

A PROGRESSIVE STUDY FOR THE AUTOMATIC ANALYSIS OF
INTERNAL LAYERS FROM POLAR RADAR IMAGERY

Jerome E. Mitchell

Submitted to the faculty of the University Graduate School
in partial fulfillment of the requirements
for the degree
Doctor of Philosophy
in the School of Informatics, Computing, and Engineering,
Indiana University
May 2018

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Doctoral Committee

Geoffrey C. Fox, Ph.D

David J. Crandall, Ph.D

Minje Kim, Ph.D

Judy Qiu, Ph.D

John D. Paden, Ph.D

Date of Defense: 05/04/2018

Copyright © 2018

Jerome E. Mitchell

To my mothers, Annie and Linda

ACKNOWLEDGEMENTS

I would like to express the deepest appreciation and admiration for my committee chair, Prof. Geoffrey C. Fox, who has the attitude and substance of a genius: he continually and convincingly conveyed a spirit of adventure in regard to research and scholarship. In addition, Prof. David Crandall has provided a lot of insights into computer vision and machine learning approaches. Without their guidance and persistent help, this dissertation would not have been possible.

I would also like to thank my remaining committee members: Prof. Judy Qiu, Prof. Minje Kim, and Dr. John Paden for their insightful comments and encouragement.

A sincere thank you to Dr. Linda Hayden and Dr. Prasad Gogeneni who introduced and provided me an opportunity to study remote sensing with application to polar science.

I thank my past and present Digital Science Center colleagues for their stimulating discussions and friendship.

Jerome E. Mitchell

A Progressive Study for the Automatic Analysis of Internal Layers from Polar Radar Imagery

Projections of global sea level critically depend on reliable mass balance estimates of the polar ice sheets. Recent satellite observations show rapid thinning of ice sheet margins, speed-up of several outlet glaciers in Greenland, the disintegration of ice shelves in West Antarctica, and the speed up of glaciers buttressed by the disintegrated ice shelves. In order to better understand the processes causing the speedup and disintegration of ice shelves, the polar science community has developed radars to observe a number of internal reflections – “layers”, but efforts to map them have been limited; this is because identifying layers is time consuming and difficult to perform with current, manual techniques. Analysis of a few internal layers has produced heightened awareness into ice sheet behavior motivating the need for additional interpretations; therefore, this dissertation provides a progressive approach to automatically detecting multiple, internal layers from polar radar imagery.

The first contribution presents a semi-automated approach, an active contours model, for identifying high intensity edges while imposing constraints, such as smoothness of layer depth and parallelism among layers; however, model tuning and a fixed, global parameter specifying the number of layers limit the performance for automatic analysis. As an effort to advance layer detection techniques, the problem is represented as a probabilistic graphical model with the second contribution utilizing a Markov Chain Monte Carlo (MCMC) approach to improve accuracy and confidence for detected layer boundaries; this work eliminated the need to adapt

parameters for each layer but focused on ice thickness estimates enabling multiple layers from other datasets to be undetected. The dissertation concludes with an automated approach by extending MCMC to span models and estimate parameters simultaneously for the best number of layers at varying resolutions given a radar image. With this technique, it is possible to identify internal layers over distances of hundred kilometers with no manual intervention, and as a result, allow for higher resolution climate models to forecast climate change.

Geoffrey C. Fox, Ph.D

David J. Crandall, Ph.D

Minje Kim, Ph.D

Judy Qiu, Ph.D

John D. Paden, Ph.D

TABLE OF CONTENTS

Acknowledgements	v
Abstract	vi
List of Tables	xi
List of Figures	xii
Chapter 1: Introduction	1
1.1 Background and Motivation	1
1.2 Interpreting Echograms from Polar Radar Imagery	2
1.2.1 The Echogram	3
1.3 Objectives and Novel Contributions	5
1.4 Structure of Dissertation	8
1.5 Data Disclaimer	8
Chapter 2: State of Art in the Development of Automated Techniques for the Analysis of Polar Radar Imagery	9
2.1 Literature Review	9
Chapter 3: A Semi Automated Approach for Estimating Near Surface Internal Layers from Snow Radar Imagery	12
3.1 Introduction	12

3.2	Methodology	13
3.2.1	Edge Detection	14
3.2.2	Curve Point Classification	15
3.2.3	Active Contours	16
3.3	Results	19
3.4	Conclusion	20
Chapter 4: Probabilistic Graphical Models		22
4.0.1	Probabilistic Graphical Models	22
4.1	Introduction	22
4.1.1	Markov Random Fields	23
4.2	Inference	24
4.2.1	Loopy Belief Propagation	25
4.2.2	Markov Chain Monte Carlo	25
Chapter 5: Estimating Bedrock and Surface Layer Boundaries and Confidence Intervals in Ice Sheet Radar Imagery using MCMC		28
5.1	Introduction	28
5.2	Methodology	29
5.2.1	Modeling layer boundaries	30
5.2.2	Statistical inference	33
5.3	Results	34
5.4	Conclusion	36

Chapter 6: A Sampling-based Approach to the Automatic Analysis of Layers from Polar Radar Imagery	38
6.1 Introduction	38
6.2 Methodology	39
6.2.1 Birth and Death	40
6.2.2 Update	40
6.3 Simulated Annealing	41
6.4 Results	42
6.5 Conclusion	46
Chapter 7: Conclusion and Future Work	47
7.1 Summary and Future Directions	47
References	49

LIST OF TABLES

5.1	Evaluation of our method on the test set. Error is measured in terms of absolute column-wise difference compared to ground truth, summarized with average mean deviation and median mean deviation across images, in pixels.	36
-----	--	----

LIST OF FIGURES

1.1	Illustration depicting the history of polar science data acquisition strategies. Information about the ice subsurface was traditionally collected using ice core drilling practices and evolved to radar remote sensing techniques producing a visualization of the internal structure.	2
1.2	An echogram (left) is a visualization of an ice sheet’s scattering properties. The A-scope (right) shows a radar signal, where the magnitude measured in decibels represents layer characteristics as row coordinates while columns depict its depth.	3
1.3	An echogram highlighting reflection intensities ranging from strong (dark) to weak (light). The surface and bedrock interfaces are the strongest reflectors while internal layers between the surface and bedrock layers and can visualized and spatially identified.	4
1.4	Different CReSIS Radar Systems	7
3.1	Original Snow Echogram	14
3.2	A Canny Edge Detection to Determine Ice Surface	15
3.3	Curve Point Classification	16
3.4	Initialization of Template shown by Yellow Curve	18
3.5	Detected Layers (green) and Maximum Curve Points (blue asterisks)	19
3.6	Detected Layers from Original Snow Echogram	19
3.7	Sample results of our approach on three snow radar echograms.	21
4.1	In image restoration, an image is corrupted with noise (middlle) and is restored (right) similarly to its original (left)	24

4.2	Different CReSIS Radar Systems	26
4.3	Figure The general algorithm for Gibbs sampling iteratively draws new values for random variables conditioned both on observed variables and other unknowns.	27
5.1	Original Echogram with Two Layer Boundaries	29
5.2	Illustration depicting a model for two layers, where each row represents a layer through the echogram. The horizontal layer constrain enforce smoothness while the vertical constraint encourage repulsiveness between layers boundaries. The image likelihood are based on local echogram characteristics	30
5.3	Detected Layers of Improved Approach with Confidence Intervals	35
5.4	Sample results of approach on two MCORDS radar echograms.	37
6.1	Illustration Depicting the Birth of a Layer	40
6.2	Illustration Depicting the Death of a Layer	41
6.3	Illustration Depicting the Update of a Layer	41
6.4	Sample results of the automated approach on three 2014 snow radar echograms acquired from Antarctica	43
6.5	Sample results of the automated approach for 2014 accumulation radar echograms acquired from Antarctica	43
6.7	Intermediate accepted results during RJMCMC iterations. In a birth or update move, the red line indicates a layer being added or modified, which allows the current configuration l^* to be different from the previous iteration. In a death move, a red layer is being removed from l^t and the remaining layers constitute l^*	45

CHAPTER 1

INTRODUCTION

This Chapter introduces the dissertation. In particular, it motivates climate change research and a need for analysis tools to better understand the rapidly changing polar ice sheets. Also, radar remote sensing is discussed from the context of data acquisition and detailed subsurface information. Afterwards, it describes challenges related to analyzing remotely sensed data, which provides a foundation for the study and its novel contributions. It concludes with an outline detailing how each chapter provides a progressive approach to the automatic analysis of internal layers from polar radar imagery.

1.1 Background and Motivation

A number of investigations show Greenland and Antarctica rapidly losing mass, but the knowledge to completely understand the processes influencing the polar ice sheets is incomplete. To address the problem, information on ice thickness, internal stratigraphy, subglacial topography, and basal conditions is critically needed as parameters to climate models for better forecasting climate change.

To better understand the polar ice sheets, traditional methods, such as ice core drilling provide detailed information on subsurface features, but it is time consuming and limited in geographical coverage. Other alternatives have employed remote sensing techniques, such as radar, for improved spatial and temporal coverage in order to infer properties of the subsurface. Figure 1.1 shows an illustration detailing past to present data acquisition strategies. With the advent of new collection strategies, the growth in data has caused significant analysis challenges. The task of identifying subsurface features from polar radar imagery is performed by domain experts, requiring significant resources to complete a

single data file consisting of thousands of measurements. Also, the process is subjective, leading to inconsistencies. Given the volume of data acquired in the past and its growth each year, an automated approach to identify subsurface features is necessary to provide results for the scientific community in a timely manner.

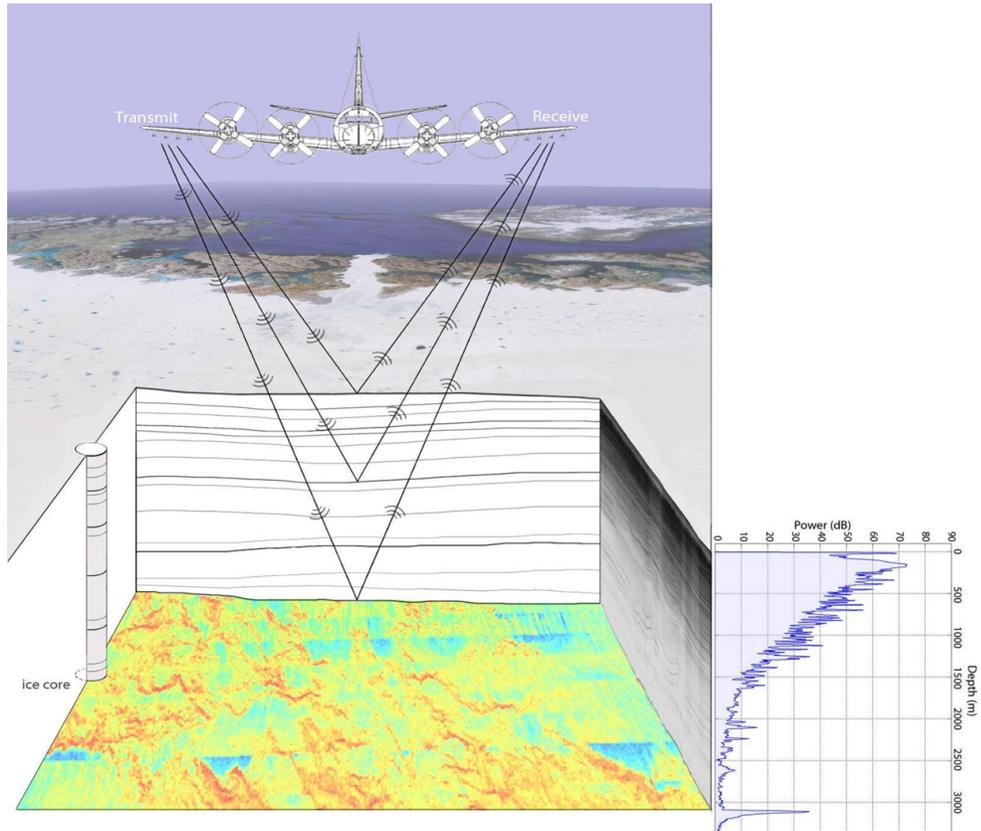


Figure 1.1: Illustration depicting the history of polar science data acquisition strategies. Information about the ice subsurface was traditionally collected using ice core drilling practices and evolved to radar remote sensing techniques producing a visualization of the internal structure.

1.2 Interpreting Echograms from Polar Radar Imagery

Analyzing subsurface features from data is an effective way to study the polar ice sheets. Radar images, or echograms (shown in Figure 1.3), provide an understanding of the ice subsurface, such as ice thickness estimations, a determination of a rough or smooth bedrock, if internal layers are flat or affected by a climate process, which are important

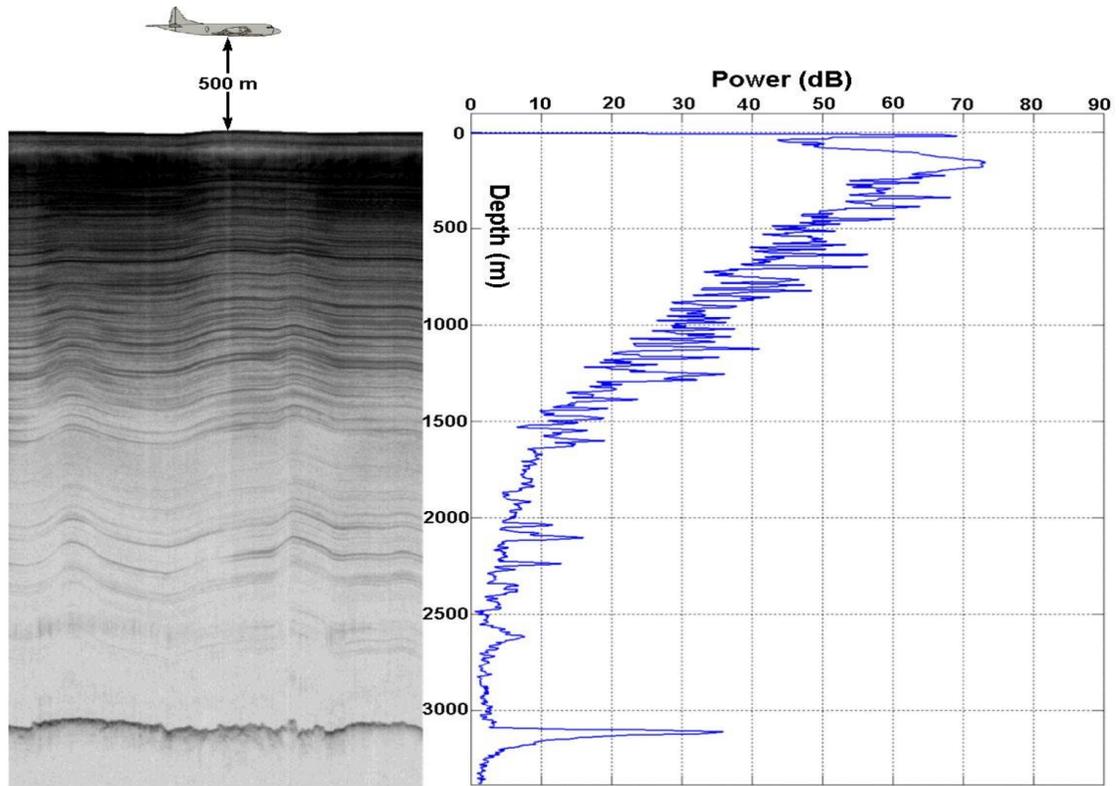


Figure 1.2: An echogram (left) is a visualization of an ice sheet’s scattering properties. The A-scope (right) shows a radar signal, where the magnitude measured in decibels represents layer characteristics as row coordinates while columns depict its depth.

characteristics for developing accurate models. Also, understanding an echogram is important for determining false positives, such as artifacts from either acquisition or processing, and identifying layers for analysis.

