point y, based on a normal distribution, see Equation 7.6. After taking the summation of the normal distribution values of all observation values at point y, one gets the Kernel density estimate on point y. When carrying this out for all points y, the total density estimate is computed.

Figure 7.6 shows two results of the density estimate of the snow reflectivity. The left plot shows the estimate using a smaller bandwidth than the optimal bandwidth for optimal distributions, which is in this case the optimal bandwidth divided by four (i.e. h becomes $0.0423/4 = 0.0106$). It clearly demonstrates that the level of detail is high enough to observe the gaps between some reflectivity values. The right plot shows the estimate using the nominal bandwidth (i.e. 0.0423). It demonstrates that the function has smoothed the estimate to such a degree that the gaps disappeared. For this example this is preferred, because one would expect the reflectivity values to approach a more or less normal distribution, instead of gaps appearing due to possible random variations in the data. The disadvantage of this approach is that information is lost and possibly important clusters that result from other influences than random variations (for example fresh and old snow) are not visible anymore.

The insets of the plots show that the Kernel density estimate is a superposition of the standard normal distribution of the observations estimated by the Kernel function. Especially the left plot clearly demonstrates the relation between the normal distributions and the density estimate. For example, the normal distribution of only one value at 1.02 causes the density estimate to increase locally. The right plot shows that the larger bandwidth makes the normal density functions become increasingly spread out and consequently have more overlap with each other. As a result, the Kernel density estimate becomes smoother and shows less detail.

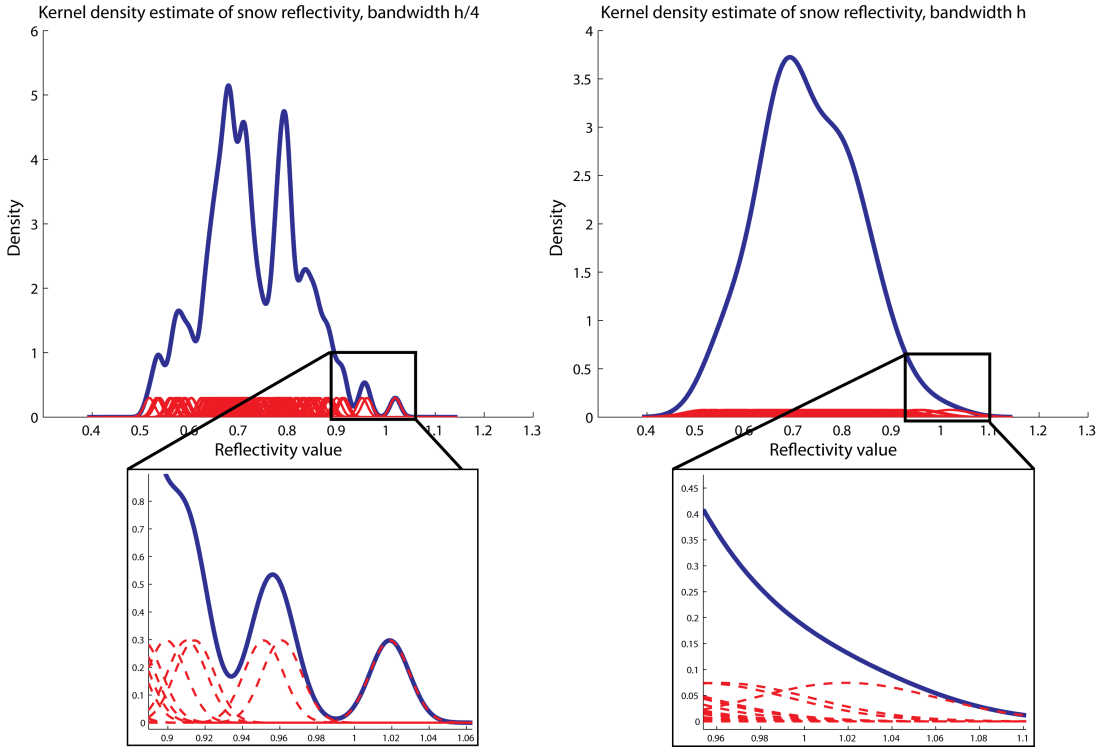

Figure 7.6: The Kernel density estimates of the snow reflectivity, displayed in blue. The left plot shows the result using the nominal bandwidth divided by four, the right plot shows the result using the nominal bandwidth. The insets show that the density estimate (blue) is composed of the Kernel function estimated normal distributions (red).

Chapter 8

Constructing the classification algorithm

In the previous chapters, the preparations were made to analyze the influences of the various land cover types on ICESat's attributes. This chapter will combine all the results from the previous chapters in order to construct the classification algorithm based on ICESat's full waveform signals. Since the data product derived attributes and the waveform derived attributes are available for analysis and the classification methodology has been explained, the classification algorithm can be constructed.

In Section 8.1, several case studies of the Dry Valleys and Lake Victoria are used to demonstrate the attribute's potential for classification. In addition to the potential for classification, in some cases the attributes show irregular values that can be related to various influences. In Section 8.2, these irregularities are demonstrated with additional case studies; it is explained that these irregular values can be related to cloud presence. In section 5.4, the proposed filtering criterion for clouds was already mentioned, but in this chapter the underlying analysis will be presented. With this knowledge, (assumed) cloud affected shots are filtered out and a detailed overview of the number of shots that are used for classification analysis can be given. This overview is discussed in Section 8.3. With the selected points, the final classification analysis can be carried out. In Section 8.4, the Kernel smoothed probability density estimates, hereafter referred to as Kernel smoothed distributions, are used to discuss the effects of land cover types on various attributes. With this information, the final classification criteria are presented and discussed in Section 8.5. These criteria together form the classification algorithm.

To avoid confusion regarding the attributes, Table 8.1 provides a list of attributes that are used for classification analysis. It combines the relevant attributes that have been discussed in Chapter 4 and Chapter 5. Four self-computed attributes that are explained in Chapter 5 have been considered for classification analysis, but are not included in the detailed analysis in this chapter. First of all, $FWHM$ shows virtually no difference from *width* of the waveform and to reduce the redundancy in the coming discussions this attribute has

been omitted. Furthermore, *maximum*, *summation* and *mean* of the waveform have been omitted as well, because these attributes do not show any contribution to the classification potential of the receiver gain, *i_gval_rcv*, and reflectivity, *i_reflctUncorr*. Moreover, since the receiver gain does not adjust properly in certain cases, see Section 8.1, the *maximum*, *summation* and *mean* are considered to be unreliable for accurate classification analysis.

Attribute ID	Relevant section	Frequency	Attribute name
i40_g_bcsc	4.5	40Hz	532nm merged attenuated backscatter profile from 10.3km to -1.0km [1/(m·sr)]
i40_ir_bscs	4.5	40Hz	1064nm attenuated backscatter profile from 10.3km to -1.0km [1/(m·sr)]
i_reflctUncorr	5.1	40Hz	reflectivity not corrected for atmospheric effects [-]
i_gval_rcv	5.1	40Hz	gain value used for received pulse [counts]
i_satNrgCorr	5.1	40Hz	saturation energy correction [0.1fJ]
i_kurt2	5.5	40Hz	kurtosis of the received echo (standard) [-]
i_skew2	5.5	40Hz	skewness of the received echo (standard) [-]
Width	5.5	40Hz	Waveform width [ns]
Risetime	5.5	40Hz	Waveform risetime [ns]
Number of peaks	5.5	40Hz	Waveform number of peaks [-]

Table 8.1: The attributes that are used classification analysis with a reference to the corresponding sections. The frequency represents the rate per second of parameter occurrence in the specified data product. Units of the parameters are indicated between brackets.

8.1 Attributes potential for classification

To demonstrate the potential of attributes for land cover type classification, several case studies over the Dry Valleys and Lake Victoria will be discussed. To show the effect of varying transmit energy on the waveform attributes, for each case study the attribute's values are demonstrated and discussed for high and low transmit energy. In this section a discussion will be given why reflectivity has not been corrected for saturation as well.

The case studies show the values of all attributes of Table 5.1 and the values of the surface elevation, *i_elev*. The reason to include case studies in this chapter is because the land cover types and surface topography can be directly related to the attributes. The disadvantage is that only a small portion of the shots can be presented. The attribute values of all shots per land cover type will be discussed in Section 8.4. The following paragraph will give a brief overview of the attribute descriptions and introduce the attributes for classification analysis.

The receiver gain is a measure for the amplification of the signal in order to record the waveform properly. As the case studies will show, the gain behavior changes dramatically with changing land cover types and surface topography. In addition, it is sensitive to varying transmit energy. The reflectivity is related to the gain because both depend on the returned energy. However, in contrast to the gain, the reflectivity is normalized

with respect to the transmit energy and the range. The reason that both the gain and reflectivity values are considered is to show that both can act as complementary rather than mutually redundant attributes. The saturation energy correction is used to show whether shots are saturated or not. This is of particular importance for understanding the limitations of the receiver and its effect on other attributes. The kurtosis and skewness are data product derived attributes and are two measures of the waveform's shape. The kurtosis indicates the peakedness of the waveform, the skewness indicates the asymmetry of the waveform. The three self-computed attributes that are taken into account are the waveform *width*, *risetime* and *number of peaks*. These three attributes have been computed from beam co-elevation angle corrected waveforms. Together with the kurtosis and skewness, these attributes indicate the shape of the waveform. Finally, the ground returns of the $532nm$ and $1064nm$ attenuated backscatter profiles are included in the analysis as well. These values, referred to as green and infrared ground return signal respectively, show the backscatter intensity per spatial column unit. The green ground return values are particularly interesting for classification purposes. However, as will be demonstrated in this chapter, in practice these signals do not show significant potential for classification improvement.

In the next paragraphs, several case studies of the Dry Valleys and of Lake Victoria are shown and discussed. These case studies serve a special purpose by showing how different land cover types and varying transmit energy affects the attributes. Four case studies are included for the Dry Valleys, and one case study for Lake Victoria.

8.1.1 Dry Valleys area

For each campaign, several repeating tracks overlay the Dry Valleys. To demonstrate the potential of the attributes, one of these repeating tracks is selected and some parts are highlighted for detailed analysis. In Figure 8.1, the tracks of campaigns L3B and L3J that cover the Dry Valleys are shown. The track that is bounded by the green rectangle is used to demonstrate the classification potential of the attributes. The reason for selecting the campaigns L3B and L3J is the significant difference in transmit energy. The mean transmit energy of the L3B and L3J track is $58.5mJ$ and $16.7mJ$ respectively, see Figure 4.3. This is important, because some of ICESat's measurements are affected by this drop in energy over campaigns, in particular the gain. Thus, the case studies are used to demonstrate both the potential of using the attributes for classification and the variability of attributes with varying transmit energy. For proper comparison, the same track and the same bounding boxes are analyzed for both campaigns. The track of L3B was flown on the 15th of March 2005 and the track of L3J on the 12th of March 2008.

The track runs from the upper left to the lower right, as indicated by the arrow. The upper (first) box covers part of Taylor Glacier, the second box covers part of Wright Valley and Lake Vanda, the third box covers Barwick Valley and Mckelvey Valley, and finally the lower (last) box covers a higher elevated snow covered plane. All boxes serve a special

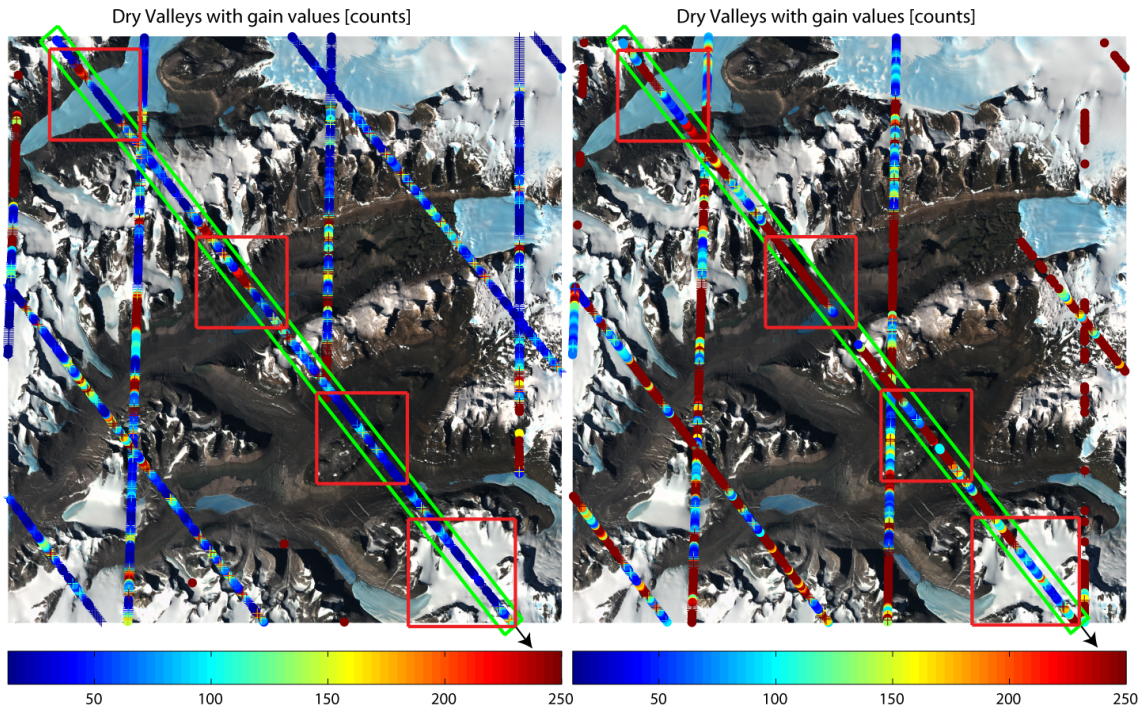

Figure 8.1: All tracks covering the Dry Valleys of laser campaign L3B (left) and campaign L3J (right). The green rectangle and the red squares in each plot highlight the track and the close-ups that are used for comparative analysis.

purpose for demonstrating the land cover type and land topography influence on the gain and several other waveform attributes. For each close-up, the results of the L3B and L3J attributes are shown on the same page to make easy comparisons per box and to demonstrate the effect of transmit energy.

Occasionally, gaps in the tracks occur, especially occurring at very high sloped ground surfaces. These gaps represent invalid values and are generally speaking caused by invalid geolocation values (latitude, longitude and elevation). Usually, invalid geolocation values indicate failed elevation estimates due to incorrect waveform recordings. However, in some cases a proper waveform is recorded and a location of the echo peak has been found. The exact reason why the geolocation values are still invalid is not known, but the figures indicate that these occurrences relate to highly inclined surfaces. A shot is considered to be saturated if the saturation energy correction is not equal to zero. This implies that a shot having invalid saturation energy correction, is considered to be saturated as well.

The vertical axes of the subplots (i.e. the ten plots on the right) have been scaled to the minimum and maximum values such that the variability can be clearly observed. However, it also implies that the axes between the case studies differ and consequently one should pay attention to these differences. The horizontal axes show the shot number from

beginning until end of the track in the direction of the arrow. After the case studies, the most important observations will be summarized in a separate paragraph.

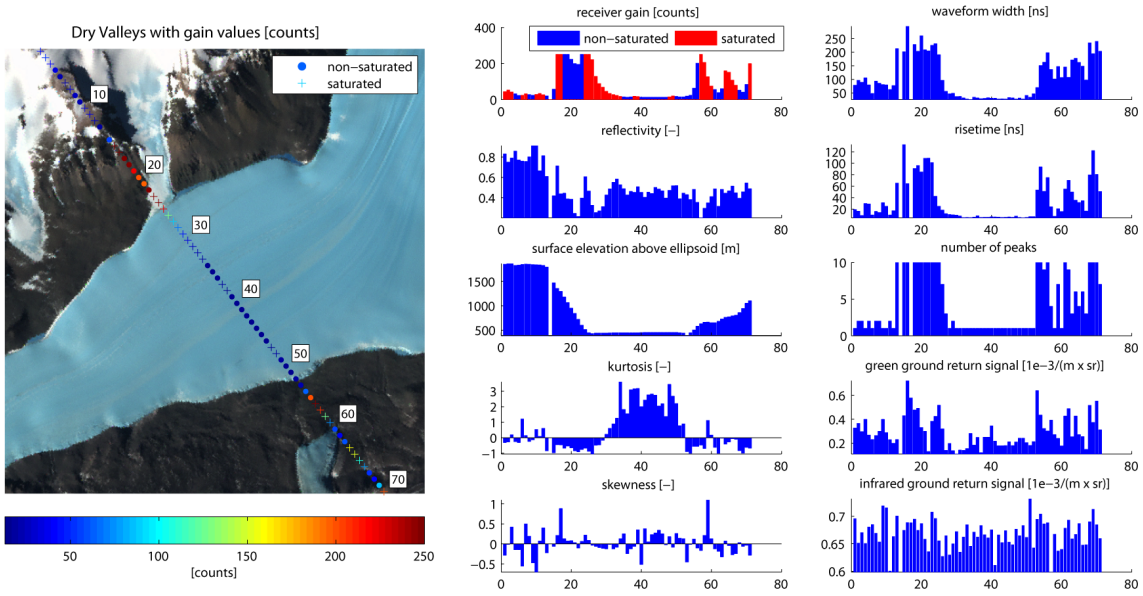

Figure 8.2: Case study 1: the close-up of the selected track of campaign L3B covering Taylor Glacier. The arrow indicates the direction of the track.

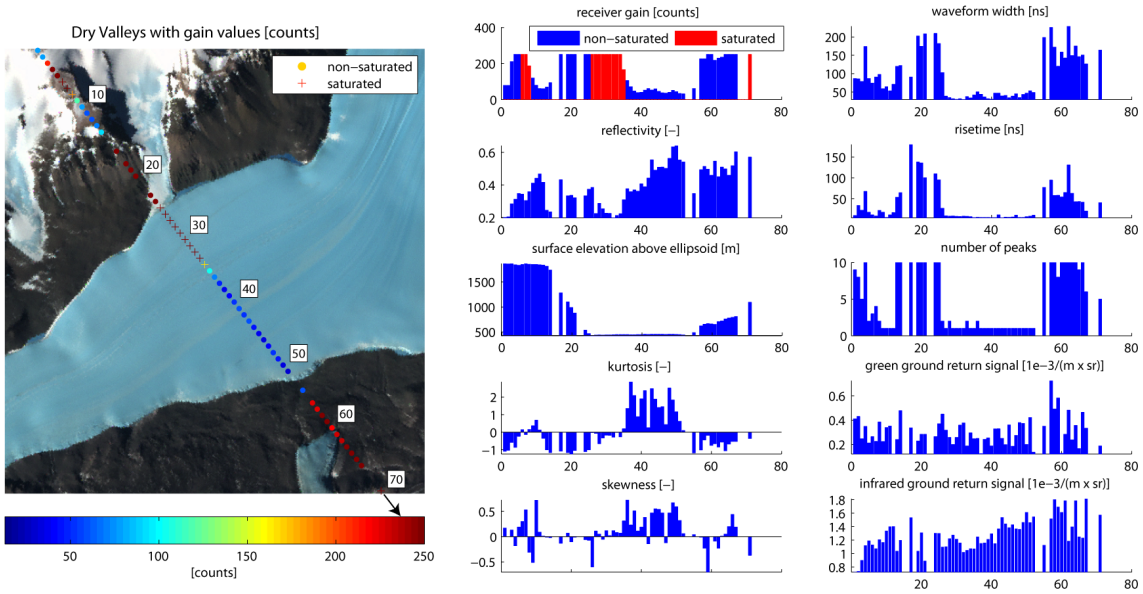

Figure 8.3: Case study 1: the close-up of the selected track of campaign L3J covering Taylor Glacier. The arrow indicates the direction of the track.

Case study 1

Figure 8.2 and Figure 8.3 show the close-ups of Taylor Glacier. Both figure show interesting gain and saturation values. Based on the elevation and waveform width, the high and low gain values are related to sloped and flat areas, respectively. The first 15 observations show extremely low gain and cover a rather flat area. Whether this part is actually rock or snow cannot be validated, however, the reflectivity shows high values (higher than 0.6) for these first observations and remain generally lower further on. Since the elevation and width plots show that these higher reflectivity values cover flat areas and the flat glacier gives lower reflectivity values, the reflectivity could be an indication of land cover type.

After the elevated flat area (around shot number 20), the gain reaches maximum values (250 counts). Based on the elevation and width values, these shots cover a highly sloped surface just before the glacier begins. Also, the reflectivity drops and the moment the shots cover the small inlet glacier just before Taylor Glacier the shots become saturated. This could be an indication that highly sloped snow or glacier surfaces return more energy than highly sloped rock surfaces. The saturation occurs due to increased energy being received by the instrument and the limited reaction time of the gain. As the figure indicates, it takes more than ten shots before the gain is adjusted properly and shots become non-saturated again. Over the glacier, the gain values remain low and the reflectivity values stay close to 0.5. The saturation and gain observations show another interesting characteristic as well; from high gain to low gain, the shots become saturated and this lasts for more than ten shots. However, from low gain to high gain, for example over rock surfaces, the gain is correctly adjusted fairly quickly, i.e. in about five shots.

The kurtosis plot shows increased values over the glacier part. Positive kurtosis means that the waveform shape is more peaked. Also interesting about the kurtosis is that it increases only after the shots become non-saturated (after shot number 30). This can be explained by the waveform shape being distorted due to the non-optimal recording (i.e. the waveform is clipped). The waveform width and risetime, on the other hand, change the moment the shots cover the flat glacier. The reason for this is that width and risetime are less affected when the waveform top is omitted. In addition, the number of peaks vary strongly between the sloped parts and the flat glacier part, from ten peaks for the rock surfaces to only one peak for the glacier surface. The infrared ground return signals do not show much variation along the track. The green ground return signals show a weak sign of changing land characteristics. More important about these signals is the magnitude, since this changes significantly with lower transmit energy.

The same part of the track taken in campaign L3J is shown in Figure 8.3. When comparing the observations of L3J to L3B, some differences come forward. First of all, the reflectivity plot shows significantly lower values for the first 15 observations compared to the L3B campaign. This is in contrast to the waveform width and risetime which show similar values for the two campaigns. The number of peaks, however, shows different values as well (5-10 instead of only single peaks). An explanation for the smaller number of peaks

could be the saturation of waveforms of L3B. Still, the lower reflectivity values of L3J are not explained by this fact, because saturated reflectivity values are underestimated rather than overestimated. Consequently, saturation would make the L3B values even higher. The exact relationship between saturated shots and reflectivity, and why the reflectivity is not corrected for saturation is explained in more detail later in this case study discussion.

Also noticeable is that the gain of the L3J shots remains high and saturated over glacier and becomes non-saturated again only after about ten shots. The big difference with L3B is that the L3B gain values gradually decrease and the L3J values drop only just before the shots get non-saturated (at shot number 35). Further on, when the shots cover the rocky area (after shot number 55), the gain reaches again near-maximum values in contrary to the same shots of L3B. The reflectivity values, however, are similar for both campaigns. A plausible reason for the difference in gain values is that the transmitted energy is more than three times as low for the L3J campaign. Here too, it appears that the gain adjusts faster from low to high gain than the other way around. Another interesting observation is the big difference between the low gain over glacier ice and the high gain over rock surface after the glacier, while the reflectivity is almost the same. From this, it appears that even for a slight decrease in returned energy, the gain has to be increased to its maximum to record proper waveforms. Apparently, the dynamic range of the gain is too limited with low transmit energy, i.e. the difference in gain is disproportional compared to the difference in reflectivity.

Looking at the kurtosis, similar results are observed over glacier for the two campaigns: the moment the shots are not saturated anymore the kurtosis becomes positive (at shot number 35). The skewness appears to increase locally for L3J, although this is not that evident for L3B. The risetime shows similar trends as the width: high values on the rock surfaces and low values on the glacier. The infrared ground return signals show a similar trend as the reflectivity, but are about twice as high compared to L3B. The next paragraph will explain how the reflectivity can be corrected for saturation and why this has not been applied, using Figure 8.3 as an example.

The reflectivity parameter shows noticeable results, as was mentioned before. The most odd reflectivity values occur in Figure 8.3: as the measurements over rock enter the glacier area the shots become saturated, due to limited response time of the gain. For at least ten shots, this saturation continues until the gain is lowered and non-saturated waveforms are recorded. The reflectivity also changes significantly during this adjustment period. Only after the saturation becomes zero, does the reflectivity over glacier appears to increase and approach the correct values. Especially during the saturated shots, the reflectivity is much lower and this is most likely due to the incorrect measurement of received energy (i.e. lower energy recorded than actually the case). On the GLAS data release description page on the NSIDC website [39], an overview of all improvements and bug fixes are explained for each new data release. One of the known problems is the effect of saturation on the reflectivity. As was shown in Section 4.2 with a detailed equation, the reflectivity is basically the received energy over the transmitted energy, both scaled by gain. The reflectivity has been

calibrated, amongst others, for gain non-linearity (i.e. is gain corrected), but this only holds for non-saturated shots. In case of saturated shots, the original reflectivity values are not computed correctly. NSIDC proposes to apply a simple ratio factor to correct for this, see Equation 8.1

$$sat.\ corrected\ reflectivity = i\_reflctUncorr \cdot \frac{i\_RecNrgAll + i\_satNrgCorr}{i\_RecNrgAll} \tag{8.1}$$

Here, *i_reflctUncorr* is the original reflectivity parameter, *i_satNrgCorr* is the saturation energy correction parameter and *i_RecNrgAll* is the laser received energy from signal beginning to signal end. However, NSIDC also notices some concerns with this approach, since the saturation corrected reflectivity values have not been fully verified. First of all, some saturation energy correction values are large compared to the received energy values, which potentially results in reflectivity values larger than one. Also, since the saturation energy corrections are determined using narrow waveforms, this correction may not be valid for wide over-land waveforms. Finally, NSIDC mentions that about 20% of the saturation energy correction values over land are not valid at all (i.e. invalid values are returned).

