


Towards an Antarctic meteorite hotspot map

Veronica Tollenaar



Towards an Antarctic meteorite hotspot map

by

Veronica Tollenaar

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday February 25, 2020 at 13:30.

Student number: 4162889
Project duration: June 17, 2019 – February 25, 2020
Thesis committee: Dr. H. Zekollari, TU Delft
Dr. S. L. M. Lhermitte, TU Delft
Dr. D. M. J. Tax, TU Delft
Dr. R. C. Lindenberg, TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Summary

Meteorites contain information on the formation and evolution of the Solar System. Antarctica is the most productive region for collecting meteorites, as the visually contrasting meteorites are easily detectable and tend to concentrate at specific areas exposing blue ice. Blue ice areas act as meteorite stranding surfaces (MSSs) if the flow of the ice sheet and specific geographical and climatological settings combine favorably. Previously, possible meteorite stranding surfaces were identified by chance or through visual examination of remote sensing data, which have limitations in discovering new locations for future meteorite searching campaigns.

In this study, datasets are combined in a novel machine learning approach in order to estimate the likelihood of a blue ice area to be a meteorite stranding surface. Input data consists of positive and unlabeled observations. The ca. 2,500 positive observations are defined as the centers of regularly spaced grid cells containing one or more meteorite finds. The ca. 2,000,000 unlabeled observations, for which the presence of meteorites is unknown, are defined as the centers of regularly spaced grid cells overlaying blue ice areas. The size of a grid cell is 450 by 450 meter. Features of the observations, such as the surface velocity, the surface temperature, and the ice thickness, are extracted from geospatial datasets. Individual features and correlations between features indicate that positive observations differ from unlabeled observations.

The unlabeled observations are classified as MSS or non-MSS by training a classifier with the nontraditional training set consisting of positive and unlabeled data. The obtained classification is validated and evaluated quantitatively with positive and negative observations, where the latter are defined after investigating fieldwork reports. With an estimated accuracy of 80 %, the classification shows promising results. The influence of the different features on the classification does confirm the current, qualitative, understanding of the meteorite concentration mechanism and provides a quantification of how individual features affect the meteorite concentration. In the visualization of the as MSS-classified observations, the probabilistic character of the obtained results is considered by using a color scale ranging from yellow to red. These colors indicate how likely it is to find meteorites at a MSS-classified observation (i.e. the precision of the classification). This leaves the interpretation of the obtained meteorite hotspot map to the user.

Contents

List of Figures	vii
List of Tables	xi
1 Introduction	1
2 Theoretical background: concentration mechanisms	3
2.1 Concentration model	3
2.1.1 Sublimation	4
2.1.2 Meteorite infall and fragmentation.	4
2.1.3 Temperature	5
2.1.4 Wind.	5
2.1.5 Glacial flow	5
2.2 Concentration model settings.	6
2.2.1 Concentration model setting 1: stranded ice.	6
2.2.2 Concentration model setting 2: closed BIA.	6
2.2.3 Concentration model setting 3: open BIA	6
2.2.4 Concentration model setting 4: lee side of barrier	6
2.3 Categorization of meteorite stranding surfaces (MSSs)	7
3 Methods	9
3.1 Observation definition	9
3.1.1 Positive observations	10
3.1.2 Negative observations	11
3.1.3 Unlabeled observations	11
3.2 Classification algorithm.	12
3.2.1 Principal component analysis (PCA) and Parzen density estimation.	12
3.2.2 Novelty detection	13
3.2.3 Standard classification.	13
3.2.4 Positive and unlabeled learning (PUL)	14
3.3 Receiver operating characteristic (ROC) curve	15
4 Data	19
4.1 Feature: Radar backscatter	19
4.2 Feature: Surface velocity	21
4.3 Feature: Ice thickness	22
4.4 Feature: Surface slope.	23
4.5 Feature: Yield stress.	24
4.6 Feature: Distance to outcrops.	25
4.7 Feature: Surface temperature	26
4.8 Feature: Change of ice thickness along the flow line	27
4.9 Comparison to the literature	28
5 Results	31
5.1 Feature selection	32
5.1.1 Sequential forward selection (SFS).	33
5.1.2 Sequential backward selection (SBS).	34
5.1.3 Comparison between SFS and SBS.	34
5.2 Principal component analysis (PCA)	35
5.3 Sensitivity analysis	37
5.3.1 Meteorite density	37
5.3.2 Meteorite finding locations	38

5.4	Classification	39
5.4.1	Selection of the threshold	39
5.4.2	Visualization	40
5.4.3	Evaluation	42
6	Discussion	43
6.1	Representativeness of locations of labeled and unlabeled observations	43
6.2	Quantitative features versus qualitative descriptions in the literature	44
6.3	Classification with positive and unlabeled data	45
6.4	Use of the meteorite hotspot map.	46
7	Conclusion and recommendations	47
	Bibliography	49
	Appendices	55
A	Meteoritical Bulletin dataset	57
B	Negative examples	65
C	Test data	69
C.1	Negative test data	69
C.2	Positive test data	70
D	Additional Figures	71
D.1	ROC curves sequential forward selection (SFS) and sequential backward selection (SBS)	71
D.1.1	Sequential forward selection (SFS)	71
D.1.2	Sequential backward selection (SBS)	73
D.2	Two dimensional histograms of principal components (PCs)	75
E	Meteorite hotspot map	79

List of Figures

2.1	Schematic sketches of four different settings of the concentration model (not to scale). Bedrock (both subglacial and exposed nunataks) is shown in brown. The color blue represent ice and white represents snow. In setting 2, 3 and 4, accumulating snow is displayed as well as the direction of katabatic winds enhancing the ablation. The red arrows indicate the occurrence of ablation. Thin gray arrows display the flow of the ice, in which meteorites are embedded. Black dots represent meteorites.	4
3.1	Overall framework for data handling, classification through machine learning and evaluation.	9
3.2	Snapshot in the area of Elephant Moraine (76°17'S, 157°20'E), illustrating the definition of observations. Unlabeled observations are the centers of regularly spaced grid cells overlaying BIAs. The size of a grid cell is 450 meter by 450 meter. When one or more meteorites are found within a grid cell, the observation is labeled positive. The labeled negative observations are extracted from the set of unlabeled observations, using information retrieved from fieldwork reports.	10
3.3	Histogram of transformed and standardized surface velocities. The surface velocity is transformed by taking the logarithm of the surface velocity, as the largest part of the velocities are low (i.e. the distribution is skewed). After that, the values are standardized, with as result that the mean is equal to zero and the standard deviation equal to one. Different scales for the y-axes are used, as the number of labeled observations is much lower than the number of unlabeled observations.	12
3.4	Normalized histogram of standardized surface velocities and the corresponding estimated density distributions $p(x s=0)$ and $p(x s=1)$. Part of the unlabeled observations are positive, consequently, there is partial overlap between the two distributions.	13
3.5	ROC curve of novelty detection (one class classification). Three different operating points are marked, and the corresponding thresholds are indicated. The threshold equals the estimated value for the probability density of $p(x s=1)$ (Figure 3.4), i.e. all test observations that score above the threshold when evaluating the estimated density on the observations, are classified as MSS. The resulting true positive rate and false positive rate, corresponding to a certain threshold, are shown as operating point on the ROC curve. For the threshold of 0.8, for example, all observations with a standardized surface velocity between roughly -1.5 and -0.75 (Figure 3.4) are classified as MSS. Negative test observations with velocities within this range contribute to the false positive rate, whereas positive test observations with velocities within this range contribute to the true positive rate.	14
3.6	Histograms of standardized surface velocities and the corresponding estimated density distributions $p(s=1 x) \times p(x)$ and $p(s=0 x) \times p(x)$. $p(s=1 x) \times p(x)$ is always smaller than $p(s=0 x) \times p(x)$. The histograms are plotted on the same scale (left y-axis), showing clearly that the number of unlabeled observations is much smaller than the number of labeled observations. $p(s=1 x)$ is always smaller than $p(s=0 x)$, as $p(s=1)$ is much smaller than $p(s=0)$. This indicates that all observations will be classified as belonging to class $s=0$	14
3.7	Estimated distributions of $p(y=1 x)$ ($\approx p(s=1 x) \times \lambda \times p(x)$) and $p(y=0 x)$ ($= 1 - p(y=1 x)$), using Bayes' minimum risk decision rule. The value for λ is chosen arbitrarily as 200.	15
3.8	Illustration of the theoretical false positive- and true positive rates. The distributions of the positive and negative test data are shown. When an observation from the positive distribution is classified as positive, it contributes to the true positive rate. When an observation from the negative distribution is classified as positive, it contributes to the false positive rate. The threshold determines the true positive and false positive rates. By varying the threshold, different true positive- and false positive rates are obtained, which are combined in a ROC curve (Figure 3.9).	16