1.2.1 The Echogram

When a radar transmits and receives energy from an antenna, changes in dielectric properties of the polar ice sheets cause the signal to reflect and refract. Larger changes in dielectric properties cause strong reflection intensities allowing targets, such as internal layering with a strong surface and bedrock to be visualized in a radar image.

Each measurement, a radar trace (shown in Figure 1.2), consists of a signal representing received energy over time. In an echogram, a trace is a column of pixels, with each pixel representing a depth. A survey segment is a collection of traces representing all columns

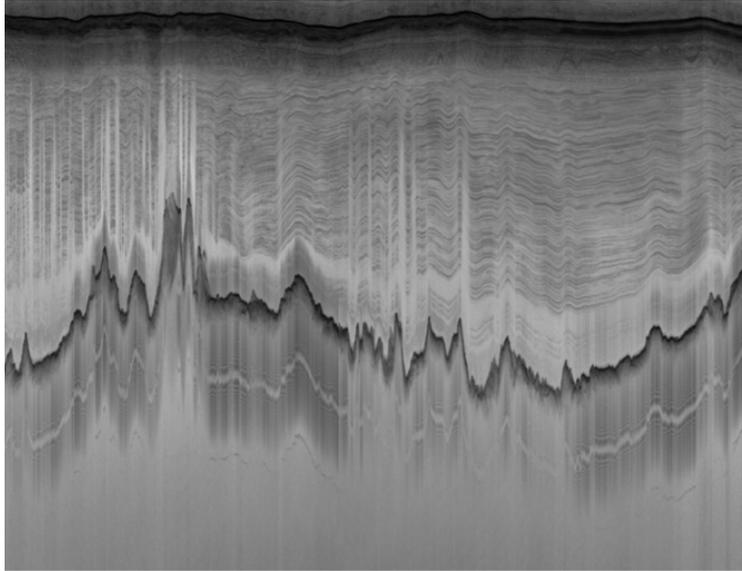


Figure 1.3: An echogram highlighting reflection intensities ranging from strong (dark) to weak (light). The surface and bedrock interfaces are the strongest reflectors while internal layers between the surface and bedrock layers and can visualized and spatially identified.

of a radar image, from the beginning of the segment (left side of the echogram) to the end (right side of the echogram) during data collection. A pixels width represents a uniform geographic distance between traces, depending on the sensing platform's speed.

Bedrock

The importance of detecting bedrock layers from polar radar imagery stems from determining if it is lubricated by the presence of water at the pressure melting point, either formed by channels, as either a liquid film or water filled sediments. Water between the ice sheet and bedrock interfaces decreases the friction, which controls the ice flow velocity.

The differences in reflectivity between a frozen and thawed bedrock is significant and can be used as an indicator for bedrock conditions. Understanding the bedrock reflectivity directly is not possible since the returned echo is attenuated. Therefore, using the bedrock reflectivity to map the presence of water requires knowledge of its local attenuation rate. Also, the presence of water near the bedrock can be determined by observing the roughness, which can be extracted from its shape. The fall-off of the bedrock reflection is significantly

faster for smooth surfaces compared to rougher bedrocks.

Internal

Internal layers is attributed to constraints in density, conductivity, and fabric. Layers of different ice density, such as melt or ash layers, can cause a partial reflection due to the discontinuity in the permittivity of the ice. Density variations are considered to be a dominant source of internal layers in the upper region of the ice sheet. With depth, high pressure air bubbles transitions into crystals, removing the density contrasts.

1.3 Objectives and Novel Contributions

With the availability of large remotely sensed datasets and the expected growth of data, there is a need for novel automatic techniques for the analysis of echograms acquired from the polar ice sheets. There are several advantages for using automatic techniques as opposed to traditional, manual approaches.

- Automatic techniques rely on computational resources for information extraction. By using advances in technology, such as machine learning and computer vision techniques, automatically processing remotely sensed data allows exploitation of large, available data to aid knowledge discovery.
- Automatic techniques are objective, which ensures consistent repeatable experiments; this is important for a coherent and quantitative analysis of the ice subsurface.

Automated analysis techniques for polar radar imagery has been marginally addressed by the scientific community. Most techniques have focused on analyzing subsurface features from echograms acquired by terrestrial datasets. Although the subsurface features from planetary and terrestrial datasets have common characteristics these echograms operate on different platforms, having different instrument

parameters and investigate different types of subsurfaces. For these reasons, the coverage and resolution, both in the vertical and horizontal directions, and the patterns highlighted are different. Therefore, automatic techniques for the analysis of terrestrial datasets cannot be applied to ice sheet subsurfaces. Automated techniques for analyzing ice sheet data encourage the development of techniques tuned to the uniqueness of remotely sensed ice sheet data. Automatic techniques can support a better use of remote sensing data for a better understanding of ice sheet subsurface.

In this dissertation, three novel contributions allow for the progressive development of an analysis tool supporting polar radar imagery. The techniques address challenges related to automatically detecting layer boundaries, which are important to the study of polar ice sheets.

In order to provide a study on the automatic detection of subsurface features, the following works are contributed:

1. A Semi Automated Approach for Estimating Near Surface Internal Layers from Snow Radar Imagery
2. Estimating Bedrock and Surface Layer Boundaries and Confidence Intervals in Ice Sheet Radar Imagery using MCMC
3. A Sampling-based Approach for the Automatic Analysis of Layers from Polar Radar Imagery

The techniques are summarized:

1. The first contribution focuses on a specialized, semi-automatically approach for estimating near surface internal layers from snow radar echograms acquired in Antarctica; the solution requires a user to estimate a global parameter for determining a number of visible layers while an active contour (snakes) model se-

quentially finds those high-intensity edges likely to correspond to layer boundaries while simultaneously imposing constraints on smoothness of layer depth and parallelism among layers

2. The second contribution improves the accuracy and utility of an automated approach designed for a fixed number of layer boundaries; this approach uses Gibbs sampling, which allows the model to solve for layer boundaries simultaneously, yielding automatic layer detection results significantly better than performed by the state of art technique. Moreover, the Gibbs sampler produces confidence intervals by estimating bands of uncertainty for layer boundaries. Since noise and ambiguity in echograms are inevitable, this ability would provide confidence in climate models.
3. The third contribution proposes an automatic approach for inferring a number of layers by using an advanced reversible jumps MCMC sampling approach, which samples candidate hypotheses to evaluate them against image observations to determine the optimal configuration explaining the echogram. The approach is used for two different echogram sources, snow and accumulation radar imageries.

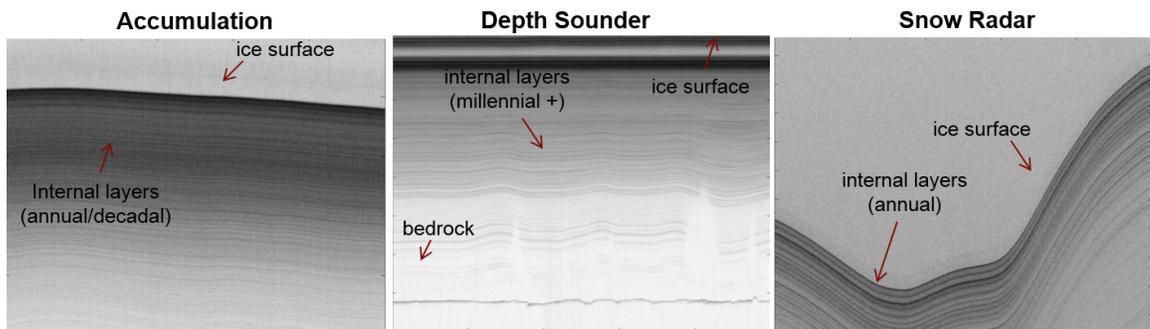


Figure 1.4: Different CReSIS Radar Systems

1.4 Structure of Dissertation

This document is composed of 7 chapters. Chapter 1 motivates the problems of climate change and a need for tools to better understand polar radar imagery. Chapter 2 presents an overview of literature related to the automatic analysis of remotely sensed data. In Chapter 3, a semi-automated approach is discussed, which uses an active contour model along with “off-the-shelf” computer vision techniques to identify near surface internal layers from polar radar imagery. While Chapter 4 introduces probabilistic graphical models, Chapter 5 shows an improved automatic analysis for the identification of bedrock and surface layers from polar radar imagery. Chapter 6 describes an automatic technique for the estimation of internal layers by extending MCMC for the detection of layers at various scales. Each Chapter from 3 to 6, excluding 4 is made up of more sections, each presenting a detailed introduction, methods, results obtained by applying to data acquired from the polar ice sheets, and related conclusions. Finally, the conclusion of the dissertation is discussed along with proposals for future research and development.

1.5 Data Disclaimer

For this study, data, specifically, products acquired from the multichannel coherent radar depth sounder (MCORDS), accumulation and snow radars were collected by the Center for Remote Sensing of Ice Sheets (CReSIS). These radars MCORDS measure ice thickness and image the ice-bed interface while snow and accumulation radars provide internal layers at different resolutions. Examples of data from the selected radars are shown in Figure 1.4). The data is publicly available and provide ground-truth for the MCORDS and little to no ground-truth for snow and accumulation radars.

CHAPTER 2

STATE OF ART IN THE DEVELOPMENT OF AUTOMATED TECHNIQUES FOR THE ANALYSIS OF POLAR RADAR IMAGERY

This Chapter provides an overview to the automatic analysis of remotely sensed data.

2.1 Literature Review

Because reflection amplitudes from the surface and subsurface can be noisy, analyzing echograms is a difficult task; these interferences may either partially or completely weaken the signal, leading to a wrongful interpretation of information. However, despite the challenges developing an automated approach for remotely sensed data, many efforts have explored semi-automated techniques. For example, [1] detected Holocene's depth in Greenland. The technique uses image processing concepts based on histogram analysis and surface fitting to identify the transition region between the Holocene and glacial ice. In [2], the authors use phase information from internal reflections to map internal layer slopes, which allow layer dips, the phase of the reflected signal to change smoothly along the track. The first method is based on the along track gradient while the second method, a Fourier transform, uses complex radar data for the along track. The central peak of this spectrum is the along track gradient of the phase, which is converted to a layer slope, as with the first method. These works focus on understanding the ice stratigraphy, which is needed for either ice flow modeling [3] or isochronous characterization of the ice [4] and have motivated, [5], [6], [7], [8], [9], to develop analysis techniques for remotely sensed data.

Also, most related work has focused on identifying layers from terrestrial datasets. The subsurface of the North Poles on Mars is characterized by layers of dry and wet ice. Therefore, patterns shown in echograms are due to echoes from various ice layers and underlying basement. Such features have been detected by automatic analyses. In [10] and

[11] , a technique detects shallow linear features; the authors use collaborative filtering to reduce noise and highlight returns from the target while linear feature extraction is exploited using Steger's filter. Another work [12], which focuses on the Martian subsurface, and in particular, the detection of basal returns; the authors use several theoretical models to characterize the statistical properties of remote sensing signals. In order to isolate returns of the basement from other echoes, the algorithm exploits a region growing technique, which combines results from the geometrical properties of subsurface features with its statistical analysis. Automatic techniques for terrestrial data analysis acquired on Mars cannot be applied to ice sheet subsurfaces from the polar ice sheets. Nevertheless, such works along with advances in the processing of ice sheet remote sensing data, represent a reliable starting point for the development of novel advanced methods for the investigation of echograms of the ice sheets.

Machine learning allows algorithms to either learn or model behavior for either data or an environment. Learning is determined by automatically extracting information from data and adjusting either the output, adapting behavior or structure over time. By enabling algorithms to learn and automatically update control parameters and behaviors, the need for human intervention can be eliminated.

Machine learning, however, has been limited in applications involving ice sheet and subsurface data. Internal layers have been used to predict the depth and thickness of certain layers. For example, efforts to predict the depth and thickness of the Eemian layer in Greenland utilized a Monte Carlo Inversion of the flow model to estimate unknown parameters constrained by internal layers [13] [14]. Similar work classified the presence of bottom crevasses and estimated their height using radar data of the Ross Ice Shelf in Antarctica [15]. A physical model was developed by studying the radar data and long echo tails, which were characteristic of bottom crevasses. That work studied the data to determine power reflection coefficients, but it did not incorporate either regression or machine learning methods. Efforts to classify subsurface layers using ground penetrating radar are

more abundant. Coupled with seismic surveying, ground penetrating radar allows subsurface properties to guide decisions of where to drill for either oil or natural gas, as well as help explain observed surface changes. Radar reflection patterns can distinguish echo returns in different regions, which provides machine learning methods to associate patterns with distinct classes. In [16], a neural network was utilized to learn and predict patterns of radar reflections in ground penetrating radar data collected from a river in Canada. In this work, higher classification accuracies were obtained from attributes, which incorporated spatial arrangement of reflections, such as spatial covariance, as opposed to amplitude. Ground penetrating radar signal interpretation is one of the more studied applications in subsurface classification, as interpretation of subsurface materials requires either coring or drilling to determine ground truth. Domain experts manually investigated the subsurface materials to segment into distinct rock or object classes. These data were used by machine learning methods to develop a model between GPR return signals and the excavated ground truth. Since ground penetrating radar has been used for many applications, machine learning techniques can aid in automating these tasks by learning a model, interpolating between measurement sites, and characterize the subsurface at other unknown locations. Machine learning algorithms, such as support vector machines have been used in many signal interpretation applications, including roadbed degradation recognition [17]. [18] From the literature review, the use of machine learning to create models for polar radar data analysis is not only feasible but also provides a high level of accuracy while offering a significant increase in efficiency.

CHAPTER 3

A SEMI AUTOMATED APPROACH FOR ESTIMATING NEAR SURFACE INTERNAL LAYERS FROM SNOW RADAR IMAGERY

Near surface layer signatures in polar firn are preserved from the glaciological behaviors of past climate and are important to understanding the rapidly changing polar ice sheets. Identifying near surface internal layers in snow radar imagery can be used to produce high-resolution accumulation maps. This Chapter presents an approach for semi-automatically estimating near surface internal layers and have applied it to snow radar echograms acquired from Antarctica; the solution requires a user to estimate a global parameter for determining a number of visible layers while an active contour (snakes) model finds those high-intensity edges likely to correspond to layer boundaries while simultaneously imposing constraints on smoothness of layer depth and parallelism among layers.