For this research, this correction on reflectivity has been considered. However, it has not been applied due to the following reasons, taking Figure 8.3 as an example. As the first of the NSIDC concerns already describes, some reflectivity values become unexpectedly large (i.e. even larger than 1.5) for shots that are similar to neighboring shots. It appears the correction introduces an unreliable and unpredicted effect. In addition, the fact that the saturation energy correction has not been verified for wide over-land waveforms makes the correction more unreliable. However, the most important reason not to apply the correction is the high percentage of invalid values. For Figure 8.3, more than 90% of the saturation energy correction values are invalid. And these occur not only over the sloped rock surfaces but to the same degree over the flat glacier parts. As a result, more than 90% of the reflectivity values would have invalid values after the correction (i.e the saturation energy cannot be applied), which makes the correction an inconvenience rather than an improvement for the analysis of reflectivity values over the Dry Valleys. In fact, for all shots that cover the Dry Valleys, 67% of the saturation values are invalid. The exact reason for the high percentage of invalid values is neither known nor analyzed for this research. It is not likely that it is the result from failed elevation estimates, because consequently the other attributes would have been invalid as well. A plausible reason is that the receiver cannot properly record unexpectedly high returned energy values. If in the future the saturation energy correction is validated for wide-overland waveforms and an alternative for invalid values is provided, the correction should be looked at again and possibly implemented. However, for this research this correction is not applied. The downside of not correcting the reflectivity is that the reflectivity of about 29% of all shots covering the Dry Valleys and 17% of Lake Victoria is not computed correctly. The direct effect on the reflectivity is that the values are underestimated and should actually be higher.

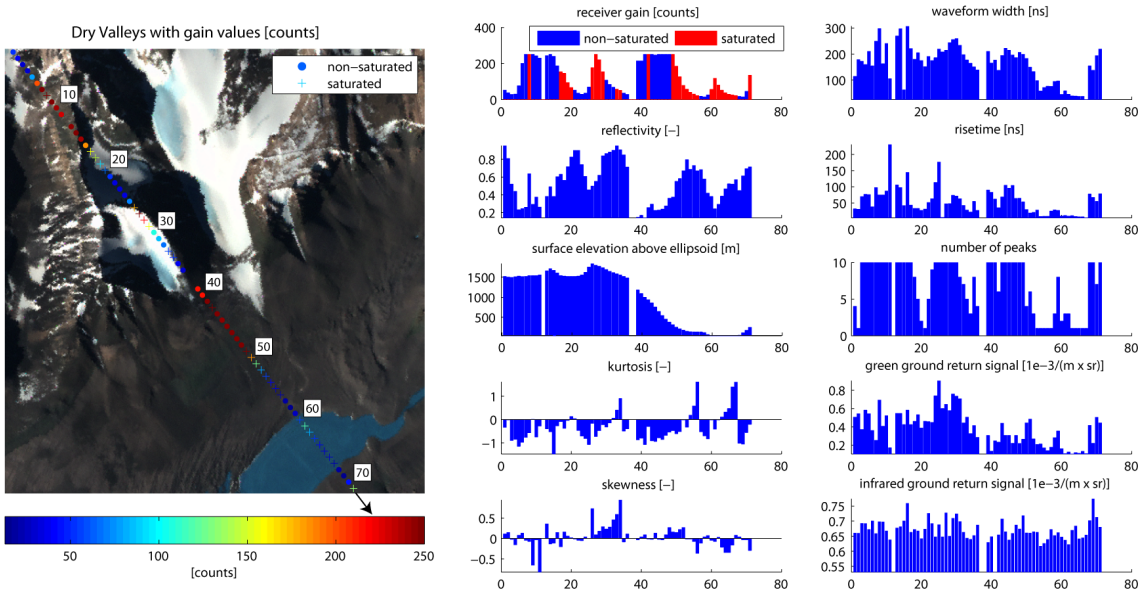

Figure 8.4: Case study 2: the close-up of the selected track of campaign L3B covering Wright Valley and Lake Vanda. The arrow indicates the direction of the measurements in the track.

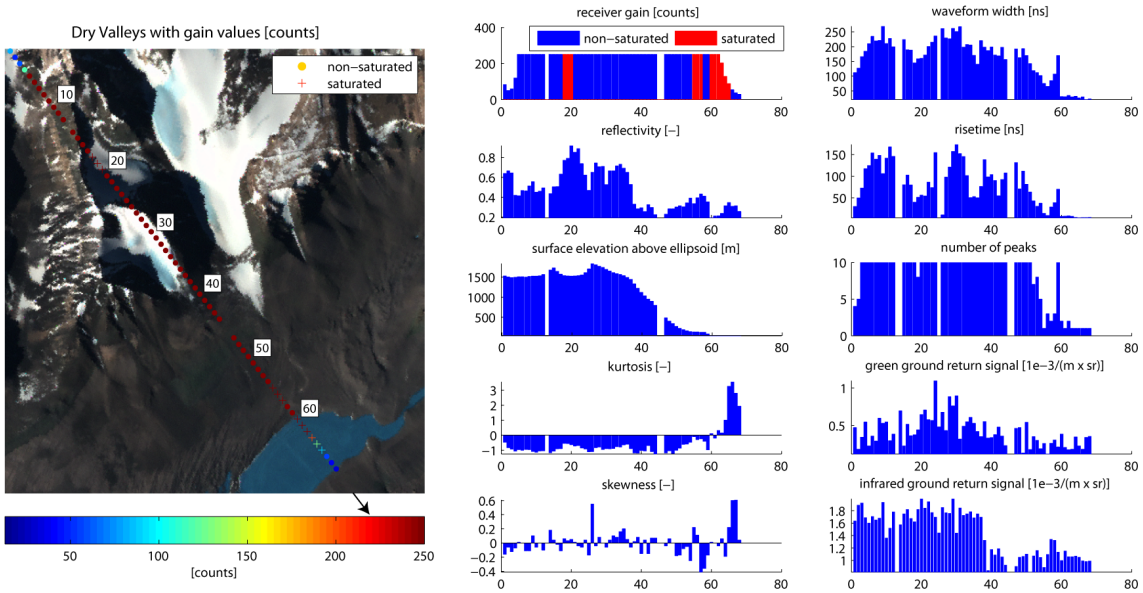

Figure 8.5: Case study 2: the close-up of the selected track of campaign L3J covering Wright Valley and Lake Vanda. The arrow indicates the direction of the measurements in the track.

Case study 2

Figure 8.4 and Figure 8.5 show observations over the dynamic mountainous area between Taylor Glacier and Wright Valley, over Wright Valley and finally over Lake Vanda. For L3B, the reflectivity values range from 0.2 to 0.9 and width values from $40ns$ to $300ns$. Interesting to notice is that the width values do not follow the same trends as reflectivity values. For example, when the shots cover the steep rock surface (between shot numbers 40 and 50), the width and reflectivity values are around $200ns$ and 0.3, respectively. When comparing these to the shots just before (between shot numbers 30 and 40), the width is about $200ns$ as well, but the reflectivity approaches almost 1. A possible reason for this is that sloped snow surfaces give high reflectivity values, while sloped rock surfaces give low reflectivity values (and consequently high gain values). The flat rock surface just before Lake Vanda (between shot numbers 50 and 60) shows low width values but high reflectivity values. However, LTER data acquired at Lake Vanda station shows that almost $4cm$ of snow fell in the end of February and still $3cm$ covered the site on the L3B track acquisition date (15th of March). In this case the temporal difference between the acquisition date of the LIMA image and the acquisition date of ICESat's tracks plays an important role. Since no LANDSAT images are available at the acquisition date of the L3B track, it cannot be fully verified which areas were covered by snow. This shows that the validation data has its shortcomings and conclusions have to be drawn with caution.

Over Lake Vanda, the width and risetime values become small (approximately $10ns$ and $5ns$ respectively) and most of the time only one peaked waveforms are detected. Since almost all shots are saturated over Lake Vanda, the kurtosis and skewness values are biased and consequently do not represent accurate figures. Similarly to Taylor Glacier, it takes about ten shots before the gain is forced down to proper values. The gain adjustment from low to high gain appears to go faster.

The first difference between L3J and L3B is the higher gain values of L3J, which is most likely due to the lower transmit energy in L3J. The reflectivity values are somewhat similar until the flat rock surface just before Lake Vanda (between shot numbers 50 and 60). Here, the L3J reflectivity values are around 0.3 (compared to around 0.6 for L3B). Based on the assumption that high reflectivity values indicate snow, this could mean that indeed snow was covering this part in L3B but not in L3J. Unfortunately at the time of writing the LTER data of 2008 had not been processed yet, so no validation can be carried out for L3J. Interesting as well, is that far less L3J shots are saturated, which means that the waveform derived attributes are more reliable. Most probably, the lower transmit energy is the main cause of less saturation. As for the kurtosis of L3J, more consistent values are observed and it appears that the higher sloped areas give a negative kurtosis (more bluntly shaped) and Lake Vanda, after the saturated shots, gives positive kurtosis (more sharply peaked).

Similar to L3B, for L3J it appears that the gain value can increase rapidly (in two to three shots for the first couple of shots) but needs more time to decrease. Even though the gain is close to maximum for a large part of the track, the reflectivity values show differ-

ent trends. Almost all shots between the shot numbers 20 and 35 have reflectivity values higher than 0.6 and further on in the track the have lower values of around 0.3, including the reflectivity values over Lake Vanda. The number of peaks shows high correlation with width, risetime and kurtosis. When looking at the green ground return signals it shows somewhat higher values than L3B but there is no evident relation with the land features. The infrared ground return signals show a slight linear relation with the reflectivity values, but more important to notice is that the values have become about three times as large compared to L3B. An exact reason for this is not known, but it could be that the lower transmit energy makes the normalized lidar signal higher because of the transmit energy normalization and consequently the values of the infrared attenuated backscatter profiles become larger. This could also hold for the green ground return signal, but the reason why these values are not three times higher is unknown. Perhaps the very low SNR plays a role.

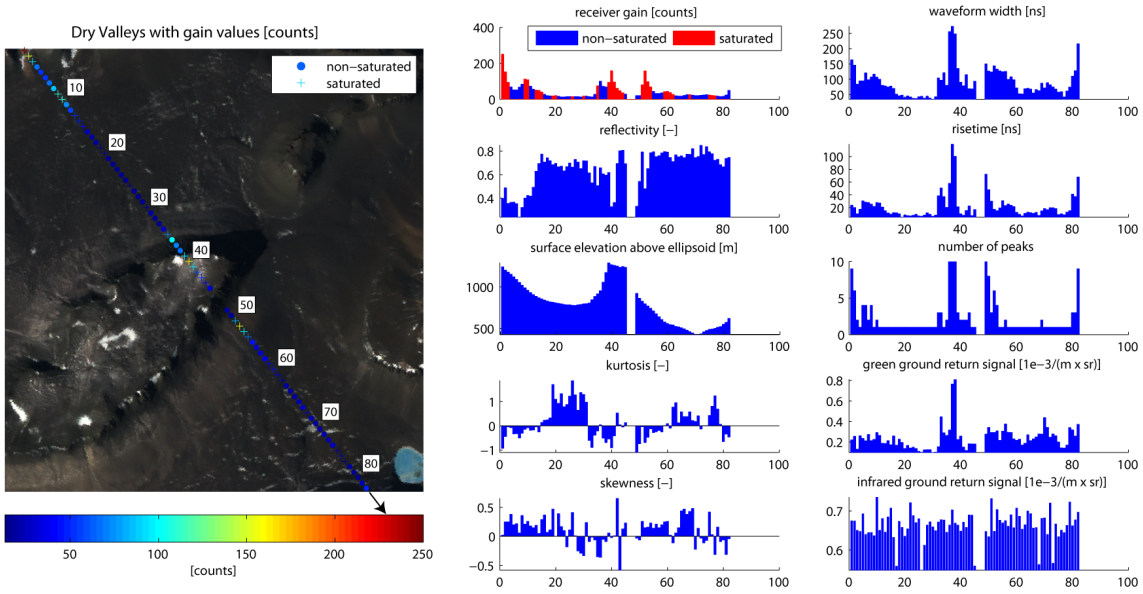

Figure 8.6: Case study 3: the close-up of the selected track of campaign L3B covering Barwick Valley and Mckelvey Valley. The arrow indicates the direction of the track.

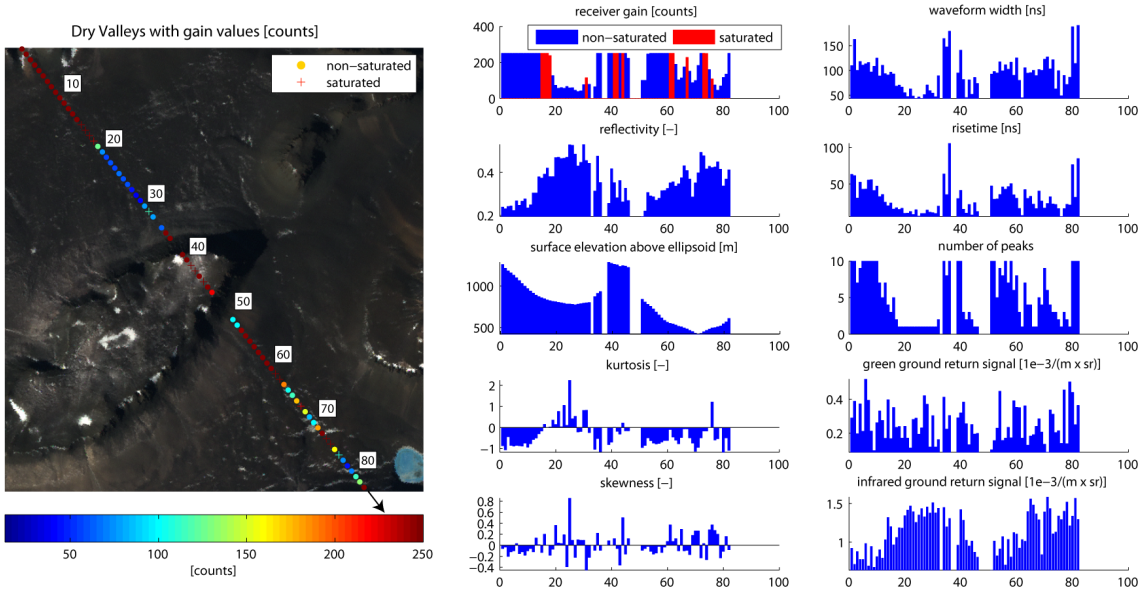

Figure 8.7: Case study 3: the close-up of the selected track of campaign L3J covering Barwick Valley and Mckelvey Valley. The arrow indicates the direction of the track.

Case study 3

Figure 8.6 and Figure 8.7 show the observations over Barwick Valley and Mckelvey Valley and are included in the analysis to show what the influence of topography is on the attributes. As the elevation, width and risetime values indicate, the area includes two flat valleys (between shot numbers 15 and 50, and between 60 and 70) and higher sloped surfaces near the elevated plateau (towards shot number 40). Interesting to notice are the high number of low gain values and the number of high reflectivity values for L3B. Even though the gain values are mostly small, still 41% of the shots is saturated. This indicates that the occurance of saturation is high even when the gain is brought down. Apparently, in a dynamic environment where the receiver applies low gain, it is not able to deal with individual shots having high returned energy. It demonstrates the low dynamic range of the receiver recording capabilities. Nonetheless, most attributes are still relevant for analysis. As was suggested before, high reflectivity values could indicate snow. Also in this case the real validation poses a problem. The LTER data of Lake Vida (which is about $8km$ away, see Section 3.2.2) also indicates that more than $4cm$ of snow fell at the end of February and the same amount of snow was there at the acquisition date of the track. Therefore, it is plausible that these valleys were covered with snow. The kurtosis shows positive values for the flat areas and slightly negative values for the sloped areas. The number of peaks show a large increase at the sloped areas and more steady values in between. Finally, the green ground return signals show values that appear to follow the risetime and width values.

The gain values of L3J are generally significantly higher and the reflectivity values lower. Saturation occurs less often as well: only 26% of the shots are saturated. One reason is that the transmit energy is lower and as a result less energy is reflected back, limiting required the dynamic range. In addition, the lower reflectivity values indicate that less energy is reflected back and consequently the dynamic range of the receiver are not pushed to its limits. The reflectivity values show that these may relate to surface topography characteristics as well. From the width and risetime values, it appears that the more sloped rock surfaces (between shot number 1 and 10, and towards shot number 40) show lower reflectivity values than the flat rock. And, similar to L3B, the kurtosis values show negative values for sloped and positive values for flat parts. The infrared ground return signals are about twice as large as for L3B and show similar trends to the reflectivity. The green ground return signals show no signs of any relationship with the surface characteristics at all.

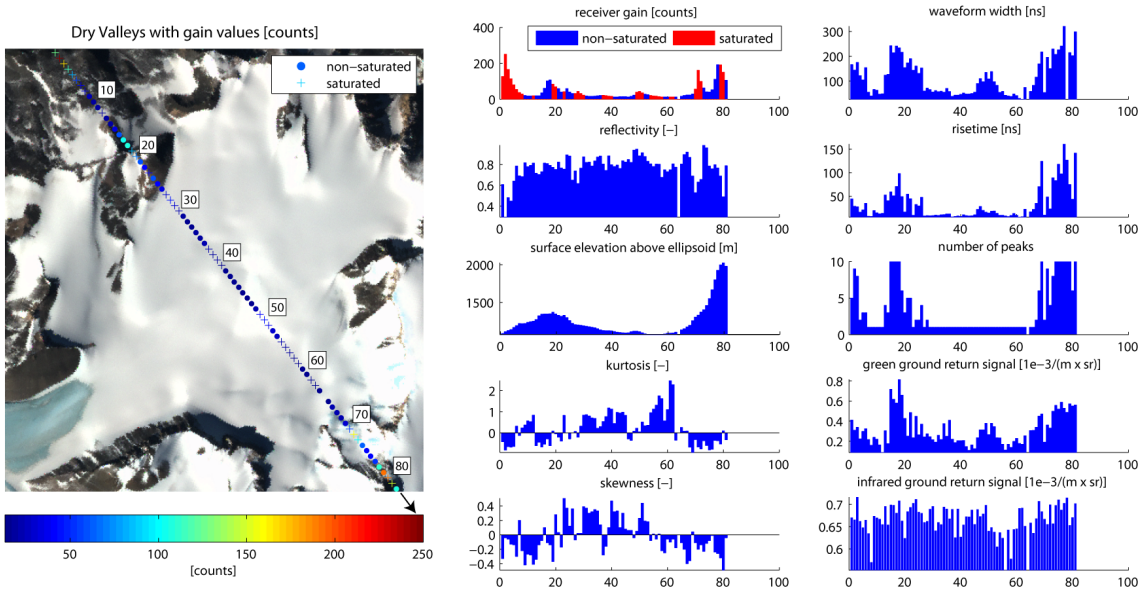

Figure 8.8: Case study 4: the close-up of the selected track of campaign L3B covering the snow plain north of Barwick Valley. The arrow indicates the direction of the track.

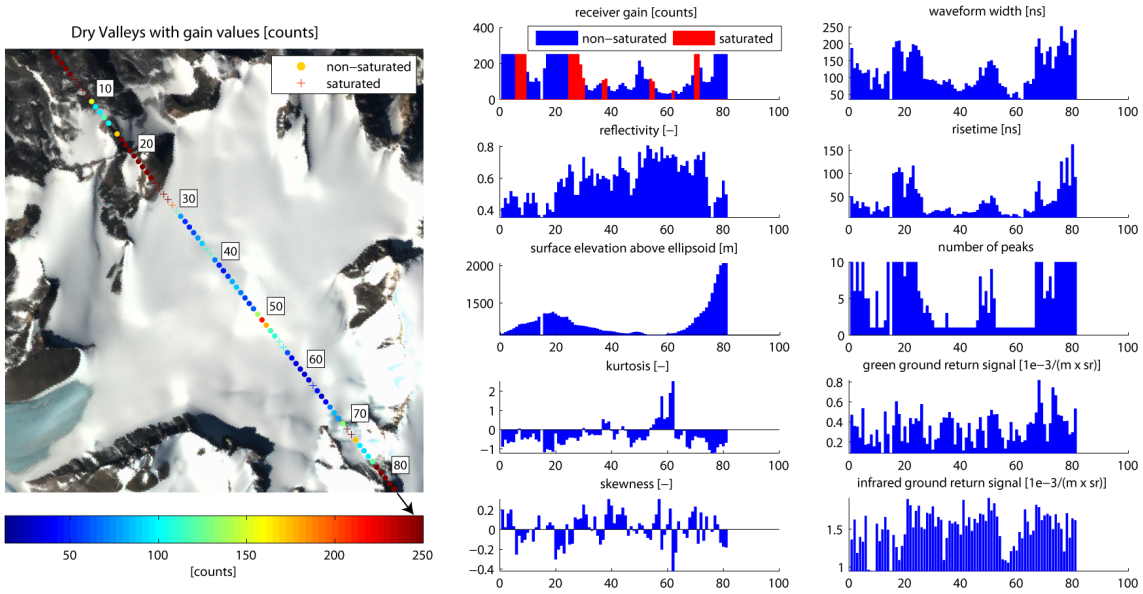

Figure 8.9: Case study 4: the close-up of the selected track of campaign L3J covering the snow plain north of Barwick Valley. The arrow indicates the direction of the track.

Case study 4

Figure 8.8 and Figure 8.9 show the final close-ups of the track of both campaigns. These figures show a plateau covered by snow and surrounded by somewhat more mountainous areas. It is assumed that the major part of the plain contains snow during both campaigns. The outer mountains showing rock surface are more doubtful since the snow extent is more susceptible to temporal differences.

The most apparent observation of the attribute values of L3B is the high reflectivity over snow. The first part of the track shows lower reflectivity values, which also happen to coincide with the (not verified) rock surfaces. Based on the elevation, width and risetime values, the main part is fairly flat except for the edges of the track. For example, the highly sloped surfaces near the last observations show high width and risetime values, but also high reflectivity values. From this information, it appears that snow under high slope angles will also result in high reflectivity values (higher than 0.6). This is in contrast to the rock under high slope that has much lower reflectivity values (around 0.3).

The number of peaks increases close to the more sloped parts and the kurtosis shows positive values for the flat areas and negative values for the sloped areas. The ground return signals show some relationship with the risetime and width, which could indicate that indeed the green laser illuminated the surface and these were received properly. Similar to the L3B observations of previous areas, saturation still occurs, even though the gain is decreased to its minimum.

The observations of L3F over the snow covered region show quite similar reflectivity values, but with higher gain values. Especially the early observations show low reflectivity and high gain values. After the gain has been decreased and the track covers the center area of the snow plain (at shot number 60), the reflectivity shows high values and the kurtosis becomes positive. Compared to L3B, the snow covered shots have higher gain and slightly lower reflectivity values. It even occurs that in the middle of the snow area (at shot number 50), where the elevations show a small bump, the gain becomes higher than 200 while the reflectivity remains the same. Also the number of peaks, width and risetime show locally higher values. This demonstrates that, even though the gain can become high (most likely due to a local sloped surface), the reflectivity value can remain the same. In addition, the observations show that not all snow covered waveforms have only one peak. The green ground signal returns show no clear relationship with land characteristics. The infrared ground signal returns are slightly correlated with the reflectivity.

Case study conclusions

In the previous paragraphs the observations of each case study were discussed and analyzed. The most important observations are summarized. First of all, the gain changes significantly between the two campaigns. This difference is most probably due to the difference in transmit energy. Even though high gain values usually occur over sloped rock

surfaces, no simple thresholding rule can be applied because of the limited reaction time of the gain. Especially from high gain to low gain, it takes about ten shots before the shots get non-saturated and the right gain level is applied. In the case of the L3J track over Taylor Glacier for example, a thresholding criterion to classify sloped rock based on gain (i.e. gain larger than 200 counts) will also include the saturated shots over the glacier. Reflectivity as a simple threshold forms a problem as well, because no saturation correction is applied. On the other hand, when separating saturated and non-saturated shots, criteria based on the attribute's values can be applied with more confidence. When taking this into account, reflectivity of non-saturated shots shows that values above 0.6 are very likely to come from snow. Unfortunately, no clear difference in reflectivity was observed between rock, glacier and lake ice.

Another useful attribute is the waveform kurtosis, showing large positive values over very flat areas like glacier ice, lake ice and flat rock. The width, risetime and number of peaks are good indicators of flat terrains as well. The skewness does not give promising results for classification, similar to the ground return signals of the atmospheric channels. When these signals do show classification potential, they generally follow the same trends as reflectivity and consequently become redundant. In Section 8.4, a more general discussion will be given on the the attribute values taking into account all shots over the Dry Valleys and Lake Victoria. Based on the current section and Section 8.4, the classification criteria will be constructed in Section 8.5.

In the case of saturated shots, the classification criteria will become even more complicated. Leaving saturated shots out would indeed give a simpler classification criteria scheme, but it would also have severe effects on the number of shots that are left, especially for glacier ice and lake ice. Since these land cover types contain many saturated shots, for lake ice even about 50% of all the lake ice covering shots, the results would no longer be representative for practical applications. In most cases of saturation, the gain, reflectivity and kurtosis are biased. However this does not mean these cannot be used at all. For reflectivity, saturation results in the underestimation of reflectivity values. This means that, when saturated reflectivity values are higher than 0.6, in reality these values are even higher and consequently could be used for snow classification. One should be cautious however, when interpreting reflectivity values below 0.6, because these could also be higher than 0.6 in reality.

Saturation information could be used to detect changing land cover types as well. For example, after sloped rock surfaces with high gain values, subsequent saturated shots could indicate glacier ice, lake ice or snow, even though the gain values are still high. The case is different for saturated kurtosis values. For example, saturated kurtosis values over glacier ice are usually negative, but when the shots become non-saturated over the same area, the kurtosis values become positive. In other words, saturated shots give unreliable kurtosis values. The waveform width, risetime and number of peaks appear to be less affected by saturation and consequently can be used for classification criteria with more confidence. Width and risetime could be used for detecting whether saturated shots cover areas like

sloped or rough surfaces that could indicate rock. The number of peaks could contribute to this detection as well, but should be interpreted with caution because part of saturated waveforms is omitted.

These observations form an important reference for the classification analysis that is carried out in Section 8.4. From the results of that section in combination with the observations of this section, the final classification criteria are constructed.

8.1.2 Lake Victoria area

In the previous paragraphs, the observations over several Dry Valleys areas were analyzed. In the coming paragraphs, one case study is used to show the effects of water on the measurement attributes. Similarly to the Dry Valleys, two examples have been highlighted to observe the effect of high and low transmit energy. The shots were taken in the same track and a nominal beam co-elevation angle was applied. The first figure includes measurements of L3B, acquisition date the 1st of March 2005, and has a mean transmit energy of $64.3 mJ$. The second figure includes measurements of L3H, acquisition date the 22nd of March 2007, and has a mean transmit energy of $24.7 mJ$.