3.9	Three ROC curves obtained by applying different thresholds to the distributions of test data. Depending on the separability of the test data, the obtained ROC curve is better (left) or worse (right). The color scheme of the upper panels corresponds to the color scheme of Figure 3.8, i.e. positive test data in green, negative test data in red.	16
3.10	Approximate locations of the nine most productive field sites of the dataset with usable coordinates as listed in Table 3.1.	17
4.1	Overview of features displaying a snapshot of the geospatial data in the area of Elephant Moraine (76°17'S, 157°20'E). Meteorite finding locations (Section 3.1.1) and the buffered BIA extent are shown in the lower right panel. References of the displayed datasets are given in the corresponding paragraphs (Section 4.1 to Section 4.8). The background image is the Center-Filled LIMA of the Landsat Image Mosaic of Antarctica (LIMA) Project.	20
4.2	Histogram of estimated radar backscatter intensities of unlabeled observations (left axis) and of labeled observations (right axis).	20
4.3	Two dimensional histogram of radar backscatter versus surface temperature (Section 4.7).	21
4.4	Histogram of estimated surface velocities of unlabeled observations (left axis) and of labeled observations (right axis).	22
4.5	Histogram of estimated ice thicknesses of unlabeled observations (left axis) and of labeled observations (right axis).	22
4.6	Histogram of average surface slopes of unlabeled observations (left axis) and of labeled observations (right axis).	23
4.7	Two dimensional histogram of ice thickness versus surface slope.	24
4.8	Histogram of estimated yield stress of unlabeled observations (left axis) and of labeled observations (right axis).	25
4.9	Histogram of estimated distances to outcrops of unlabeled observations (left axis) and of labeled observations (right axis).	26
4.10	Histogram of estimated mean surface temperatures of unlabeled observations (left axis) and of labeled observations (right axis).	27
4.11	Illustration of the calculation of the change of ice thickness along the flow line. The arrow indicates the direction of the flow line, estimated using the maximum slope of the average filtered surface elevation data, for which the 15 km circular kernel is displayed in pink. The ice thickness at the two observation locations, displayed as black dots, is also filtered with an averaging filter with circular kernel of 14.5 km, displayed with diagonal shading. From the difference between the two average thicknesses, the change of ice thickness along the flow line is computed.	27
4.12	Histogram of estimated yield stresses of unlabeled observations (left axis) and of labeled observations (right axis).	28
5.1	Antarctic meteorite hotspot map. A high resolution overview map, as well as local maps are provided in Appendix E	31
5.2	Qualitative summary of features, illustrating observation uncertainties and the link to the physical phenomenon of MSSs.	32
5.3	ROC curves when training the classifiers with a single feature. The eight curves are labeled with the name of the single feature.	33
5.4	ROC curves when training the classifiers with an increasing amount of features. The three curves are labeled with the number of the iteration and the additional feature in that iteration.	33
5.5	ROC curves when training the classifiers with a decreasing number of features. The three curves are labeled with the number of the iteration and the discarded feature in that iteration.	34
5.6	Subsets of features selected by the SFS and SBS method. The number of features in each iteration is indicated. The methods converge, but do not result in the same subset of features.	34
5.7	ROC curves after the sequential forward selection (SFS) of features and the sequential backward selection (SBS) of features. ROC curves of individual sites (thin lines) from which the weighted average ROC curve (bold line) is computed are plotted for comparison.	35
5.8	Explained variance by the principal components.	36
5.9	Two dimensional histogram of the first principal component and the second principal component, clearly indicating that the unlabeled data differs from the positive labeled data.	36

5.10	ROC curves when training the classifiers with a different number of principal components. Subsets of PCs are always defined so that the selected PCs explain the maximum variance (Figure 5.8). The ROC curves of the left panel are indicated in the right panel using dashed lines, and the other way around.	37
5.11	ROC curves when training the classifier with the first five PCs obtained in a PCA on six features (Section 5.1). Alternately, the observations at a single field site (as listed in the legend) are kept apart from the training data and used as positive validation data (Section 3.3). The ROC curve labeled with 'other classification' displays the weighted average ROC curve of the right panel in the left panel, and the other way around.	37
5.12	ROC curves when training the classifier with the almost the same training data as in Figure 5.11, with as difference that positive observations with a single meteorite in a grid cell are not used for training, but only for validation. The curves labeled with 'reference' are the weighted averages as displayed in Figure 5.11.	38
5.13	ROC curves when training the classifier with the exact meteorite finding locations as positive observations. The positive validation data also consists of the exact meteorite finding locations. The curves labeled with 'reference' are the weighted averages as displayed in Figure 5.11.	38
5.14	Metrics obtained by the number of true positives, false positives and false negatives. The F1 score represents the harmonic mean between precision and recall. The maximum value of the F1 score is indicated with an x.	39
5.15	Snapshot of the meteorite hotspot map in the area of the Allan Hills and Elephant Moraine. The classification values correspond to the classification values of Figure 5.16. The area over which the classification is performed is indicated with black outlines, labeled as 'buffered BIAS' (Section 6.1).	40
5.16	Histogram of classification values (upper panel) and related performance metrics (lower panel). The color scheme of the upper panel corresponds to the classification values displayed on the maps in Appendix E.	41
6.1	Snapshot of the meteorite hotspot map, displaying the Far Western Icefield in the area of the Allan Hills (77°S, 157°E). The positive and the negative observations used for validation are displayed. The obtained classification shows good accordance with the validation data, except in the Northern part of the area absent of meteorites.	44
D.1	ROC curves when training the classifiers with a single feature. The eight curves are labeled with the name of the single feature. <i>Figure is identical to Figure 5.3.</i>	71
D.2	ROC curves when training the classifiers with two features: the surface temperature and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on a single feature, the surface temperature, and is identical to the light blue curve in Figure D.1.	72
D.3	ROC curves when training the classifiers with three features: the surface temperature, the surface velocity, and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on two features, surface temperature and surface velocity, and is identical to the green curve in Figure D.2.	72
D.4	ROC curves when training the classifiers with four features: the surface temperature, the surface velocity, the radar backscatter, and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on three features, surface temperature, surface velocity, and radar backscatter, and is identical to the red curve in Figure D.3.	73
D.5	ROC curves when training the classifiers with seven features, all features but one. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on all features. The first five principal components used for the classifications with seven features explain between 87.6 and 93.1% of the variance of the data. The first five principal components used for the classifications with all eight features explain 86.9% of the variance of the data.	73

D.6	ROC curves when training the classifiers with six features, all features but the change of ice thickness and another feature. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on seven features, all features but the change of ice thickness, and is identical to the pink curve in Figure D.5. The first five principal components used for the classifications with six features explain between 94.6 and 97.0% of the variance of the data. The first five principal components used for the classifications with seven features explain 92.8% of the variance of the data.	74
D.7	ROC curves when training the classifiers with five features, all features but the change of ice thickness, the yield stress, and another feature. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on six features, all features but the change of ice thickness and the yield stress, and is identical to the orange curve in Figure D.6. The first five principal components used for the classifications with five features explain 100% of the variance of the data, as there are only five features considered. The first five principal components used for the classifications with six features explain 96.6% of the variance of the data.	74
D.8	Two dimensional histograms of the principal components (continued on next page).	75
D.8	Two dimensional histograms of the principal components (continued on next page).	76
D.8	Two dimensional histograms of the principal components.	77

List of Tables

2.1	Examples of MSSs in different concentration model settings.	7
3.1	Nine most productive field sites of the dataset with usable coordinates (Section 3.1.1). The numbers of usable coordinates is used as weights for computing the weighted average of the cross-validated ROC curves. The number of positive observations equals the number of grid cells containing one or more meteorite (Section 3.1.1), as used for training and validating the classifications. The colors represent geographic regions, indicating that field sites QUE, MIL, LEW and MAC are (relatively) close by each other, as well as that EET and ALH are close by each other (Figure 3.10, Appendix A).	17
4.1	Datasets used for the definition of features. The resolution refers to the used resolution, some data is available at higher resolution.	19
4.2	Summary statistics (mean, median, standard deviation (std dev), and range) of the eight features. The set of unlabeled observations contains both MSSs and non-MSSs, whereas the set of labeled observations contains only MSSs. Units are indicated in the first column. Corresponding histograms are shown in the corresponding Sections.	29
5.1	Linear combinations of original features from which the principal components are computed. .	36
5.2	Equations of classification performance metrics, where TP denotes the number of true positives, FP the number of false positives, FN the number of false negatives, and TN the number of true negatives.	40
5.3	Performance metrics estimated using the validation data.	40
5.4	Confusion matrix obtained by using the independent test data.	42
5.5	Performance metrics estimated using the independent test data.	42
A.1	Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].	58
A.1	Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].	59
A.1	Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].	60
B.1	Overview of negative examples as extracted from the literature.	65
B.1	Overview of negative examples as extracted from the literature.	66
B.1	Overview of negative examples as extracted from the literature.	67
C.1	Overview of additional negative examples used for testing.	69
C.1	Overview of additional negative examples used for testing.	70
C.2	Overview of additional positive examples used for testing	70

Introduction

Meteorites contain essential information for the understanding of the formation and evolution of the Solar System. This understanding leads to advancements in studies on e.g., the history of life on Earth, early Mars' habitable environments, and the components of the solar nebula [e.g. Corrigan et al., 2014; Dehant et al., 2019]. Until the mid-twentieth century, the geographic distribution of meteorite recoveries was primarily correlated to the human population density [e.g. Brown, 1960; Hughes, 1981; Wetherill, 1976]. In the 1950s and 1960s, a handful of meteorites were recovered from Antarctica, and the big continent was at that time not considered to be of particular interest for the collection of meteorites [Harvey, 2003]. This changed when in December 1969 a survey team of the tenth Japanese Antarctic Research Expedition (JARE-10) found nine meteorites at a blue ice field close to Yamato Mountains in Queen Maud Land [Harvey, 2003; Yoshida et al., 1971]. Two years later, the members of the inland team of the JARE-10 expedition, Yoshida et al. [1971], proposed a meteorite concentrating model, based on their field experience in Antarctica. This led to further research and subsequent searches for meteorites in Antarctica [Corrigan et al., 2014; Yoshida, 2010]. In the following years, meteorite search programs were established such as the Japanese Antarctic Research Expedition program (JARE), the U.S. Antarctic Search for Meteorites program (ANSMET) [Marvin, 2014], the European Meteorite program (EUROMET) [Corrigan et al., 2014], and the Chinese National Antarctic Research Expedition program (CHINARE) [Miao et al., 2018]. Nowadays, over 50 years later, almost 44,000 Antarctic meteorites have been collected, as listed in the database of the Meteoritical Society [a]. The proposed concentration model of Yoshida et al. [1971], following their meteorite recoveries, turned out to be of similar importance for the field of planetary science as the Apollo missions that brought in the same year, 1969, the first humans to the moon [Marvin, 2014; Yoshida, 2010].

As a result, despite the remoteness and desolation of Antarctica, 64 % [Meteoritical Society, a] of all recovered meteorites on Earth are found on this large continent, which is almost completely covered by an ice sheet. This high percentage can be mainly attributed to the meteorite concentration mechanism. When the flow of the ice sheet and specific geographical and climatological settings combine favourably, regions can act as so-called *meteorite stranding surfaces* (MSSs) [Harvey, 2003]. These MSSs typically expose old, blue colored ice [Cassidy et al., 1992], which is also interesting from a paleoclimatic perspective, as it bears information on the history of the ice sheet and on the past climate [Bintanja, 1999; Cassidy et al., 1992; Zekollari et al., 2019]. Besides the concentration mechanism there are other factors that contribute to the high number of recovered meteorites from Antarctica. An important element resides in the fact that the rates of meteorite weathering are low in the cold and dry environment of Antarctica, as compared to other places at the Earth's surface, where the meteorites are further from chemical equilibrium [Cassidy et al., 1992]. Another important aspect contributing to Antarctica's high meteorite retrieval number is the fact that it is easy to recognize a meteorite in a contrasting area of light-colored ice by human eye, as compared to a background composed of terrestrial materials [Cassidy et al., 1992; Harvey et al., 2015; Zekollari et al., 2019].