3.1 Introduction

Global sea level rise has been attributed to the melting of glaciers along with other causes. To assess the contribution of the polar ice sheets to sea level rise, there is a need to determine the mass balance. A key variable in assessing the mass balance of an ice sheet using this method is the accumulation rate. Other required measurements are surface elevation, ice thickness, aerial extent, intensity of summer melt, wind-generated surface roughness, ice temperature, and absolute velocities.

Traditionally, accumulation rates are determined by ice cores and pits. These data are sparsely distributed ice cores and pits. A more efficient way to estimate the accumulation rate, which would benefit the polar science community in their modeling of the ice sheets

and mass balance estimation, is to map a continuous profile of layers from the polar ice sheets. These changes can be measured remotely using radar allowing snow and ice properties, such as density, acidity, and crystal orientation fabrics, to influence the visibility of the internal reflection.

With an increase interest to identify internal layers from polar radar imagery, a solution proposes to identify those layers using image processing techniques by applying high-level constraints, such as how the ice-air boundary should be most prominent and how snow layers should be modestly parallel.

3.2 Methodology

Observations about how domain experts detect layer boundaries were used to develop a semi-automated algorithm to mimic these behaviors. As shown in Figure 3.1 and as is typical for our experimental images, the surface reflection is very strong and near surface layer intensity generally decreases as depth increases. Also, near surface layers are approximately parallel, but may have modest changes in slope both to one another and to the ice surface. A technique is proposed, which attempts to find the prominent surface reflection and searches for similar (but invariably weaker) layer structures below the surface. Each layer is used as an estimate of the appearance for the layer below it and an active contours (snakes) model snaps the correct layer structure given this estimate. The process of detecting the surface, estimating layer location using curve point classification and refining the use of snakes are discussed in subsections 3.2.1, 3.2.2, and 3.2.3, respectively, and use Figure 3.1 as a demonstration of the proposed approach.

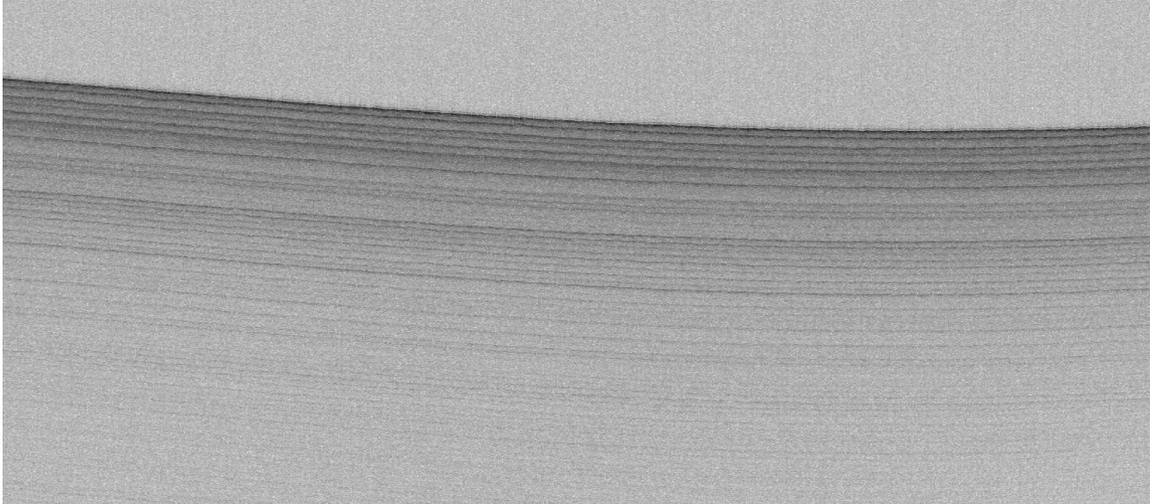


Figure 3.1: Original Snow Echogram

3.2.1 Edge Detection

To find the location of the surface boundary, which is typically the most prominent edge in the echogram, a Canny edge detector [19] was used because of its performance in detecting strong intensity contrasts for the near surface layer dataset (see Figure 3.2). Based on this criteria, the Canny edge detector smoothes the echogram to eliminate noise. It then finds the image gradient to highlight regions with high spatial derivatives. The algorithm then tracks along these regions and suppress any pixel, which is not at the maximum (non-maximum suppression). The gradient array is now further reduced by hysteresis. Hysteresis is used for determining the remaining pixels, which were not suppressed. Canny edge follows two threshold values. If the magnitude is below the first threshold, it is set to zero otherwise it is an edge. And, if the magnitude is between the two threshold, then it is set to zero unless there is a path from this pixel to a pixel with a gradient about the lower threshold value. In detecting this initial ice surface, the following fixed Canny parameters were used: a sigma of 2 for the standard deviation of the Gaussian filter and a low and high thresholds of 0.7 and 1.8, respectively. Since the ice surface is symmetrical to subsequent layers, it provides a good starting template.

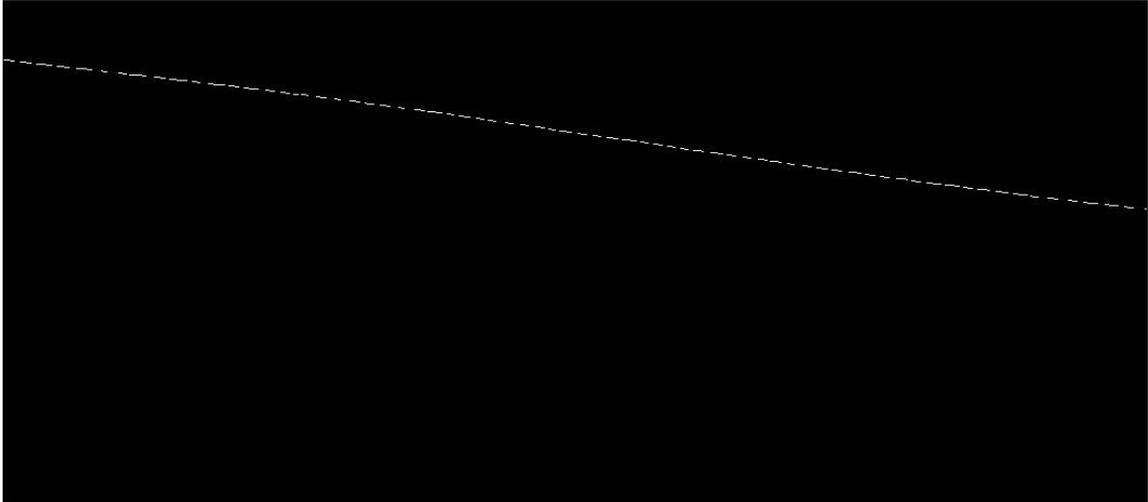


Figure 3.2: A Canny Edge Detection to Determine Ice Surface

3.2.2 Curve Point Classification

While the ice surface can be readily detected by edge detection, using it for near surface internal layers is not possible because of the very weak layer boundaries and noise inherent in echograms. As a consequence, Stegers approach [20] was used to identify points in an echogram, which were likely to be part of curvilinear structures. In short, this approach computes statistics on gradient structure within local image patches, in particular determining areas with prominent gradients in a coherent direction. identified peaks for scores computed by Steger (shown as blue asterisks in Figure 3.3 and used these to suggest initial curve positions for estimating near surface internal layers. For the first layer, the ice surface was used to estimated previously and shifted it down,(in the y direction) so it intersected the first maximum point. This process was repeated until the number of near surface internal layers specified by the user has been found and gave initial estimates of layer positions and shapes, which we refined in the next step.

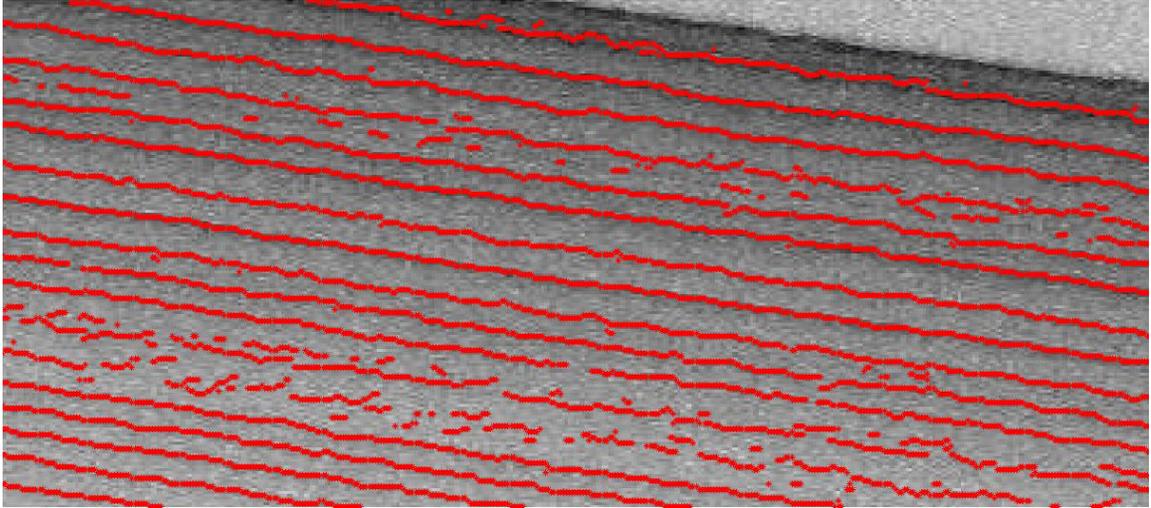


Figure 3.3: Curve Point Classification

3.2.3 Active Contours

Snakes [21], also known as deformable models, gravitate to an object's boundary by minimizing an energy function. A snake conforms to the boundary of an object by adjusting its shape and moving through the spatial domain of the image until it reaches a location, where the energy functional E_{snake} is at a minimum. The energy function is composed of two terms: internal and external. The internal energy is calculated from the shape and location of the snake and serves to preserve its continuity and smoothness. The external energy term, is composed of two terms: the external image energy and an external E_{ext} energy term. The external image energy, E_{Image} , is calculated using image information, typically the negative magnitude of the image gradient and is used to moving the snake towards image features, such as lines, edges. The external constraint E_{Con} , is an optional term, which is responsible for forcing the snake to or from a particular image feature. The following discusses a snake's energy terms, which influence its behavior.

Internal Energy

The internal energy, E_{int} , of the contour is composed of two terms. The first term allows the snake to resist stretching and enabling it to be continuous while the second term causes

the snake to resist bending as a smoothness constraint. The coefficients α and β , known as tension and rigidity parameters, are scaling factors to control influence of the smoothness and continuous terms. α determines the extent to which the snake can stretch at a point s on the snake while β determines how the snake can bend at a point s . As α increases, the “tension” reduces the length of the snake while increasing the value of β increases the bending “rigidity” of the snake, which allows the snake to be smoother and less flexible.

External Energy

As was previously mentioned, E_{ext} the external energy, is calculated from the image, which contains the target boundary; this is accomplished by using the gradient image.

Since there is a good approximation of individual layers using the maximum eigenvalue to denote the strength of the curve, an active contours (snakes) model was used to refine it. Active contours is a cost minimization approach, which encourages the curve to align with high-gradient edge pixels, but to discourage it from having either discontinuities or sharp bends. These two goals are often in stiffness and tension, and the energy minimization function is used to find the curve with the best trade off between them.

In this methodology, the contour adapts until it either reaches equilibrium or it has gone through a number of iterations. A window, which computes the local stiffness and smoothness characteristics during the contour’s adjustment, is maintained at either upward or downward until it fits the near surface layer. The local stiffness costs are computed using the distance between the current pixel location and mean of the current pixel location of the contour in neighboring pixels. The pixel location corresponding to the closest distance to the mean of neighboring pixels receives the lowest cost while the furthest distance receives the highest cost for a particular pixel. Therefore, the contour is encouraged to smoothly fit to the echogram and bridge gaps in the near surface internal layer.

A stiffness weight was experimentally found to be robust to image noise and other artifacts. Figure 3.4 illustrates how the contour actively adapts to the echogram as processing progresses. The contour typically fits the data at various places, and the stiffness cost pulls the remainder of the contour either up or down to a reasonable location.

The pixel location corresponding to the minimum total cost for the contour at each pixel is selected as the contours starting configuration for the next iteration. If the configuration does not change between either iterations or until a specified number of iterations, the contour is determined to have reached equilibrium for a near surface layer. Figure 3.6 shows the automatically selected surface and bottom layers for all traces of the example echogram.

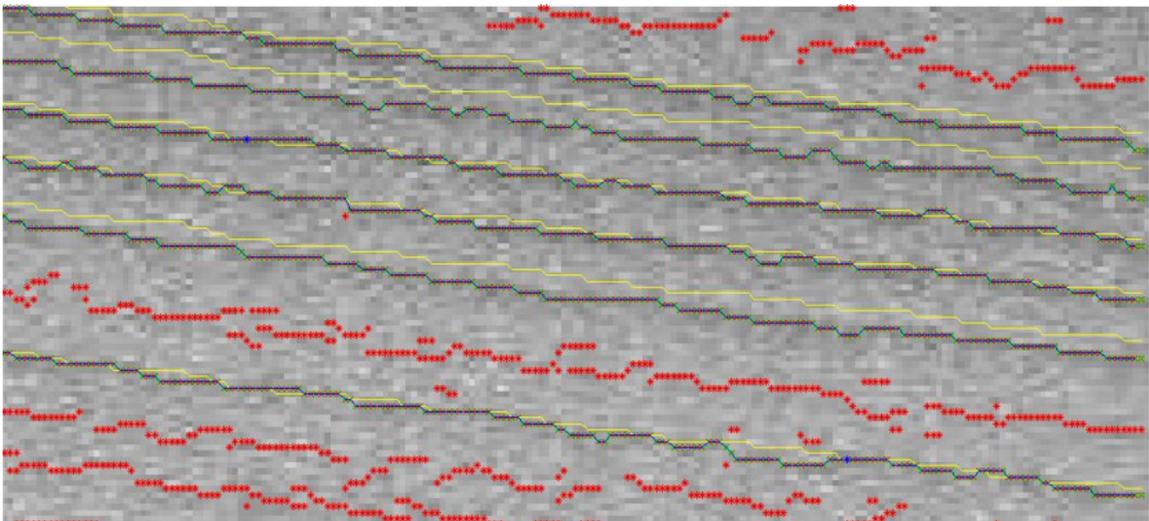


Figure 3.4: Initialization of Template shown by Yellow Curve

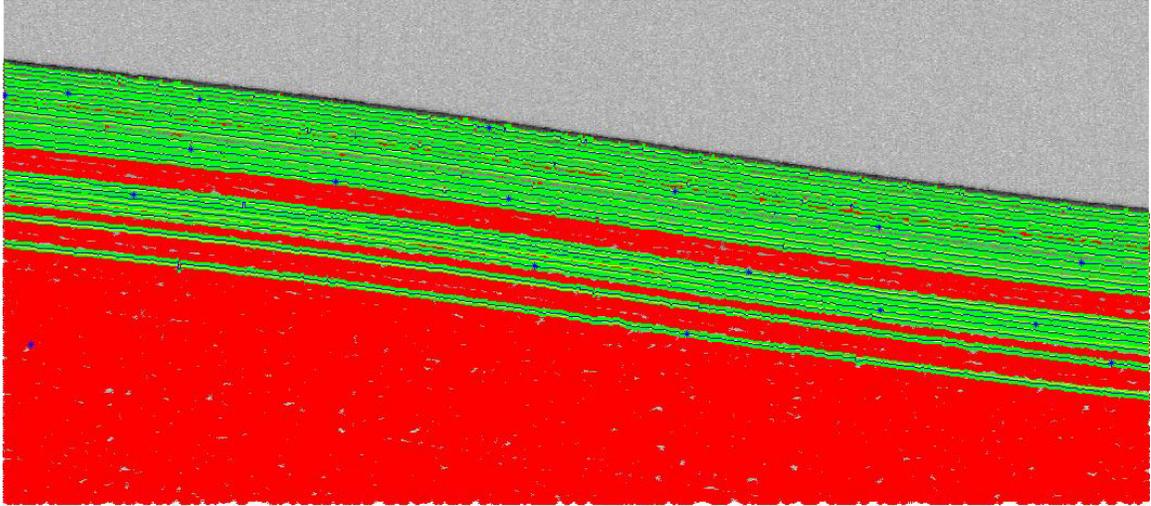


Figure 3.5: Detected Layers (green) and Maximum Curve Points (blue asterisks)

3.3 Results

Figure 3.6 shows the result of our approach for Figure 3.1. We observe it has successfully found over a dozen layers correctly, although it misses some of the very faint layers towards the bottom of the echogram. Figure 5.4 shows results for three additional echograms. While the algorithm works quite well for layers near the surface, it does miss or incorrectly identify some of the deeper layers (such as the discontinuities in Figure 3.3 in which the estimates skip from one layer boundary to another).