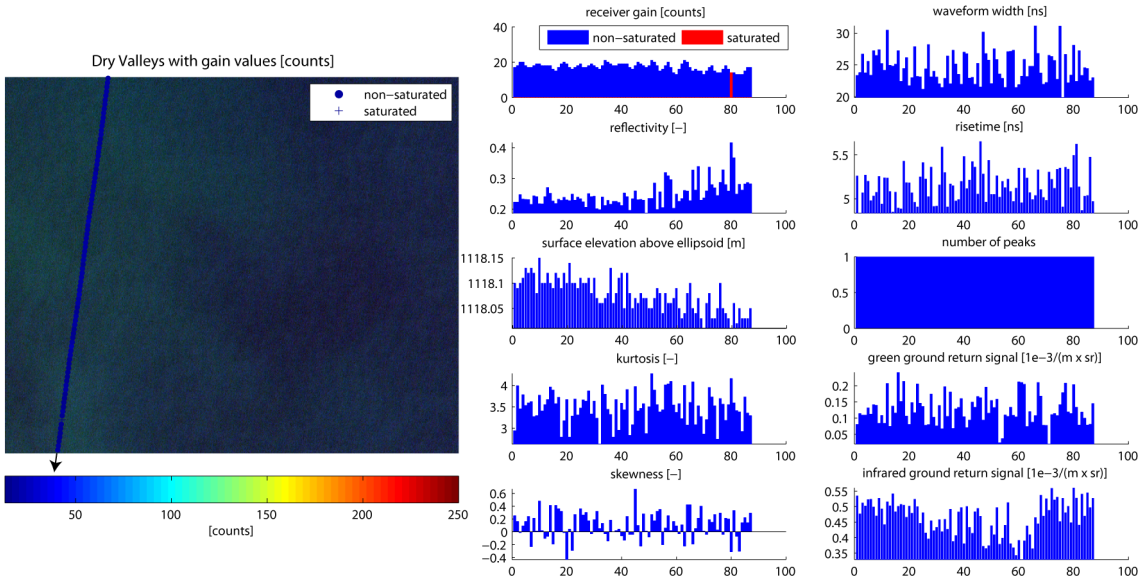

Figure 8.10: The close-up of a selected track of campaign L3B covering Lake Victoria. The arrow indicates the direction of the track.

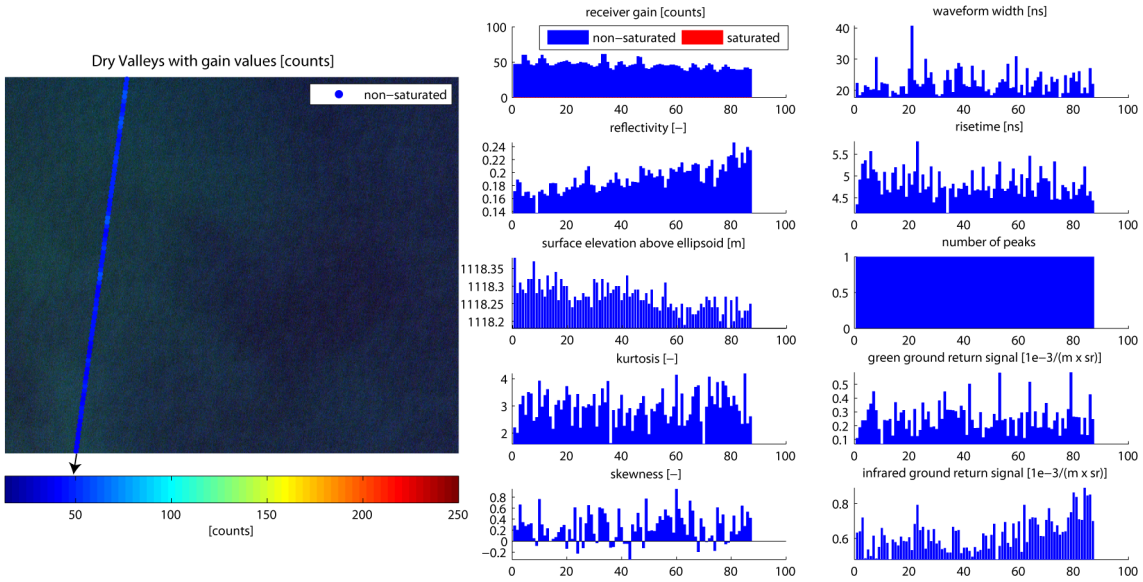

Figure 8.11: The close-up of a selected track of campaign L3H covering Lake Victoria. The arrow indicates the direction of the track.

Figure 8.10 shows the measurements of L3B and the first interesting observations that come forward are the virtual absence of saturated shots and the small elevations differences. The reflectivity values stay around 0.3 and do not show significant variation along the track. There appears to be a slight increase in the reflectivity though, which might relate to the smaller surface elevation values. Despite the slight increase in reflectivity, the values are still smaller compared to most reflectivity values over the Dry Valleys. The gain is low and fairly consistent as well; it is close to the minimum value of 13 counts. In addition, the figure confirms the assumption stated before, that saturation dominantly occurs in highly dynamic environments.

The kurtosis indicates that the waveforms are sharply peaked and show even higher values than the kurtosis over glacier ice and lake ice. The skewness does not show any consistent values along the track. The width is small and comparable to the width over flat glacier ice, lake ice and snow. Also the risetime and number of peaks show that the waveforms consist of one narrow peak and that virtually no variation occurs. The green ground return signals are small compared to most L3B values over the Dry Valleys. The infrared ground return signals show consistent values and are somewhat lower than the values of L3B over the Dry Valleys. These signals can be directly related to the reflectivity.

The observations of L3H are shown in Figure 8.11. The gain values are about twice as high compared to the L3B gain values and no saturated shots occur. The reflectivity values are similar to the L3B campaign and show a slight increase along the track as well. The elevations appear to slightly decrease as well. The kurtosis values are highly positive compared to most of the Dry Valleys values, but are slightly smaller than the values of L3B. The width is similar to L3B, although occasionally high values occur. The values of risetime and number of peaks are virtually identical to L3B. The green ground return signals are similar with the exception of some local spikes. The infrared ground return signals are somewhat higher than for L3B and also show a slight increase, similar to the reflectivity.

In general, the two figures mainly show similarities, in particular the reflectivity, waveform width, risetime and number of peaks. The observations differ slightly in gain, kurtosis, and infrared ground return signals. When comparing these observations with the Dry Valleys observations, already some potential classification criteria and some difficulties can be identified. First of all, the kurtosis of Lake Victoria shows higher values compared to the Dry Valleys, although (non-saturated) shots covering glacier ice and lake ice can give high kurtosis values as well. The reflectivity, waveform width, risetime and number of peaks in many cases is comparable to glacier ice and lake ice observations. In some cases, the attribute values over flat rock are similar to the water values.

8.2 Attributes irregularities analysis

In the previous section the effect of land cover type and topography on the attributes have been analyzed with case studies. The potentials of using attributes for classification have been identified. In this section, two case studies are highlighted that show inconsistent values with respect to the case studies of the previous section. These inconsistent values will be related to cloud presence. See Section 5.4 for the technique based on gain values that is currently applied for cloud filtering. Similar to the previous section, the Dry Valleys and Lake Victoria are used for analysis.

8.2.1 Dry Valleys area

The first case concerns part of a L3C track over the glacier invading Wright Valley, see Figure 8.12. It demonstrates high gain values for a land cover type and topography that does not comply with the observations of the previous section. As the figure shows, the gain is constantly 250 counts and none of the shots is saturated. The fact that this occurs over glacier, makes these observations particularly interesting when comparing these to the observations over Taylor Glacier of the previous section for example. Even if the temporal difference between the original background LIMA image and campaign L3C resulted in snow accumulation, the gain values would still be smaller than 50 counts and the reflectivity would be close to 0.6, as was seen in the previous section.

Possible explanations are a high beam co-elevation angle, a thin layer of melted ice, abnormal solar background radiation or the presence of clouds. To find out what explanation could cause the high gain values, Figure 8.13 is used for comparative analysis. This figure shows a different track over the same glacier taken in the same campaign. Most importantly, the comparison with Figure 8.12 shows that both tracks were run under generally the same conditions. The beam co-elevation angle of Figure 8.12 is 0.33° and of Figure 8.13 this angle is 0.32°. Also the terrain topography is similar, looking at the elevation profile of both tracks. As will be shown later on, it is likely there was snow accumulation in between the tracks. However, possible topography change due to snow fall would definitely not result in such low reflectivity values. Furthermore, the solar radiation circumstances are identical as well, since for both campaigns the Sun does not come above the horizon. Approximately the same transmit energy is used, for Figure 8.12 this is about $47mJ$ and for Figure 8.13 about $45mJ$. By no means did this difference result in the dramatic increase in gain. For example, the gain values over Taylor Glacier of L3J (see previous section) were subject to even lower transmit energy while having significantly lower gain values.

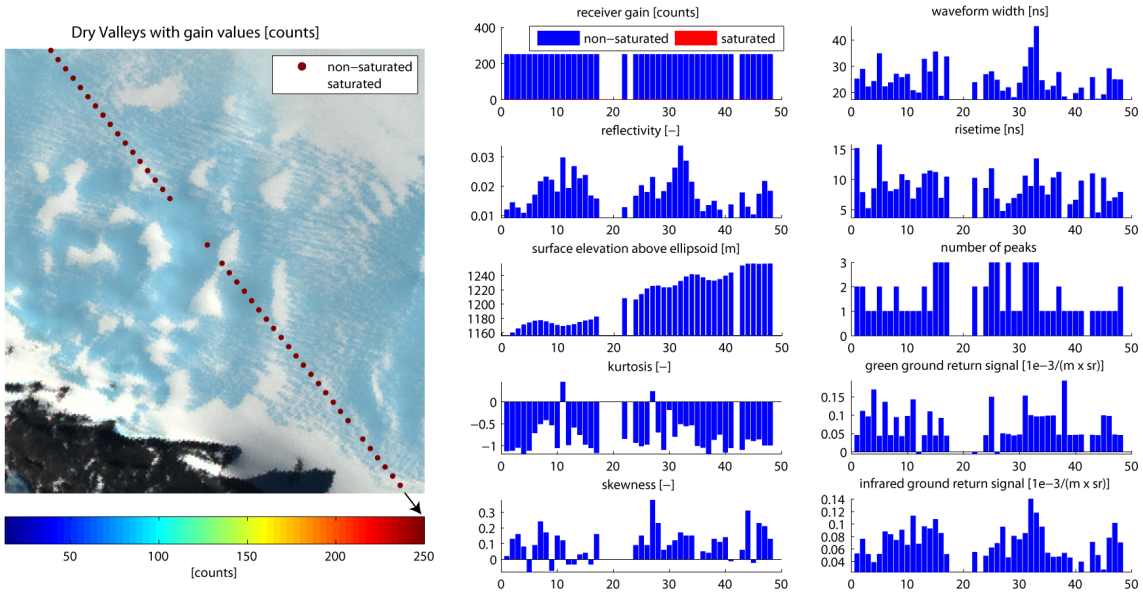

Figure 8.12: Close-up of a selected track of campaign L3C covering the glacier invading Wright Valley, demonstrating inconsistent gain values. The arrow indicates the direction of the track.

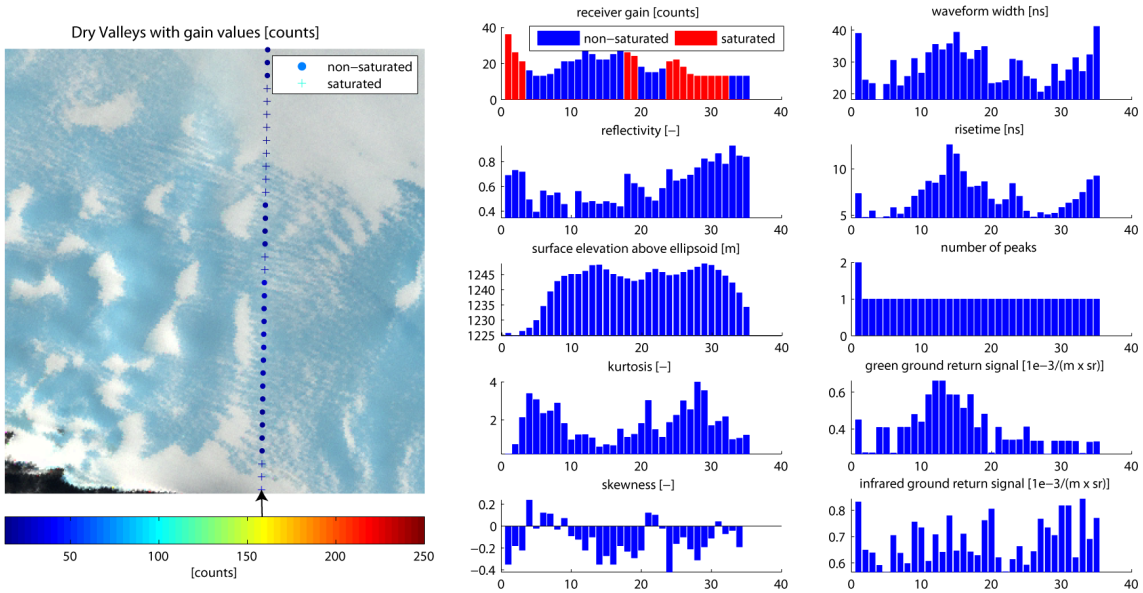

Figure 8.13: Close-up of a selected track of campaign L3C covering the glacier invading Wright Valley, demonstrating expected gain values. The arrow indicates the direction of the track.

Since both tracks were flown in mid-winter, a thin layer of water cannot have been the cause either. The left plot of Figure 8.14 confirms this assumption by showing the air temperatures that were recorded at the closest weather station, Lake Vanda (at about $25km$ distance from the track), over the entire L3C campaign. The dashed lines indicate when both tracks were flown, showing that the temperatures were well below zero degrees Celsius (i.e. lower than -15°). The right plot shows that about $8cm$ of snow was accumulated in one day at Lake Vanda. Even though the amount of snow rapidly declines (most probably due to rapid sublimation), there is still snow left at the time the second track was flown (i.e. a couple of centimeter). Even though Lake Vanda is located about $25km$ from the ground tracks, it indicates the temperatures were most definitely below zero and heavy snow fall occurred. It might give a hint of the cloudiness at the time of the high gain track, because of the heavy snow fall a couple of days later. In this discussion it is not that important to know whether there was snow accumulation or not at the locations of the tracks. More interesting is the presence of clouds, and although the weather station is about $25km$ distant from the locations of the track, the snow accumulation does make it more probable that clouds were present at the time of measurements. However, this could not be verified and should be interpreted with caution.

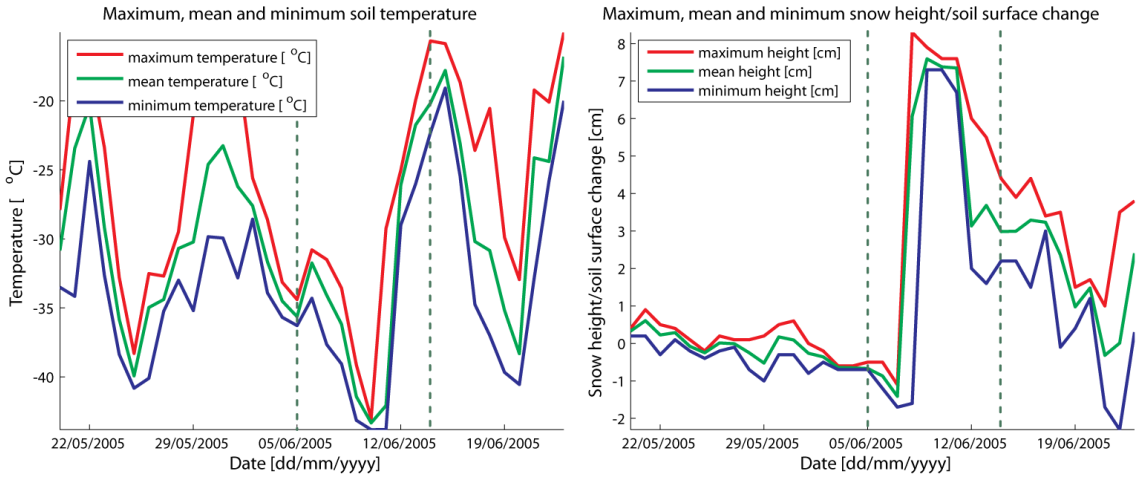

Figure 8.14: The soil temperature and soil surface change or snow height for both L3C tracks taken at Lake Vanda. The left dashed lines of both plots indicate the time that the track of Figure 8.12 was flown, the second dashed lines indicate the time that the track of Figure 8.13 was flown. Data taken and adapted from LTER, [17].

Since beam co-elevation angle, solar radiation and a thin layer of water are most probably not the causes of the high gain values, and heavy snowfall occurred close to the acquisition date of the track, the most plausible reason is the presence of clouds at the time of measurements. In the discussion about clouds, also other forms of severe atmospheric influences are implied by this term. It is beyond the scope of this research to analyze the severity of each atmospheric disturbance on waveform attributes, but at least some kind of relevant filter criteria will be proposed to reduce the probability of cloud affected

attributes. A disadvantage of the fact that both tracks were flown in night conditions, is that no LANDSAT images are available and no on-site solar radiation levels of the LTER stations can be used.

The subplots with the attribute values of Figure 8.12 show odd observations. First of all, the reflectivity values are extremely low (i.e. lower than 0.4), generally even more than an order of magnitude lower compared to Figure 8.13. Even in cases of highest gain and lowest reflectivity discussed in the previous section, for example the sloped rock surfaces of Figure 8.5, the reflectivity was still about ten times higher. In other words, the laser pulse is attenuated or reflected such that almost no energy is returned. Additionally, the sign of the kurtosis is reversed compared to Figure 8.13, which means that the raw waveform shape is changed from sharp peaked to blunt peaked. The ground return signals of the atmospheric channels are considerably lower as well and most of the signals differ an order of magnitude. This means that between $300m$ and -$300m$ above the ground surface, the maximum peak of both signals is not that different from the neighboring signal returns, i.e. the ground return signals have similar values as the atmospheric return signals.

Despite the fact that almost all attributes of Figure 8.12 show completely different results than one would expect based on the previous section and Figure 8.13, the elevation values are similar to the elevation of Figure 8.13. The most plausible explanation is that even though the received signal is extremely low, the infrared signal still touched the ground surface and surface elevations were successfully estimated. Perhaps fog was present at the time of measurement, or the infrared signal is strong enough to penetrate heavier clouds as well. Next to the abnormal gain values, also some gaps occur in the data. This is again due to invalid values of the geolocation parameters, resulting from failed elevation estimates. Perhaps for some shots, the attenuation was that strong that no usable waveform was reflected back and no Gaussian peaks could be fitted in.

Although there are some more tracks with similar unexpected gain values over the Dry Valleys that could be highlighted, it is decided not to discuss these since the essence is the same. Instead, another figure (Figure 8.15) is included that relates to the analysis of Figure 8.12 as well. It concerns a similar track of L3D that runs parallel to the track of Figure 8.12 and is only $200m$ next to the track of Figure 8.12. The L3D figure shows how the gain values in this case are more or less as expected. The most apparent differences with Figure 8.12 are again related to the reflectivity (0.02 for L3C compared to 0.3 for L3D), kurtosis (-1 compared to 3), atmospheric green channel ground returns (around 0.1 compared to 0.3) and the atmospheric infrared ground return (0.06 compared to 0.9). Thus, quite some important attributes differ considerably for approximately the same track. On the other hand, elevation, width and risetime show almost similar results. In particular the elevation is almost identical, indicating that the infrared laser signal of Figure 8.12 certainly hit the ground and proper elevations were estimated.

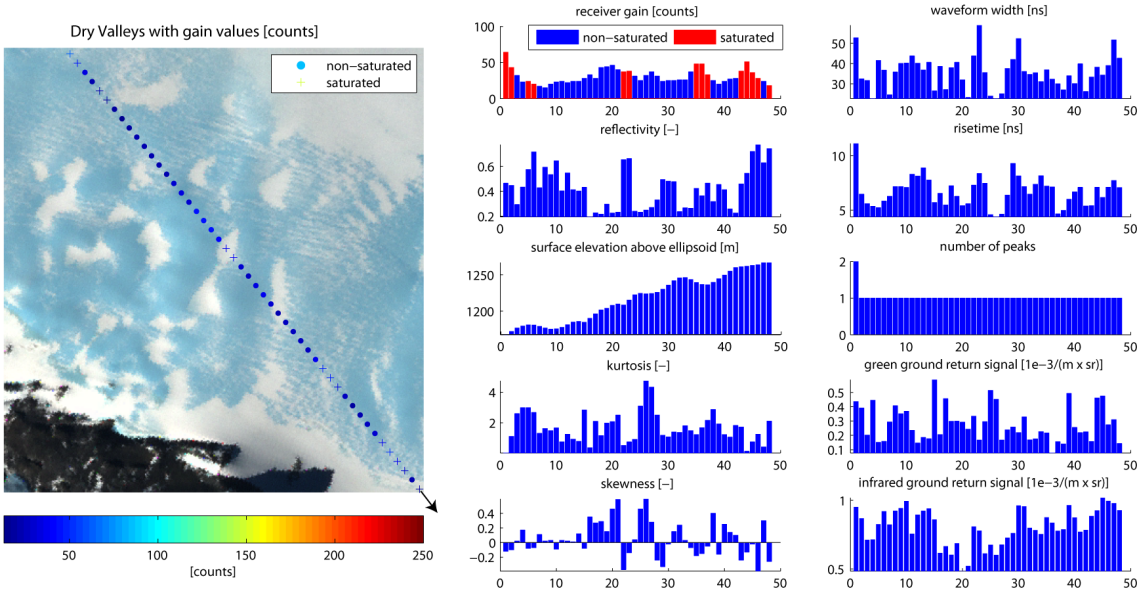

Figure 8.15: Close-up of a selected track of campaign L3D covering the glacier invading Wright Valley, demonstrating expected gain values. The arrow indicates the direction of the track.

8.2.2 Lake Victoria area

For Lake Victoria one track has been highlighted as well to show some gain irregularities that are likely related to cloud presence. The observations were taken on the 31st of May 2005 with a beam co-elevation angle of 0.37° and a mean transmit energy of 47.6mJ. Although there are LANDSAT images available that were acquired about a week before and after the acquisition date of the track which indeed show cloud presence, it is decided not to include these for cloud validation, because of the high variability in cloud cover over time.

Figure 8.16 shows a track of L3C with unexpected low gain values. As demonstrated in the previous section, generally speaking the waveform attributes are approximately constant over water. However, this figure shows such varying gain, reflectivity, kurtosis and infrared ground return signals that there must be some kind of disturbing factor causing these irregularities. The most probable factor is cloud interference; this figure contributes to the assumption that extremely low reflectivity values are associated with cloud presence as well. Especially the first 50 shots of the track shows irregular values with low reflectivity values and high gain values. In the remainder of the shots, the attributes show values that are similar to the attribute values over Lake Victoria in the previous section. Cloud forming at locations where the shots have high gain and low reflectivity could be a plausible explanation. The reason that occasionally the number of peaks is relatively high can be explained by either high waves or the effect of low SNR. This effect of SNR can be explained as follows. Since the reflectivity (and consequently the received energy) of these shots is extremely low, the returned waveform is closer to the noise threshold. Con-

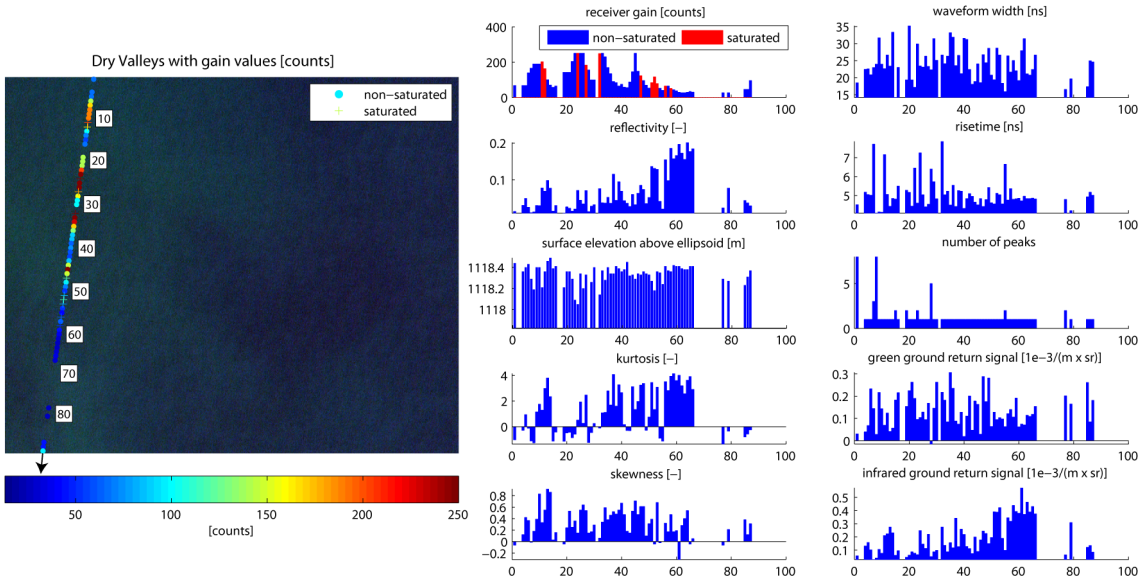

Figure 8.16: Close-up of a selected track of campaign L3C covering Lake Victoria. The arrow indicates the direction of the track.

sequently, the noise affects the waveform shape in such a way that the number of peaks may locally become higher.

8.3 Data points used for classification

Based on the previous discussions, a simple method to filter out the most irregular shots will be proposed in this section. This filtering is a very rough estimate, but is meant to filter out those shots that are affected severely by (assumed) cloud presence. Additionally, this section will give an overview of the data points used for final classification analysis.