Blue ice areas (BIAs) that might contain a meteorite concentration are usually identified by examining maps and images such as Landsat imagery, Advanced Very High Resolution Radiometer (AVHRR) mosaics, Radarsat imagery, MODIS imagery and licensed high resolution satellite imagery [Harvey et al., 2015]. During a field season these areas are visited by snowmobile, aircraft or helicopter, to investigate whether a mete-

orite concentration is present [Harvey et al., 2015]. The discovery of meteorite concentrations is thus partly depending on the expertise and experience of the persons examining maps and imagery, and partly on the costly field reconnaissance visits. Because of this big human factor in the reconnaissance approach it is likely that (some) MSSs remain undiscovered. The goal of this research is to reveal these undiscovered MSSs, in other words, to map locations that are likely to be a MSSs in Antarctica (i.e. an Antarctic meteorite hotspot map). Producing such a map will also lead to a better understanding of MSSs and the concentrating mechanism.

The Antarctic meteorite hotspot map presented in this thesis is created using a big data approach. As a first step in this approach, quantitative information is extracted from geospatial datasets at locations where meteorites are found, as well as at locations that will be classified. This data serves as features of the so called labeled- and unlabeled observations, respectively. The definition of the features is based on the current understanding of the concentration mechanism. Defining thresholds on the calculated features of the positive observations in order to perform a simple overlay analysis as classification is not feasible, as the blue ice areas containing a concentration of meteorites show different settings of specific glaciological and geographical characteristics [Harvey et al., 2015]. Therefore, a binary classifier is trained. As useful negative observations for training the classifier are not available, the classifier is trained with positive and unlabeled observations. Training a classifier with this nontraditional training set is referred to as positive and unlabeled learning (PUL). The classification is validated (i.e. optimized) using observations of MSSs and non-MSSs, and evaluated (i.e. tested) with independent test data. The observations classified as MSSs are visualized in order to obtain an Antarctic meteorite hotspot map.

Through the various steps leading to the creation of a map of potential MSSs in Antarctica, the following research questions are addressed:

1. What are representative locations of labeled and unlabeled observations?
2. Which quantitative features do represent the characteristics of MSSs that are to date described qualitatively in literature? And do the feature-values of the labeled observations confirm the qualitative descriptions in the literature?
3. To what degree is a classification feasible given the lack of negative observations and the expected diversity of unlabeled observations?
4. In what way is the obtained classification to be used in meteorite recovery missions?

The structure of the report is as follows. In Chapter 2, the mechanism responsible for the meteorite concentrations is described. In Chapter 3, details on the observation definition, on the classification algorithm, and on the receiver operating characteristic (ROC) curve used for the validation of the classification are given. Chapter 4 lists the features used to represent the characteristics of MSSs. The results are presented in Chapter 5 and discussed centered around the four research questions in Chapter 6. A brief conclusion and recommendations on what further steps should be taken towards an Antarctic meteorite hotspot map are given in Chapter 7.

2

Theoretical background: concentration mechanisms

Meteorites fall regularly on Earth, distributed fairly equally across various latitudes [Halliday, 2001; Halliday and Griffin, 1982]. When meteorites fall on the Antarctic ice sheet, they can become entrapped in the ice sheet during the snow compaction process, as falling snow compacts under its own weight first to firn and later to ice. This snow compaction process takes place in the accumulation area of Antarctica, which makes up almost the entire ice sheet surface of Antarctica. Only ca. 1% of the ice sheet shows net ablation, namely at some areas at the coast, where temperatures in summer cause melting, and some spots inland, where abrasion and sublimation is observed [Bintanja, 1999]. The ice flows under gravitational forces towards the margins of the ice sheet, where various loss mechanisms (calving, sublimation, abrasion and melt) cause ablation. This flow of ice brings the embedded meteorites to other locations along with the ice. Although most of the meteorites end up in the ocean, at specific sites a small fraction of these englacially transported meteorites, embedded in the ice, are brought to the surface of the ice sheet. These sites, where a concentration of meteorites can be found, are called *meteorite stranding surfaces* (MSSs). The complex mechanism responsible for concentrating meteorites has been subject to diverse studies. Several models have been proposed that all describe a specific interplay between bedrock topography causing the ice flow to be redirected towards the surface, and ablation causing exposure of the embedded meteorites. The terminology used in the literature is not consistent and terms such as "concentration mechanisms" [Cassidy et al., 1992], "stranding surface models" [Cassidy et al., 1992], "ice flow models" [Corti et al., 2003; Folco et al., 2006], "mechanisms behind meteorite traps" [Zekollari et al., 2019], and "conveyor-belt model" [Harvey, 2003] are used inconsistently in various studies. In this study, the **conceptual model** of the concentration mechanism is called the *concentration model*. Descriptions of mechanisms resulting from studies on specific sites can be summarized into four cases, to which I refer as *concentration model settings*.

2.1. Concentration model

A large part of the meteorites at MSSs have been embedded in the ice, before they reappear at the surface. This means that the older ice in which they were embedded also appears at the surface. Consequently, MSSs typically expose old ice that has a blue color that sharply contrasts with white snow [Bintanja, 1999]. Areas where blue colored ice is exposed are called *blue ice areas* (BIAs) [Bintanja, 1999]. Bintanja [1999] defines a BIA as a region where the surface mass balance is negative, but hardly influenced by melt. The ablation, caused by sublimation, abrasion, and/or snowdrift, is thus larger than accumulation by precipitation [Bintanja, 1999]. Blue ice is also exposed at areas where melt does occur, typically at lower elevations in coastal regions, where summer temperatures can be relatively high [Bintanja, 1999]. Several authors [e.g. Hui et al., 2014; Winther and Liston, 2001] expand Bintanja [1999]'s definition of BIAs to include also these melt-driven areas, and also in this research the looser definition of BIAs is followed (i.e. including melt driven low elevation sites). BIAs occur scattered widely over Antarctica [Bintanja, 1999], but presumably only a minor part of these BIAs act as MSS [Harvey et al., 2001]. The known physical processes influencing the concentration of meteorite samples are listed here.




Figure 2.1: Schematic sketches of four different settings of the concentration model (not to scale). Bedrock (both subglacial and exposed nunataks) is shown in brown. The color blue represent ice and white represents snow. In setting 2, 3 and 4, accumulating snow is displayed as well as the direction of katabatic winds enhancing the ablation. The red arrows indicate the occurrence of ablation. Thin gray arrows display the flow of the ice, in which meteorites are embedded. Black dots represent meteorites.

2.1.1. Sublimation

Sublimation is an ablative mechanism where ice is converted into atmospheric moisture. Sublimation takes place with energy from solar radiation, and the low albedo (reflectivity of solar radiation) of blue ice as compared to snow increases the energy available [Bintanja, 1999; Harvey, 2003]. Sublimation is thus mainly occurring in summer. Removal of ice through sublimation, with typical rates of a few centimeters per year [e.g. Schultz and Annexstad, 1984], and exceptional values of up to five centimeter per day [Harvey, 2003], exposes meteorites that were entrapped in the removed layers. Sublimation is thought to be the dominant ablative mechanism at MSSs¹ [Bintanja, 1999]. As MSSs are a subset of [Bintanja, 1999]. Other ablative mechanisms such as melt and abrasion are discussed in Section 2.1.3 and Section 2.1.4, respectively.

2.1.2. Meteorite infall and fragmentation

If a meteorite falls directly on a MSS, it augments the meteorite concentration at the MSS. The contribution of direct infalls to the total number of meteorites at a MSS depends in part on the rate of infall. This infall rate is fairly evenly distributed over the surface of the earth [Halliday, 2001; Halliday and Griffin, 1982], indicating that rates are similar over BIAs. Direct infall is thus not a concentrating mechanism, but rather a contributor to the number of specimens. Zolensky et al. [2006] estimates today's influx of meteorites larger than 10 gram on the total surface of the Earth as ~ 42000 infalls/ 5.11×10^8 km²/yr. If this number would be constant over time this reflects ~ 82 infalls/km²/Myr.

¹Bintanja [1999] states that sublimation is thought to be the ablative mechanism at BIAs, where BIAs are defined as regions where melt is (almost) absent. MSSs are a subset of the BIAs as defined by Bintanja [1999] and therefore the statement holds.

The specimens found at MSSs can often be paired. This is because collision of the meteorite and the Earth's surface can cause the meteorite to fall apart into fragments [Marvin, 2014; Zolensky et al., 2006]. Moreover, meteorites often fall in groups of multiple fragments, so called 'meteorite showers', also resulting in samples that can be paired [Lindstrom and Score, 1994]. Estimates of the factor for the pairing of meteorites vary from 2 to 6, meaning that a single meteorite fall can result in 2 to 6 specimens [Lindstrom and Score, 1994; Zolensky et al., 2006]. As many identified pairing groups consist of samples that have a complete fusion crust, which is created while falling, it is likely that the fragmentation during falling or in later stages in the glacial transport are of minor importance as compared to atmospheric breakups resulting in showers [Corrigan et al., 2014]. Cassidy et al. [1992] mentions that if the specimens fall onto snow accumulation surfaces, the fragmentation is likely to be much lower, as the snow provides a softer landing surface compared to ice. However, this effect is not quantified.

2.1.3. Temperature

Low temperatures are beneficial for the preservation of meteorites exposed at MSSs for two reasons [Harvey, 2003]. (i) In the first instance, low temperatures prevent contact with liquid water, which can cause chemical weathering as well as freeze-thaw fracturing, as liquid water in the meteorite expands during refreezing [Cassidy et al., 1992; Harvey, 2003]. (ii) Also, low temperatures prevent the sinking of specimens. When temperatures rise, meteorites can become warmer than the surrounding air, causing the surrounding ice to melt and the specimen to move downward in the ice, creating a so called cryoconite hole [Cassidy et al., 1992; Sinisalo and Moore, 2010]. This process is not only dependent on air and ice temperatures, but also on solar radiation, weight and thermal conductivity of the meteorite and sublimation rate [Evatt et al., 2016; Harvey, 2003].

As temperatures are negatively correlated to elevation, favorable conditions are typically found at higher elevation, above 1500 m altitude and melt is rare at MSSs [Harvey, 2003; Harvey et al., 2001].