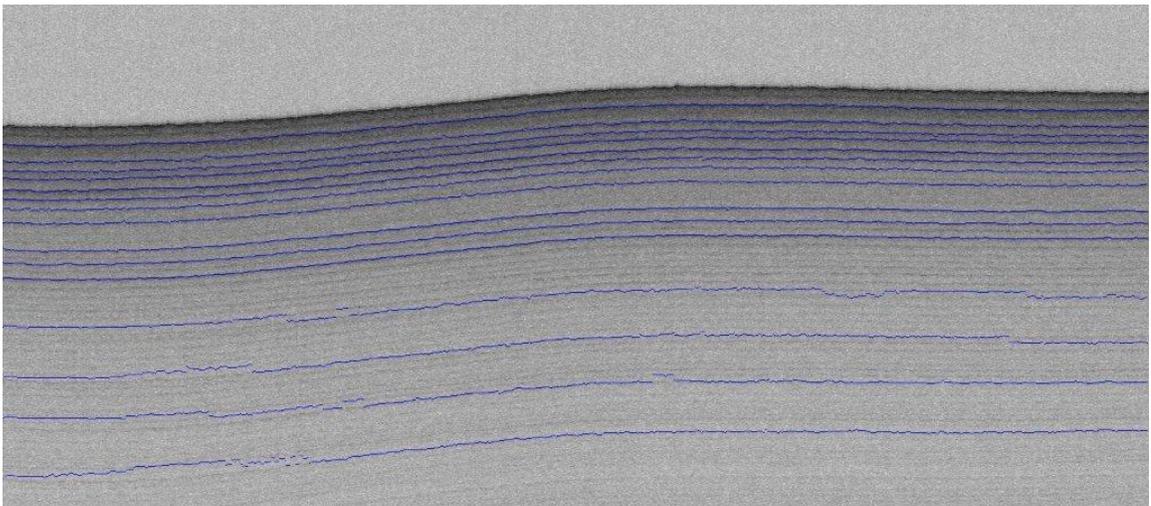


Figure 3.6: Detected Layers from Original Snow Echogram

3.4 Conclusion

A semi-automated approach for estimating near surface internal layers in snow radar imagery was developed. It utilizes an active contour model in addition to edge detection and Stegers curve classification. The technique is a step towards the ultimate goal of unburdening domain experts from the task of dense hand selection. The proposed approach requires variable control parameters per image for a particular dataset and assumes the ice surface is parallel to subsequent internal layers, which is incorrect for most images since the lower portion of the ice sheet typically follows the bed topography. Therefore, innovative methods are needed to refine the solution of detecting internal layers.

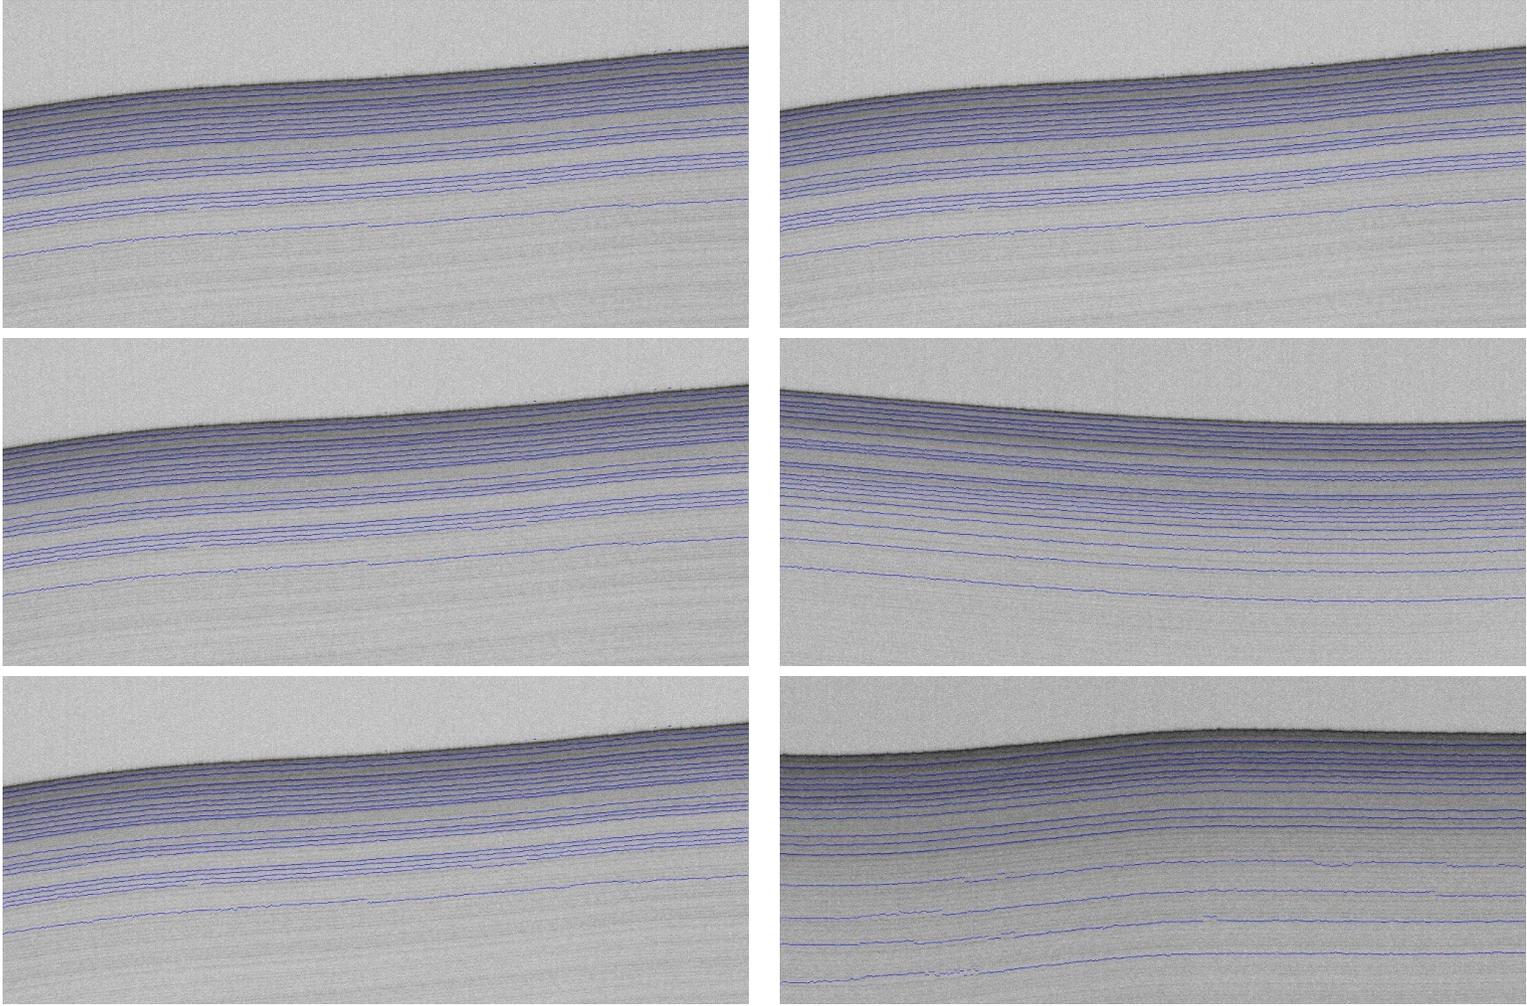


Figure 3.7: Sample results of our approach on three snow radar echograms.

CHAPTER 4

PROBABILISTIC GRAPHICAL MODELS

Echograms have a large degree of uncertainty caused by unpredictable signal scatter and reflection patterns, which cause difficulty to interpret subsurface features. In this chapter, probabilistic graphical models are discussed, specifically Markov Random Fields, as the main model with a focus on inference techniques

4.0.1 Probabilistic Graphical Models

4.1 Introduction

Probabilistic graphical models[22], [23], [24], [25] represent the intersection between graph and probability theories. They provide a tool for dealing with two problems occurring in real world systems: uncertainty and complexity. As a modeling and inference tool, probabilistic graphical models use a graphical structure to represent probability distributions of multiple variables. The nodes in the graph correspond to random variables while the links represent the statistical relationships between nodes. The graph structure encode conditional dependence and independence among random variables. Given a probabilistic graphical model, the joint (or conditional) probability distribution can be factored as a product of a set of functions defined from a subset of random variables; this factorization simplifies the modeling multivariate joint probability distributions. Based on probability theory, algorithms have exploited the structure of graphs in order to efficiently determine either the marginal or conditional probabilities.

There are two main graphical models: directed and undirected. These models represent and encode different conditional independence relationships on random variables. The undirected graphical model are used for specifying spatial constraints between variables

while directed models express casual relationships. Each probabilistic graphical model is represented by a graph $G = (V, E)$, where V denotes the vertices (nodes) in the graph and E , the edges (links) connecting these vertices. For undirected graphs, E only consists of undirected links as opposed to directed links for directed models. The joint probability distribution, which is represented by the model, can be factored as the product of local potentials functions or local conditional probabilities, which simplifies the estimation for performing probabilistic inference.

4.1.1 Markov Random Fields

Undirected graphical models use an undirected graph to model relationships between random variables. Markov random fields (MRF), a type of undirected graph, has a wide range of applications such as image processing (image denoising and restoration), object recognition, image segmentation, and texture synthesis. A MRF models the joint probability distribution of labels X and observations Y based on interactions among a set of random variables within a selected neighborhood. In an undirected graph, a node X_j is a neighbor of node X_i if and only if there is an undirected link between them; this is modeled as a Markovian prior probability distribution on random variables. The Markov property in MRFs suggest a random variable X_i is conditionally independent of another random variable, given random variables in its neighborhood.

For example, in the application of image restoration

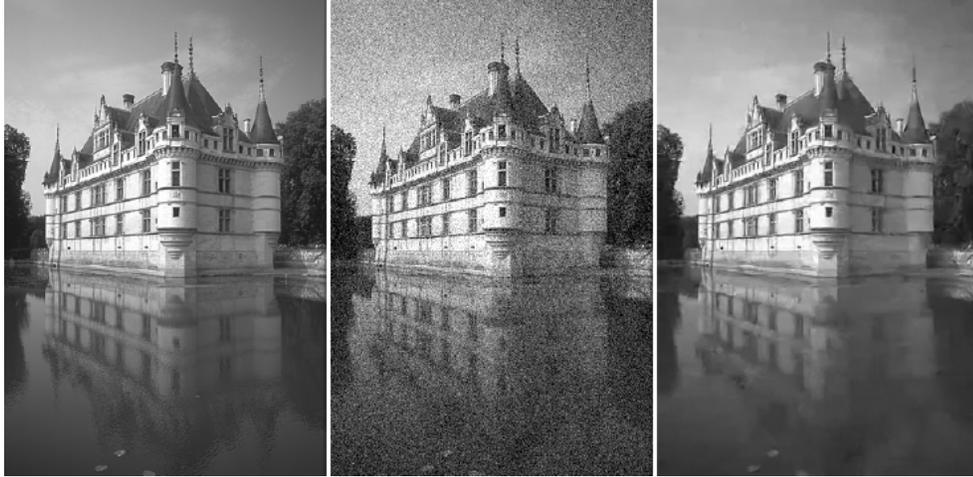


Figure 4.1: In image restoration, an image is corrupted with noise (middle) and is restored (right) similarly to its original (left)

4.2 Inference

Inference aims to calculate the marginal probability, inferring the maximum a posterior (MAP) solution. For example, in a MRF model, the MAP inference estimates the best states of labels X given observations Y :

$$X^* = \operatorname{argmin}_x P(X|Y) = \operatorname{argmin}_x \frac{P(X,Y)}{P(Y)} = \operatorname{argmin}_x \frac{P(X|Y)P(X)}{P(Y)}$$

where $P(Y|X)$ is the likelihood model and $P(X)$ is the prior probability distribution of labels.

Inference algorithms for probabilistic graphical models can be divided into exact inference and approximate inference methods. For single connected tree structures, belief propagation (BP) [24] can find the exact solution based on local message-passing. In this algorithm, belief propagation iteratively updates messages between adjacent nodes until convergence and infers the best states for either a single node or set of nodes. However, this method cannot guarantee an exact solution for a probability graphical model with undirected loops. For a probabilistic graphical model with undirected loops, the junction tree algorithm [26] [27] provides an exact solution by converting it into a tree structure by grouping a set of nodes. Exact inference can be performed on this converted probabilistic

graphical model using the belief propagation method.

For approximate inference techniques, these methods include loopy belief propagation (LBP), the mean field method, and variational algorithms [28]. The loopy belief propagation method (briefly described in 4.2.1 shows the approach for applied to layer boundaries 4.2 applies the belief propagation principle to a graphical model with undirected loops. It can produce an approximate solution, which may not guarantee convergence for the message-passing process.

Besides the methods discussed, there are other approaches to perform inference for graphical models. The Markov Chain Monte Carlo is a simulation technique. Either gibbs or Metropolis Hastings are used to extract samples in the MCMC algorithm.