As was analyzed in Section 8.2, the reflectivity that occurs at the irregular gain values, is more than ten times lower than any reflectivity value of the nominal gain values. Even though some other attributes like kurtosis, gain and infrared ground return signal also show significantly different results, only the reflectivity is used for cloud filtering. The reason that kurtosis and gain are not taken into account for filtering is because kurtosis occasionally changes from large negative to slightly positive values for the irregular case studies and the gain does not show consistent irregular values due to varying transmit energy. Additionally, the gain is already accounted for in the reflectivity and the infrared ground return signals show in many cases the same trends as the reflectivity.

As demonstrated by all figures in Section 8.1, shots having nominal gain values all have reflectivity values higher than 0.1. For the irregular gain values, the reflectivity drops to

about 0.03 for Figure 8.12 and 0.01 to 0.06 for Figure 8.16. As mentioned in Section 8.2, these extremely low values are that different from expected values, that these are assumed to be affected by clouds. As such, a threshold filtering is proposed to **exclude all shots having reflectivity value lower than 0.1**. One should keep in mind that this cloud filtering is based on comparison between expected and irregular gain values and has not been confirmed by validation data.

Besides the cloud filtering, some other criteria were formulated to select specific shots for classification analysis as well, see Section 7.1. The two criteria that involve data filtering are the Dry Valleys mask and the selection of only those campaigns that are taken in the same season in which the LIMA image was taken (i.e. austral spring). As a result only the L2A, L3A, L3D, L3G and L3I campaigns are used for the Dry Valleys. For Lake Victoria all campaigns that overlap the specific subset are used, which are L2A, L2B, L3A, L3B, L3C, L3D, L3E, L3F and L3H.

After all filtering steps have been taken, the result is that 10% of the total shots covering the Dry Valleys subset is left. In case of Lake Victoria this is 74%. Table 8.2 shows the impact of the mask, campaign and cloud filtering over the Dry Valleys and Lake Victoria. Since no mask or campaign filtering is used for Lake Victoria, only cloud filtering and waveform selection make a difference.

Limiting criterion	Dry Valleys	Lake Victoria
Total over subset	24408	1804
Mask selection	5864	1804
Campaigns selection	2733	1804
Cloud selection	2375	1334
Waveform selection	2326	1059

Table 8.2: The number of points that are left after each filtering criterion for both validation areas. The campaign and mask selection criteria are based on temporal and spatial preferences, the cloud filtering is based on attribute thresholding, and the waveform selection is based on the availability of valid waveforms.

Table 8.2 shows that 76% of the Dry Valleys shots are filtered out at the mask filtering step. The campaign selection accounts for filtering out 53% of the shots covering the mask, resulting in 11% of the total number of shots. These two filtering procedures may seem very strict, but the big advantage is that the analysis of the resulting classified footprint becomes more reliable. The cloud criterion filters out 13% and 26% of the shots that are left for the Dry Valleys and Lake Victoria, respectively.

After the filtering has been carried out, all shots over the Dry Valleys are divided into land cover classes according to the supervised classification map described in Section 7.1. For Lake Victoria, all shots that comply with the filtering criteria are assumed to cover water. Figure 8.17 shows the classified shots over the Dry Valleys and Lake Victoria.

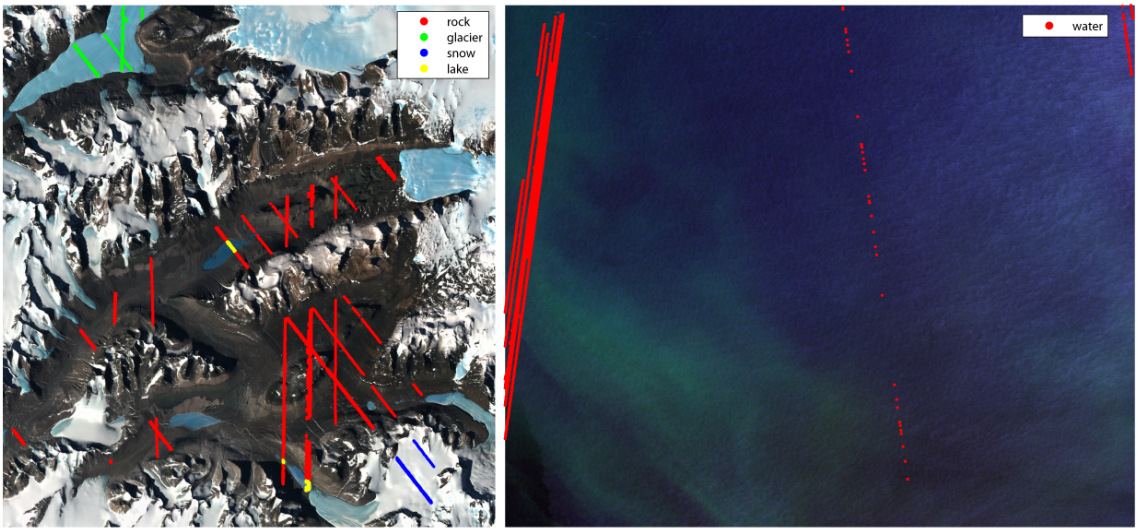

Figure 8.17: The shots that are used for classification analysis over the Dry Valleys (left) and over Lake Victoria (right). Over the Dry Valleys, the shots are classified based on the classification map.

Of all the shots that are classified with the Dry Valleys classification map, 1861 cover rock, 183 cover glacier, 195 cover snow and 67 cover lake, which adds up to 2306 points. The reason that this is not the same as 2326 as indicated by Table 8.2 is because some footprints do not cover one land cover type for 95% or more. In case of Lake Victoria, 1059 shots are used for classification analysis.

8.4 Attributes distributions

In the previous sections, various attributes were analyzed for their potential for classification. Using Kernel smoothed distributions per land cover type, the distinguishing power of each attribute will be determined.

Ideally, the Kernel smoothed distribution of each land cover type is significantly different from the other land cover types, because a certain attribute interval could then exclusively be assigned to each land cover type. In the following paragraphs the Kernel smoothed distributions will be given of the attributes of Table 8.1. With the distributions, the classification potential of the attributes will be analyzed. No final classification criteria will be given in the coming discussions. The classification criteria will be constructed in Section 8.5 based on the Kernel smoothed distributions and the findings of Section 8.1. For the percentages and distributions that are mentioned in the following discussions, unless otherwise indicated, the following order of land cover type is used: glacier ice, rock, snow, lake ice and water (the same order as the legends of the figures).

8.4.1 Gain and saturation energy correction

As the case studies in Section 8.1 demonstrated, the receiver gain is dependent on land cover type and topography as well as on transmit energy. Over sloped rock surfaces, the gain is mostly close to the maximum value of 250. However, since reflectivity accounts for changing transmit energy by normalization, the reflectivity is a better alternative than gain in classifying land cover types. The case studies showed that changing land cover type from rock surface to high reflectivity surfaces like snow or ice for example, the shots become saturated because the gain is not adjusted properly. And even though in many of these cases the reflectivity values will be biased, this saturation information can also be used for classification.

Of all shots that cover glacier ice, 74% is non-saturated; of all rock covered shots 79% is non-saturated; of all snow covered shots 68% is non-saturated; of all lake ice covered shots, 55% is non-saturated; of all water covered shots, 75% is non-saturated. The reason that the percentage for lake ice is low, is because the Dry Valleys lakes are surrounded by rock surface . Therefore the gain values are moderate to high just before the shots cover the lake. Because the lakes are relatively small and the gain adjusts slowly, many shots over lake are saturated.

Some attribute values are significantly different when saturation occurs (for example kurtosis). Consequently, the first important decision in the classification is separating saturated and non-saturated shots. **This means that saturated shots will be subject to other classification criteria than the non-saturated shots.** Consequently, in the following discussions a clear distinction is made between distributions of saturated and non-saturated shots. The first distribution that is analyzed is the gain per land cover type.

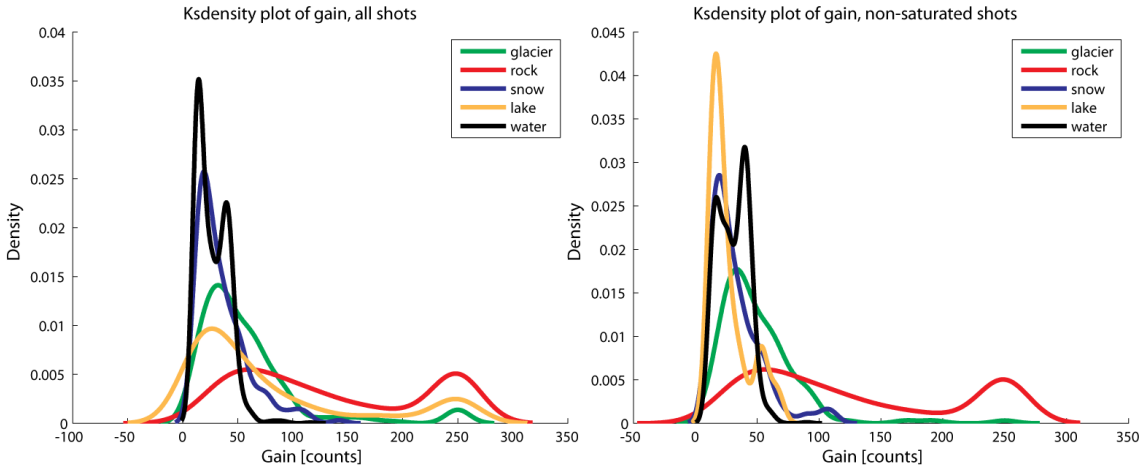

Figure 8.18: The Kernel smoothed probability density of gain per land cover type. The left plot shows the estimate for all shots, the right plot shows the estimate for all non-saturated shots.

Figure 8.18 shows the Kernel smoothed probability density estimates of gain per land cover type. The left plot shows the estimate of all shots and the right plot shows the estimate of the non-saturated shots. The reason that the gain distributions of saturated shots are not shown is because these do not present any relevant information. For many saturated shots, the same gain values are observed over different land cover types, because of the limited response time of the gain.

The right plot indicates that the majority of all non-saturated shots having a gain larger than 100 counts, belong to rock: of all shots that are classified as rock, 39% is non-saturated and has a gain value higher than 100 counts. For the shots that are classified as glacier ice, snow, lake ice and water these percentages are 2.2%, 2.1%, 0% and 0%, respectively. Consequently, a substantial part of all rock covering shots (i.e. 39%) could be classified as rock, with only minimal losses for the other land cover types.

8.4.2 Reflectivity

The reflectivity is the most important attribute for the classification algorithm to distinguish snow. The case studies in Section 8.1 already indicated that shots covering snow typically have reflectivity values higher than 0.6. This is also confirmed by Figure 8.19, where the Kernel smoothed distributions of reflectivity are shown per land cover type.

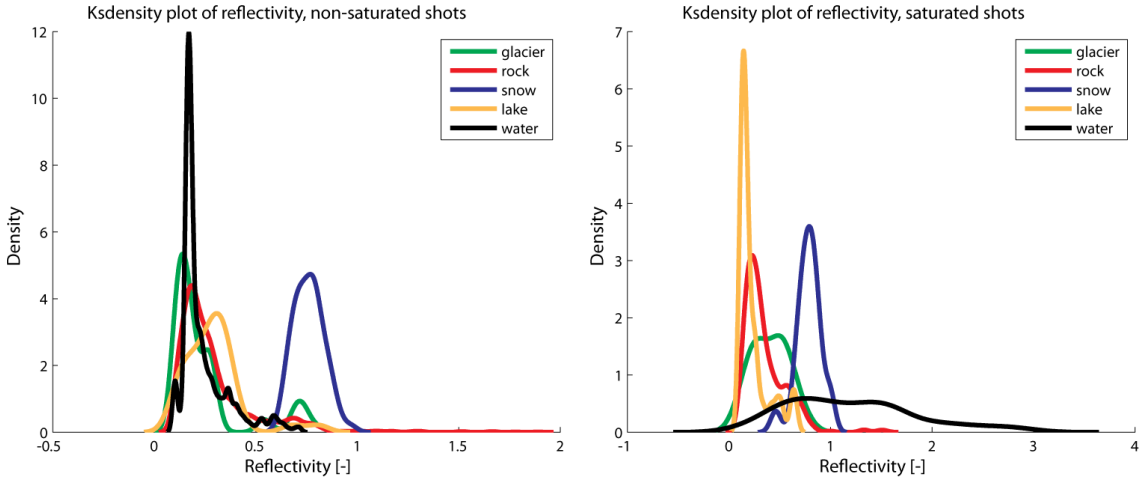

Figure 8.19: The Kernel smoothed probability density of reflectivity per land cover type.

For the non-saturated shots, 100% of the snow covering shots have reflectivity values higher than 0.6. For glacier ice this is 11%, for rock 9.0%, for lake ice 5.6% and for water 2.5%. Based on the percentages of non-saturated shots over all shots (i.e. 74%, 79%, 68%, 55% and 75%, see the previous paragraph) this means that 68% of all snow covered shots can be classified as snow. However, setting this classification criterion also means that 8.2%, 7.1%, 3.0% and 1.9% of all shots covering glacier ice, rock, lake ice and water, respectively,

are incorrectly classified as snow. For the shots covering the Dry Valleys, this is according to what is assumed to be 'true' from the supervised classification map. The validation of the clusters in the left plot (for example the glacier cluster at around 0.6) poses a problem because of the lack of consistent validation data. Since the classification of the shots are based on one image taken in November 2001, temporal variation of land cover type can give incorrect classification results. For some cases, no proper validation data is available to validate the exact land cover type at the moment of measurement. LANDSAT images are mostly not available for the exact date of interest and the LTER data set usually does not cover the area of interest exactly. Although it might be possible that the supervised classification map resulted the incorrect exclusion of shots classified as snow, this research adheres still to the classification results from the supervised classification map. Also, the left plot shows that the ice and rock observations having reflectivity values higher than 0.6 form small isolated clusters. The fact that these clusters have low density values and are well separated from their main peak, indicates that they represent irregular observations from the main peak. For example for glacier, the main peak is around 0.2 and one small cluster is in the 0.6 region.

When looking at the reflectivity distribution of saturated shots, again the snow is generally higher than the other Dry Valleys land cover types. However, more overlap occurs and especially the saturated shots over water show a wide distribution and even have a mean and maximum reflectivity of 1.0 and 2.1 respectively. As a consequence, if the same simple criterion, classifying all shots with reflectivity values higher than 0.6 as snow, would be used, 12%, 11%, 95%, 6.5% and 81% of the saturated shots covering glacier ice, rock, snow, lake ice and water would be classified as snow. This comes down to 3.3%, 2.4%, 31%, 3.0% and 20% of all shots covering the corresponding land cover types. Especially for water, this would make the classification accuracy significantly worse. Ignoring the lower percentages (i.e. glacier ice, rock and lake ice), the two most important land cover types for reflectivity value higher than 0.6 would be snow and water.

8.4.3 Kurtosis

The kurtosis is the most important attribute for classifying water. As the case studies of Section 8.1 already demonstrated, the shots over water usually give high positive kurtosis values. Positive kurtosis indicates peak shaped waveforms. Figure 8.20 shows the Kernel smoothed probability density estimates of kurtosis per land cover type.

The left plot shows that the kurtosis can be used for non-saturated shots to distinguish between water and other land cover types, though glacier ice and lake ice have more overlap which makes it more difficult. For the saturated shots, the distributions show increased overlap and consequently it is harder to find sound classification criteria.

The percentages of non-saturated shots having a kurtosis higher than 2.5 are 17%, 0.48%, 0%, 22% and 89% for glacier ice, rock, snow, lake ice and water respectively. Taking into account the percentages of non-saturated shots over all shots, the result is that

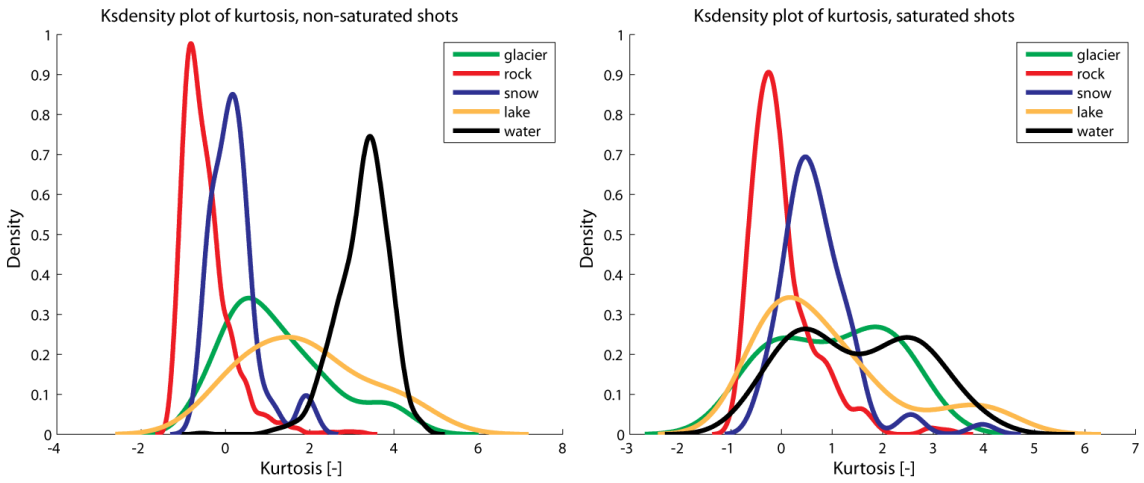

Figure 8.20: The Kernel smoothed probability density of kurtosis per land cover type.

13% of all glacier ice covered shots have kurtosis higher than 2.5, 0.38% for rock, 0% for snow, 12% for lake ice and 66% for water. Using this threshold as a classification criterion will result in the majority of all water covering shots to be correctly classified. However, it also results in 13% and 12% of all glacier ice and lake ice shots to be incorrectly classified as water. For two of these cases, some extra research is carried out to check whether it is possible that some melting occurred.

The first case concerns shots over Taylor Glacier taken in L3D on the 22nd November 2005. The second case concerns the shots over Lake Vanda taken in L3G on the 17th of November 2006. Figure 8.21 shows the ice and soil temperatures of Taylor Glacier and Lake Vanda during the specific campaigns with an indication of the specific dates of measurements. For the Taylor Glacier observations, the maximum, mean and minimum temperatures are -2.6°C, -7.0°C and -13°C respectively. And even though the day before the maximum temperature was +1°C, it does not convincingly prove that there was a thin layer of ice on top of the glacier. For the Lake Vanda observations, the maximum, mean and minimum temperatures are +2.5°C, -3.1°C and -7°C respectively. The day before the maximum temperature was as high as +5.8°C. These temperatures indicate there might have been a small layer of water on top of the lake during the time of measurement. However, since the observations are not supported sufficiently and it would be too labor intensive to check all of these cases, this research yet again adheres to the supervised classification map results. Again, this decision might give the accuracy assessment a more pessimistic result, however the LTER observations can provide several of the incorrect classification results with explanations.

The percentages that correspond with the kurtosis criterion (i.e. kurtosis larger than 2.5) show that for non-saturated shots, water can be clearly separated from rock and snow. However, it also results in incorrect inclusion of a part of glacier ice and lake ice. But

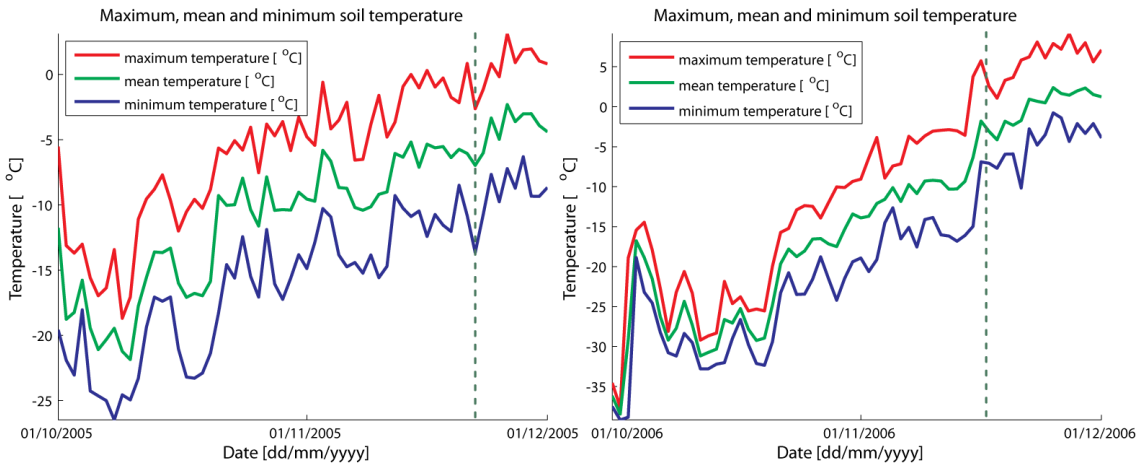

Figure 8.21: The ice temperature during L3D taken by the Taylor Glacier weather station (left plot) and the soil temperature during L3G taken by the Lake Vanda weather station data (right plot). The dashed lines indicate the specific dates of measurements. Data taken and adapted from LTER, [17].

because there is no other attribute more optimal for distinguishing water from the other land cover types and because some of the misclassifications might be explained by LTER data, it is decided to use kurtosis higher than 2.5 for classifying non-saturated shots as water.

For saturated shots, the classification is even more difficult. As the right plot of Figure 8.20 shows, part of the saturated water shots have kurtosis values close to zero. As discussed earlier in Section 8.1, saturation influences the kurtosis values due to the distorted waveforms (part is clipped off). Since a significant part of the shots is saturated (i.e. 26%, 21%, 32%, 45% and 25% for glacier ice, rock, snow, lake ice and water respectively) it is still important to make sure these percentages are classified as correctly as possible, especially in case of lake ice. As already mentioned in the paragraph about reflectivity, the reflectivity can be used to filter out most of the snow and water covering shots.

8.4.4 Waveform width

The waveform width is an important attribute to classify between ice and rock and water. The previous paragraphs show that reflectivity only helps in separating between snow and other land cover types for non-saturated shots and between -snow and water- and -rock and ice- for saturated shots. Kurtosis is useful for non-saturated shots only, to separate between water and other land cover types. However for non-saturated shots the distinction between ice and rock, and for saturated shots between water, snow, ice and rock the width is the only attribute that offers reasonable classification potential.

In Chapter 6, the relationship between width and surface topography is analyzed and

discussed. Even though theoretically the beam incidence angle as well as roughness affect the width, it was shown in this chapter that most dominantly the incidence angle has influence on the width. Since the incidence angle combines beam co-elevation angle with surface slope angle, the uncorrected width is also dependent on the beam geometry. To correct for this, the beam co-elevation correction was applied to the width. The waveform is compressed according to the extra width the angle induces. This is particularly important when carrying out comparative analysis between tracks that do not have identical beam co-elevation angles and when looking at the Kernel smoothed distribution of the width. As a result, the width plots of the figures in Section 8.1 are beam co-elevation corrected as well as the Kernel smoothed distribution of Figure 8.22.

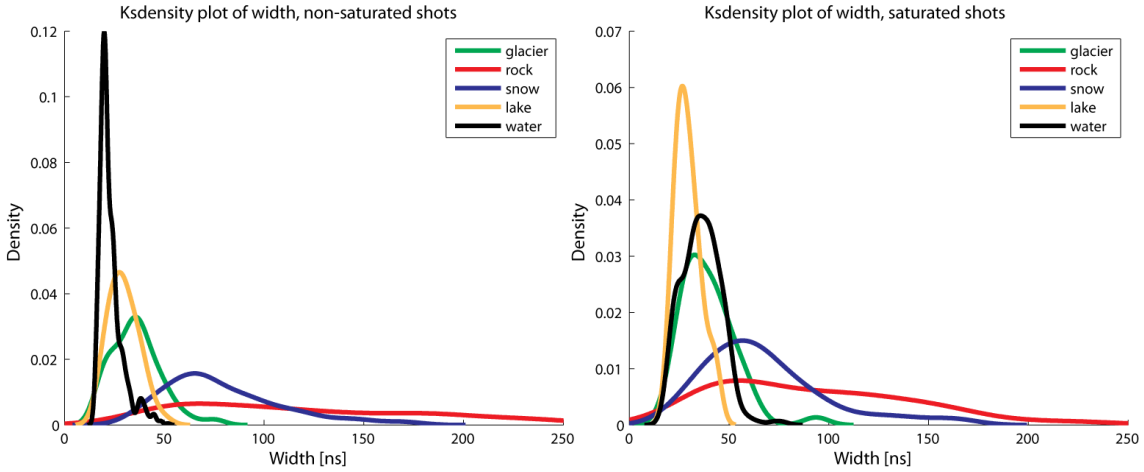

Figure 8.22: The Kernel smoothed probability density of width per land cover type. The plots are truncated at $250ns$ for visualization purposes, resulting in the exclusion of 3.5% (left plot) and 2.7% (right plot) of the rock distributions.

The left plot shows that the width can be used for non-saturated shots to distinguish rock from glacier ice and lake ice after water is separated by use of kurtosis. For the saturated shots, two cases can be identified based on the previous sections. First of all, when reflectivity can separate snow and water from the other land cover types, the width can separate snow from water. Secondly, after reflectivity isolates rock, glacier ice and lake ice from snow and water, width can be used to distinguish between ice and rock. Taking into account the large percentage of shots over lake ice that are saturated, this forms potential for discrimination between rock and lake ice in particular.

For non-saturated shots as well as for saturated shots, a width threshold is used of $50ns$, which corresponds to approximately 3° surface slope. For non-saturated shots, 89%, 12%, 14%, 100% and 100% of the corresponding land cover types has width smaller than or equal to $50ns$. When only considering the non-saturated shots having width values smaller than or equal to $50ns$, 66%, 9.5%, 9.7%, 54% and 75% of all shots would be classified as one land cover type. Since this does not result in optimal classification, other classification

criteria will precede the width criterion, see Section 8.5. For saturated shots, 83%, 24%, 33%, 100% and 94% of the corresponding land cover types has width smaller than or equal to 50ns. In this case, when one would assign these shots to one land cover type, the result would be that 21%, 5%, 11%, 46% and 24% of all shots would be classified. Again, this criterion is not applied without considering preceding classification criteria (see Section 8.5).

8.4.5 Other attributes

In this paragraph, the attributes that do not contribute to the classification criteria are briefly analyzed and discussed. Only those attributes that are also incorporated in the figures in Section 8.1 are taken into account, which are skewness, waveform risetime, waveform number of peaks and the atmospheric ground return signals. For these attributes the Kernel smoothed distributions are again based on separating non-saturated from saturated shots, and were applicable correlation coefficients are used to analyze the redundancy with other attributes.