2.1.4. Wind

Wind can contribute in various ways to the concentration of meteorites. (i) Wind can redistribute smaller meteorites with weights up to ~200 g, with velocities ranging from 1 m/year for heavier stones (~100 g to ~170 g) to 10 m/year for lighter ones (~25 g) [Folco et al., 2002]. On the one hand, this wind transport can remove meteorites from an icefield by driving them into accumulating snowfields, moraines or crevasses. On the other hand, meteorites transported by wind can also be trapped in surface textures of the BIA, in the rough firn on the downwind side of the BIA, or in a downwind border of morainal material [Harvey, 2003]. Aeolian transport can thus reduce the number of samples by removal or it can change the site of concentration by redistribution. (ii) Strong katabatic winds at MSSs, which are often located in regions where those winds converge, do cause newly accumulated snow to erode, preventing accumulation [Bintanja, 1999; Harvey, 2003]. This ensures that the exposed meteorites remain exposed. However, these strong winds can also cause loss of meteorites at the surface through enhanced erosion. (iii) Wind does also contribute to abrasion, the process of ice loss through scouring of wind-blown particles of snow and ice [Harvey, 2003]. This process increases the removal of ice at the surface and thus leaves behind the meteorites that were entrapped in the removed layers. The contribution of abrasion to the total amount of ablation is not well understood, and some scientists suggest that the process is of minor importance [Bintanja, 1999; Harvey, 2003]. (iv) Lastly, turbulent winds in the lee of mountains increase sublimation rates (Section 2.1.1) [Bintanja, 1999]. These turbulent winds thus indirectly contribute to the exposure of samples in the ice.

2.1.5. Glacial flow

Convergent flow patterns and slow flowing ice are beneficial for the meteorite concentration at MSSs. Convergence of flow results in upward movement, bringing older layers of ice to the surface when ice is removed by other mechanisms such as sublimation [Harvey, 2003]. A slow glacial flow at MSSs ensures that the exposed meteorites reside for a long period of time at the MSS. Both the retardation of flow as well as the convergence of flow are related to the bedrock topography. Namely, bedrock obstacles have a buttressing effect on the ice, resulting in the lifting up of internal layers [Corti et al., 2008]. Outflow of ice from a MSS can bring along exposed meteorites and favor a fast export from a MSS.

2.2. Concentration model settings

Known specific geographic settings in which the discussed processes combine favorably to a meteorite concentration are described as concentration model settings. The distinction between the different model settings, which show some overlap, is in practice not always clear. Sometimes a MSS shows features of different settings, and it can best be described as a combination of two or more settings.

2.2.1. Concentration model setting 1: stranded ice

At some sites horizontal ice velocities are extremely low. In- and outflow of ice is limited here, as the body of ice is isolated from faster flowing ice. This situation can occur in a setting next to moving ice or in a downhill setting, where the ice flows into a (quasi) dead-end [Harvey, 2003]. This is sketched as the left BIA and right BIA in Figure 2.1a, respectively. A diffuse shear zone separates the stranded ice from the flowing ice [Harvey, 2003]. When the ablation rates equal the rate of inflow, there is a static system where the sublimation and abrasion of the upper layers results in exposure of the meteorites contained in the "stranded" ice, resulting in a significant concentration of meteorites [Cassidy et al., 1992; Harvey, 2003; Harvey et al., 2001]. Direct inflow of meteorites also contributes to the concentration, as the infalling meteorites do not become embedded in the ice and transport at the surface is very low. The ice at these locations can be very old, as there is a very limited inflow of younger ice. Consequently the meteorites can have terrestrial ages that are old, exceeding 3 million years [Harvey, 2003]. This model setting is illustrated in Figure 2.1a.

2.2.2. Concentration model setting 2: closed BIA

An emerged barrier can cause the ice flow to the margins to stagnate. If the barrier is perpendicular to the flow and long enough, not all the ice can flow around it. The central part of the ice will basically be unable to leave, causing the ice to lift up, a so called buttressing effect [Cassidy et al., 1992; Corti et al., 2008; Harvey, 2003]. The reduced ice flow velocities in combination with enhanced ablation by wind, exposes older ice in which meteorites are trapped [Cassidy et al., 1992; Corti et al., 2008; Harvey, 2003]. This situation is sketched in Figure 2.1b. Harvey [2003] calls this process a "conveyor-belt" process, as the ablated ice is replaced by inflowing ice that is lifted up again and ablated again, delivering trapped meteorites as a conveyor-belt [Cassidy et al., 1992; Harvey, 2003]. The type of BIA where this MSS occurs is called a 'closed type BIA', a categorization according to its flow regime [Sinisalo and Moore, 2010], hence the name of the concentration model setting.

2.2.3. Concentration model setting 3: open BIA

Similar to what happens in the closed BIA model setting, a submerged barrier can slow the ice flow and can cause the ice flow to be directed towards the surface due to a buttressing effect [Corti et al., 2008]. Enhanced ablation by wind then exposes the older ice in which meteorites are trapped [Corti et al., 2008]. The ablation by wind is enhanced, because the surface topography shows an increased inclination as a result of the bedrock topography, resulting in an increased speed of the katabatic wind [Takahashi et al., 1992]. This strong katabatic wind causes more drifting snow transport, enhancing abrasion and removing newly accumulated snow [Takahashi et al., 1992]. This type of BIA, categorized according to its flow regime, is called 'open type BIA' [Sinisalo and Moore, 2010]. The low velocity rates in these open BIAs leave meteorites exposed at the surface for long periods of time, which can be quantified with an estimated mean surface residence time (MSRT) [Zekollari et al., 2019]. For example, at the Nansen B icefield, Zekollari et al. [2019] uses the surface velocities to estimate a MSRT of over 5000 years. Meteorites eventually reach an accumulation zone at the edge of the BIA and become embedded in the ice (again) and transported away. This is illustrated in Figure 2.1c.

2.2.4. Concentration model setting 4: lee side of barrier

Meteorite stranding surfaces are also found downstream of major topographic barriers (nunataks). At the upstream side of an emerged barrier, the flow trajectories diverge to move around the obstacle, resulting in a shear zone [Corti et al., 2003]. At the lee side of the barrier, ice curves around the sub-ice extensions of the barrier and forms eddy currents, resulting in an area with low flow velocities [Cassidy et al., 1992]. When ice passes the barrier on two sides, two minor ice flows at the lee side converge [Corti et al., 2003]. Submerged bedrock topography at the lee side of the barrier induces uplifting of deep ice [Corti et al., 2003; Folco et al., 2002]. Together with wind induced ablation this results in exposed blue ice [Corti et al., 2003; Folco et al., 2002]. The ablation rates at the downwind side of the emerged bedrock obstacle can be very high, as the barrier causes turbulent wind falls of the katabatic wind [Corti et al., 2003], and the exposed bedrock can increase the total radiation energy available for sublimation [Folco et al., 2002]. This enhanced ablation

results in exhumation of meteorites [Cassidy et al., 1992]. The ice moves out of the blue ice area very slowly, so exhumed meteorites remain on the ice surface for a long time (long MSRT), before they escape from the MSS [Corti et al., 2003]. For example, at the MSS at Frontier Mountain Corti et al. [2003] estimates that meteorites remain exposed at the surface for over 60 000 years. This setting contains elements of Setting 1 and Setting 3, but distinguishes itself by the turbulent wind falls that induce ablation, as illustrated in Figure 2.1d.

2.3. Categorization of meteorite stranding surfaces (MSSs)

The description of the concentration model and its settings is most likely not all-encompassing. A limitation of the proposed model and its settings lies in the assumption that the mechanism is a stagnant system. The meteorites exposed at MSSs today have fallen over a long period (the oldest terrestrial age of a meteorite is estimated to be over 2 million years [Jull, 2001; Scherer et al., 1997; Welten et al., 1997]) and may have been embedded in the ice and transported over long distances [Zolensky et al., 2006]. The ice flow has not been constant over time, and this can be a factor influencing MSSs to a unknown degree. However, as the processes on the Antarctic ice sheet are slow and MSSs are operational for a long time [Cassidy et al., 1992], this possible temporal factor is not considered.

Categorizing a MSS as one of the four settings based on descriptions is not an unambiguous task. For instance, the MSS at Frontier Mountain, where two minor ice flows converge at the lee side of the barrier, is seen as Setting 4 by Corti et al. [2003]. However, Folco et al. [2002], considers the MSS at Frontier Mountain as a special case of Setting 3, claiming that it is the bedrock topography that causes the ice to move upwards behind the barrier. In the case of Yamato, the MSS is described as a very large MSS (about 4000 km²), consisting of multiple large BIAs and many small blue ice patches [Cassidy et al., 1992]. The existence of the MSS is related to the Yamato Mountains barrier and a presumably very wide subsurface barrier [Cassidy et al., 1992]. This results in a categorization into both setting 2 (closed BIA) and setting 3 (open BIA). Table 2.1 provides examples of MSSs of all four settings, derived from an interpretation of descriptions in literature. Although this categorization is a very rough simplification, it serves as illustration of the different model settings and can serve as guidance in the evaluation of the results. The four settings of MSSs as proposed can serve as template for studies on other MSSs and can be supplemented as the understanding on the mechanism develops.

Model setting	Icefield	Field site (Appendix A)	Literature	Remarks
1: stranded ice (Figure 2.1a)	Lewis Cliff Ice Tongue	Lewis Cliff (LEW)	Harvey [2003]	-
	Mare Meteoriticus	Queen Alexandra Range (QUE)	Harvey [2003]	-
	Tail's End Icefield	Queen Alexandra Range (QUE)	Harvey [2003]	-
2: closed BIA (Figure 2.1b)	Main	Pecora Escarpment (PCA)	Cassidy et al. [1992], Schutt [1989]	-
	Foggy Bottom Icefield	Queen Alexandra Range (QUE)	Harvey [2003]	-
	Yamato	Yamato (Y)	Cassidy et al. [1992]	Also 3: open BIA
3: open BIA (Figure 2.1c)	Far Western	Allan Hills (ALH)	Cassidy et al. [1992]; Harvey [2003]	-
	Near Western Icefield	Allan Hills (ALH)	Cassidy et al. [1992]	-
	Nansen icefields	Asuka (A)	Zekollari et al. [2019]	-
	Reckling Peak ice field	Reckling Peak (RKP)	Cassidy and Rancitelli [1982]	-
	Yamato	Yamato (Y)	Cassidy et al. [1992]	Also 2: closed BIA
4: lee side of barrier (Figure 2.1d)	Frontier Mountain icefield	Frontier Mountain (FRO)	Cassidy et al. [1992]; Corti et al. [2003]	-
	Johannessen Nunataks	Johannessen Nunataks (JOH)	Corti et al. [2003]	-
	Miller Butte	Miller Butte (MIB)	Corti et al. [2003]	-
	Mount Walton	Mount Walton (WAL)	Corti et al. [2003]	-
	Moulton Escarpment	Thiel Mountains (TIL)	Cassidy et al. [1992]	-

Table 2.1: Examples of MSSs in different concentration model settings.