4.2.1 Loopy Belief Propagation

Belief propagation is an iterative algorithm, which propagates local messages along the nodes of a MRF. For each iteration, every node sends outgoing messages to its neighboring nodes and accepts incoming messages from each one of its neighboring nodes. Sending and receiving new messages for the next iteration is based on the existing set of messages at the current iteration; this procedure is repeated until all messages have converged.

The main idea of belief propagation is the concept of messages. From each iteration, every node must send one message per label to each one of its neighbors. Intuitively, the message sent from node A to node B about label x_A indicates how likely node A thinks of node B should be assigned that specific label. After stabilization of the messages, each node chooses a label, which has the maximum support based on all incoming messages at the node.

4.2.2 Markov Chain Monte Carlo

Inference is the process of For parameter estimation, the posterior distribution cannot be derived analytically and a numerical analysis is not feasible. The long standing problem

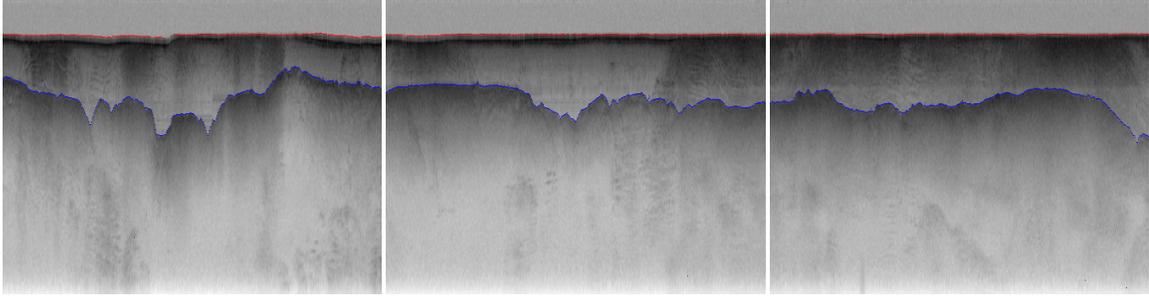


Figure 4.2: Different CReSIS Radar Systems

was solved by Markov Chain Monte Carlo. MCMC outputs a sequence of parameter sets (Markov Chain) whose empirical distribution approximates (converges to) the posterior distribution)

Gibbs Sampling

Gibbs sampling as our inference algorithm of choice. Gibbs sampling is a Markov Chain Monte Carlo (MCMC) method, which is capable of producing samples X^1, \dots, X^J from a distribution $f(x)$ without requiring the ability to directly sample or even know the form of $f(x)$; this is accomplished by iteratively sampling each variable conditioned on the remaining variables. The algorithm for gibbs sampling is shown in Figure 4.3. This sampler provides a flexible framework for generating samples from a complex distribution, assuming samples can be taken from simpler full conditionals. At run-time the mJ samples must be drawn from the full conditionals where m is the number of variables and J is the number of iterations, which depend on the sampling from the full conditionals. As the gibbs sampler requires a initialization of the model, many early samples may not reflect the true distribution and are discarded as a "burn-in" samples.

With these samples from the joint distribution, the functionals from the samples can be computed. For example, the mean of the samples approximates the expectation of the joint. In Chapter 5, confidence intervals for the improved model prediction, producing estimates of error in addition to the inference process

```

1: Initialize  $X^0 = \{x_1, \dots, x_m\}$ ;
2:  $j = 1$ ;
3: while  $j < J$  do
4:    $X^{(j)} = X^{(j-1)}$ ;
5:   for all  $x_i$  in  $X^{(j)}$  do
6:      $x_i^{(j)} \sim P(x_i | X^{(j)} - \{x_i^{(j)}\})$ ;
7:   end for
8:    $j = j + 1$ ;
9: end while

```

Figure 4.3: Figure The general algorithm for Gibbs sampling iteratively draws new values for random variables conditioned both on observed variables and other unknowns.

Reversible Jumps MCMC

The reversible jump MCMC sampler, an extension of Metropolis-Hastings, jumps between models and parameter spaces of different dimensions in the sampling process. The RJMCMC moves represent either model changes or an update to current parameters. The move, which requires an update to parameters requires no dimension changes or change in parameter spaces. Proposal for either a new model or to update parameters are evaluated based on an acceptance probability.

The reversible jump sampler has two types of moves: updating a current model's parameters and transitioning to different models. For the specific model, M_k , a move is chosen to either updating its parameters or move to a different model, M_{k+1} , which is attainable from M_k in one iteration. For an update move, a proposal for parameter values are obtained from the proposal distribution. A new set of parameters is sampled from the parameter space for M_k , and the probability of acceptance is calculated to be compared with a random value from a uniform distribution, and if it is less than the transition to M_{k+1} , it is accepted. Otherwise, the model remains as M_k . The process continues convergence. In Chapter 6, a detailed approach to RJMCMC is discussed with application to layer detection.

CHAPTER 5

ESTIMATING BEDROCK AND SURFACE LAYER BOUNDARIES AND CONFIDENCE INTERVALS IN ICE SHEET RADAR IMAGERY USING MCMC

In this Chapter, identifying surface and bedrock layers is posed as an inference problem on a probabilistic graphical model similar to work introduced by Crandall et al [29]. However, several important contributions are introduced to improve its accuracy and utility; this approach uses Gibbs sampling to perform inference instead of the dynamic programming-based solver, which allows the model to solve for both layer simultaneously, yielding automatic layer detection results significantly better. Moreover, the Gibbs sampler produces explicit confidence intervals by estimating bands of uncertainty for layer boundaries. Since noise and ambiguity in echograms are inevitable, this ability would provide confidence in climate models.

5.1 Introduction

Recent glaciological surveys have revealed a significant increase in ice discharge from the polar ice caps. There is a need for ice flow models to better reproduce current changes and forecast future behavior of the polar ice sheets. However, these models require important parameters, such as surface topography and bedrock elevation, which control the dynamics and ensuring the overall mass balance of the polar ice sheets.

Radar systems have evolved to improve data collection strategies and have increased the number of datasets capable of determining bed and surface topographies. In addition to advances in radar instrumentation and algorithm signal processing, analysis approaches to estimating ice thickness have also provided a solution to the obsoleting traditional methods; for this method, a probabilistic graphical model is posed by incorporating local and global features for layer detection, but although the approach is completely automated, it

sometimes solve for layers incorrectly since it uses a simplistic model and solved for them incrementally producing a single answer for an ice layer. In order to supplement the limitations in accuracy, the model is improved by solving for layer boundaries simultaneously and providing confidence intervals. The method is evaluated using Multichannel Coherent Depth Sounder (MCORDS) data acquired from Antarctica.

5.2 Methodology

Observations about how domain experts detect layer boundaries were used to improve the automated algorithm. As shown in Figure 5.1 and typical for the MCORDS dataset, the surface reflection is strong while the bedrock layer is rigid with weaker edges. The process of modeling and determining the location of layer boundaries are discussed in subsections 5.2.1 and 5.2.2, respectively. Figure 5.1 is used as a demonstration for the proposed approach.

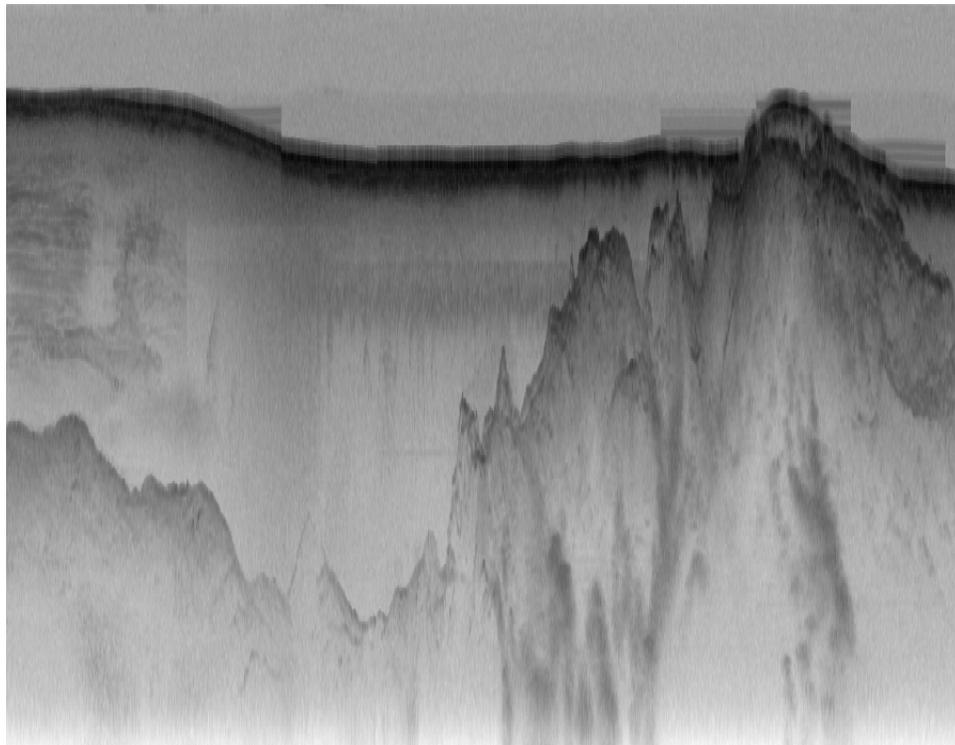


Figure 5.1: Original Echogram with Two Layer Boundaries

5.2.1 Modeling layer boundaries

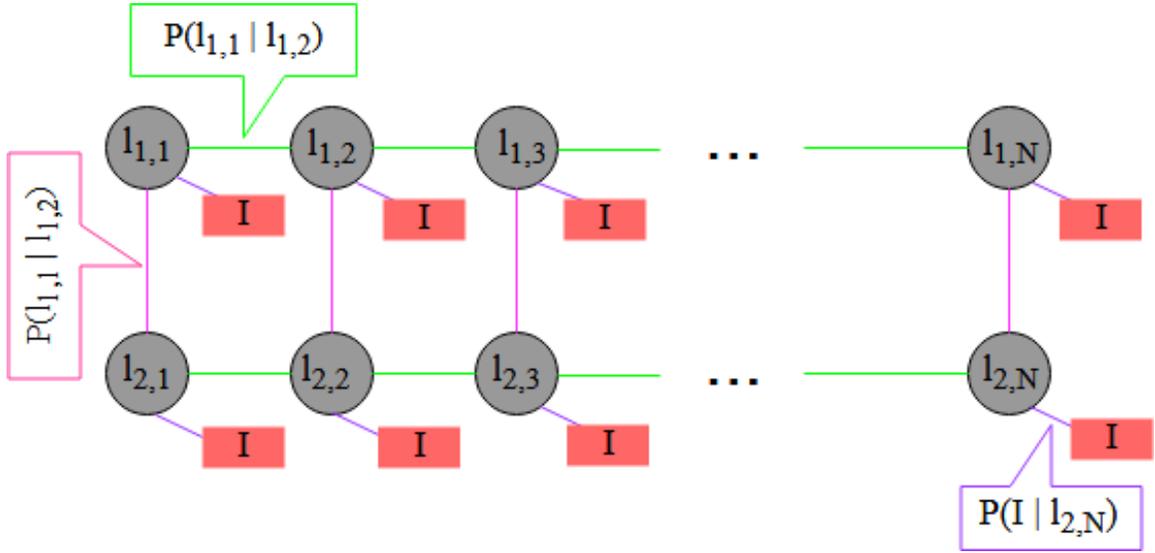


Figure 5.2: Illustration depicting a model for two layers, where each row represents a layer through the echogram. The horizontal layer constrain enforce smoothness while the vertical constraint encourage repulsiveness between layers boundaries. The image likelihood are based on local echogram characteristics

In this work, an echogram has k layer boundaries (with $k=2$ in our case). Given an echogram I of dimension $m \times n$, the goal is to estimate unknown variables $L = \{L_1, \dots, L_k\}$, where $L_i = \{l_{i1}, \dots, l_{in}\}$ and l_{ij} denotes the row coordinate of layer i in column j .

The structure of this problem allows it to be posed as a grid-shaped probabilistic graphical model. In this framework, the goal is to estimate $P(L_1, \dots, L_k | I)$, the joint probability over the layer boundaries given the echogram. Figure 5.2 shows as graphical representation of a two layer model for an echogram of width N . For each layer, each column n of the image is associated with random variables $l_{1,n}$ and $l_{2,n}$ corresponding to the positions of each layer in that column. These variables are connected via dependency relationships to the image I and their immediate neighbors (both along a layer and between layers). In Figure 5.2, the node representing l_{11} is highlighted and the conditional distributions with respect to l_{11} are written.

Unfortunately, this distribution has an alarming dimension of order $O(m^{kn})$, so com-

putation and storage is intractable even for small images. To address this problem, three assumptions are followed:

- 1 all echograms are equally likely
- 2 image characteristics are determined by local layer boundaries; and
- 3 variables in L exhibit a Markov property with respect to their local neighbors

Under the first assumption, the joint distribution can be factored into a product according to Bayes' Law,

$$P(L_1, \dots, L_k | I) \propto P(I | L_1, \dots, L_k) P(L_1, \dots, L_k). \quad (5.1)$$

This decomposition reduces the full joint into two intuitive distributions: $P(I | L_1, \dots, L_k)$ captures how well the image data can be explained by a set of layers L_1, \dots, L_k , and $P(L_1, \dots, L_k)$ captures prior knowledge about the boundaries, like that they are smooth and do not intersect.

The second assumption implies parts of the image not near the layer boundaries are generated by noise, so only model pixels near boundaries. Therefore, it can be factored $P(I | L_1, \dots, L_k)$ into a product over layers and columns,

$$P(I | L_1, \dots, L_k) = \prod_{i=1}^k \prod_{j=1}^n P(I | l_{i,j}). \quad (5.2)$$

Since boundaries are dark edges, the right hand term is modeled as a product of gradient magnitude and image intensity,

$$P(I | l_{i,j}) \propto |\nabla I(l_{i,j}, j)| \cdot (1 - I(l_{i,j}, j)), \quad (5.3)$$

where $|\nabla I(x, y)|$ is the gradient magnitude at coordinate (x, y) of the image, and we assume that pixel values have been scaled such that $I(x, y) \in [0, 1]$. The gradient magnitude

was approximated using finite differences on a 5×5 window.