Skewness

As the case studies in Section 8.1 indicate, the skewness does not offer significant classification potential. Figure 8.23 shows the Kernel smoothed distributions per land cover type. It shows that for non-saturated shots the distributions of all land cover types overlap, except for a part of lake ice. If the threshold would be set on -0.3 for non-saturated shots, the result would be that 1.5%, 9.2%, 3.0%, 50% and 1.8% of the non-saturated shots covering the corresponding land cover types would be classified as lake ice. However, this research does not use skewness as a classification attribute because the classification criteria that are used for identifying lake ice (i.e. a combination of saturation, reflectivity and width) offer better results. One of the main reasons for this is that the skewness cannot provide a high percentage (i.e. more than 50%) for lake ice alone without also including large portions of other land cover types. In fact, more than 50% of the non-saturated shots covering lake ice still have skewness values that completely coincide with the distributions of other land cover types. In addition, a physical explanation why lake ice would have smaller skewness values than flat rock or water cannot be provided as well. Since only 67 shots cover lake ice (compared to many more shots over other land cover types, see Section

Risetime

The risetime is computed using the beam co-elevation corrected waveform as well. The waveform attribute is useful to separate land cover types having narrow peaked waveforms like lake ice and water from other land cover types. Figure 8.24 shows the risetime distributions per land cover type. As the figure indicates, the risetime for non-saturated shots covering water is small compared to snow and rock. The water distribution shows overlap with the glacier ice and lake ice for the non-saturated shots. The rock distributions in both

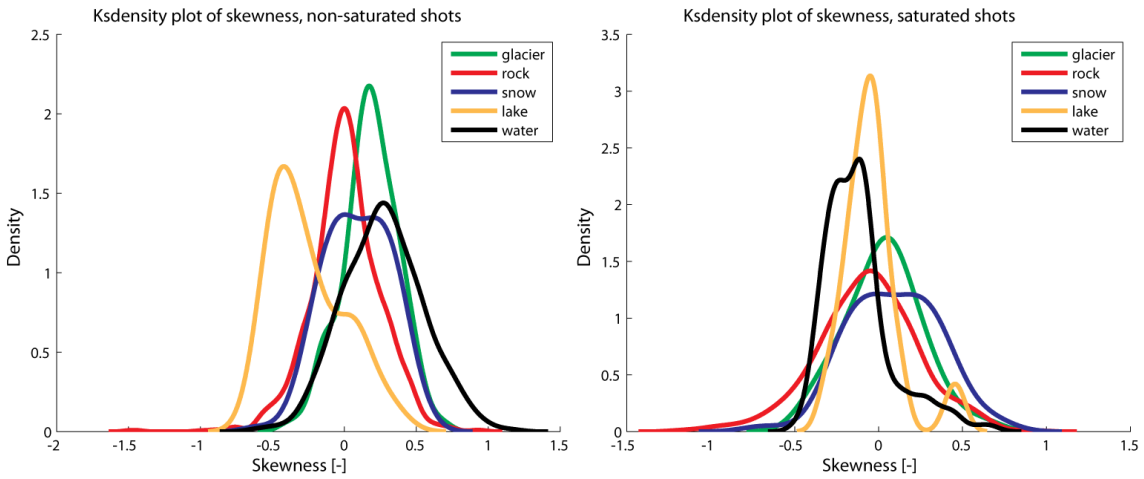

Figure 8.23: The Kernel smoothed probability density of skewness per land cover type.

plots are low and broad, indicating that the risetime values of rock is dispersed.

In case of saturated shots, the distributions become more dispersed and consequently have more overlap. In addition, some clusters are shown in the right plot for water. The most probable reason for this is that the computation of risetime is affected by the clipped part of saturated waveforms. Since the risetime is the time difference between 10% and 90% of the maximum peak but for saturated waveform this peak might be clipped off, the risetime might be incorrectly computed. So especially for saturated waveforms, the waveform width is more reliable than risetime.

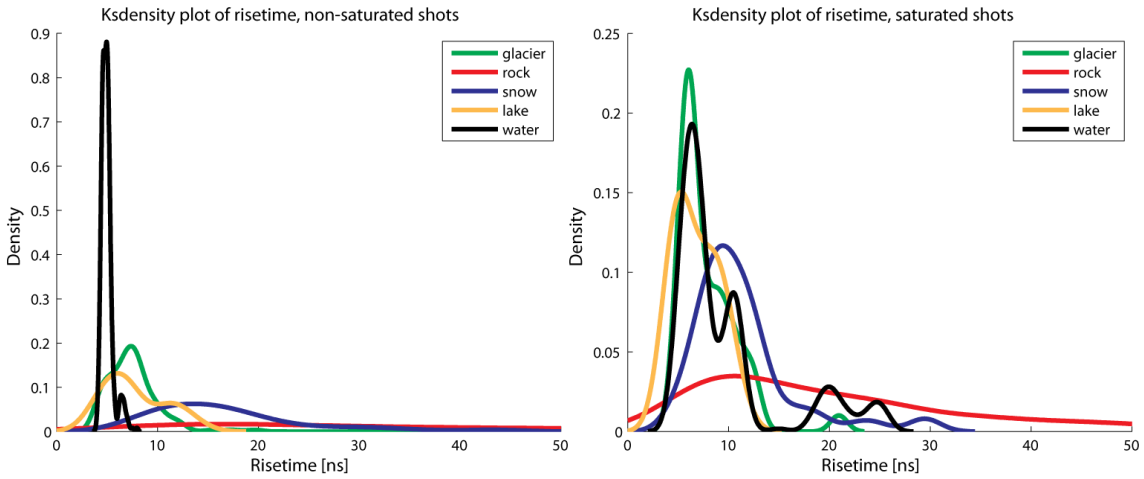

Figure 8.24: The Kernel smoothed probability density of risetime per land cover type. The plots are truncated at $50ns$ for visualization purposes, resulting in the exclusion of 38% (left plot) and 9.3% (right plot) of the rock distributions.

For non-saturated shots, the risetime is a good alternative for width. However for saturated shots, the risetime values are likely to be biased by incorrect peak identification and using the width becomes more favorable. As the case studies in Section 8.1 demonstrated, the risetime mainly follows the same trends as the waveform width. The correlation coefficient between width and risetime is 0.83, indicating the two are considerably correlated. The part of risetime that is not correlated with the width does not contribute to better classification results. Because of this correlation with width; the lack of contribution of the non-correlated part to the classification; the biased values for non-saturated shots; and to keep the classification algorithm as simple as possible, this research does not include the risetime as a classification attribute.

Number of peaks

The number of peaks is also an alternative for waveform width, albeit less optimal. Figure 8.25 shows the Kernel smoothed distribution of the number of peaks. It shows that shots covering water and lake ice mostly have only one peaked waveforms for both non-saturated and saturated shots. The correlation coefficient with waveform width is 0.82, similar to the correlation coefficient of risetime with width. The main reason that the number of peaks offers a less optimal classification potential compared to the width is because it concerns discrete values and the overlap between the land cover types is significantly larger. Of all non-saturated shots, 89%, 22%, 81%, 100% and 100% of the corresponding land cover types have a number of peaks equal to one. For saturated shots these percentages are 76%, 36%, 83%, 94% and 64%. That 36% of all saturated shots covering water have a number of peaks larger than one could indicate that waves or surface roughness are the cause of saturation. On the other hand, for saturated shots the computation of peaks is also affected by the incorrect recording of the waveform. The main reasons that the number of peaks are not incorporated in the classification algorithm are the high correlation to waveform width (which is a better alternative) and the high overlap of distributions per land cover type.

Atmospheric channels ground return signals

The green and infrared ground return signals do not show classification potential based on the case studies in Section 8.1 and suffer from highly undesired effects caused by changing transmit energy. The correlation coefficient between the two attributes is 0.77; between the green infrared signals and reflectivity 0.34; and between the infrared signals and reflectivity 0.60. All three values are not sufficiently convincing to conclude that one attribute makes the other attribute redundant, although the correlation between the two atmospheric attributes can be interpreted as interesting. This correlation is most probably because both signals are recorded similarly and are influenced by the similar attenuation and scattering factors. The reason both are not entirely correlated is likely because the transmit energy is much lower for the green channel and the green signal is more sensitive for atmospheric

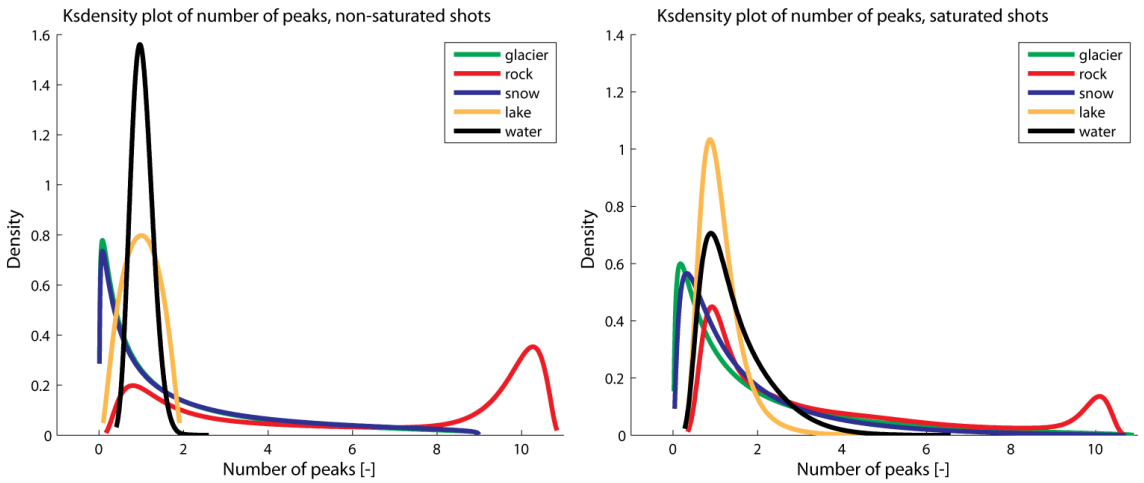

Figure 8.25: The Kernel smoothed probability density of the number of peaks per land cover type.

effects than the infrared signal. Figure 8.26 and Figure 8.27 show the distributions of the green and infrared ground return signals, respectively.

For green signals, none of the plots show dominant distribution clusters that can be used for classification purposes. One of the interesting features of both plots is the small individual clusters of lake ice in the left plot and of water in the right plot. These clusters are a direct consequence of the varying transmit energy. For lake ice, only a few non-saturated shots cover this land cover type each campaign. As a result, the differences between campaigns is clearly visible and show the effect of transmit energy on the backscatter values. About five main peaks can be identified in the lake ice distribution. As demonstrated by the case studies, the backscatter values increase with decreasing transmit energy.

Similarly, for the infrared signals no clear separation can be made between the land cover types. The non-saturated distributions show that about half of the values for snow are somewhat higher than the other land cover types. This can be related to the higher reflectivity values of snow compared to the other land cover types. In the left plot, the distribution of lake ice shows more homogeneous results than is the case for the green signals. This can be explained by the smaller transmit energy differences of the infrared signals throughout the campaigns and by the smaller sensitivity of infrared signals to atmospheric effects.

The main reasons that the green and infrared signals of the atmospheric channels are not used for classification purposes are the reliability issues with the signals caused by the transmit energy variations (in particular for the green channel) and the poor potential for making clear distinctions between the land cover types. If the same attributes would be available for the green signals as are available for the infrared signals (i.e. similar waveform recording, etc.) and the transmit energy would be sufficiently high and constant, the green

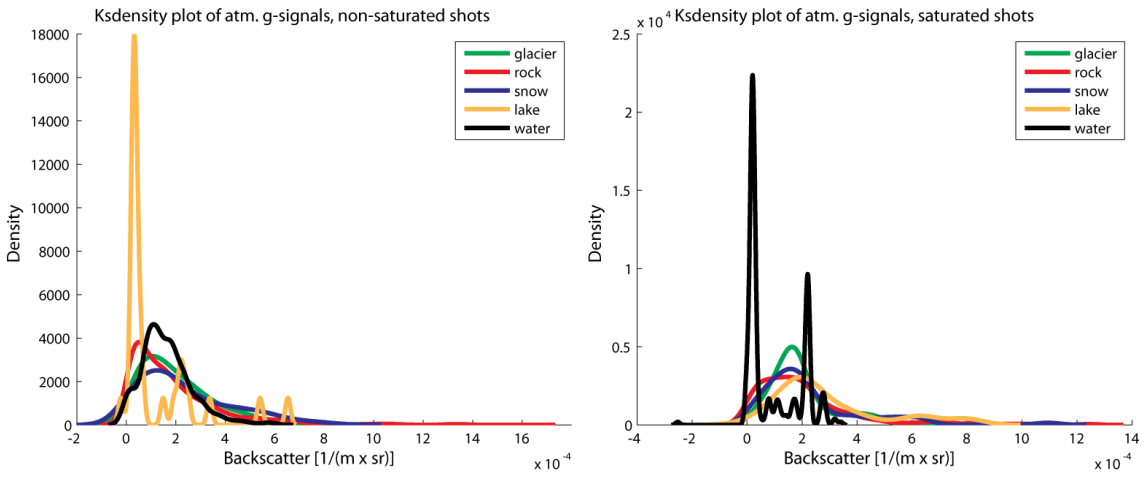

Figure 8.26: The Kernel smoothed probability density of the green ground return signal using the atmospheric channel per land cover type.

signal may offer significant improvements to the classification algorithm. In particular, the expected reflectivity of green signals in combination with the reflectivity of infrared signals could result in a spectral signature that would significantly improve the classification. For an example of what difference green light reflectance makes compared to infrared light reflectance, see Figure 4.4 in Section 4.5.

8.5 Classification criteria

In the previous section, the distributions of the attributes per land cover type and their classification potentials were analyzed and discussed. Based on that information and the case studies, the classification criteria are determined. In this section these criteria will be explained by means of a decision tree. For each decision, two branches are defined and the impact of each decision on the further branches will be analyzed with percentages. Finally, each end branch leads to a classification result. These results are analyzed with percentages as well. The exact methodology and functioning of the tree and its assessment will become clear during the process of classification. The intention of this research is to make the decision tree with all its classification criteria as simple as possible, while keeping the classification as accurate as possible, such that the classification algorithm becomes generally applicable.

To accompany the entire classification process, the decision tree is shown in Figure 8.28. The attributes shown in the top are required for the classification after which the classification is started. For each branch, the percentage per land cover type is shown with respect to the previous decision. Finally, the diamonds show the final classification decisions. The percentages shown below these diamonds are the final percentages of the

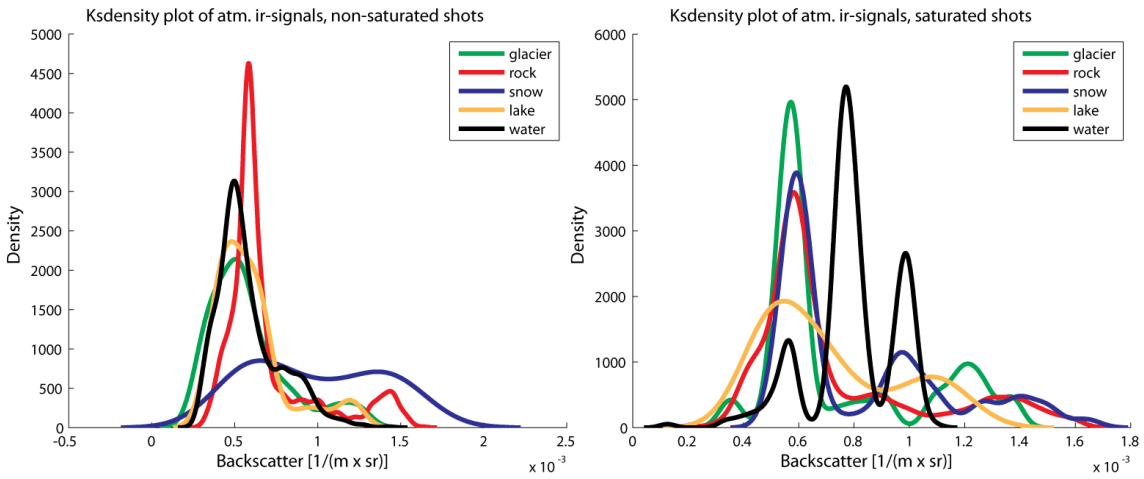

Figure 8.27: The Kernel smoothed probability density of the infrared ground return signal using the atmospheric channel per land cover type.

land cover types with respect to all shots. All percentages will be explained in the following discussions. The branches are discussed from top to bottom and from left to right. The two main branches, the non-saturated and saturated branches, are indicated individually in the coming discussions. Please keep in mind that following order of land cover type is used, unless stated otherwise: glacier ice, rock, snow, lake ice and water (the same order as the decision tree).

The classification criteria are generally based on particular attribute values. The values have been defined such that the classification algorithm is defined in the most optimal way. However, the values have been rounded to a easily interpretable number as well. For example, the criteria that are related to the waveform width are based on a $50ns$ value. A slightly higher value (e.g. $50.12ns$) might result in an improvement of a tenth of a percentage in classification accuracy, but it negatively affects the simplicity of the algorithm. Thus, for each criterion a balance is found between an optimal algorithm and interpretable values.

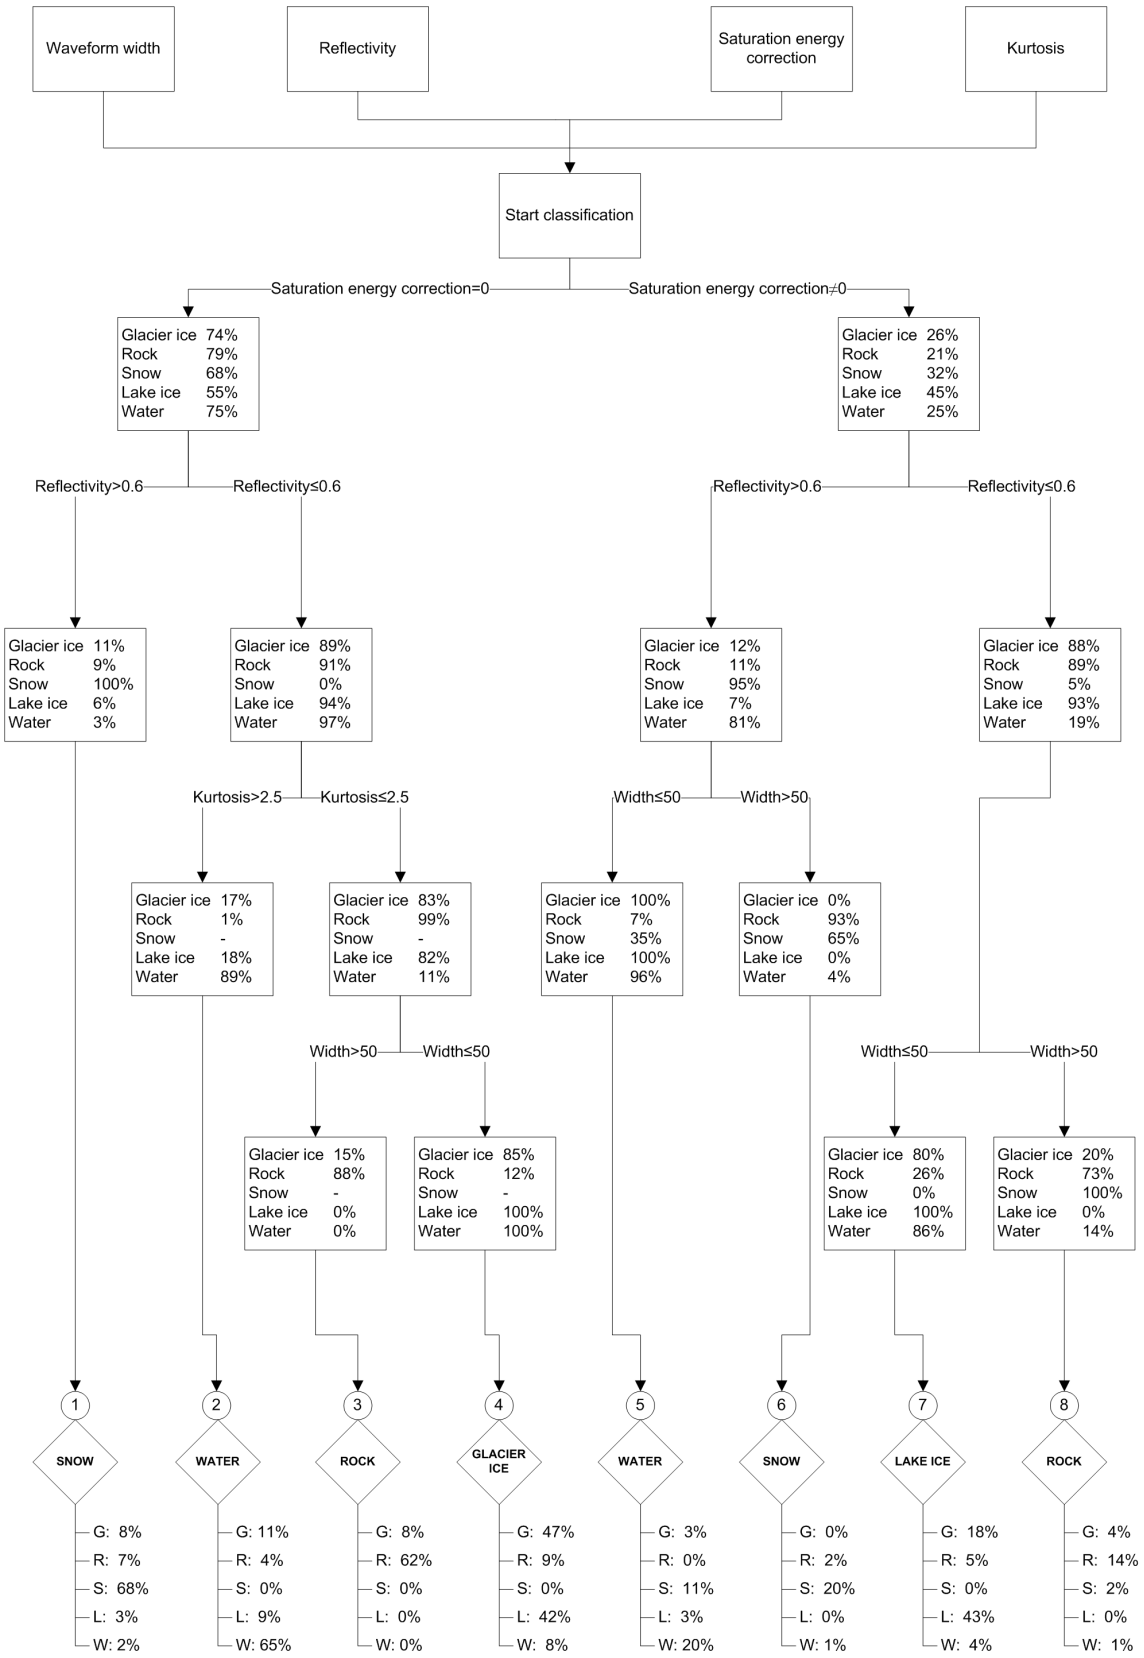

Figure 8.28: Classification decision tree based on ICESat's attributes. The squares indicate classification criteria, the diamonds indicate the decisions and the circles indicate the numbers corresponding to the decisions. Please note that the glacier ice and lake ice will be merged as will be explained further in this section.

8.5.1 Non-saturated shots

Of all shots covering glacier ice, rock, snow, lake ice and water, 74%, 79%, 68%, 55% and 75% respectively are non-saturated, see Section 8.4. As mentioned before, the reflectivity is the most important attribute to distinguish snow from the other land cover types. Consequently, the decision after the isolation of the non-saturated shots distinguish the shots based on the reflectivity criterion of 0.6. For the branch reflecivity higher than 0.6, the percentages indicate the number of shots having reflectivity values higher than 0.6 with respect to the non-saturated shots. So, 11%, 9.0%, 100%, 5.6% and 2.5% of the non-saturated shots covering the corresponding land cover types have reflectivity values higher than 0.6. Directly after this branch the first decision is made: **if a shot is non-saturated and has reflectivity value higher than 0.6, the shot is classified as snow**. The result is that 8.2%, 7.1%, 68%, 3.0% and 1.9% of all shots covering the corresponding land cover types is classified as snow.

For the non-saturated shots having reflectivity values lower than or equal to 0.6, more complex classification criteria are required. As the decision tree indicates, 89%, 91%, 0%, 94% and 97% of all non-saturated shots comply with this reflectivity criterion. To filter out the shots covering water, the criterion of kurtosis higher than 2.5 is used. As a result, 17%, 0.53%, -, 18% and 89% of all non-saturated shots having reflectivity lower than of equal to 0.6 covering the corresponding land cover types have kurtosis higher than 2.5. The '-' indicates that no snow shots are left for this criterion, because of the reflectivity criterion earlier. After the kurtosis criterion the second decision is made: **if a shot is non-saturated, has a reflectivity value lower than or equal to 0.6 and has kurtosis value higher than 2.5, the shot is classified as water**. The result is that 11%, 3.8%, 0%, 9.0% and 65% of all shots covering the corresponding land cover types is classified as water.

For the shots having kurtosis lower than or equal to 2.5 and comply with the previous criteria, 83%, 99%, -, 82% and 11% of the corresponding land cover types are left. As these percentages together with the percentages of the preceding criteria indicate, a large portion of glacier ice, rock and lake ice has to be properly distinguished. Section 8.4 already indicated that the distinction between these land cover types is problematic, because none of the attributes offer a clear separation potential between these types. The attribute that offers the best solution is the waveform width. Basically, this width criterion is a measure for the flatness of the surface. For the Dry Valleys this criterion holds reasonably well: 85%, 12%, -, 100%, 100% of all shots that comply with the previous criteria have width smaller than or equal to $50ns$. This comes down to the surface having a slope smaller than 3°. For the glacier and lakes of the Dry Valleys this indeed is true. However, when glacier ice has high slopes or large roughness, this width criterion may pose a problem. This is one of the major pitfalls of this classification scheme. However, since attributes like reflectivity, kurtosis or skewness do not offer classification potential, and waveform attributes like risetime and number of peaks are heavily correlated and related to the same problem, it is chosen to use the width for classification criterion. Consequently, after the kurtosis

criterion the third decision is made: **if a shot is non-saturated, has reflectivity value lower than or equal to 0.6, has kurtosis value lower than or equal to 2.5 and has width value larger than 50*ns* the shot is classified as rock**. The result is that 8.2%, 62%, 0%, 0% and 0% of all shots covering the corresponding land cover types is classified as rock. Also the fourth decision can be stated: **if a shot is non-saturated, has reflectivity value lower than or equal to 0.6, has kurtosis value lower than or equal to 2.5 and has width value smaller than or equal to 50*ns* the shot is classified as glacier ice**. The result is that 47%, 8.7%, 0%, 42% and 8.3% of all shots covering the corresponding land cover types is classified as glacier ice.