3

Methods

In order to create an Antarctic meteorite hotspot map, various datasets are combined into a machine learning algorithm. The process of going from observations to the estimation of the likeliness of a location to be a meteorite stranding surface, is illustrated in Figure 3.1. The process is iterative, shown with the back-looping arrows. An evaluation of the results can lead to redefining the observations or changing the set of input features. All elements of this process are discussed in this Chapter, except for the data preparation, which is discussed in Chapter 4.

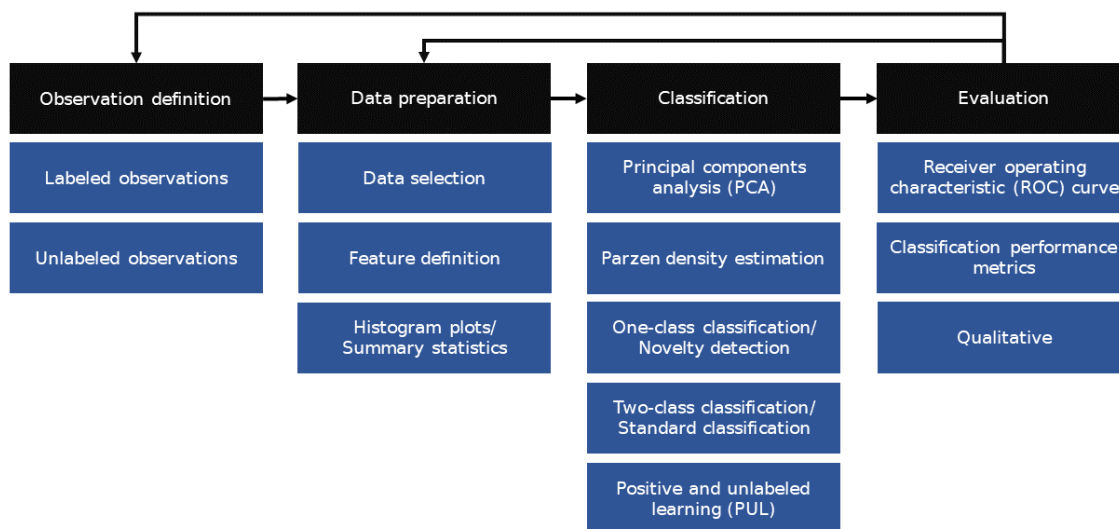


Figure 3.1: Overall framework for data handling, classification through machine learning and evaluation.

3.1. Observation definition

For classifying unlabeled observations as meteorite stranding surfaces (MSSs) or non-MSSs, observations of known MSSs and known non-MSSs are needed. Observations consist of values of features that are extracted from geospatial datasets using location data (Chapter 4). The locations of known MSSs are referred to as positive observations. The locations of non-MSSs are referred to as negative observations. Negative observations are only used for the evaluation of the classification, and not for the classification itself. This is because information on sites absent of meteorites is deduced from fieldwork reports. Consequently, this data is of limited reliability and not representative for all locations absent of meteorites, as only potential MSSs have been

searched (Section 3.1.2). The observations that need to be classified as MSS or non-MSS are referred to as unlabeled observations. Because of the absence of usable negative observations, the set of unlabeled observations, together with the set of positive observations, is used to train the classifier. Training a classifier with this nontraditional training set is referred to as positive and unlabeled learning (PUL) [Elkan and Noto, 2008; Li et al., 2011]. PUL is successfully applied in classification problems occurring in real-world domains, such as land surface classifications using remote-sensing data, as addressed by Li et al. [2011], or classification of genes or proteins in molecular biology databases, as addressed by Elkan and Noto [2008].

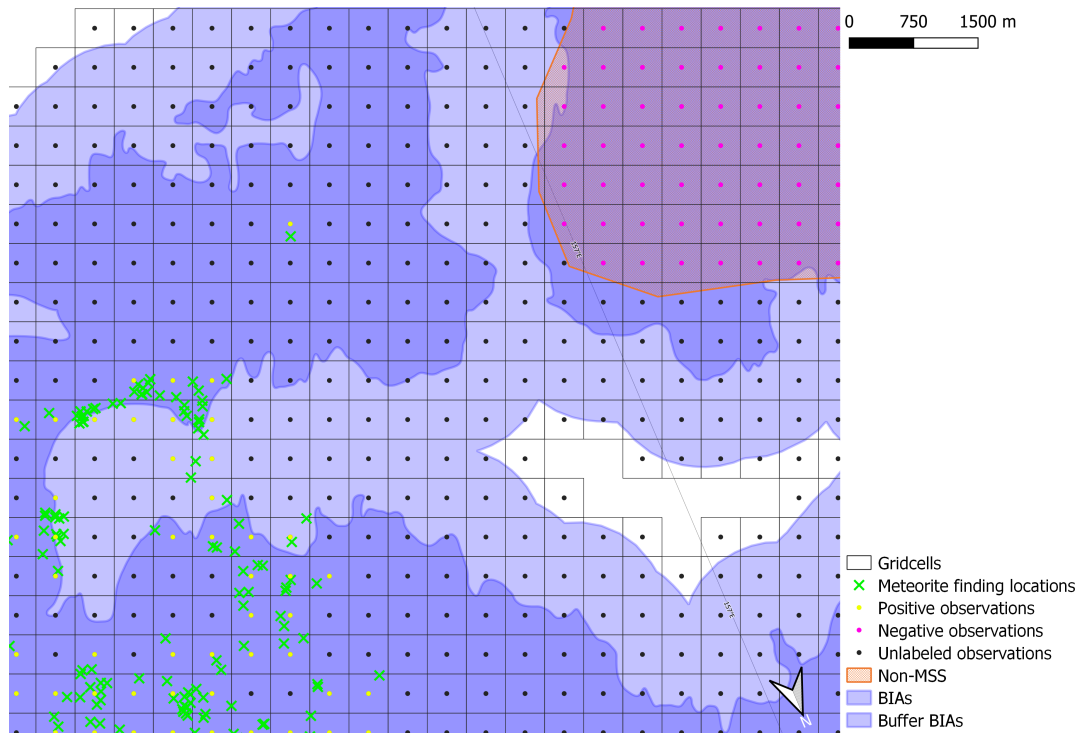


Figure 3.2: Snapshot in the area of Elephant Moraine (76°17'S, 157°20'E), illustrating the definition of observations. Unlabeled observations are the centers of regularly spaced grid cells overlaying BIAs. The size of a grid cell is 450 meter by 450 meter. When one or more meteorites are found within a grid cell, the observation is labeled positive. The labeled negative observations are extracted from the set of unlabeled observations, using information retrieved from fieldwork reports.

3.1.1. Positive observations

The Meteoritical Society [a] provides a large online database of collected meteorites with authoritative information on meteorite names, together with basic information, such as type, and place and year of discovery. The dataset, as consulted on 05/07/2019, contains information on 43,988 Antarctic meteorites, collected over around 50 years. Because of the large number of different missions to collect meteorites, and thus the large number of individuals involved, the data is neither complete nor consistent. Nevertheless, the dataset provides the best overview of meteorite finds. The dataset is presented in Appendix A, providing a table and several maps covering all meteorite recovery locations.

The literature complements the dataset with descriptive information on specific sites. Site descriptions, however, are often not detailed and not clearly demarcated, and consequently difficult to interpret. Also, naming of ice fields is inconsistent, and this makes it sometimes difficult to link descriptions to the data provided by the Meteoritical Society [a]. Therefore, only the data from the Meteoritical Society [a] is considered for positive observations.

Of the 43,988 meteorites in the database, 12,906 entries have usable location information [Meteoritical Society, a]. These finding locations of meteorites have not been used directly as positive observations (Figure 3.2), as there are two limitations to this approach. (i) Firstly, some meteorite finding locations are very close to each other, resulting in very similar values of the features of these observations. (ii) Secondly, the

observations are sensitive to the precision of the location data. Until 1990, meteorite location data of the Antarctic Search for Meteorites program of the United States (ANSMET), which has been carried out yearly from 1976 onwards [Marvin, 2014], was for instance collected using crude surveying methods [Schutt et al., 1993]. Only starting from 1990, GPS receivers were used, with reported measurement accuracy's ranging from 2 m to 40 m [Cassidy and Schutt, 1991; Schutt et al., 1993]. This was a considerable improvement, as the older methods were subject to biases introduced by accumulating errors in locating base stations [Schutt et al., 1993].

To overcome these two limitations, grid cells of a regular spaced raster are labeled positively when one or multiple meteorites are found within the grid cell. Using the same 450 meter resolution raster as for the unlabeled observations (Section 3.1.3), this approach results in a number of 2,554 observations. This reduction of observations implies a reduction of the number of features that can be used in the classification from six to five¹ [Silverman, 1986]. The influence of the reduced sample size is investigated in Section 5.3 by using the 12,906 finding locations directly as positive observations. The influence of the location precision is regarded negligible, as nearly three-quarters of the 12,906 meteorite findings occurred after 1990, when the use of GPS receivers became widespread [Meteoritical Society, a]. The reported errors of these locations are small compared to the resolution of the grid cells, and also as compared to the used datasets (Table 4.1).

3.1.2. Negative observations

Information on sites that are known for an absence of meteorites is collected by investigating field work reports of meteorite recovery missions. The information on these sites is scattered, inconsistent and tenuous. The location is often not clear, and also details of the search are not always indicated. It is not very certain that meteorites are absent at these locations, as the search might not have been extensive or the circumstances can have been unfavorable. For example, it is very often reported that a temporary snow cover, or sastrugi, can have hidden specimens [e.g. Cassidy et al., 1992; Corrigan et al., 2014; Marvin, 1982, 2014; Zekollari et al., 2019; Zolensky et al., 2006].