The third assumption simplifies the problem by assuming the graphical model has the property that each node $l_{i,j}$ is independent of the remaining variables in L given its immediate neighbors in the graph. Under this assumption:

$$P(L_1, \dots, L_k) \propto \prod_{i=1}^k \prod_{j=1}^n P(l_{i,j} | N(l_{i,j})) \quad (5.4)$$

where $N(l_{i,j})$ is the set of directly connected nodes in the graph (i.e. $N(l_{i,j}) = \{l_{a,b} \mid 1 = |a - i| \text{ and } 1 = |b - j|\}$). $P(l_{i,j} | N(l_{i,j}))$ is defined as the product of independent vertical and horizontal components. Along the same layer, l_i 's are encouraged to be smooth by a zero-mean Gaussian which is truncated to zero outside a fixed interval,

$$P(l_{i,j} | l_{i,j-1}) \propto \begin{cases} \mathcal{N}(l_{i,j} - l_{i,j-1}; 0, \sigma) & |l_{i,j} - l_{i,j-1}| < \phi_H \\ 0 & \text{otherwise,} \end{cases} \quad (5.5)$$

while a step function encourages layers not to overlap,

$$P(l_{i,j} | l_{i-1,j}) \propto \begin{cases} 0 & l_{i,j} \leq l_{i-1,j} \\ 0.1 & l_{i,j} - l_{i-1,j} < \phi_V \\ 1 & \text{otherwise.} \end{cases} \quad (5.6)$$

This model is similar to [29] but with important improvements. In [29], the vertical pairwise potentials are zero at and above intersection points and uniform elsewhere. But it is common in this data to see radar reflections of the surface layer directly below the actual surface, so a fixed-width low probability region is added directly below to reduce false bedrock detections on these reflections. Perhaps more importantly, the model in [29] breaks these vertical constraints in order to simplify inference by greedily solving each layer conditioned on the previous one. The experiments show a holistic inference approach offers substantial improvement in accuracy.

5.2.2 Statistical inference

The model defined by equations (5.1), (5.2), and (5.4) is a first-order Markov Random Field. Unfortunately, finding the values of L that maximizes equation (5.1) is NP-hard in the general case [30]. Rather than trying to solve this as an optimization problem, an attempt is to estimate functionals of the full joint distribution by using gibbs sampling (a Markov Chain Monte Carlo technique discussed in sec.)

It can be shown by Bayes Law and the independence assumptions in equations (5.2) and (5.4) that the full conditionals for each l_{ij} can be computed easily,

$$P(l_{ij}|I, N(l_{ij})) = P(I|l_{ij})P(l_{ij}|N(l_{ij})). \quad (5.7)$$

As the domain of l_{ij} is discrete and finite, sampling from this conditional is well-defined. As an additional optimization, the vertical and horizontal thresholds in equation (5.6) were used to sparsify the computation of $P(l_{ij}|N(l_{ij}))$, since most entries are known to be zero. The Gibbs sampler is applied to generate a sequence of samples $L^{(B)}, \dots, L^{(J)}$ where B is a burn-in time during which samples are discarded. This is a common practice with MCMC methods to reduce sensitivity to initial values. To predict the layer locations, the sample means $M = J - B$ samples, which approximates the expectation of the joint distribution for large M ,

$$E [P(L_1, \dots, L_k|I)] = \lim_{M \rightarrow \infty} \frac{1}{M} \sum L^{(i)}.$$

To produce confidence intervals around this mean, the marginal distribution of a variable is estimated by discarding other variables in the sample and yielding 2.5% and 97.5% quantiles.

5.3 Results

Layer identification was evaluated using a set of 826 publicly available radar echograms from the 2009 NASA Operation Ice Bridge program, collected with the airborne Multi-channel Coherent Radar Depth Sounder system of [31]. Each echogram has a resolution of 700 by 900 pixels (where 900 pixels represents about 30km of data on the horizontal axis, and 700 pixels corresponds to 0-4km of ice thickness on the vertical axis). This dataset was used by [29] and can directly compare the accuracy of the two techniques.

The images have ground-truth labels produced by human annotators, but these labels are often quite noisy. For instance, sometimes the annotators could not find a reasonable layer boundary and simply “gave up” by not marking anything at all. To decouple the error in the ground-truth from the method evaluation, echograms with incomplete ground truth are removed (including those with partially defined layers and those with less than two layer boundaries). The proposed method ran on the remaining 560 images. For each image, 10,000 samples were collected after a burn-in period of $B=20,000$ iterations.

Figure 5.3 shows detected surface and bedrock boundaries in addition to confidence intervals. Figure 5.4 highlights more results of the proposed technique (including the confidence interval) as well as the ground truth. Also, the confidence intervals were quantified by computing the percentage of ground truth layer points that are contained within the estimated intervals. Based on the proposed work, 94.7% of the surface boundaries and 78.1% of the bedrock boundaries are within the intervals, for an overall percentage of 86.4%; this number is close to but less than 95% and reflects how the framework is a good but not perfect for layers in echograms.

A quantitative performance metrics is shown in Table 5.1. Accuracy was measured by viewing ground truth and estimated layer boundaries as 1-D signals, and computing the mean absolute deviation (in pixels) between the two. The following two summary statis-

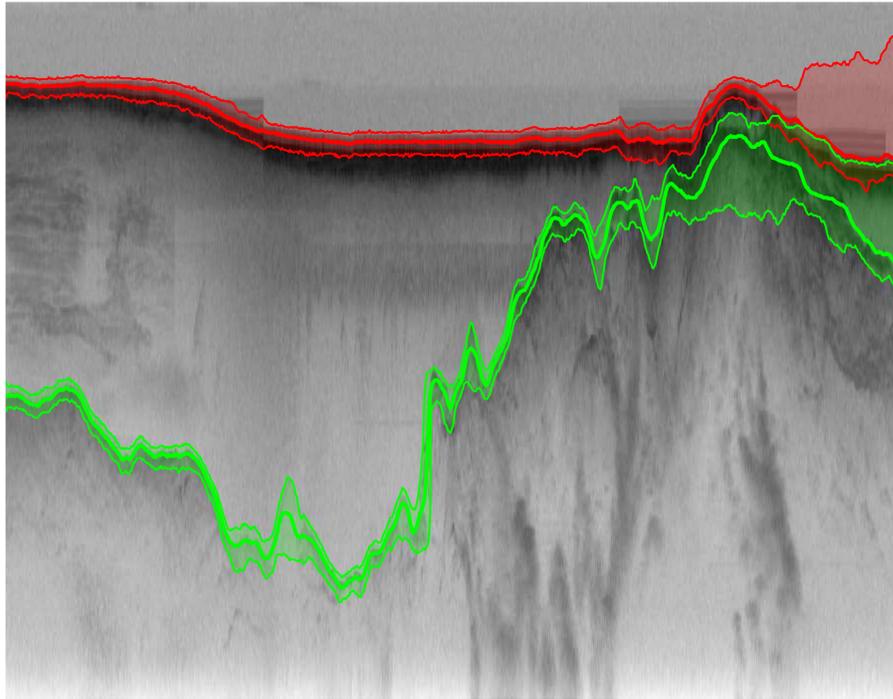


Figure 5.3: Detected Layers of Improved Approach with Confidence Intervals

tics: mean column-wise absolute error over all images and the median of the column-wise mean absolute errors across images. The first measures how well predicted layers match the ground truth, treating columns within an image as uncorrelated, while the later metric recognizes that high error in one column in an image is highly correlated with the error in the remaining columns and looks at error from a per-image viewpoint. Under both metrics, the improved approach outperforms the method of [29] significantly, by decreasing the error rate by about 44.3% for surface boundaries and 48.3% for bedrock. Our technique is slower than [29] (about 17 seconds per image), but since layer finding is trivially parallelizable across images, the accuracy is much more important than compute time.

Approach	Mean Error		Median Mean Error	
	Surface	Bedrock	Surface	Bedrock
[29]	22.3	43.1	10.6	14.4
Ours	9.3	37.4	5.9	9.1

Table 5.1: Evaluation of our method on the test set. Error is measured in terms of absolute column-wise difference compared to ground truth, summarized with average mean deviation and median mean deviation across images, in pixels.

5.4 Conclusion

In this chapter, an automated approach to estimate bedrock and surface layers in multi-channel coherent radar imagery and demonstrated its effectiveness on a real-world dataset against the state-of-the-art. The proposed technique also produces confidence interval estimates and was evaluated for correctness. Layer detection methods provide confidences, which could improve climate models by quantifying error for the input data. Although the approach improved layer accuracy, it requires a global parameter for specifying the number of layers; therefore limiting it to understanding ice thickness estimates.

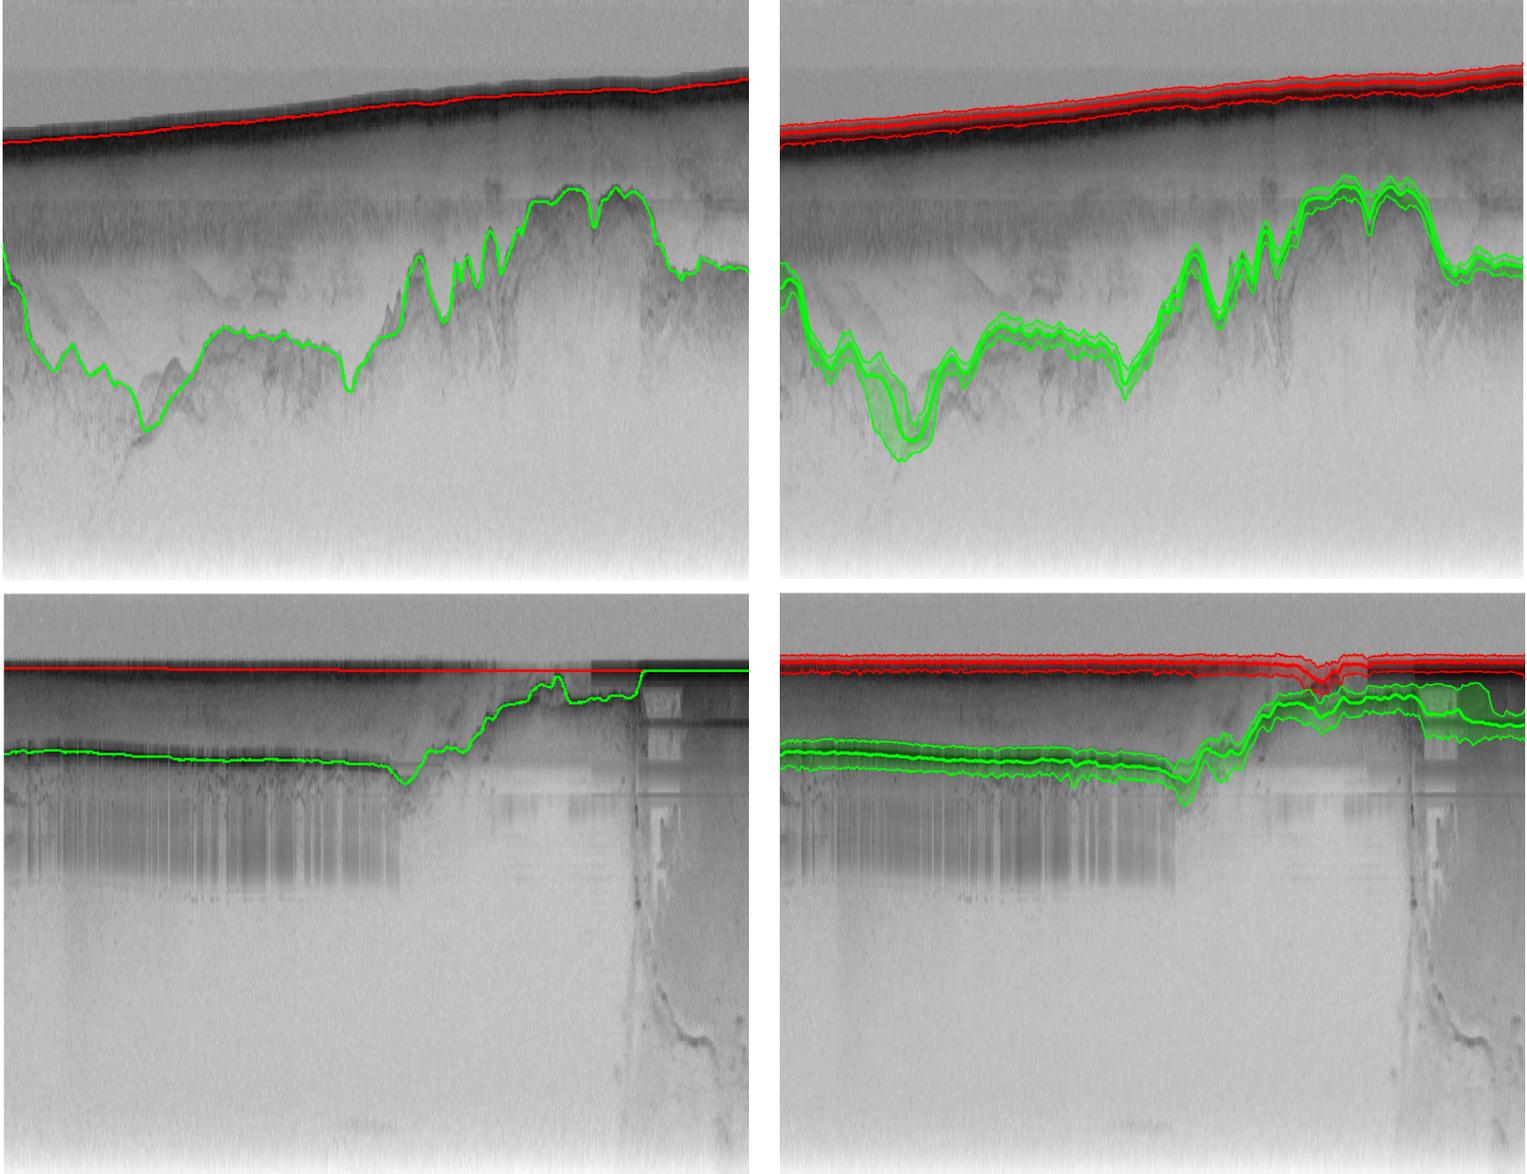


Figure 5.4: Sample results of approach on two MCORDS radar echograms.

CHAPTER 6

A SAMPLING-BASED APPROACH TO THE AUTOMATIC ANALYSIS OF LAYERS FROM POLAR RADAR IMAGERY

A number of studies have investigated the detection of internal layers from polar radar imagery. Most research have largely developed semi-automated techniques, but the need for an automated approach to support varying layer resolutions would greatly contribute to understanding global climate change. In this Chapter, an automatic approach is proposed for inferring a number of layers by using an advanced reversible jumps MCMC sampling approach, which samples candidate hypotheses to evaluate them against image observations to determine the optimal configuration explaining the echogram. The approach is used for two different sensors, snow radar and the accumulation radar.

6.1 Introduction

Ice penetrating radars have been collecting large volumes of data from the polar ice sheets, which reveal internal ice structures for many years. Analysis of these reflections highlight many aspects of ice sheet dynamics, including the strain history and their responses to climate and subglacial forces. As a result, a lot of datasets exist, which few or no internal reflections have been interpreted; this situation has motivated the development of an automated approach for the detection of internal layers from a variety of data sources.

Radar data acquired from snow and accumulation radars can have a number of internal layers; therefore, an automated approach need to provide good approximations to a variety of model structures. Although most MCMC samplers address parameter estimation, such as in Chapter 5, many cannot address the problem of choosing the best model for a specific problem domain. An automated method, reversible jumps MCMC, an extension of

Metropolis-Hastings, proposes to detect internal layer boundaries by estimating parameters and determining the appropriate model, simultaneously.