8.5.2 Saturated shots

Of all shots covering glacier ice, rock, snow, lake ice and water, 26%, 21%, 32%, 45% and 25% respectively are saturated, see Section 8.4. Again, the first criteria involves the reflectivity to separate the largest portion of snow from the other land cover types. The result is that 12%, 11%, 95%, 6.5% and 81% of the saturated shots covering the corresponding land cover types have reflectivity higher than 0.6. This means that a substantial portion of the saturated water shots are also selected, see also Figure 8.19. Since the kurtosis is not reliable for saturated shots, it is decided to use width to distinctguish between snow and water. Setting the width criterion to 50*ns*, results in 100%, 6.7%, 35%, 100% and 96% of all shots that comply with the previous criteria have a width value smaller than 50*ns*. From these percentages alone, it would make sense to classify these shots as glacier ice or lake ice. However, when taking into account the percentages of the criterion before (i.e. the reflectivity criterion), the larger effective percentage belongs to water. So the fifth decision is made: **if a shot is saturated, has reflectivity higher than 0.6 and has width smaller than or equal to 50*ns* the shot is classified as water**. As a result, 3.3%, 0.16%, 11%, 3.0% and 20% of all shots covering the corresponding land cover types is classified as water. This means that still a significant part of the shots covering snow is classified as water. It shows that saturated shots give different results than non-saturated shots. Also the sixth decision can be stated: **if a shot is saturated, has reflectivity value higher than 0.6 and has width value larger than 50*ns* the shot is classified as snow**. The result is that 0%, 2.3%, 20%, 0% and 0.76% of all shots covering the corresponding land cover types is classified as snow.

For the non-saturated shots with reflectivity values lower than or equal to 0.6, again the width is used to separate between glacier ice and lake ice, and rock. Consequently, 80%, 26%, 0%, 100% and 86% of all shots that comply with the previous criteria have width smaller than or equal to 50*ns*. This also yields the seventh decision: **if a shot is saturated, has reflectivity value smaller than or equal to 0.6 and has width value smaller than or equal to 50*ns* the shot is classified as lake ice**. As a result, 18%, 4.9%, 0%, 43% and 4.1% of all shots covering the corresponding land cover types is classified as lake ice. Also the eighth decision can be stated: **if a shot is saturated, has reflectivity value smaller than or equal to 0.6 and has width value smaller**

larger than 50*ns* the shot is classified as rock. As a result, 4.4%, 14%, 1.5%, 0% and 0.66% of all shots covering the corresponding land cover types is classified as rock.

When looking at the percentages under each decision, it appears that for many decisions the classification tree performs reasonably well. In the next chapter, the exact classification accuracy assessment is carried out. However from these percentages one noticeable pitfall in particular becomes apparent, namely the lack of distinction between glacier ice and lake ice. This was also indicated in Section 8.4 and with the observations in this section it is decided to merge the two land cover types. When also considering the practical implication, merging the two is reasonable since lake ice is very similar to glacier ice. The difference between the two is more a matter of slope angle and roughness than a difference in physical substance. In addition, flat glacier ice in particular is practically indiscernible from lake ice, merging the two would prevent unnecessary misclassifications. In addition, when considering the research question, the emphasis is on finding a classification algorithm between ice, rock, snow and water, rather than between different ice forms. For improved mass balance estimations, computing the difference between the ice forms does not make a significant contribution, in contrary to the other land cover types. Therefore, because of the limited relevancy and the difficulty of making proper distinction between the different ice forms, the classification of glacier ice and lake ice are merged. The combination is simply called 'ice' hereafter. The reason that glacier ice and lake ice were separated in the first place, is because it was relatively easy to implement in the supervised classification map (see Section 7.1). As a consequence, it was possible to analyze the potential of ICESat's signals to discriminate between the two ice types.

Since the classification algorithm has been determined, it can be applied on test areas. In Chapter 9 several the classification results on various areas will be discussed. By comparing the classification results of the Dry Valleys and Lake Victoria with the supervised classification map, the accuracy can be determined. Furthermore, in Chapter 9 the classification results of Jakobshavn in Greenland and of Greenland itself will be discussed as well.

Chapter 9

Classification results

In this chapter, the classification algorithm will be applied on several test areas. First of all, in Section 9.1 the classification results of the Dry Valleys and Lake Victoria will be compared with the supervised classification map. By comparing the two, the accuracy of the classification algorithm can be determined. In Section 9.2, the classification algorithm is tested on the Jakobshavn area in Greenland. Finally, in Section 9.3, the entire area of Greenland is used as test area. The resulting classification maps serve as practical examples of the algorithm.

For the Dry Valleys and Lake Victoria, the same shots are used with which the classification criteria have been determined. Consequently, the classification results of Section 9.1 can be directly compared to Figure 8.17. For the Jakobshavn area, three individual campaigns are classified and plotted on corresponding LANDSAT images. For entire Greenland, two campaigns are classified. The campaigns that are selected for the Jakobshavn area were flown in the end of the winter and in spring, such that the areas are not only covered by snow and sea ice, but also include bare rock, open water and glacier ice. Similarly, for the Greenland area, the two selected campaigns were flown in winter and spring to show the difference in classification; it makes the classification more challenging.

9.1 Classification accuracy

With the defined classification criteria, the shots that were classified with the supervised classification map are considered to be true and act as reference data. This data can be used to assess the classification accuracy by feeding these into the classification decision tree and compare the resulting computed classes to the reference classes. With this information, a confusion matrix, the overall accuracy and $\hat{\kappa}$-accuracy can be computed. The accuracy assessment is carried out for all available reference data as well as a random subset of the reference data. First of all, the classification results of all shots plotted on the Dry Valleys and Lake Victoria images will be discussed.

9.1.1 Dry Valleys and Lake Victoria classification results

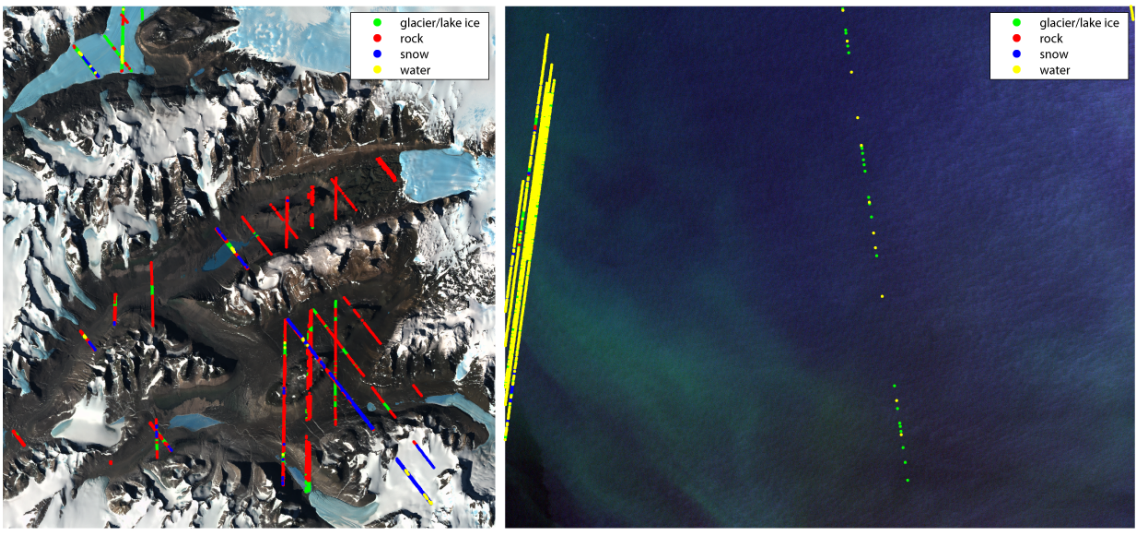

Figure 9.1: The classification results over the Dry Valleys (left) and over Lake Victoria (right).

Figure 9.1 shows the classification results of all shots that were classified by the supervised classification map, i.e. all reference data available over the Dry Valleys and Lake Victoria. As the left image shows, the classification of shots generally correspond to the land cover of the LIMA image. As expected, the misclassification occurs mostly in the flat part of valleys, where shots are classified as glacier but in reality cover rock. In addition, some misclassification occurs at Taylor Glacier, where rock and water classified shots are mixed with ice classified shots. Finally, one particular track shows snow classified shots over the entire region. Because some tracks overlap in the Dry Valleys, some close-ups of this region are included, see Figure 9.2. Before discussing these, first some comments on Lake Victoria will be given. As the right image of Figure 9.1 shows, the shots are mostly classified correctly. As the confusion matrix of Table 9.1 will indicate, 84% of all shots covering water have been correctly classified as water. In particular the isolated track shows that the majority of the shots have been classified as ice.

The four close-ups of Figure 9.2 are given to provide some detailed analysis of the misclassified shots over the Dry Valleys. First of all, the shots over Taylor Glacier show a combination of correct and incorrect classification results. In most misclassified cases the shots are classified as water, and these occur mainly on the flat parts. The reason for this is that these parts give high kurtosis values and are consequently classified as water. The shots that are misclassified as rock mainly occur at the rougher parts of the glacier (i.e. the upper shots). These parts give high waveform width values and consequently they are not classified as ice. Finally, some consecutive shots were classified as snow. These shots be-

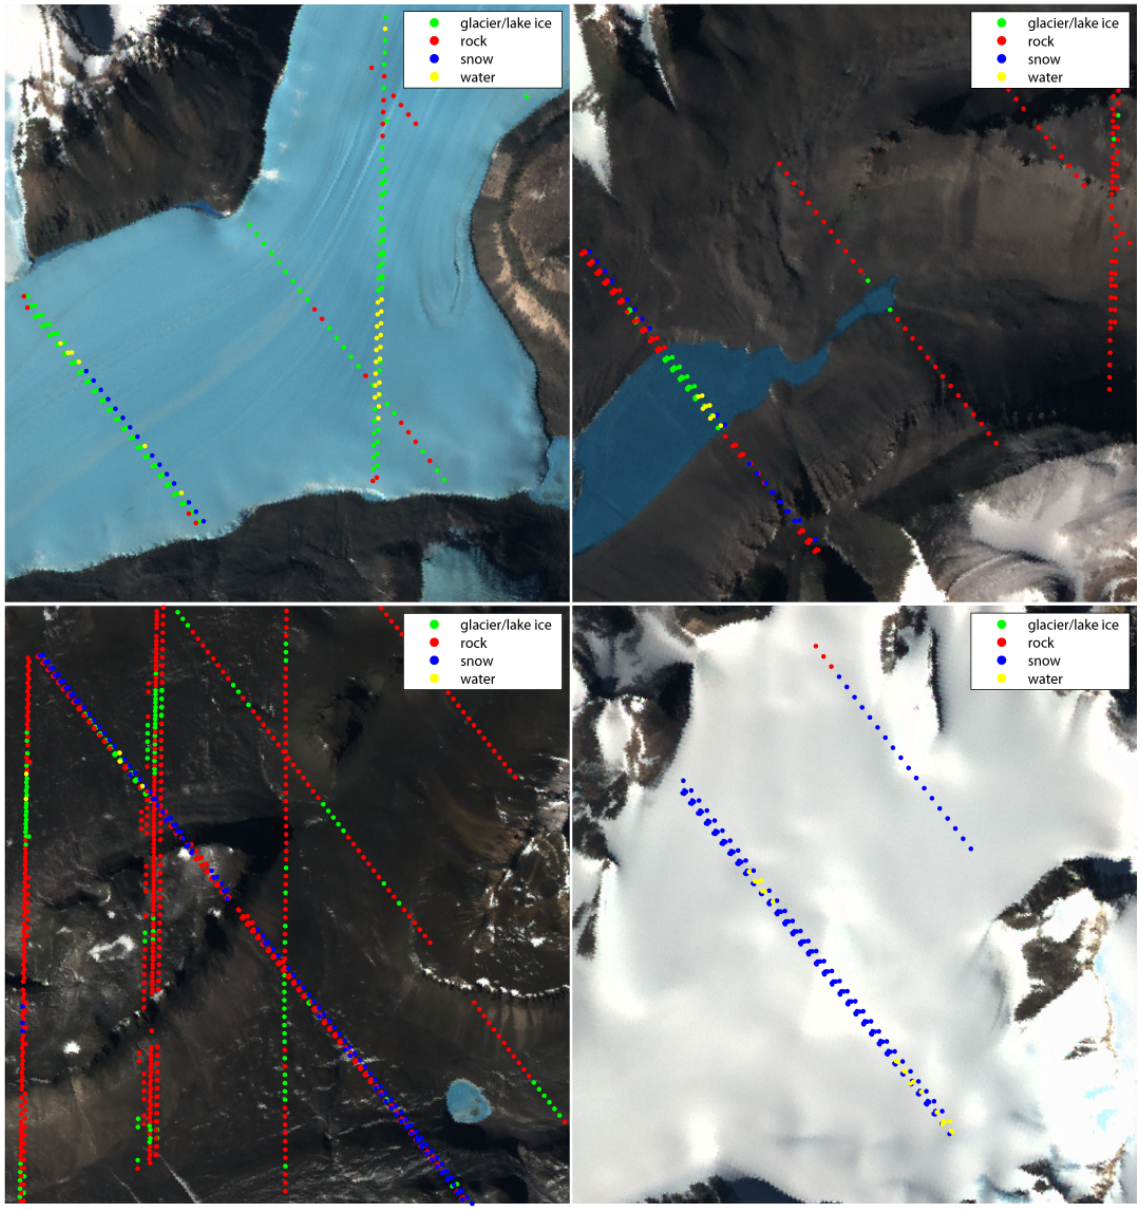

Figure 9.2: Detailed images of the classification results over the Dry Valleys. The areas of the images starting at the upper left in clockwise direction are: Taylor Glacier; Lake Vanda; snow elevated plain; and Barwick Valley and Mckelvey Valley.

long to one track over the Dry Valleys that contains a high number of snow classified shots.

For the Lake Vanda region, the shots have been classified mostly as rock over the rock area and ice over the lake. In addition, some shots that have been classified as water occur on the lake. Noticeable is the clear transition between where the shots classified as rock

stop and where the shots classified as ice begin. The shots that were classified as snow mostly occur in the same track that show snow classified shots over Taylor Glacier.

For the snow elevated plain, by far most shots have been correctly classified as snow and only a few have been classified as water or rock. The water classified shots appear to occur at specific places, at the most flat parts of the plain where the kurtosis values are high. Three rock classified shots occur closer to the boundary of the plain, which might indicate the surface is more tilted.

The fourth and final image shows the Barwick Valley and Mckelvey Valley region. Most tracks show a mix of ice and rock classified shots. Especially in the flat valleys parts, the shots are misclassified as ice. This is mainly due to the low waveform width values that occur on these locations. Next to the ice and rock classified shots, some shots were classified as snow as well. Most of these shots belong to the same track mentioned before. This could indicate that during the acquisition date of this track, some snow covered the Dry Valleys or that anomalies occurred during the entire track. Finally, the image shows that virtually none of the shots were classified as water. This is most probably due to the low kurtosis values.

Generally, the shots were classified correctly. However, some persistent misclassifications occurred at specific sites. The shots that were misclassified as ice mostly occur at the flat valley areas; the shots that were misclassified as rock mostly occur at rougher glacier areas; the shots that were misclassified as snow mostly belong to one particular track; and the shots that were misclassified as water mostly occur at flat areas like flat glacier parts, lake ice and flat snow parts.

9.1.2 Confusion and accuracy matrix, all shots

Table 9.1 shows the confusion matrix of the shots used in the previous discussion (i.e. all shots over the Dry Valleys and Lake Victoria). The equations that are used to make the confusion matrix have already been explained in Section 7.1. A short description of what the matrix represents is briefly summarized below using Table 9.1 as an example. The classes per column indicate the land cover types that have been classified with the classification map. These are considered to be reference data, or ground truth. The classes per row indicate the land cover types that have been classified with the classification algorithm based on ICESat's attributes (i.e. the decision tree). The 'total' is the sum of the row entries and represents the number of ground truth shots that are fed into the classification algorithm. So, ideally on the diagonal of the inner matrix these total numbers should be present, i.e. 250 on the intersection 'ice' and 'ice', 1861 on the intersection 'rock' and 'rock', etc. If this would be the case, all omission and commission errors would be 0% and the overall and $\hat{\kappa}$-accuracy would be 100%, i.e. perfect classification. However, the classification is not perfect because some shots have been misclassified.

The misclassified shots should be interpreted as follows. Since the columns represent the ground truth, all off-diagonal entries in the columns indicate the number of shots that are classified as the land cover type indicated by the row, but should be classified as the land cover type indicated by the column. For example for the column ice, 176 shots were classified correctly, but 23 were classified as rock, 17 were classified as snow and 34 were classified as water. The omission error then indicates the percentage that is omitted from correct classification. For ice, $(23 + 17 + 34)/250 \cdot 100\% = 30\%$ is classified as other land cover types instead of ice. The commission error indicates the percentage that is incorrectly included in the land cover type. For ice, $(252 + 0 + 131)/559 \cdot 100\% = 69\%$ is classified as ice, but these should have been classified as rock and water respectively. The total error is the sum of all off-diagonal entries of the inner matrix divided by the total number of classified shots, and is equal to 100% subtracted by the sum of the diagonal entries of the inner matrix divided by the total number of classified shots: $(252 + 0 + 131 + 23 + 3 + 7 + 17 + 174 + 28 + 34 + 10 + 21)/3365 \cdot 100\% = 100\% - (176 + 1425 + 171 + 893)/3365 \cdot 100\% = 21\%$.

		Ground truth					
Classification	**Class**	**Ice**	**Rock**	**Snow**	**Water**	**Total**	**Commission**
	Ice	176	252	0	131	559	69%
	Rock	23	1425	3	7	1458	2.3%
	Snow	17	174	171	28	390	56%
	Water	34	10	21	893	958	6.8%
	Total	250	1861	195	1059	3365	
	Omission	30%	23%	12%	16%		21%

Table 9.1: The confusion matrix for the for the classification algorithm. The table is expressed in number of pixels and the corresponding error percentages.

Table 9.2 shows the accuracy matrix that corresponds to the confusion matrix of Table 9.1. The producer's accuracy is identical to 100% minus the commission error and the user's accuracy is identical to 100% minus the omission error. The producer's accuracy shows that particularly the classification of rock and water performs well since it includes few shots that belong to other classes. The user's accuracy shows that the classification of all land cover types performs well by not excluding many shots that belong to the particular class. The overall accuracy is equal to 100% minus the overall error and is 79%. The $\hat{\kappa}$-accuracy basically compensates the accuracy for random classification and is always lower than the overall accuracy. See Section 7.1 for a detailed explanation about the $\hat{\kappa}$-accuracy. In this case the $\hat{\kappa}$-accuracy is considerably lower than the overall accuracy because of relatively high commission errors of ice and snow.

Briefly explained, the overall accuracy indicates that the classification algorithm correctly classifies 79% of all shots. The $\hat{\kappa}$-accuracy indicates that the classification algorithm correctly classifies 68% of all shots taking into account chance agreement. However, these percentages should be interpreted with great caution, because of the large differences in available shots per true land cover type. For example, for rock, 1861 shots are available while for snow only 195 shots are available, more than a factor nine difference. To normalize these differences, another confusion matrix and accuracy matrix are constructed.

These matrices use only a subset of the shots available, see Paragraph 9.1.4. First of all, the misclassifications of Table 9.1 will be discussed in the next paragraph.

Class	Producer's accuracy	User's accuracy
Ice	31%	70%
Rock	98%	77%
Snow	44%	88%
Water	93%	84%
Overall accuracy	79%	
$\hat{\kappa}$**-accuracy**	68%	

Table 9.2: The accuracy matrix corresponding to Table 9.1.

9.1.3 Misclassification analysis

The confusion matrix gives a clear overview where the classification algorithm performs correctly and where misclassification occurs. It shows the direct impact of misclassified shots on the resulting error percentages. Four primary misclassification results are highlighted and possible explanations are given.

First of all, ice has a very high commission error. One of the main causes relate to the 252 shots covering rock which are incorrectly classified as ice. The shots were classified as ice, because of the width criterion, i.e. when the width is smaller than or equal to $50ns$ the shots are classified as ice, otherwise the shots are classified as rock. As explained before, since there is no other attribute that offers a clear distinction between the two, the best possibility is to classify the two as correctly as possible by using the width criterion, which is highly correlated with the surface slope. As a result, the shots covering rock surface with very little surface slope (i.e. in the midst of the Valleys), are incorrectly classified as ice. Because there are many more shots available that cover rock compared to the number of shots that cover ice, the impact on the commission error is significant. The argument that in reality some of the misclassified rock shots are indeed ice is not very plausible since the temporal difference between the supervised classification map is more likely to be dominated by snowfall, resulting in higher reflectivity values. The hypothesis that in reality a melting snow layer became a small ice layer at the time of measurement might contribute to the misclassification. However, this cannot be verified and is much less likely than the argument of the width criterion.

The second misclassification result that warrants attention relates to the 131 shots covering water that are incorrectly classified as ice. Similar to before, the reflectivity does not offer clear separation capabilities between ice and water. As the previous section and the decision tree indicate 8.3% and 4.1% of all shots covering water are classified as ice corresponding to the fourth and seventh decision respectively. For both decisions the preceding criteria include waveform derived attributes (i.e. kurtosis and width). Since the width distributions of the two land cover types have some overlap, the misclassification is a understandable consequence. Apparently, some of the shots covering water have larger width

values and smaller kurtosis values than usual and are consequently classified as ice. Waves and roughness of the water surface could cause this effect, however this is not verified. About 12% of all water shots are incorrectly classified as ice. The much higher number of shots that cover water compared to the number of shots that cover ice contributes to the high comission error.

The third misclassification result that warrants attention relates to the 174 shots covering rock but have been classified as snow. As mentioned in the previous section, it may be possible that part of these shots are indeed snow and are consequently correctly classified as snow. However, because of difficulties with validation this cannot be confirmed. Although it is not desired to use attribute findings for hypothesis testing, the reflectivity of snow is one of the most reliable classification criteria of the decision tree. Since these particular shots covering rock have been classified as snow based on the reflectivity criterion, this somehow indicates that the shots might be snow in reality. Although this cannot be accepted as an argument, it indicates the possible deviation from the results based on the supervised classification map. Additionally, many of these shots were taken in the same unique track (i.e. on the same day), which makes it more probable that on that particular day there was some snow covering the rock surfaces.

The fourth and final case concerns the 34 shots that cover ice but are classified as water. Again, the waveform derived attributes play a major role in the classification criteria. For flat ice, the kurtosis and width values may overlap with the values assigned to water. Another possibility is that a layer of water was in fact present on the glacier and lake at the time of measurement, as was already discussed in the previous section.

9.1.4 Confusion and accuracy matrix, randomly selected shots

Table 9.3 shows the confusion matrix based on a random subset of all available shots. For each land cover type a random subset is taken such that the total number of reference shots per land cover types is equal. Since the number of shots covering water is the smallest among the other total numbers, for each land cover type 195 shots are randomly selected for classification. As a result, for snow all the available shots are taken into account and for the other land cover types a random subset of 195 shots is taken.

As the table shows, the error percentages are more balanced amongst one another than before. For example, the commission error of ice is now 25% instead of 67% and the comission error for rock is 13% instead of 2.3%. The reason for these changes is that for each land cover type the same number of shots is used for classification, and consequently no more high numbers of incorrectly classified shots are compared to low total numbers and vice versa. Table 9.4 shows the corresponding accuracy percentages. The overall error and consequently the overall accuracy are almost identical as before. Also the user's accuracies are very similar, however the producer's accuracies show less variation. The reason that the producer's accuracies are more equalized is because the commission errors are more

	Ground truth						
Classification	**Class**	**Ice**	**Rock**	**Snow**	**Water**	**Total**	**Commission**
	Ice	138	26	0	20	184	25%
	Rock	20	152	2	0	174	13%
	Snow	13	15	167	7	202	17%
	Water	24	2	26	168	220	24%
	Total	195	195	195	195	780	
	Omission	29%	22%	14%	14%		20%

Table 9.3: The confusion matrix comparing classification map results (ground truth) with the classification algorithm results based on ICESat's attributes. Here a random subset of shots is used. The table is expressed in number of shots and the corresponding error percentages.

sensitive to variation in total numbers than the omission errors are. Consequently, the $\hat{\kappa}$-accuracy is improved to 74%, mainly because of the lower commission errors. Since the subsets are taken randomly, the confusion matrix changes with other subsets. The overall accuracy ranges between 79% and 82%, the $\hat{\kappa}$-accuracy ranges between 71% and 76%.

Class	**Producer's accuracy**	**User's accuracy**
Ice	75%	71%
Rock	87%	78%
Snow	83%	86%
Water	76%	86%
Overall accuracy	80%	
$\hat{\kappa}$-accuracy	74%	

Table 9.4: The accuracy matrix corresponding to Table 9.3.

9.2 Jakobshavn area classification results

The second area on which the classification algorithm has been applied is the Jakobshavn area in Greenland. Three campaigns have been selected for analysis. One in the late winter, two in late spring.

Figure 9.3 shows the location of the Jakobshavn area on Greenland. In fact, Ilullissat is the Greenlandic name of Jakobshavn and is also shown on the map. The area of interest covers Jakobshavn, some ice fjords and islands. In the winter the sea is covered with ice and the coastline is covered with snow. In the spring and summer the sea ice melts and changes to open water, and the coastline shows bare rock and ice. More inland, the land is covered with snow. The figures on the following pages show the classification results of several campaigns plotted on top of the corresponding LANDSAT images. Without any doubt, temporal differences occur between the ICESat acquisition dates and the LANDSAT acquisition dates. However, each LANDSAT image can give an impression of what the conditions were like during each ICESat's campaign. For some images, black stripes run through the data. These data gaps are cause by the faulty SLC of the LANDSAT satellite, see Section 3.2.1 for more information. The blue parts of the images represent

Figure 9.3: A map of Greenland showing the stable ice cap region and indicating the Jakobshavn area. Image taken and adapted from [47]. Image courtesy Philippe Rekacewicz, UNEP/GRID-Arendal.

mainly ice, the red parts represent bare rock, the white parts represent snow and the black parts represent water.