In Appendix A, a table is provided with an overview of sites containing no (concentration of) meteorites as mentioned in the consulted literature. The information from the fieldwork reports is translated into polygons, using imagery provided by Google Earth Pro (v. 7.3.2.5776). These polygons are used to remove observations from the set of unlabeled observations (Section 3.1.3) and label them as non-MSS (Figure 3.2). This results in a total of 8,870 negative observations.

These negative observations can not be used as input to the classification algorithm for two reasons. (i) Firstly, as mentioned before, it is not certain whether the data is actually negative. (ii) Secondly, these areas are a very selective subset of all areas that do not contain meteorites. Many sites have never been searched, because people did not expect to find meteorites there. The collected negative data is thus not representative for all negative data. Nevertheless, negative observations are essential for the evaluation of the classification. The fact that some negative observations might actually be positive results in a conservative evaluation of the classification, but the fact that the data is not representative for all negative data, might result in an optimistic evaluation of the classification (Section 6.3).

3.1.3. Unlabeled observations

Unlabeled observations are defined as regular spaced points over classified BIAs. The sampled points over the BIAs are regular spaced with a resolution of 450 meter (Figure 3.2). They are aligned with the velocity data (Section 4.2), as the ice velocity is a very important process influencing the meteorite concentration (Section 2.1.5). Locations that do not expose blue ice can be ignored, as MSSs always expose blue ice (Chapter 2). To select locations that expose blue ice, a classification of BIAs is used, provided by the Norwegian Polar Institute's Quantarctica QGIS package [Hui et al., 2014]. To reduce the number of unlabeled observations, BIAs located above 200 meter are considered², as meteorite finds at lower elevations are very unlikely (Section 2.1.3). This results in a total of approximately 2,000,000 unlabeled observations. The used BIA-classification is reported to be validated using other studies' estimates of the size of BIAs [Hui et al., 2014].

¹This figure is based on a table provided by Silverman [1986], which lists the required sample sizes for Parzen density estimations (Section 3.2.1) in different dimensionalities.

²For the estimation of the elevation of the BIAs the Reference Elevation Model of Antarctica on 200 meter resolution is used [Howat et al., 2019].

Despite the fact that this validation of selected individual sites shows good results [Hui et al., 2014], it does not eliminate the possibility that some (parts of) BIAs are missed by the classification. This concern is confirmed by a part of the Yamato BIA, which contains meteorites, being not classified. To evaluate the results of the performed MSS-classification qualitatively, the dataset of BIAs is expanded manually in order to cover the entire Yamato BIA. Also, a 1 km buffer has been drawn around the polygons delineating BIAs. Despite this measure, the possible missed BIAs in the used BIA dataset impose a limitation to the results obtained in this research.

3.2. Classification algorithm

The classification aims to label locations as MSS ($y = 1$) or as non-MSS ($y = 0$), where the variable y denotes the label. The training data for the classifier consists of positive observations, which are MSSs ($y = 1$) (Section 3.1.1), and of observations at locations where the presence of meteorites is unknown. The latter can be either MSS or non-MSS, and are referred to as unlabeled observations (Section 3.1.3). Observations of locations that are non-MSSs ($y = 0$) are not usable as training data, as the lack of meteorites is not certain and the sites are not representative for all negative data (Section 3.1.2).

To illustrate the approach to training a classifier with only positive and unlabeled data, the variable s is introduced, indicating whether the data is labeled ($s = 1$), or unlabeled ($s = 0$) [Elkan and Noto, 2008]. The probability that a labeled observation $\langle x, s = 1 \rangle$ is a MSS ($y = 1$), is equal to one:

$$p(y = 1 | x, s = 1) = 1. \quad (3.1)$$

The probability that an unlabeled observation $\langle x, s = 0 \rangle$ is a MSS, $p(y = 1 | x, s = 0)$, is estimated in the classification.

As it is not certain whether the unlabeled observations do contribute to the separability of the classes (Section 3.3), two types of classifications are performed. The first one is a novelty detection (or one-class classification), where only the labeled training data is used. The second one is a standard binary classifier (or two-class classification), using the labeled and unlabeled training data. The algorithm for both classifications is described here, and is illustrated with a one-dimensional example.

3.2.1. Principal component analysis (PCA) and Parzen density estimation

As a first step in both classifications, the features are transformed and standardized (Figure 3.3). With standardizing the features, the variance is set to unity, eliminating the influence of the arbitrarily selected units on the principal component analysis (PCA) that is performed consecutively. The PCA is performed in order

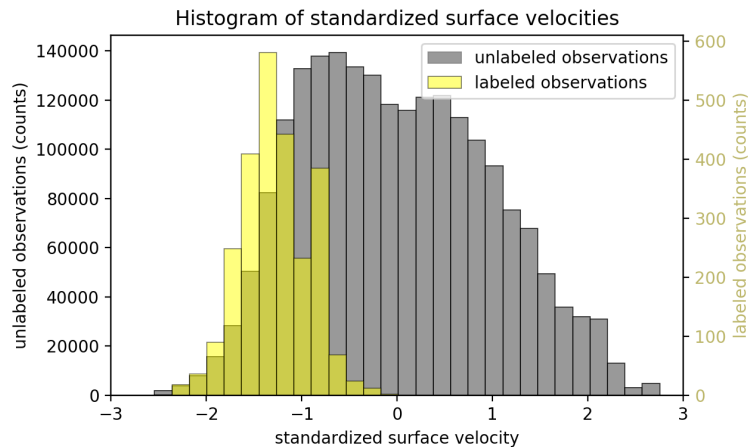


Figure 3.3: Histogram of transformed and standardized surface velocities. The surface velocity is transformed by taking the logarithm of the surface velocity, as the largest part of the velocities are low (i.e. the distribution is skewed). After that, the values are standardized, with as result that the mean is equal to zero and the standard deviation equal to one. Different scales for the y-axes are used, as the number of labeled observations is much lower than the number of unlabeled observations.

to reduce the dimensionality of the data. The dimensionality of the data needs to be reduced, as the number of features is eight (Chapter 4), while to obtain a reliable estimate of the density using the limited number of positive observations, the dimension of the data should be reduced to five (Section 3.1.1). The densities are estimated using a Parzen density estimation (also known as kernel density estimation) with a Gaussian kernel. The density is estimated both for the set of labeled observations, and for the set of unlabeled observations (Figure 3.4). These Parzen density estimations lead to estimated probabilities $p(x|s=1)$ and $p(x|s=0)$,

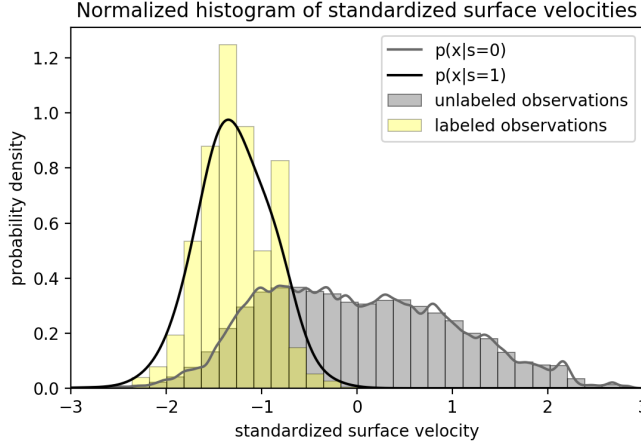


Figure 3.4: Normalized histogram of standardized surface velocities and the corresponding estimated density distributions $p(x|s=0)$ and $p(x|s=1)$. Part of the unlabeled observations are positive, consequently, there is partial overlap between the two distributions.

respectively. The bandwidth of the kernel, accounting for the variability in the distribution, is set by means of a 10-fold cross validation. For the Parzen density estimation of the unlabeled observations a randomly selected subset of 10,000 observations is used in order to reduce computational costs.

3.2.2. Novelty detection

With the estimated density model of the labeled observations, $p(x|s=1)$, a one class classification, or novelty detection, is performed. The estimated density model is evaluated on the unlabeled observations. Observations that score high in this evaluation likely belong to the positive class. A threshold to the score can be set arbitrarily. Selecting different thresholds and applying these to the testing data results in different values for the true positive rate and false positive rate. The combinations of these values can be shown as operating points, which form the so called receiver operating characteristic (ROC) curve (Section 3.3, Figure 3.5).

3.2.3. Standard classification

For the two-class classification, or standard classification, both the estimated density model of the labeled observations, $p(x|s=1)$, and the density model of the unlabeled observations, $p(x|s=0)$, are evaluated on the unlabeled observations. The resulting scores need to be transformed to $p(y=1|x)$ and $p(y=0|x)$ in order to decide whether an unlabeled observation, $\langle x, s=0 \rangle$, belongs to the MSS-class ($y=1$) or the non-MSS-class ($y=0$).

In a regular binary classification problem these *a posteriori* probabilities are estimated using Bayes rule, and the law to total probability:

$$p(y=1|x) = \frac{p(x|y=1) \times p(y=1)}{p(x)} = \frac{p(x|y=1) \times p(y=1)}{p(x|y=1)p(y=1) + p(x|y=0)p(y=0)} \quad (3.2)$$

and

$$p(y=0|x) = 1 - p(y=1|x). \quad (3.3)$$

If the *a posteriori* probability of the observation $\langle x, s=0 \rangle$ being a MSS is larger than the *a posteriori* probability of the observation belonging to the non-MSS-class, the observation can be classified as MSS. In formula:

$$\text{if } p(y=1|x) > p(y=0|x), \text{ then } x \in (y=1). \quad (3.4)$$

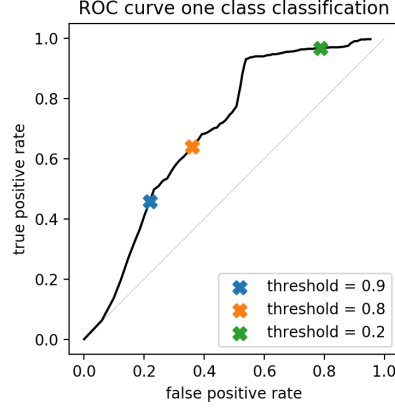


Figure 3.5: ROC curve of novelty detection (one class classification). Three different operating points are marked, and the corresponding thresholds are indicated. The threshold equals the estimated value for the probability density of $p(x|s = 1)$ (Figure 3.4), i.e. all test observations that score above the threshold when evaluating the estimated density on the observations, are classified as MSS. The resulting true positive rate and false positive rate, corresponding to a certain threshold, are shown as operating point on the ROC curve. For the threshold of 0.8, for example, all observations with a standardized surface velocity between roughly -1.5 and -0.75 (Figure 3.4) are classified as MSS. Negative test observations with velocities within this range contribute to the false positive rate, whereas positive test observations with velocities within this range contribute to the true positive rate.