6.2 Methodology

The Markov Chain Monte Carlo method generate samples from complicated distributions by constructing a Markov Chain with the target distribution as its equilibrium distribution [32]. The Metropolis-Hastings algorithm simulates the desired Markov Chain by starting with an initial state and proposing a new state θ^* from a proposal distribution, which depends on the current state of θ . The proposed state is accepted as the next state in the Markov Chain with accepted probability

Reversible Jump Markov chain Monte Carlo algorithm (RJCMCMC) [33] explores configurations of different dimensionality, in this case, the number of internal layers. To generate configuration hypotheses, three proposals have been designed: birth, death, and update. Among the three, birth and death are a pair of reversible jump moves because adding or deleting layers from the current configuration causes a dimension change. Denote $L^t = l_1^t, \dots, l_n^t$ as the current state at iteration t , which consists of n layers. A new configuration L^* is proposed by first choosing a move from all three candidates according to a move probability distribution, which is a uniform distribution, but adapted to the current configuration for better accuracy. For example, if the current layer count is zero ($n = 0$), only the birth move is allowed. The RJCMCMC can be described as:

- Initialize the RJCMCMC sampler by choosing random state for the Markov chain.
- Begin with the starting state X^* of the Markov Chain
- Choose a move type from a set of possible move types
- Apply the chosen move. This involves proposing a new configuration
- Compute the acceptance ratio

- Add the n th sample to the Markov chain. If $\alpha = 1$, add the proposed configuration X^* to X_n . If not, add the proposed configuration with probability α . If the proposed configuration is rejected, add the previous configuration.

The following describes each proposal:

6.2.1 Birth and Death

To propose a new candidate layer, prior information about the appearance of the layers is used. A layer candidate is considered by sampling a row index according to the prior. The proposed configuration hypothesis is $l^* = \{l^t, l^m + 1\}$. The reverse move of birth is death; in this move, a layer is chosen from the current configuration and is removed based on the largest layer energy in the current configuration.

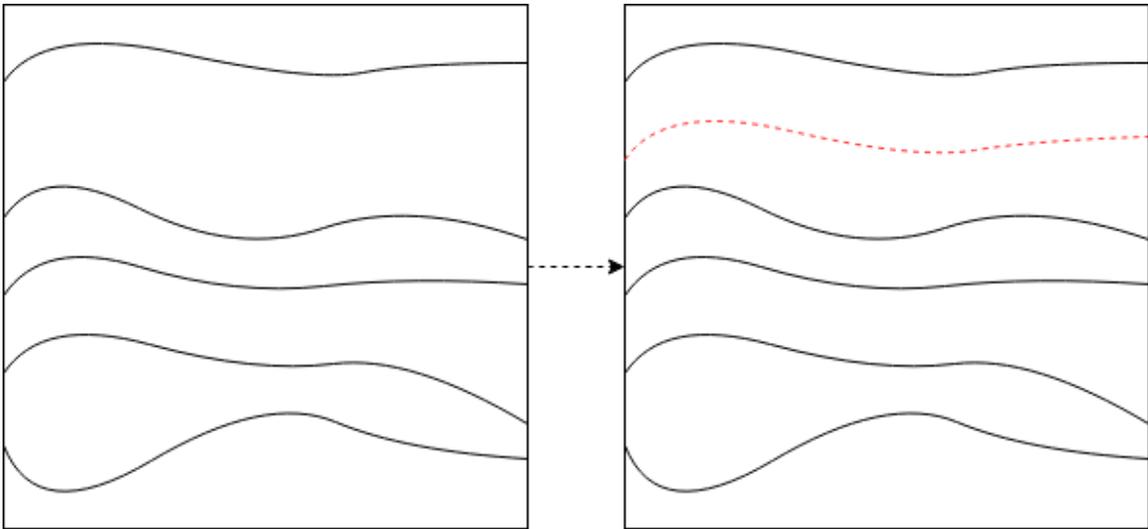


Figure 6.1: Illustration Depicting the Birth of a Layer

6.2.2 Update

The update proposal keeps the dimension of the current configuration but generates a new estimate from local adjustments to an existing layer detection. The update proposal is its own reverse move. A layer is selected for update from the current configuration at random.

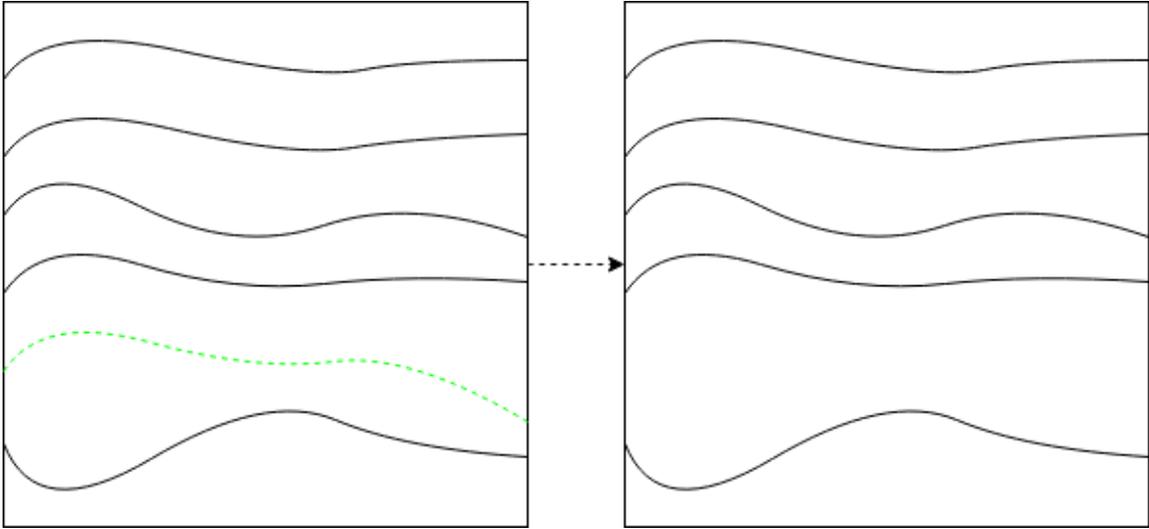


Figure 6.2: Illustration Depicting the Death of a Layer

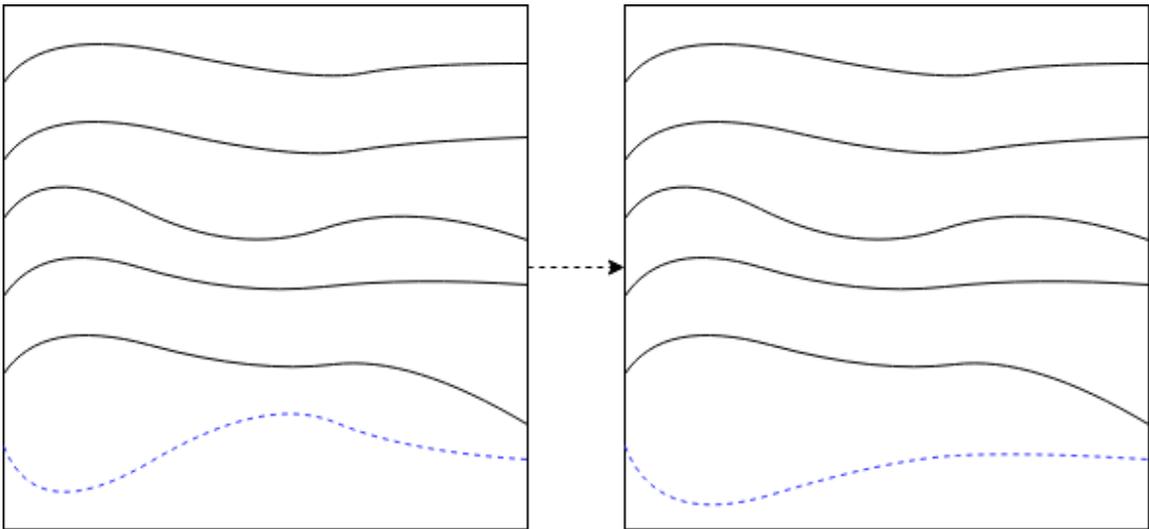


Figure 6.3: Illustration Depicting the Update of a Layer

6.3 Simulated Annealing

Simulated annealing [34], a Monte Carlo global optimization technique derived from the Metropolis Hastings algorithm, is used to ensure the convergence process. In short, simulated annealing begins selecting an initial solution and then generating a new state, randomly generating a new solution in the neighborhood of the current solution; this new state is evaluated and compared with the previous solution. If the new solution from the new

state is better than the previous, it is accepted; but if it is not, it is accepted or rejected with some probability.

The location of the current internal layer is used as an input parameter to an energy function. The proposed layer is then perturbed to yield a new fit, and the difference in energy between the two configurations is calculated. The following shows the core of the simulated annealing algorithm:

$$d_t = \begin{cases} d_{t-1} & \text{if } \Delta E \leq 0 \\ d_t & \text{if } \Delta E > 0 \text{ and } r > \exp(-\Delta/T(t)) \end{cases}$$

If the energy is lower for the new configuration d_t than d_{t-1} the new configuration is immediately accepted as a better fit (accept a). Otherwise, if changes caused the energy to be increased, there is a small chance the new configuration d_t would be accepted despite having a higher energy (accept b); this is controlled by the temperature formula

$$\frac{500}{\log(x + 1)}$$

at the current time step t , the energy difference and a uniformly distributed random variable formula. Finally, if the new fit was not accepted, the fit would be rolled back to the previous time step (reject). The b branch is included to make sure the algorithm does not get stuck in a local minimum of the energy function. As the probability decreases with temperature and time, it gradually less likely for a higher energy configuration to be accepted.

6.4 Results

For all experiments, a random number of layers were chosen as an initial configuration. If the layer could be initialized to an appropriate fit, it would converge faster. The RJMCMC sampler is coupled with simulated annealing in order to ensure convergence of the Markov

Chain by a temperature parameter. The temperature gradually decreases as the number of iterations increases; layers were detected between 50 and 100 iterations.

Figures 6.4 and 6.5 show an example of results obtained with RJMCMC. Some details illustrate how the algorithm identified layers with . However, on this example, some layers are not detected and others were. The missed and false detections can be alleviated with tuning of model parameters

In order to show the effectiveness of the algorithm, it was tested with echograms acquired from snow and accumulation radars. The results show the algorithm is able to detect most layers with varying resolutions and moderately

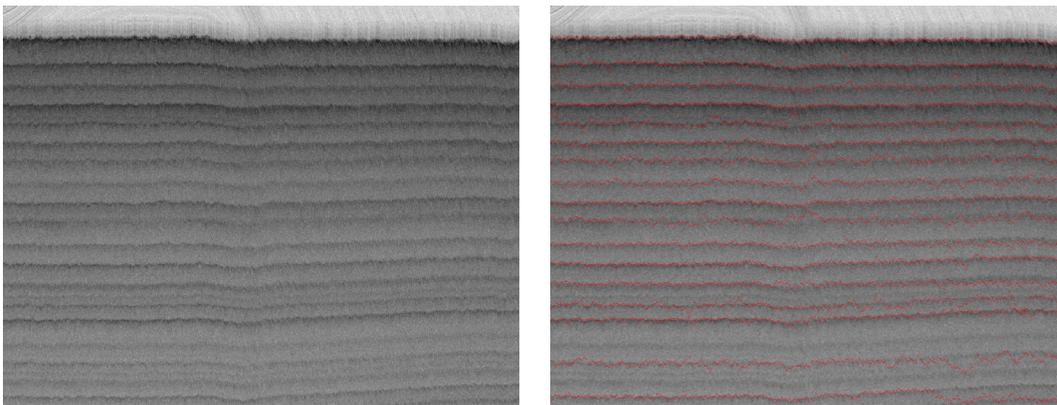


Figure 6.4: Sample results of the automated approach on three 2014 snow radar echograms acquired from Antarctica

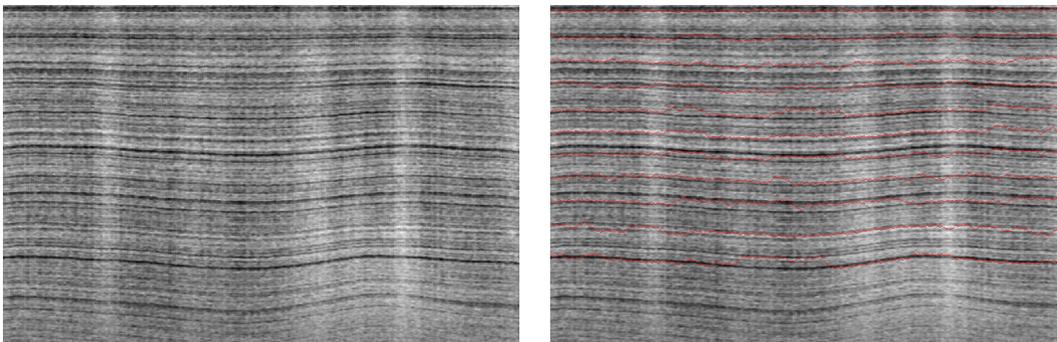
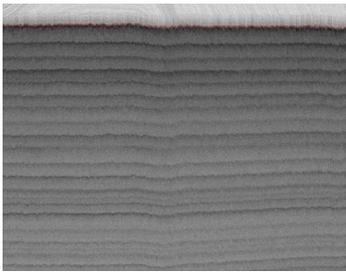
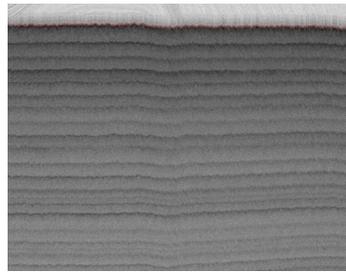


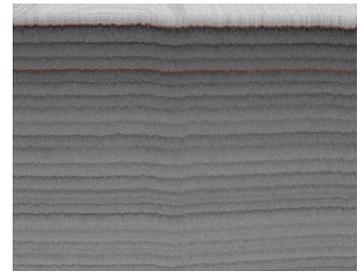
Figure 6.5: Sample results of the automated approach for 2014 accumulation radar echograms acquired from Antarctica