9.2.1 Campaign L1A

Figure 9.4 shows the classification results of L1A, during the late winter of 2003. The LANDSAT image was taken about three weeks later and consequently might show more bare rock than actually was the case during the ICESat's measurements. In addition, there might have been more ice than that the image shows. For the next discussion, the tracks are analyzed upwards. As the upper left image shows, the shots are classified mainly as snow and water and these alternate mostly according to the land cover types the LANDSAT image indicates. The close-ups show more detailed results, especially the differences between several tracks. The upper right image shows the track A that covers partly the sea and partly land. One of the most successful results this image shows is the clear separation between shots covering water (or sea ice) and land. The moment the shots cover land, the classification changes from ice and water into rock and snow. Further up, the classification fails to recognize water or ice. The separation between water (or ice) and

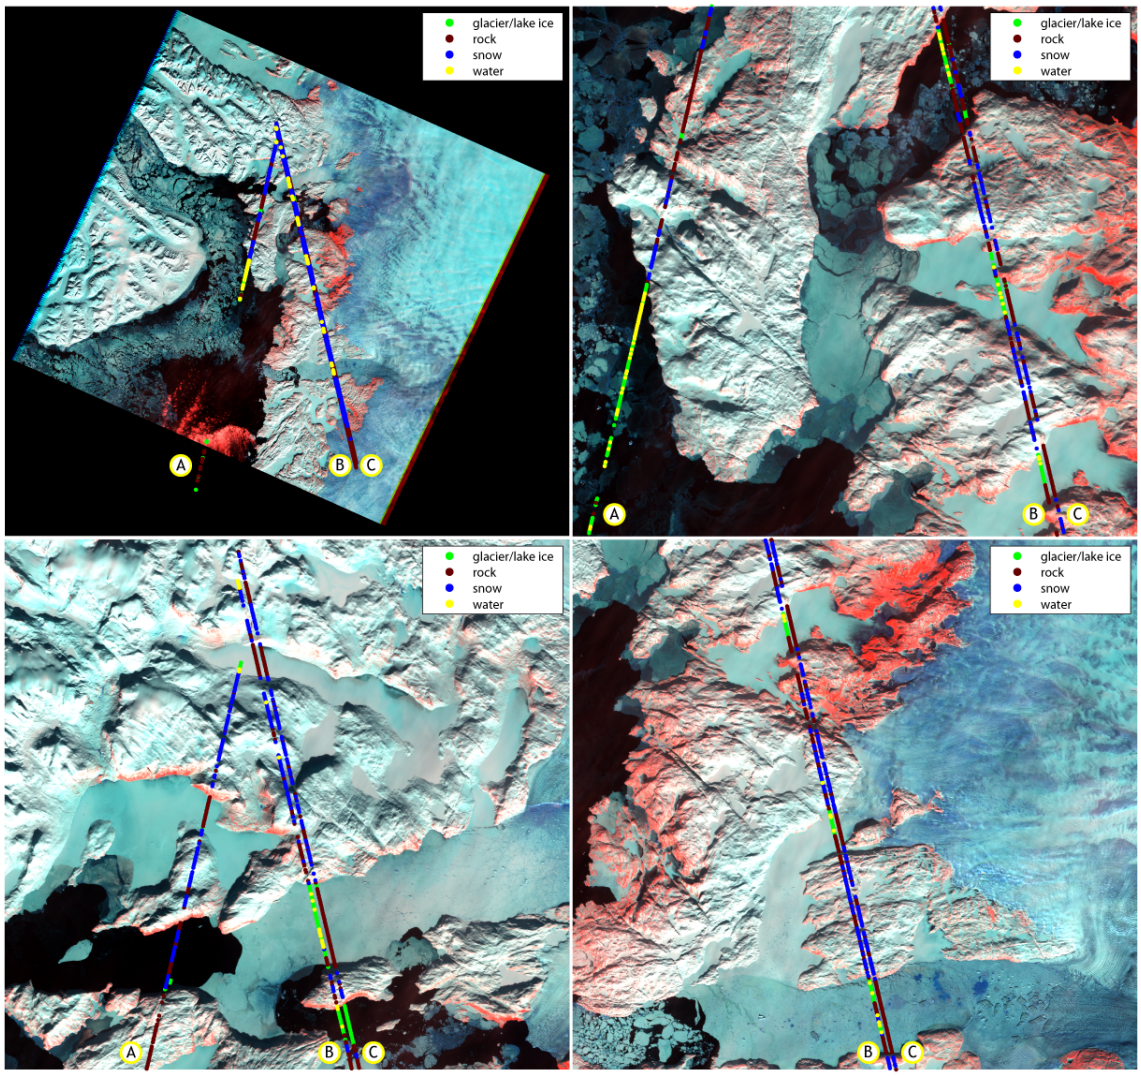

Figure 9.4: The L1A classification results over the Jakobshavn region. The shots were taken between the 21st of February 2003 and the 7th of March 2003, the LANDSAT image was taken on the 30th of March 2003.

land is less evident for the tracks B and C in the right part of the image. Track B shows acceptable results for the first part, but becomes less accurate further upwards. Track C shows the change in land cover type even less accurately. For this track, it appears that the classification confuses between rock and ice. One possible explanation for this is that the ice surface is rougher and consequently has higher width values. Perhaps the difference between tracks B and C is that for track B the ice was flat and track C the ice already broke into pieces and formed sea ice.

The lower right and lower left images show similar results to the upper right image,

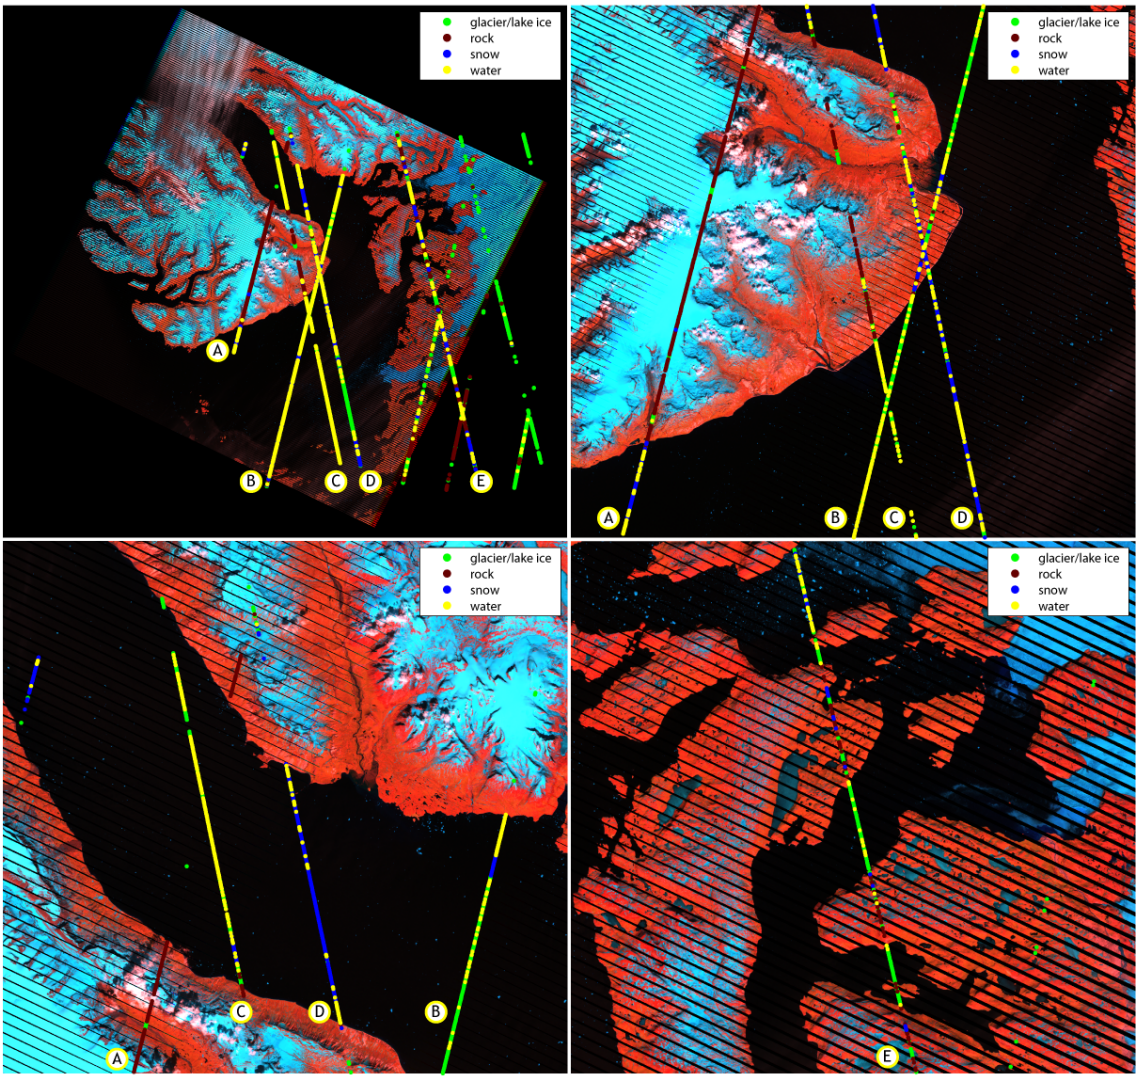

Figure 9.5: The L2C classification results over the Jakobshavn region. The shots were taken between the 19th of May 2004 and the 21st of June 2004, the LANDSAT image was taken on the 12th of May 2004.

where the shots correctly indicate differences in land and ice but with confused land cover types. Here too, the track B show better results than track C, especially downwards. Track A in the lower left image shows the similar results to track C. For all images in this figure, the shots that were classified as snow correspond most accurately with the LANDSAT image. In some cases, the classification correctly shows the transition from ice and water to land, and sometimes it does not.

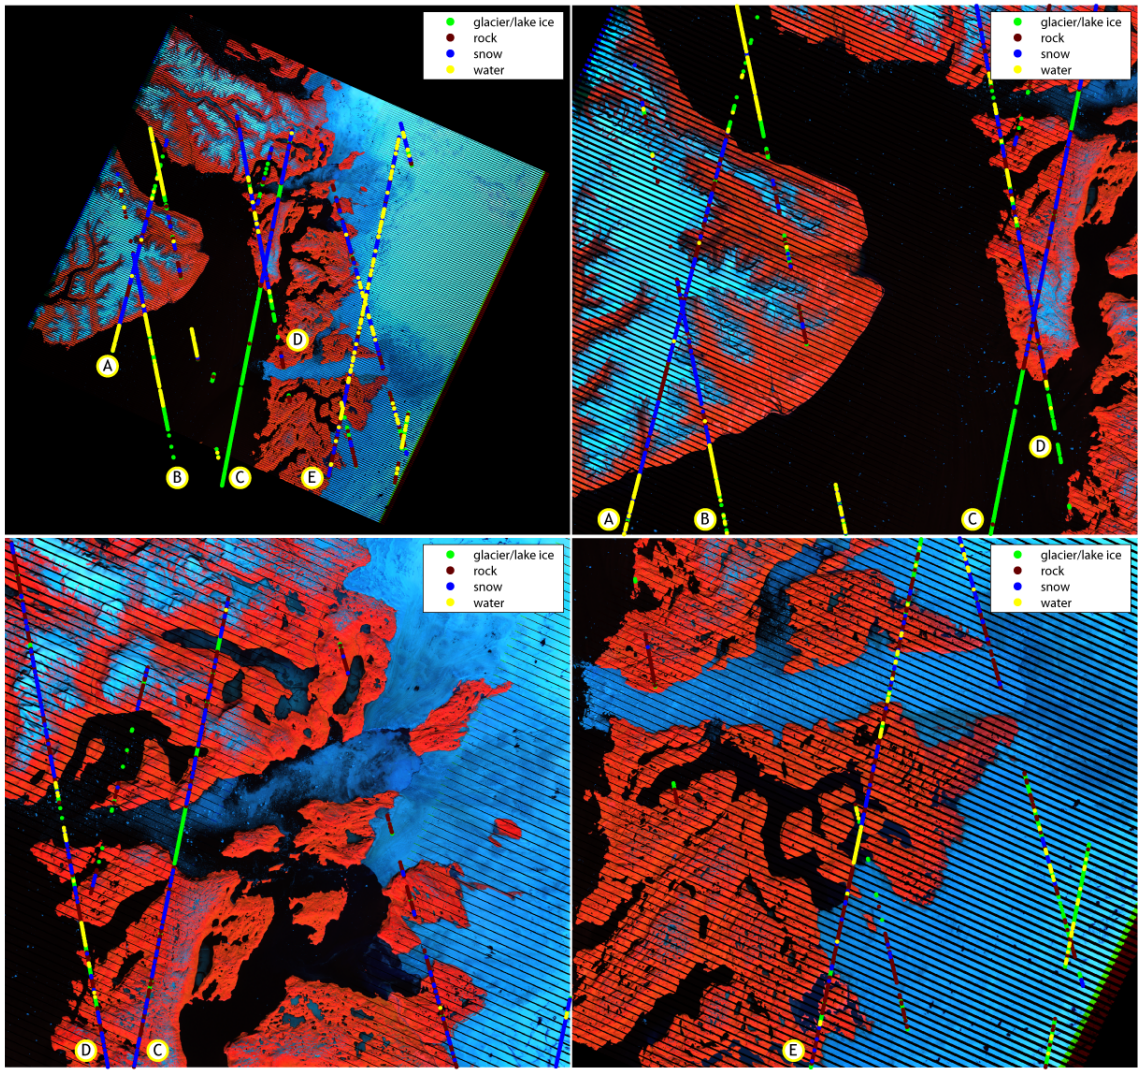

Figure 9.6: The L3F classification results over the Jakobshavn region. The shots were taken between the 25th of May 2006 and the 26th of June 2006, the LANDSAT image was taken on the 26th of June 2006.

9.2.2 Campaign L2C

Figure 9.5 shows the classification results of L2C, during the late spring of 2004. The upper left image shows that most shots have been classified as water and more to the right of the image more shots have been classified as ice. The upper right image contains tracks that cover both land and water. Track A correctly shows the transition from water to land, but fails to classify any shots over land as snow. Whether there was indeed rock at the time of measurement has not been verified. It could be that the extremely low transmit energy causes the low SNR to affect the reflectivity values. The tracks B, C and D show mixed classification results between snow, water and ice. The reason why most shots classified as

snow cover water is not clear, perhaps the water surface roughness plays a role or again the complications with the low transmit energy. The lower left image shows the misclassification of water as well. Especially track D contains many shots that have been classified as snow. The lower right image shows better results for the transitions between water and land. Especially the shots covering rock are classified correctly and clear transitions are present. Even the small island in the lower half of the image has been classified correctly in track E.

For all images in this figure, the most interesting is the misclassification of snow over (assumed) water. On the other hand, transitions in land cover types were generally accurately classified, especially for the lower right image. The extremely low transmit energy might have its effect on the classification results. In addition, the temporal differences between ICESat's measurements and the LANDSAT image could result in misinterpretation.

9.2.3 Campaign L3F

Figure 9.6 shows the classification results of L3F, during the late spring of 2006. The upper left image shows that few shots are classified as rock and most transitions in land cover types have been correctly classified. The upper right image contains tracks that cover both water and land. The tracks A and B show a clear transition between water and snow. Track A continues with shots classified as snow (and occasionally rock) upwards until the shots again cover water. Some shots of track B that cover water have been classified as ice. This occurs even more for the tracks on the right, like track C and D. For these tracks, the transition from water to land have been correctly indicated as well, only most water covering shots have been classified as ice. The over-land shots have been classified as rock and snow. More upwards of track C and D, see the lower left image, the transitions are again indicated correctly, although it is not clear whether shots were correctly classified as ice or not (due to the uncertainty in LANDSAT image). Track D shows similar results as discussed before, with shots mainly classified as snow over land and shots classified as water and ice elsewhere. Interesting to notice is that track C shows a consistent pattern in shots classified as ice and shots classified as snow. The transitions of the narrow fjords have been classified correctly as well (although it is not sure if there was ice or water). The lower right image shows similar results as discussed before as well. For some ice or water parts of track E, some shots have been incorrectly classified as rock, especially over the icy areas. This could relate to rough ice sheet or sea ice.

For all images in this figure, the transitions between land cover types mostly correspond with transitions in classification results. Whether the transition between rock and snow is correct cannot be accurately verified, however the transition between land and water or ice is more evident in the classification results. The fact that there were few shots classified as rock could be related to the time difference between LANDSAT and ICESat's measurements, meaning little rock was indeed visible. The transition from ice or water to rock or snow mostly corresponds neatly to the transitions the LANDSAT image indicates.

9.3 Greenland area classification results

The final test area covers Greenland. For this area, two campaigns are classified: a spring campaign and a winter campaign. The spring campaign, L3C, was flown from the 20th of May until the 23rd of June 2005. The winter campaign, L3E, was flown from the 22nd of February until the 28th of March 2006. For each campaign, a random subset of 500,000 points is taken to reduce computational processing loads. The two classification maps show significant differences due to changing season and give an insight into the success of the algorithm, but also into its deficiencies. A low resolution version of the coastlines is plotted on top of the maps.

9.3.1 Campaign L3C

Figure 9.7 shows the results of the L3C campaign. The shots classified as rock mostly occur at the coastlines. Although the results have not been validated, the locations of the shots that are classified as rock form plausible results. Since the tracks are more densely distributed in the north, the hegemony of classified regions can be better observed than in the south. From north to south along the east coast of Greenland, large continuous areas have been classified as rock. More to the south, the regions are more difficult to identify. Virtually all shots over the island east of Greenland, i.e. Iceland, have been classified as rock, which is very plausible during that time of year. In addition, the area west of Greenland, i.e. the east coast of Canada, shows a large portion of shots classified as rock as well.

Less successful is the classification of Greenland's inland, where relatively few shots are classified as snow and relatively many shots have been classified as ice or water. Some individual tracks show high densities of shots classified as snow, but generally the tracks contain shots classified as ice or water inland and rock at the coast. One possible explanation for the inland results is that higher temperatures in the summer cause the snow to have different properties than the snow that is used for validation in the Dry Valleys. There are also large quantities of results mixed up at the ocean areas. The reason that the shots have not been classified as water could relate to the surface roughness at open sea causing the kurtosis value of the waveforms to decrease and consequently the shots to be classified as ice. This example shows the drawback of using the waveform-derived attributes for classification. In case the reflectivity of water would have been substantially different from the other land cover types, classification based on ICESat data would be simpler and most probably more accurate. In addition, the map shows individual fragments of tracks on the inner part of Greenland that have been classified as rock. This is most probably the result from changing surface topography. When the surface becomes more sloped or rough, the waveform width increases and as a consequence, the shots will be classified as rock.

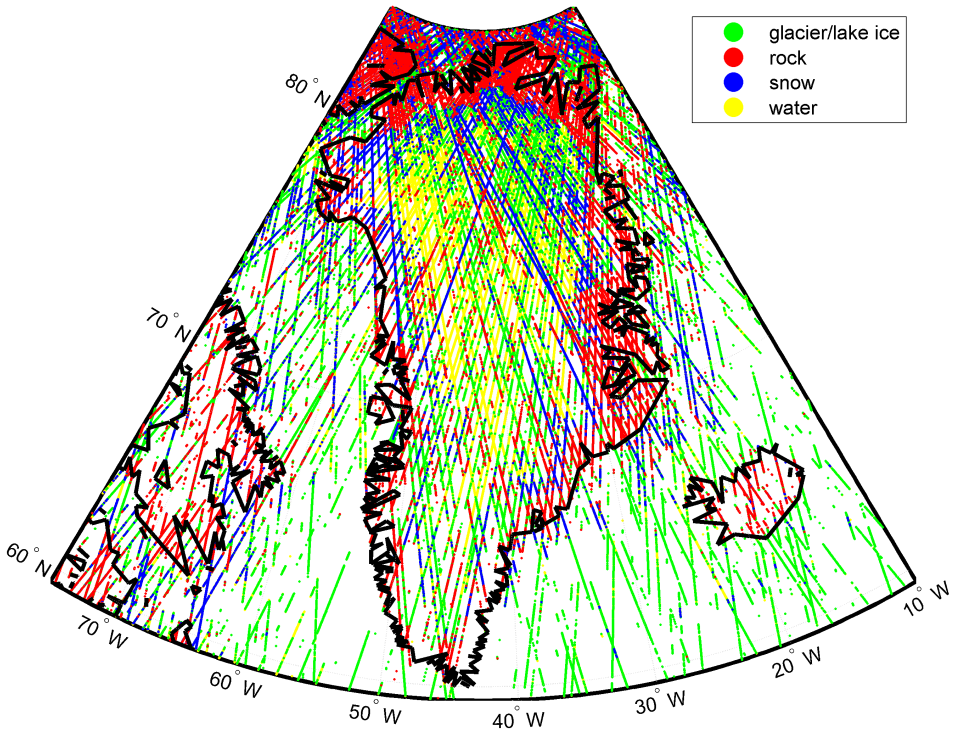

Figure 9.7: The L3C classification results over Greenland. The shots were taken between the 20th of May 2005 and the 23rd of June 2005.

9.3.2 Campaign L3E

The results of L3E show quite different results from the results of L3C. Most dominantly is the higher number of shots that are classified as snow. This is most probably due to the presence of fresh snow that corresponds better to snow in the Dry Valleys. In addition, the map shows few shots that are classified as rock, especially at the coast of Greenland. Furthermore, most shots covering the sea west of Greenland have been classified as snow. This is most likely due to a snow layer on top of sea ice.

The less successful parts of the classification map are the inner part of Greenland and the sea south-west of Greenland. The inner part is mainly covered by shots that have been classified as snow. However, quite some shots have been classified as rock or water as well. As for rock, the reason for misclassification can be found with the criteria that use waveform width, irrespective of saturation. Apparently, some areas of the inner part of Greenland have high slopes or roughness. The same type of misclassifications are present in the classification map of L3C as well. The water waveforms are most likely due to the exactly opposite reason: extreme flat parts resulting in high kurtosis values. This holds true for non-saturated shots. For saturated shots the reason lies with flat areas as well, but in combination with high reflectivity. Saturated shots having a high reflectivity value and low width are classified by the algorithm as water. This is one of the major pitfalls of the algorithm: for Lake Victoria these criteria worked most optimally, however for other areas these criteria might work less optimally or even adversely. As a result, in the recommendations, see Section 10.2, it is advised to take a closer look how to deal with saturated shots, in particular for classifying snow and water.

The second area that shows typical misclassification results of the L3E map is the sea south east of Greenland. This area shows a mixture of shots classified as ice and shots classified as water. Similarly to the L3C discussion, this could be related to the surface roughness at the time of measurements causing the kurtosis to decrease. Whether there was sea ice at the time of measurements has not been validated.

Both resulting classification maps show the potential of using ICESat data for land cover type classification. Especially the information of areas that are classified as rock can be used for better understanding of processes in the Arctic regions. Although the classification can provide classification maps with non-uniform coverage and having non-perfect accuracy, it can give estimates of which areas are covered by rock that could be extracted from mass balance estimates. Especially since current techniques apply smoothing and interpolation for the estimation of mass balance, classification maps could be used to identify which areas to exclude from the estimation. In particular further away from the poles, the density of shots decreases and fewer cross-overs are available resulting in increased interpolation errors.

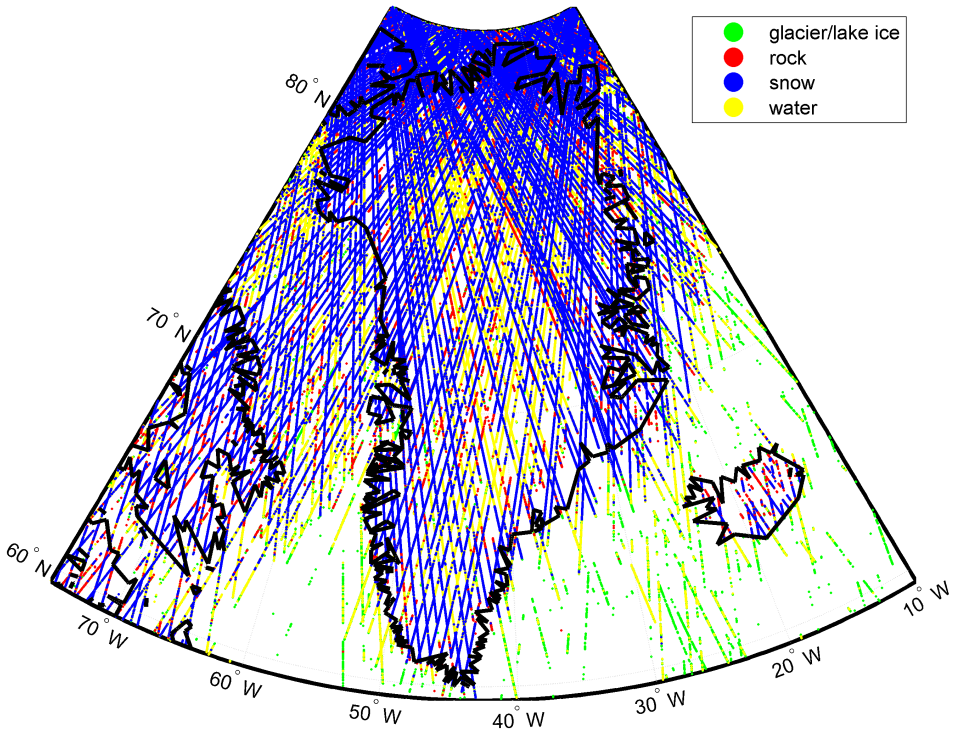

Figure 9.8: The L3E classification results over Greenland. The shots were taken between the 22nd of February 2006 and the 28th of March 2006.

Chapter 10

Conclusions and recommendations

Based on the findings of previous chapters, conclusions and recommendations can be given. Section 10.1 gives the conclusions; Section 10.2 gives the recommendations.