Because of Equation (3.3), the decision rule can be also formulated as:

$$\text{if } p(y = 1|x) > 0.5, \text{ then } x \in (y = 1). \quad (3.5)$$

3.2.4. Positive and unlabeled learning (PUL)

In the classification problem considered here, the probabilities $p(x|s = 1)$ and $p(x|s = 0)$ are estimated, and not the probabilities $p(x|y = 1)$ and $p(x|y = 0)$, as the training dataset does not contain any labeled negative observations ($y = 0$) (Figure 3.4). These probabilities are transformed to estimations of $p(s = 1|x)$ and $p(s = 0|x)$ according to Equation (3.2) and Equation (3.3), respectively, where y is replaced by s (Figure 3.6).




Figure 3.6: Histograms of standardized surface velocities and the corresponding estimated density distributions $p(s = 1|x) \times p(x)$ and $p(s = 0|x) \times p(x)$. $p(s = 1|x) \times p(x)$ is always smaller than $p(s = 0|x) \times p(x)$. The histograms are plotted on the same scale (left y-axis), showing clearly that the number of unlabeled observations is much smaller than the number of labeled observations. $p(s = 1|x)$ is always smaller than $p(s = 0|x)$, as $p(s = 1)$ is much smaller than $p(s = 0)$. This indicates that all observations will be classified as belonging to class $s = 0$.

After that, additional scaling is applied in order to estimate $p(y = 1|x)$. This additional scaling accounts for the fact that a part of the unlabeled observations is positive, which results in the partial overlap of the two distributions $p(s = 1|x)$ and $p(s = 0|x)$ (Figure 3.4). The scaling can be achieved by using Bayes' minimum risk

decision rule, which introduces a factor λ that represents the cost of assigning observation x to the MSS-class ($y = 1$), while observation x actually belongs to the non-MSS-class ($x \in (y = 0)$) (Figure 3.7). The decision rule

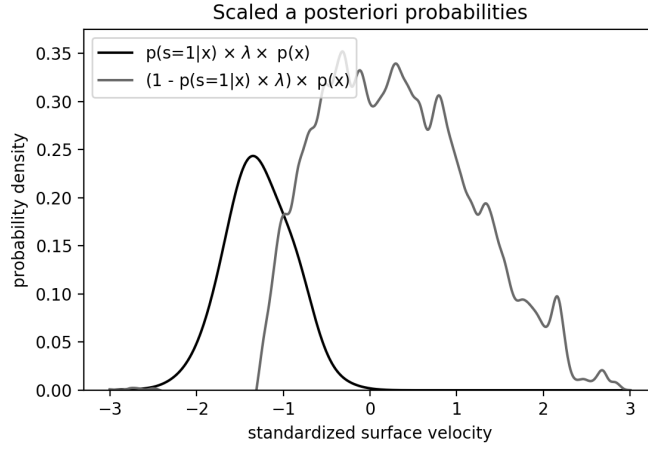


Figure 3.7: Estimated distributions of $p(y = 1|x)$ ($\approx p(s = 1|x) \times \lambda \times p(x)$) and $p(y = 0|x)$ ($= 1 - p(y = 1|x)$), using Bayes' minimum risk decision rule. The value for λ is chosen arbitrarily as 200.

then becomes:

$$\text{if } p(s = 1|x) \times \lambda > 0.5, \text{ then } x \in (y = 1). \quad (3.6)$$

Choosing different values for λ and applying these to the testing data results in different operating points on the ROC curve. With this ROC curve an optimal value for λ can be chosen (Section 3.3).

The value for λ can also be calculated for this specific problem of positive and unlabeled training data, using a theory presented by Elkan and Noto [2008]. They prove that the *a posteriori* probability $p(y = 1|x)$ can be estimated by:

$$p(y = 1|x) = p(s = 1|x)/c, \text{ where } c = p(s = 1|y = 1). \quad (3.7)$$

The value for c is in practice best estimated by e_1 :

$$e_1 = \frac{1}{n} \sum_{x \in P} p(s = 1|x), \quad (3.8)$$

where P is a subset of labeled (hence positive) observations of a validation set V [Elkan and Noto, 2008]. This theory holds only under the assumption that the labeled observations are a random subset of all positive observations [Elkan and Noto, 2008].

This specific value for the cost, $\lambda = \frac{1}{c} \approx \frac{1}{e_1}$, results in a specific operating point on the ROC curve (Section 3.3).

3.3. Receiver operating characteristic (ROC) curve

In the quantitative evaluation of the classification (Chapter 5), the receiver operating characteristic (ROC) curve is used. This curve shows the relation between the false positive rate (on the x-axis) and the true positive rate (on the y-axis), which is estimated by applying the classifier to test data. This test data contains meteorite finding locations (positive observations, Section 3.1.1), as well as locations known for their absence of meteorites (negative observations, Section 3.1.2). The classifier classifies the test observations as MSS or as non-MSS. Comparing the classified class to the actual class of the test data, results in an estimate of the true positive and the false positive rate (Figure 3.8).




Figure 3.8: Illustration of the theoretical false positive- and true positive rates. The distributions of the positive and negative test data are shown. When an observation from the positive distribution is classified as positive, it contributes to the true positive rate. When an observation from the negative distribution is classified as positive, it contributes to the false positive rate. The threshold determines the true positive and false positive rates. By varying the threshold, different true positive- and false positive rates are obtained, which are combined in a ROC curve (Figure 3.9).

A good classification has a high true positive rate, while maintaining a low false positive rate. In other words, the number of observations correctly classified as MSS has to be high, while the number of observations incorrectly classified as MSS has to be low. The threshold of the classifier can be set arbitrarily. Applying different thresholds results in different true positive- and false positive rates, from which the ROC curve is formed (Figure 3.9). Each point on the ROC curve is called an operating point. When the test data is easy to separate, the ROC curve is better. Graphically, this translates into ROC curves that are located in the upper left part of the graph (Figure 3.9).

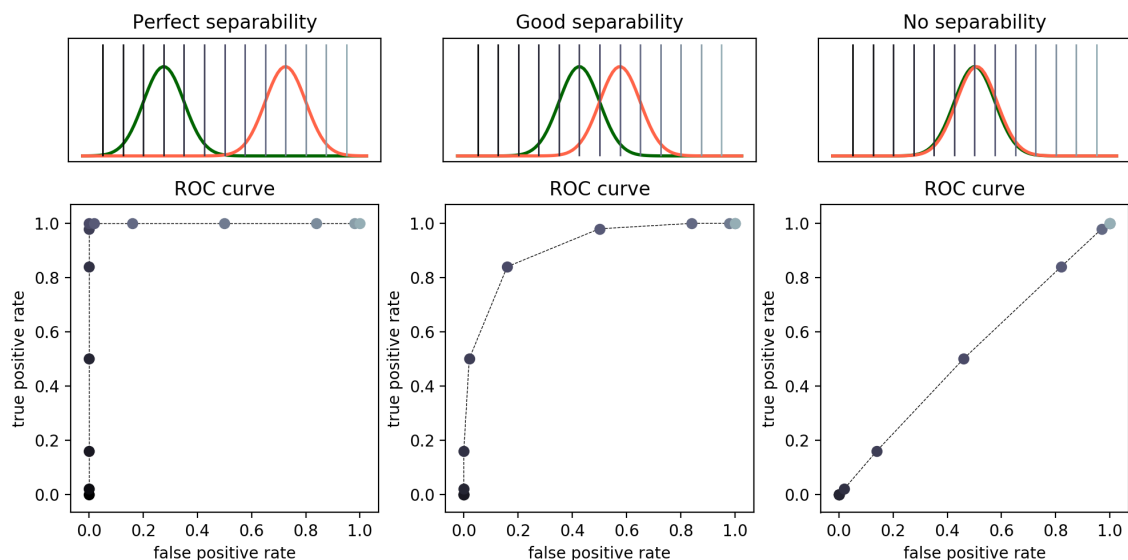


Figure 3.9: Three ROC curves obtained by applying different thresholds to the distributions of test data. Depending on the separability of the test data, the obtained ROC curve is better (left) or worse (right). The color scheme of the upper panels corresponds to the color scheme of Figure 3.8, i.e. positive test data in green, negative test data in red.

The ROC curve is used to evaluate the separability of the test data, which is influenced by the different features and the number of principal components used in the classification. Also, the difference between the two types of classifications can be investigated using the ROC curve. While evaluating these parameters,

the classification is tuned in order to have a better performance. Therefore, strictly speaking, a validation is performed (as opposed to an evaluation). Therefore, in this thesis, the testing data used to perform this validation is referred to as validation data.

The ROC curve strongly depends on the validation data. To mitigate this dependence, a cross-validation is performed. Observations of one field site are extracted from the training data and used as validation data. This is done recursively for the nine most productive field sites, and all remaining field sites combined (Table 3.1, Figure 3.10). The negative validation data always contains all negative observations to reduce the influence of possible positive examples in the negative data (Section 3.1.2). With this approach 10 ROC curves are obtained. To compare this set of 10 ROC curves to another set of 10 ROC curves, the true positive and false positive rates are combined into one ROC curve. This is done by calculating a weighted average, where the weights are defined by the number of meteorite finds on the field sites (NB this does not equal the number of observations on a field site, Table 3.1).

Field site (abbreviation)	Field site (full name)	no. of usable coordinates	no. of positive observations
QUE	Queen Alexandra Range	2564	233
MIL	Miller Range	2267	629
LEW	Lewis Cliff	1820	159
EET	Elephant Moraine	1740	317
GRV	Grove Mountains	1532	168
ALH	Allan Hills	962	440
MAC	Mac Alpine Hills	543	83
PCA	Pecora Escarpment	524	152
FRO	Frontier Mountain	492	78
rest	-	462	296

Table 3.1: Nine most productive field sites of the dataset with usable coordinates (Section 3.1.1). The numbers of usable coordinates is used as weights for computing the weighted average of the cross-validated ROC curves. The number of positive observations equals the number of grid cells containing one or more meteorite (Section 3.1.1), as used for training and validating the classifications. The colors represent geographic regions, indicating that field sites QUE, MIL, LEW and MAC are (relatively) close by each other, as well as that EET and ALH are close by each other (Figure 3.10, Appendix A).