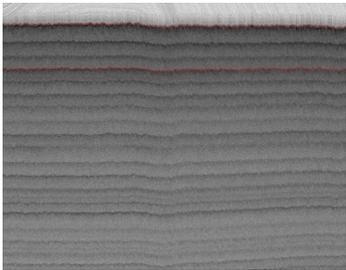
(a) Birth



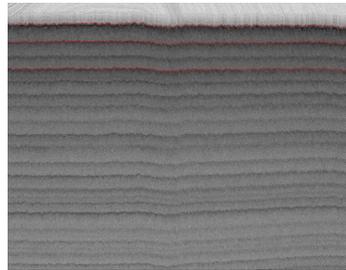
(b) Update



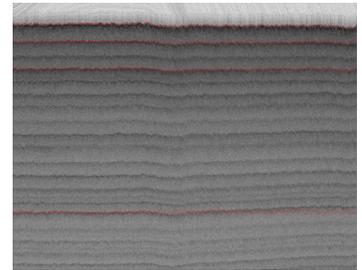
(c) Birth



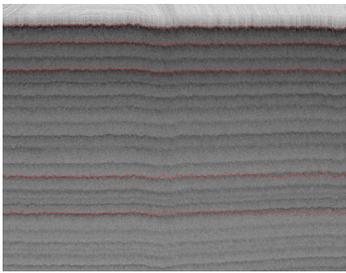
(d) Update



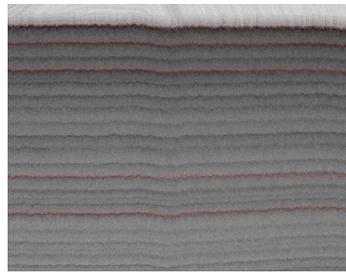
(e) Birth



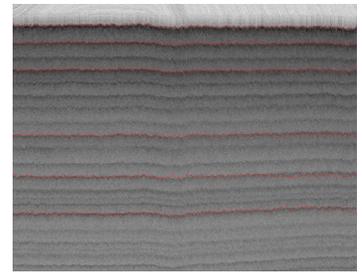
(f) Birth



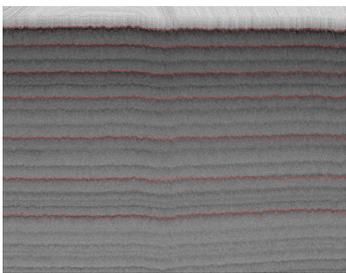
(g) Birth



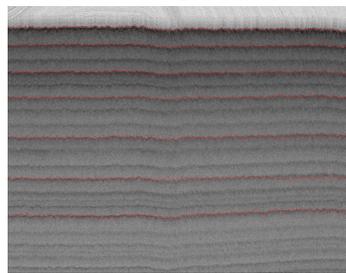
(h) Update



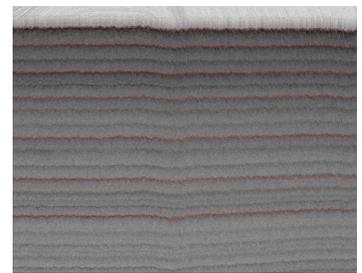
(i) Update



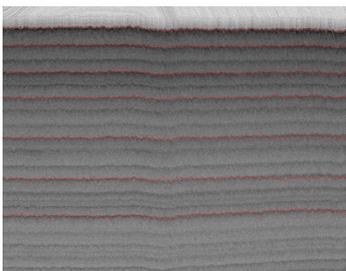
(j) Birth



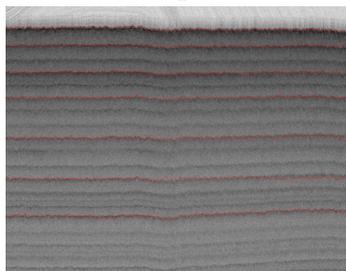
(k) Update



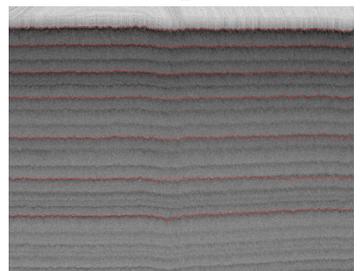
(l) Update



(m) Birth



(n) Update



(o) Update

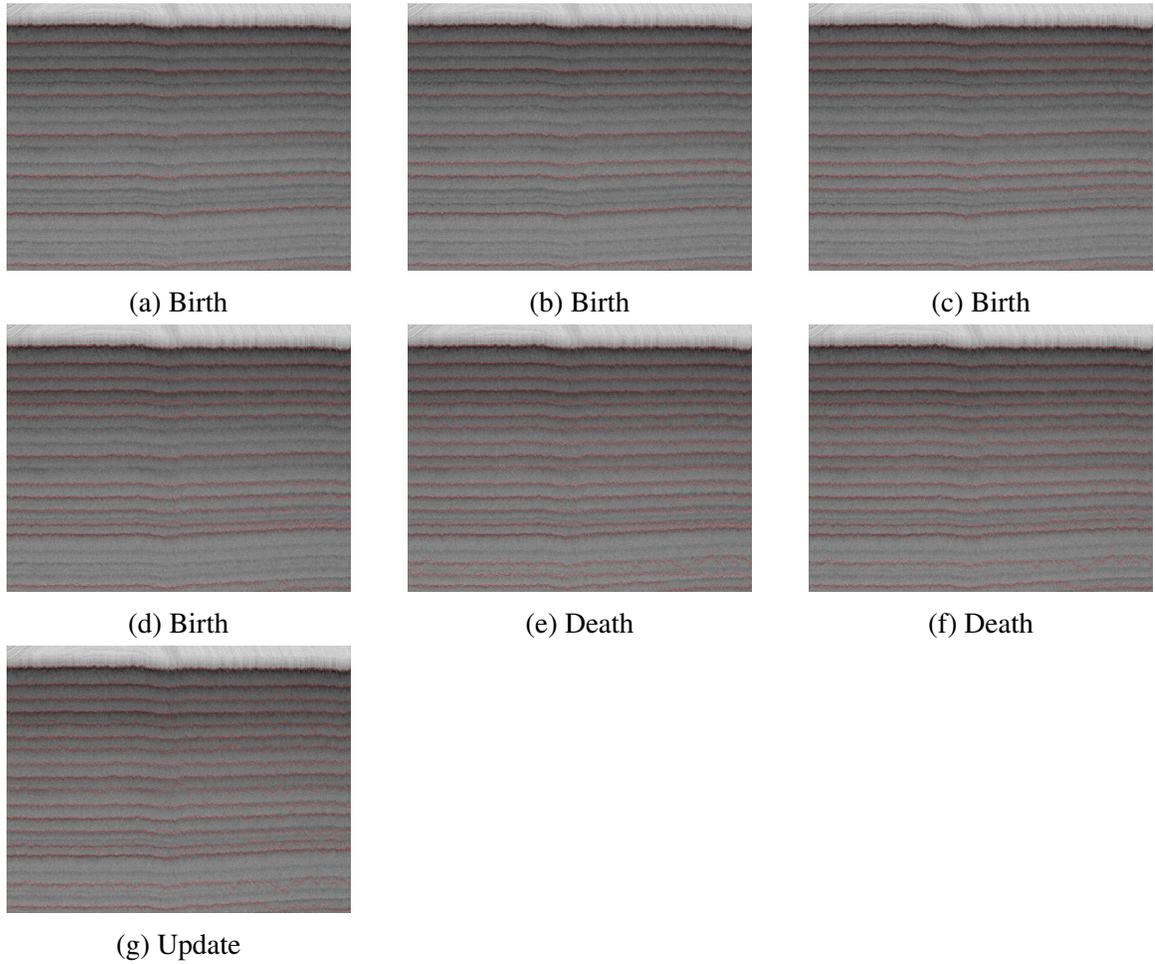


Figure 6.7: Intermediate accepted results during RJMCMC iterations. In a birth or update move, the red line indicates a layer being added or modified, which allows the current configuration l^* to be different from the previous iteration. In a death move, a red layer is being removed from l^t and the remaining layers constitute l^*

6.5 Conclusion

The development of RJMCMC for the application of layer detection has allowed for an automated approach supporting the detection of layers irrespective of radar and layer resolution. The RJMCMC sampler uses data driven proposals, such as birth, update and death, in addition to simulated annealing to ensure convergence for a number of unknown layers. The automated approach was applied to 2014 Greenland accumulation and Antarctica snow radar imagery and identified most layer boundaries.

CHAPTER 7

CONCLUSION AND FUTURE WORK

This Chapter concludes the dissertation by providing a summary of novel contributions and ideas for future developments.

7.1 Summary and Future Directions

Studying the rapidly changing polar ice sheets is important to understanding the future behavior and implications of global climate. This dissertation progressively develops an automated analysis tool to better understand the polar subsurface.

There are large amounts of radar data collected from the polar ice sheets. Due to the wide coverage and different types of information, analyzing echograms can lead to an increased understanding of the polar ice sheets. However, in past decades, echograms were interpreted manually with limited support of semi-automated techniques. Manual approaches limit scientific discovery, which could be achieved by more sophisticated techniques. The availability of radar data and projected data growth from future field campaigns support the need for novel automatic techniques for interpreting echograms. However, developing automatic techniques for polar radar imagery has been marginally addressed by the scientific community.

In this dissertation, a progressive approach to automatically identifying internal layers improves the literature by focusing on analysis techniques for polar radar imagery. In particular, the proposed methods are: a semi-automatic approach, which uses an active contour's method in addition to "off-the-shelf" image processing techniques, an improved probabilistic graphical model for automatically detecting bedrock and surface layers, and

an automated approach, which identifies layers irrespective of radar and layer resolution. These techniques address three challenges in studying polar ice sheets. The first method is a specialized approach, which tunes the algorithm based on observations from the internal layer data, the second contribution improves an existing detection approach and provides the foundation for a third method, which automatically detects internal layers from different radars and different resolutions. Future developments regarding the methods described in this dissertation should address the following:

- Further tuning and validation of the proposed techniques by applying them to datasets with different characteristics, acquired both in Antarctica and Greenland
- Definition and extraction of more discriminative features for echogram detection and development of an appropriate technique for removing the misclassified samples.

REFERENCES

- [1] Nanna B Karlsson et al. “Tracing the depth of the Holocene ice in North Greenland from radio-echo sounding data”. In: *Annals of Glaciology* 54.64 (2013), pp. 44–50.
- [2] Joseph A MacGregor et al. “Radiostratigraphy and age structure of the Greenland Ice Sheet”. In: *Journal of Geophysical Research: Earth Surface* 120.2 (2015), pp. 212–241.
- [3] Olaf Eisen. “Inference of velocity pattern from isochronous layers in firn, using an inverse method”. In: *Journal of glaciology* 54.187 (2008), pp. 613–630.
- [4] Edwin D Waddington et al. “Inference of accumulation-rate patterns from deep layers in glaciers and ice sheets”. In: *Journal of Glaciology* 53.183 (2007), pp. 694–712.
- [5] Vincent de Paul Onana et al. “A semiautomated multilayer picking algorithm for ice-sheet radar echograms applied to ground-based near-surface data”. In: *IEEE Transactions on Geoscience and Remote Sensing* 53.1 (2015), pp. 51–69.
- [6] M. Fahnestock et al. “Internal layer tracing and age-depth-accumulation relationships for the northern Greenland ice sheet”. In: *Journal of Geophysical Research* 106.D24 (2001), pp. 33789–33.
- [7] A.-M. Ilisei, A. Ferro, and L. Bruzzone. “A technique for the automatic estimation of ice thickness and bedrock properties from radar sounder data acquired at Antarctica”. In: *IEEE International Geoscience and Remote Sensing Symposium*. 2012, pp. 4457–4460.
- [8] Nanna B Karlsson et al. “A continuity-index for assessing ice-sheet dynamics from radar-sounded internal layers”. In: *Earth and Planetary Science Letters* 335 (2012), pp. 88–94.

- [9] Louise C Sime, Richard CA Hindmarsh, and Hugh Corr. “Automated Processing to derive Dip Angles of Englacial”. In: ().
- [10] Greg J Freeman, Alan C Bovik, and John W Holt. “Automated detection of near surface Martian ice layers in orbital radar data”. In: *Image Analysis & Interpretation (SSIAI), 2010 IEEE Southwest Symposium on*. IEEE. 2010, pp. 117–120.
- [11] Adamo Ferro and Lorenzo Bruzzone. “Automatic extraction and analysis of ice layering in radar sounder data”. In: *IEEE Transactions on Geoscience and Remote Sensing* 51.3 (2013), pp. 1622–1634.
- [12] Adamo Ferro and Lorenzo Bruzzone. “Analysis of radar sounder signals for the automatic detection and characterization of subsurface features”. In: *IEEE Transactions on Geoscience and Remote Sensing* 50.11 (2012), pp. 4333–4348.
- [13] Susanne L Buchardt and Dorthe Dahl-Jensen. “Estimating the basal melt rate at NorthGRIP using a Monte Carlo technique”. In: *Annals of Glaciology* 45 (2007), pp. 137–142.
- [14] SL Buchardt and D Dahl-Jensen. “Predicting the depth and thickness of the eemian layer at NEEM using a flow model and internal layers”. In: *Geofysikdag* (2007).
- [15] Matthew E Peters et al. “The distribution and classification of bottom crevasses from radar sounding of a large tabular iceberg”. In: *IEEE Geoscience and Remote Sensing Letters* 4.1 (2007), pp. 142–146.
- [16] S Moysey et al. “Quantification of reflection patterns in ground-penetrating radar data”. In: *AGU Fall Meeting Abstracts*. 2005.
- [17] Huasheng Zou and Feng Yang. “Study on signal interpretation of GPR based on support vector machines”. In: *International Conference on Life System Modeling and Simulation*. Springer. 2007, pp. 533–539.
- [18] Christian Panton. “Automated mapping of local layer slope and tracing of internal layers in radio echograms”. In: *Annals of Glaciology* 55.67 (2014), pp. 71–77.

- [19] J. Canny. “A computational approach to edge detection”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (1986), pp. 679–698.
- [20] C. Steger. “An unbiased detector of curvilinear structures”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20.2 (1998), pp. 113–125.
- [21] M. Kass, A. Witkin, and D. Terzopoulos. “Snakes: Active contour models”. In: *International Journal of Computer Vision* 1.4 (1988), pp. 321–331.
- [22] Stan Z Li. *Markov random field modeling in image analysis*. Springer Science & Business Media, 2009.
- [23] Steffen L Lauritzen. *Graphical models, volume 17 of Oxford Statistical Science Series*. 1996.
- [24] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.
- [25] Christopher M Bishop. “Graphical models”. In: *Pattern recognition and machine learning* 4 (2006), pp. 359–422.
- [26] Ross D Shachter, Stig K Andersen, and Peter Szolovits. “Global conditioning for probabilistic inference in belief networks”. In: *Uncertainty Proceedings 1994*. Elsevier, 1994, pp. 514–522.
- [27] Finn V Jensen. “Bayesian updating in causal probabilistic networks by local computations”. In: *An introduction to Bayesian networks* (1996).
- [28] Martin J Wainwright and Michael I Jordan. “11 A Variational Principle for Graphical Models”. In: (2005).
- [29] D. Crandall, G. Fox, and J. Paden. “Layer-finding in radar echograms using probabilistic graphical models”. In: *International Conference on Pattern Recognition*. 2012, pp. 1530–1533.

- [30] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT Press, 2009.
- [31] C. Allen et al. “Antarctic ice depthsounding radar instrumentation for the NASA DC-8”. In: *IEEE Aerospace and Electronic Systems Magazine* 27.3 (2012), pp. 4–20.
- [32] Peter J Green. “MCMC in image analysis”. In: *Markov chain Monte Carlo in practice* (1996), pp. 381–400.
- [33] Peter J Green and David I Hastie. “Reversible jump MCMC”. In: *Genetics* 155.3 (2009), pp. 1391–1403.
- [34] Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. “Optimization by simulated annealing”. In: *science* 220.4598 (1983), pp. 671–680.