10.1 Conclusions

The most important conclusion is that classification based on ICESat's full waveform signals and attributes essentially works. Based on comparison of the algorithm with validation data, a classification accuracy of 74% was found. The algorithm has been applied to other test areas as well, showing seasonal changes in land cover types.

The final classification algorithm requires only four attributes and is based on seven simple classification criteria. After the criteria, the shots are divided into the land cover types using eight decisions. Distinction is made between non-saturated and saturated shots to deal with the effects of ICESat's measurement system on recording saturated waveforms. To summarize the classification algorithm, the next list gives the most important criteria:

- The reflectivity is the most important attribute to identify snow for both non-saturated and saturated shots.
- The kurtosis is the most important attribute to identify water for non-saturated shots.
- The width is the most important attribute to identify rock and ice for non-saturated shots, and to distinguish water from snow, and ice from rock, for saturated shots.

The classification algorithm has been applied to several test areas. First of all, the classification results over the validation areas (Dry Valleys and Lake Victoria) were used to assess the accuracy of the classification. To make sure the accuracy of the classification algorithm was not biased by the availability of shots over different land cover types, random subsets were taken. From comparing the results of the classification algorithm with the

ground truth, a final accuracy was computed to be 74%. This number can be interpreted as the percentage of shots that are classified correctly with respect to a random classifier.

In addition, the classification has been tested on the Jakobshavn area in Greenland and Greenland itself. Both areas were not used for quantitative accuracy assessment, but rather for analyzing visually the success of classification. Over Jakobshavn, the most dominant success relates to the clear transitions in classification results over changing land forms. For example, the algorithm is capable in most cases to distinguish water and ice from rock and snow. This is especially striking for inlet bays and fjords. Furthermore, the classification between rock and snow give reasonable results with respect to the land cover types indicated by the LANDSAT images. However, accurate validation between the land cover types remains problematic due to the temporal difference between ICESat's measurements and LANDSAT's images.

For the final test case of the classification algorithm, the entire Greenland area was used. Two campaigns have been selected, with the first flown during spring 2005 and the second during winter 2006. The two resulting classification maps show clear differences in land cover types. The spring campaign shows a high number of shots classified as rock at the coastal regions of Greenland. In addition, virtually all shots over Iceland and the east coast of Canada have been classified as rock as well. The classification map of the winter campaign shows totally different results. A large portion of the sea west of Greenland have been classified as snow, most probably due to a snow layer on top of the sea ice. In addition, almost all shots that cover the mainland of Greenland have been classified as snow as well. The shots over Iceland show a mixture of snow and rock, which is plausible for that time of the year.

Besides the (presumed) success of the classification over the Greenland area, several (presumed) deficiencies can be identified as well. First of all, both classification maps show shots covering the inner part of Greenland that have been classified as water. These results most certainly do not represent reality. Furthermore, both maps show a large portion of shots covering water that have been classified as ice, especially for the spring campaign. From these two cases of misclassification it was concluded that the algorithm is not working optimally primarily due to the confusion between ice and water. A plausible cause for this can be found in the interaction of surface topography with waveform signals. Since ice and water are mainly identified in the classification algorithm based on kurtosis (i.e. peakedness) of the waveform, increased water roughness could cause the kurtosis to decrease and consequently the algorithm leads to ice. This holds for non-saturated waveforms. For saturated waveforms, ice is distinguished from water by reflectivity. Since reflectivity is biased by saturation, the algorithm performs less confidently for these shots. This could result in shots covering ice (or snow) to be classified as water.

A general note about the classification concerns the uncertainty of land cover type just below the surface. Even if the classification algorithm would perform perfectly, i.e. the accuracy is 100%, it would still give information about the top layer of the Earth's surface

only. This means, for example, that if an area that is usually covered by rock has a thin layer of snow on top, the algorithm would most probably classify the surface as snow. This problem can be solved by comparing overlapping footprints. In the recommendations, this is explained in more detail.

In addition to the construction of the classification, the research study came across several interesting aspects of ICESat as well. First of all, the data filtering and correction techniques include some unconventional approaches to preparing the data. For data filtering, the most important difference with conventional research studies is the inclusion of shots irrespective of gain value. The exclusion of shots based on gain values is widely accepted in the ICESat research community for filtering out cloud affected shots [55, 21, 46]. However, in particular for campaigns suffering from lower transmit energy, the gain can become high for certain land forms as well. In particular, shots covering sloped rock surfaces have high gain values and regularly even reach the maximum gain value of 250 counts. For campaigns with low transmit energy, the number of shots having a high gain value dramatically increases, irrespective of cloud presence. One advantage of low transmit energy, and consequently higher gain values, is that saturation occurs less often and as a result makes the data products of ICESat more reliable. To filter out cloud affected shots, a filtering criterion based on reflectivity values, instead of gain values, was proposed. All shots with a reflectivity value lower than 0.1 were filtered out. Since this criterion has not been properly verified, it should be interpreted with caution and is therefore included in the recommendations.

For correcting the data, two techniques in particular were included in the research. The first technique corrects the geolocation of the footprint of the shots for atmospheric refraction. Although the correction is not significant for small beam co-elevation angles, it becomes significant for higher angles. The technique relocates the footprint by more than $14m$ for a 5° angle. This technique has been adapted from one of the GLAS Algorithm Theoretical Basis Documents. The second technique corrects ICESat's waveforms for beam co-elevation angle. Especially for classification purposes this is required, since the waveform width is included in the classification criteria. The beam co-elevation angle has a direct effect on the waveform's shape by making it wider and this could negatively affect the classification algorithm. With the beam co-elevation correction, the waveform is compressed according to the angle-induced extra width. This technique has been developed especially for this research study.

For this research study, several ICESat signal related attributes have been analyzed. One of these attributes is the green laser signal, with which atmospheric measurements are carried out. As was discussed in Section 4.5, several land cover types show different reflectance in the infrared and in the green. The potential was demonstrated that the combination of the infrared signals with the green signals at ground level could result in a spectral signature per land cover type. Consequently, this could make the classification algorithm more accurate. However, based on analysis of the atmospheric channels, it was shown that the green signals at ground level do not provide any improvement to the classi-

fication algorithm. The main reasons for this are the extremely low laser transmit energy in the green for later campaigns, and the invalid values of the green backscatter signals for the first two campaigns.

Finally, the answer on the research question can be stated.

Can ICESat's full waveform signals be used to discriminate between land cover types over the cryosphere?

Up to an accuracy of 74%, yes.

10.2 Recommendations

In this section a list of recommendations will be given. These could be used for further research studies, but should most definitely be taken into account when using the results of this research study.

- The first and most important recommendation is the actual implementation of the results of this research on current mass balance estimation techniques and to analyze their improvements. Furthermore, analyzing the change in land cover type might give interesting results of the variability of polar regions over time. Even though the classification is not accurate enough for detailed ice sheet monitoring such as glacier retreat and advance over time, general impressions of changing snow coverage over time should be possible to obtain.
- The classification algorithm makes distinction between saturated and non-saturated shots. For the non-saturated shots, reflectivity and waveform width are used. Since the reflectivity has not been corrected for saturation, this part of the classification is the least reliable. By correcting the reflectivity and reviewing the corresponding classification criteria, the classification accuracy might increase and less confusion is made between snow and water.
- The varying transmit energy has its implications on ICESat's attributes, especially on gain and saturation energy correction. In addition, the computation of full waveform attributes like waveform width becomes less reliable with lower transmit energy because of the lower signal-to-noise ratio (SNR). Analyzing the effects of SNR on the waveform derived attributes, might result in measures to correct for negative effects or even make classification criteria dependent on transmit energy.
- Related to the previous recommendation: even though the reflectivity is based on the ratio of the received energy to (amongst other influences) the transmit energy, it was observed that the reflectivity changes inversely with transmit energy over stable areas. Although it has not been included in the report, during analysis of ICESat's attributes, several plots indicated that lower transmit energy results in

higher reflectivity values for stable areas (like Dry Valleys snow and Lake Victoria water). Perhaps the SNR plays a role in the computation of the reflectivity. By analyzing these effects, the classification criteria might be reviewed and changed.

- The shots for this research have been selected irrespective of gain values. Since the gain is commonly used for filtering out cloud affected signals [55, 21, 46], another filtering criterion was proposed in this research. Based on case study analysis, all shots were removed with a reflectivity value of 0.1 or lower. However, validation of this cloud filtering criterion has not been carried out. Analyzing the influence of clouds on ICESat's signals might improve the proposed criterion and consequently might result in a better classification algorithm.

- In addition to the atmospheric effects on ICESat's signals, solar radiation should be analyzed as well. It might affect the SNR and consequently the reflectivity values. Even though the received energy has been corrected for background noise, some quick analysis (not included in the report) indicated differences in reflectivity over the same area for night and for day conditions. Analyzing the interference of solar radiation on ICESat's signals might change the interpretation of these signals and consequently might change the classification criteria as well. In addition, solar glint of certain land cover types might influence the classification as well.

- For this research study, a fast working algorithm to detect overlapping footprints was made. However, this algorithm has not been used for constructing the classification algorithm, because no valuable information was found. Nevertheless, this algorithm still provides an effective way of detecting overlapping footprints. By identifying overlapping footprints and comparing the land cover types, a classification map can be made of the cryosphere. This map can help to identify areas with covering surfaces but are temporally covered by snow; it could very well be used for improving the classification of land cover types over time. For a detailed explanation of the overlapping footprints algorithm and the programming code, please refer to the author or to the supervisors.

- For the surface topography analysis, it was assumed that the waveform width is dominantly influenced by surface topography and beam co-elevation angle. To deconvolve the effects of topography and beam co-elevation angle from the width, the beam co-elevation angle correction was applied on the waveforms. This technique assumes a circular footprint to determine the amount of waveform compression, because the orientation of the footprint with respect to the illuminated surface slope was considered to be unknown. However, a global DEM that roughly indicates the slope direction of the beam illuminated surface could improve the beam co-elevation correction technique. Analyzing the improvement of a global DEM to the beam co-elevation technique might result in more accurate compression of waveforms. On the other hand, using a global DEM should be applied with caution especially for the polar regions because of changing land cover types. Furthermore, research could be

carried out on the actual contribution of scatters outside the footprint dimensions that are indicated by the laser profiling array (LPA). Since these dimensions are based on an intensity threshold (see Section 4.3), the laser footprint actually covers a larger region and this might influence the full waveform signals.

- In several cases the data descriptions on the NSIDC website [35] were incorrect. First of all, according to the source, the gain has a maximum value of 200 counts. Actually, the maximum value is 250 counts. This change has also not been incorporated in the descriptions of past data releases. Secondly, NSIDC proposes a technique to correct reflectivity for saturation and mentions some limitations to the correction as well. One of the limitations mentions that about 20% of the non-zero saturation energy values are invalid. However, this research found out that for one of the case studies the percentage was even 90% and over the entire Dry Valleys was 67%. Further analysis to the causes of the high percentage of invalid values might result in improved computation techniques for saturation energy correction and eventually in the successful correction of all reflectivity values of saturated shots.

- The final recommendation that relates to the data products provided by NSIDC relates to the signals of the atmospheric channels. According to the data descriptions of the GLA09 data product [35], the two following attributes should be available but in fact hold nothing but invalid values: *i_FRg_grd_sig*, 'full-resolution ground return signal at $532nm$', and *i_FRir_grd_sig*, 'full-resolution ground return signal at $1064nm$'. After correspondence with the Goddard Space Flight Center (GSFC) [41], it was concluded that most probably some kind of error in the computation algorithms causes the attributes to give nothing but invalid values. Correcting for this error will make it easier for the data user to access the ground return values of the atmospheric channels, instead of deriving the return signals with extra programming. In addition, the green backscatter signal values embedded in GLA07 (i40_g_bcsc) are invalid for the L1A and L2A (at least over the Dry Valleys).

- As was mentioned before in this section and in Section 4.5, attributes derived from green laser signals have the potential to improve land cover type classification significantly. Combining these attributes with the attributes of the infrared laser could result in a spectral signature per land cover type. This might be especially relevant for improved discrimination between ice and rock. Even though for ICESat this potential was not realized due to low transmit energy and invalid values, perhaps this recommendation can be taken into account for designing the ICESat follow-up mission: ICESat II.

- Other recommendations for the ICESat follow-up mission relate to the improved functioning with respect to ICESat. First of all, recording of returned waveforms with higher temporal resolution could improve the identification of effects along the laser's path line-of-sight. In addition, research should be carried out on implementing a two dimensional recording array of the laser pulse. Consequently, the waveform (as the time series of two dimensional images) can be analyzed for the spatial contribution

of scatters along the laser's path line-of-sight. As such, the waveform does not loose a dimension, as was explained in Section 4.3. The array could be similar to the LPA, however having higher spatial resolution and more sensitivity. Furthermore, the interaction of the receiver gain with the waveform recording should be improved by making the response time of the gain faster, such that saturation of waveforms can be prevented. In addition, the dynamic range of the gain should be increased, such that shots with extremely low reflectivity and extremely high values are recorded properly. Furthermore, the transmit pulse should be generated more constantly, not only in terms of energy but in terms of shape as well. Even though this research assumes that the variability of the transmit pulse shape does not affect the received pulse shape, even minor changes might have impact on the received pulse shape. The plot of the transmit waveforms (Figure 5.5) indicated a small bump after the main peak as well. Perhaps the waveforms were recorded incorrectly due to ringing effect, but otherwise the shape transmit pulse shape should be made more constant to minimize the effect on the received waveform. Finally, more attention should be paid to clearly describe the relevance of quality flags and their validity, especially in case the satellite might not perform as anticipated.

- The final recommendation that relates to ICESat II concerns the incorporation of a camera system on the satellite. Based on the experiences of this research study, the validation of weather conditions, land cover types and solar effects can be troublesome. Especially the temporal differences between the laser measurements and validation data could be solved by simultaneously acquiring images. In addition, the characteristics of the atmosphere and land cover types (for example solar glint) from the instrument's point of view can be analyzed more accurately with a camera system on-board.

Bibliography

[1] A. Bowman and A. Azzalini. *Applied Smooting Techniques for Data Analysis*. Oxford Science Publications, 1997.

[2] A. L. Neuenschwander. Evaluation of waveform deconvolution and decomposition retrieval algorithms for ICESat/GLAS data. Unpublished research note, Canadian Journal of Remote Sensing, Manuscript ID: CJRS-08-0028.R2, 2008.

[3] A.C. Brenner, H.J. Zwally, C.R. Bentley, B.M. Csathó, M.A. Hofton, J. Minster, L. Roberts, J.L. Saba, R.H. Thomas, and D. Yi. *Derivation of Range and Range Distributions From Laser Pulse Waveform Analysis for Surface Elevations, Roughness, Slope, and Vegetation Heights. Version 4.1*, GLAS Algorithm Theoretical Basis Document edition, September 2003. Available at http://www.csr.utexas.edu (Last visited June 2009).

[4] B. Jutzi and U. Stilla. Range determination with waveform recording laser systems using a Wiener Filter. *ISPRS Journal of Photogrammetry & Remote Sensing*, 61:97–107, September 2006.

[5] B.E. Schutz, H.J. Zwally, J. Bufton, C. R. Bently, T. Herring, J. Minster, J. Spinhirne, and R. H. Thomas. Geoscience Laser Altimeter System GLAS Science Requirements. Version 2.01. Technical report, GLAS Science Team, October 1997. Available at http://www.csr.utexas.edu (Last visited June 2009).

[6] British Antarctic Survey (BAS). BAS Natural Environment Research Council. URL, 2007. Last visited June 2009, http://www.antarctica.ac.uk.

[7] C.C. Tscherning, R. H. Rapp, and C. Goad. A Comparison of Methods for Computing Gravimetric Quantities from High Degree Spherical Harmonic Expansions. *Manuscripta Geodaetica*, 8:249–272, 1983. Available at http://www.gfy.ku.dk/~cct/publ_cct/cct174.pdf (Last visited June 2009).

[8] CSR. Center for Space Research ICESat/GLAS website. URL, 2008. Last visited June 2009, http://www.csr.utexas.edu/glas.

[9] D.J. Saul, J.M. Aislabie, C.E. Brown, L. Harris, and J.M. Foght. Hydrocarbon contamination changes the bacterial diversity of soil from around Scott Base, Antarctica. *FEMS Microbiology Ecology*, 53:141–155, 2005. Available at http:

//www3.interscience.wiley.com/cgi-bin/fulltext/118651457/PDFSTART (Last visited June 2009).

[10] G. Sun, K.J. Ranson, D.S. Kimes, J.B. Blair, and K. Kovacs. Forest vertical structure from GLAS: An evaluation using LVIS and SRTM data. *Remote Sensing of Environment*, 112 (2008):107–117, September 2008.

[11] H. Basagic, Research Assistant, Department of Geology, Portland State University. Personal communation, March 2009.

[12] H. Duong, N. Pfeifer, and R. Lindenbergh. Full waveform analysis: ICESat laser data for land cover classification. In *ISPRS Commission VII Mid-term Symposium "Remote Sensing: From Pixels to Processes"*, May 2006. Enschede, The Netherlands.

[13] H.J. Zwally and C. Shuman. ICESat Brochure. Technical report, NASA/GFSC, 2002. Available at http://icesat.gsfc.nasa.gov (Last visited June 2009).

[14] H.J. Zwally, B. Schutz, W. Abdalati, J. Abshire, C. Bentley, A. Brenner, J. Bufton, J. Dezio, D. Hancock, D. Harding, T. Herring, B. Minster, K. Quinn, S. Palm, J. Spinhirne, and R. Thomas. ICESats laser measurements of polar ice, atmosphere, ocean, and land. *Journal of Geodynamics*, 34:405–445, 2002.

[15] International Lake Environment Committee (ILEC). World Lakes Database. URL, 2007. Last visited June 2009, http://www.ilec.or.jp.

[16] J.G. Bockheim and M. McLeod. Soil distribution in the McMurdo Dry Valleys, Antarctica. *Geoderma*, 144:1914–1924, 2008. Available at http://linkinghub.elsevier.com/retrieve/pii/S0016706107002960 (Last visited June 2009).

[17] Long Term Ecological Research (LTER). McMurdo Dry Valleys LTER. URL, 2009. Data sets were provided by the MCMurdo LTER database, partnership between the Department of Geology, Portland State University, and Byrd Polar Research Center. Significant funding for collection of these data was provided by the National Science Foundation Long-Term Ecological Research program (NSF Grant number ANT-0423595). Last visited June 2009, http://www.mcmlter.org.

[18] M. Aubury and W. Luk. Binomial Filters. *Journal of VLSI Signal Processing*, i:1–8, 1995. Available at http://www.doc.ic.ac.uk/~wl/papers/bf95.pdf (Last visited June 2009).

[19] M. Kennedy. *Understanding Map Projections*, 2000. Available at https://dspace.ist.utl.pt/bitstream/2295/52750/1/Kennedy_2000_UnderstandingMapProjections.pdf (Last visited June 2009).

[20] NASA Science Missions. URL, 2008. Last visited June 2009, http://nasascience.nasa.gov/missions/icesat.

[21] N. Kurtz, T. Markus, D. Cavalieri, W. Krabill, J. Sonntag, and J. Miller. Comparison of ICESat Data With Airborne Laser Altimeter Measurements Over Arctic Sea Ice. *IEEE Transactions on Geoscience and Remote Sensing*, 46(7):1914–1924, 2008. Available at http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=04539634 (Last visited June 2009).

[22] NASA. *SRTM30 Documentation*. Available at http://icesat.gsfc.nasa.gov/tools/SRTM30_Documentation.html (Last visited June 2009).

[23] NASA. ICESat homepage. URL, 2008. Last visited June 2009, http://icesat.gsfc.nasa.gov.

[24] NASA. The Geoscience Laser Altimeter System (GLAS) Homepage. URL, 2008. Last visited June 2009, http://glas.gsfc.nasa.gov/gallery.html.

[25] NASA. NASA's Earth Observing System. URL, 2009. Last visited June 2009, http://eospso.gsfc.nasa.gov/eos_homepage/mission_profiles/index.php.

[26] NASA Earth Observatory. Perspectives from on High. URL, 2008. Last visited June 2009, http://earthobservatory.nasa.gov/Study/TimeShelf/time_shelf3.html.

[27] P. Besl, J. Birch and L. Watson. Robust Window Operators. *Machine Vision and Applications*, 2:179–191, 1989. Available at http://www.springerlink.com/content/n240238644267225/ (Last visited June 2009).

[28] P. Ditmar. Data Analysis and Mapping. Lecture notes on Fourier Analysis, AE3-E03, Faculty of Aerospace Engineering, Delft University of Technology, The Netherlands, October 2008.

[29] P.J.G. Teunissen, D.G. Simons, and C.C.J.M. Tiberius. Probability and Observation Theory. Lecture notes on Probability and Observation Theory, AE2-E01, Faculty of Aerospace Engineering, Delft University of Technology, The Netherlands, December 2007.

[30] R. Kichak. Independent GLAS Anomaly Review Board Executive Summary. Technical report, November 2003. Available at http://icesat.gsfc.nasa.gov/docs/IGARB.pdf (Last visited June 2009).

[31] S. Bae and T.J. Urban. Summary of Laser Profile Array (LPA) Parameter Estimation, Version 1.0. University of Texas, Center for Space Research, online document, October 2006.

[32] S. M. Riad. The deconvolution problem: An overview. *Proceedings of the IEEE*, 74(1):82–85, January 1986.

[33] S. Palm, W. Hart, D. Hlavka, E. Welton, A. Mahesh, and J. Spinhirne. *GLAS Atmospheric Data Products. Version 4.2*, GLAS Algorithm Theoretical Basis Document edition, October 2002. Available at http://www.csr.utexas.edu (Last visited June 2009).

[34] National Snow and Ice Data Center. *GLAS Altimetry Product Usage Guidance*, 2008. Available at http://nsidc.org/data/docs/daac/glas_altimetry/usage.html (Last visited June 2009).

[35] National Snow and Ice Data Center. GLAS/ICESat L1 and L2 Global Altimetry Data. URL, 2008. Available at http://nsidc.org/data/docs/daac/glas_icesat_l1_l2_global_altimetry.gd.html (Last visited June 2009).

[36] National Snow and Ice Data Center. GLAS/ICESat L1 and L2 Global Atmospheric Data. URL, 2008. Available at http://nsidc.org/data/docs/daac/glas_icesat_l1_l2_global_atmosphere.gd.html (Last visited June 2009).

[37] National Snow and Ice Data Center. ICESat/GLAS Frequently Asked Questions. URL, 2008. Available at http://nsidc.org/data/icesat/faq.html (Last visited June 2009).

[38] National Snow and Ice Data Center. ICESat/GLAS Order Data. URL, 2008. Available at http://nsidc.org/data/icesat/order.html (Last visited June 2009).

[39] National Snow and Ice Data Center. GLAS/ICESat Data Releases. URL, 2009. Available at http://nsidc.org/data/icesat/data_releases.html (Last visited June 2009).

[40] ITT Visual Information Solutions. IDL Data Visualization Solutions. URL, 2009. Available at http://www.ittvis.com/ProductServices/IDL.aspx (Last visited June 2009).

[41] S.P. Palm, Research Meteorologist, Goddard Space Flight Center. Personal communation, March 2009.

[42] T. Schenk, B. Csathó, Y. Ahn, T. Yoon, S. W. Shin, and K.I. Huh. DEM Generation from the Antarctic LIDAR data, SITE REPORT. Technical report, No specific institution, September 2004. Available at http://usarc.usgs.gov/lidar/lidar_pdfs/Site_reports_v5.pdf (Last visited June 2009).

[43] T.A. Herring and K. Quinn. *Atmospheric Delay Correction to GLAS Laser Altimeter Ranges. Version 2.1*, GLAS Algorithm Theoretical Basis Document edition, March 2001. Available at http://www.csr.utexas.edu (Last visited June 2009).

[44] T.J. Urban and S. Bae, Center for Space Research, The University of Texas at Austin. Personal communation, June 2008.

[45] T.J. Urban, B. Schutz, and A.L. Neuenschwander. A Survey of ICESat Coastal Altimetry Applications: Continental Coast, Open Ocean Island, and Inland River. *Terrestrial, Atmospheric and Oceanic Sciences*, 19(1-2):1–19, April 2008.

[46] T.J. Urban, R. Gutierrez, and B. Schutz. Analysis of ICESat Laser Altimetry Elevations over Ocean Surfaces: Sea State and Cloud Effects. Abstract for 2008 IEEE International Geoscience & Remote Sensing Symposium.

[47] United Nations Environment Programme (UNEP). Greenland (Denmark), ice cap, topography and bathymetry. URL, 2001. Last visited June 2009, http://maps.grida.no/go/graphic/greenland_denmark_ice_cap_topography_and_bathymetry.

[48] US Geological Survey (USGS) - CMG. Coastal and Marine Geology (CMG) InfoBank. URL, 2007. Last visited June 2009, http://walrus.wr.usgs.gov/infobank/.

[49] US Geological Survey (USGS) - EROS. Earth Resources Observation and Science (EROS). URL, 2007. Last visited June 2009, http://eros.usgs.gov.

[50] US Geological Survey (USGS) - LANDSAT. 2003 LANDSAT Updates. Technical report, USGS, 2003. Available at http://landsat.usgs.gov/about_Landsat_Updates.php (Last visited June 2009).

[51] US Geological Survey (USGS) - USARC. USGS Atlas of Antarctic Research - United States Antarctic Resource Center (USARC). URL, 2007. Last visited June 2009, http://usarc.usgs.gov/lidar_dload.shtml.

[52] US Geological Survey (USGS)- LIMA. LANDSAT Image Mosaic of Antarctica (LIMA). URL, 2007. Last visited June 2009, http://lima.usgs.gov.

[53] USGS. LANDSAT: A Global Land Imaging Project. Technical report, US Department of the Interior, US Geological Survey, October 2005. Available at http://landsat.usgs.gov/documents/Landsat_Fact_sheet.pdf (Last visited June 2009).

[54] W.G. Rees. *Physical Principles of Remote Sensing*. Cambridge University Press, second edition, 2004.

[55] Y. Donghui, H. Zwally, and X. Sun. ICESat measurement of Greenland ice sheet surface slope and roughness. *Annals of Glaciology*, 42:83–89, 2005. Available at http://icesat.gsfc.nasa.gov/publications/pubs/Yi_AnnGlac2005_SlopeRoughness.pdf (Last visited June 2009).

Index