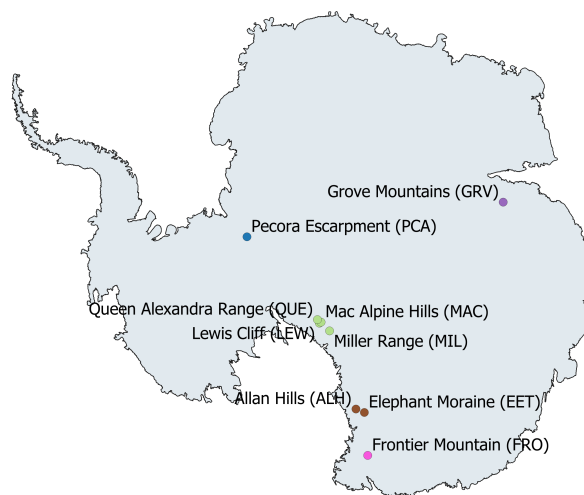


Figure 3.10: Approximate locations of the nine most productive field sites of the dataset with usable coordinates as listed in Table 3.1.

After tuning the classifier, a final estimate of the performance of the classifier is obtained by an actual evaluation, where independent, test data is used, gathered from additional field work reports. This independent testing data is presented in Appendix C.

4

Data

Meteorite stranding surfaces (MSSs) can be characterized with some typical features, despite the fact that individual MSSs show different concentration settings (Figure 2.1). The characterizing features quantify the concentration model as described in Chapter 2. Since the process of defining features can be relatively subjective, the starting point of each defined feature always lies in the current understanding of the concentration mechanism. This link between the quantitative features and the physical phenomenon is very important, as irrelevant features introduce noise in the classification. The features also need to be examined critically with regard to measurement uncertainties and/or biases that could negatively influence the classification.

The selection of features used in the classification is thus based on the concentration model and the quality of the available data. In examining the data, not only individual features are investigated, but also relevant correlations are investigated. The conclusions drawn on the features and their correlations are compared to the qualitative descriptions in the literature (Section 4.9).

The values of the features are extracted from geospatial datasets, where the original data is linearly interpolated to the locations of the labeled and unlabeled observations (except for the surface temperature, Section 4.7). All used datasets and their original resolution are summarized in Table 4.1. An overview of the eight features extracted/calculated from this data is shown in Figure 4.1.

Dataset	Resolution	Used for feature
RAMP AMM-1 SAR Image Mosaic of Antarctica, Version 2 [Jezek et al., 2013]	200 m	Radar backscatter
MEaSURES Phase-Based Antarctica Ice Velocity Map, Version 1 [Mouginot et al., 2019b]	450 m	Surface velocity
MEaSURES BedMachine Antarctica, Version 1 [Morlighem, 2019]	500 m	Ice thickness, Yield stress, Change of ice thickness along the flowline
Reference Elevation Model of Antarctica [Howat et al., 2019]	200 m	Surface slope, Yield stress, Change of ice thickness along the flowline
Rock Outcrop medium resolution v7.1 [ADD, 2019; Burton-Johnson et al., 2016]	NA	Surface slope, Distance to outcrops
MOD11A2 MODIS/Terra Land Surface Temperature Daytime, 8-Day Global, V006 [Wan et al., 2015]	1000 m	Surface temperature

Table 4.1: Datasets used for the definition of features. The resolution refers to the used resolution, some data is available at higher resolution.

4.1. Feature: Radar backscatter

BIAs where meteorites are found are located in the accumulation zone (i.e. surrounded by accumulating snowfields) and are sublimation-dominated (as opposed to melt-dominated) [Harvey, 2003; Harvey et al., 2001]. With radar backscatter, these type of BIAs can be distinguished from snow covered areas and (blue ice) areas where melt and refreezing takes place [Liu et al., 2006]. The radar backscatter at observation locations is extracted from a radar backscatter mosaic, composed from C-band synthetic aperture radar (SAR) backscatter intensity data collected in the austral winter season between September 9 and October 22, 1997 [Jezek et al., 2013]. As SAR can penetrate into thin snow layers [Jezek, 1999], temporary or seasonal snow covering BIAs [Brown and Scambos, 2004], does not influence the measurements to a large degree. As the 200 meter

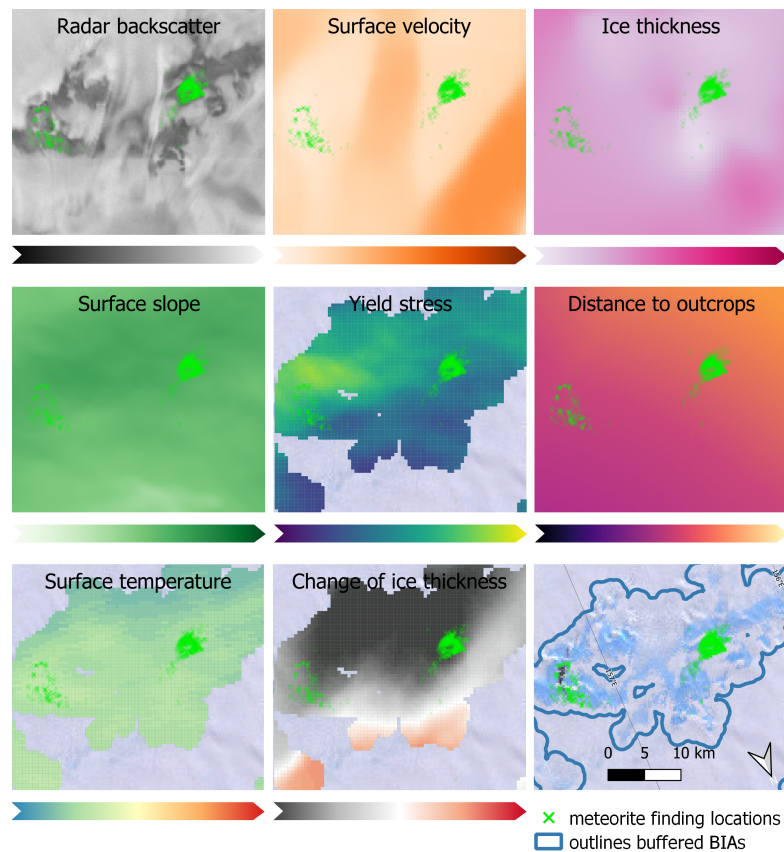


Figure 4.1: Overview of features displaying a snapshot of the geospatial data in the area of Elephant Moraine ($76^{\circ}17'S$, $157^{\circ}20'E$). Meteorite finding locations (Section 3.1.1) and the buffered BIA extent are shown in the lower right panel. References of the displayed datasets are given in the corresponding paragraphs (Section 4.1 to Section 4.8). The background image is the Center-Filled LIMA of the Landsat Image Mosaic of Antarctica (LIMA) Project.

resolution dataset contains integer gray-scale values between 0 and 255 representing the radar backscatter intensity [Jezek et al., 2013], extracting data at the grid locations results in duplicates (merely approximately 30,000 unique values). Duplicates form a problem in the cross-validation of the bandwidth of the Parzen density estimation, as simply the smallest bandwidth will result in the maximum log-likelihood (Section 3.2). To solve for this, random Gaussian noise with a standard deviation of 0.25 is added to the observations.




Figure 4.2: Histogram of estimated radar backscatter intensities of unlabeled observations (left axis) and of labeled observations (right axis).

Radar backscatter is lower at meteorite finding locations as compared to the values at all (buffered) BIAs (Figure 4.2). This can be explained by the fact that a large part of the unlabeled observations (BIAs + buffer) are located at low elevation in the percolation zone and wet snow zone of Antarctica, where surface melt occurs during summer [Liu et al., 2006]. The (re)frozen subsurface ice is effective in reflecting the signal and therefore the radar backscatter intensities are positively correlated with the surface temperatures (Figure 4.3) [Liu et al., 2006]. At colder temperatures, where melt is absent and meteorites are found, the linear relation between the surface temperature and the radar backscatter is lost. The intensity of the radar backscatter is now dominated by the surface roughness, and not by refrozen features [e.g. Campell, 1996]. Specular surfaces do not reflect the radar signal in the direction of the radar antenna. As a consequence, specular surfaces, such as BIAs, appear dark in the radar image (low backscatter), while snow in the accumulation area appears bright (high backscatter). The feature radar backscatter thus allows to distinguish between areas where melt and re-freezing takes place (high backscatter), snow (high backscatter), and specular blue ice areas (low backscatter).




Figure 4.3: Two dimensional histogram of radar backscatter versus surface temperature (Section 4.7).

4.2. Feature: Surface velocity

Due to the presence of sub- and supraglacial bedrock obstacles, ice flow rates at MSSs are typically low, ranging from centimeters per year to a few tens of meters per year, [Harvey et al., 2001]. The ice surface velocities give insight in the ice flow, as the velocities at the surface are similar to the depth-averaged velocities [Cuffey and Paterson, 2010; Zekollari et al., 2019].

The surface velocities at observation locations are extracted from a continent-wide map of Antarctic ice velocity developed by Mouginit et al. [2019a][Mouginit et al., 2019b]. For the surface velocity map, data from six different sensors, collected over a timespan of 25 years, has been used [Mouginit et al., 2019a]. The velocities in faster flowing areas are estimated using feature and speckle tracking techniques. Velocities lower than 10 m/year, covering 60 % of Antarctica, are estimated using InSAR techniques. The reported precision is 2-5 m/year for fast flowing areas, and 15-30 cm/year for slow flowing areas [Mouginit et al., 2019a]. The observations are log-transformed in order to reduce the skewness of the distribution.




Figure 4.4: Histogram of estimated surface velocities of unlabeled observations (left axis) and of labeled observations (right axis).

Surface velocity observations do confirm that the ice flow at MSSs is very low as compared to the velocities at the grid locations (Figure 4.4). The maximum observed velocity at MSSs is 12.5 m/year (Table 4.2).

4.3. Feature: Ice thickness

The interplay between subsurface topography and removal of ice through sublimation results in relatively thin ice at MSSs [Cassidy et al., 1992]. Ice thickness observations are extracted from the state-of-the-art subglacial bed topography map developed by Morlighem et al. [2020a][Morlighem, 2019]. At slow flowing areas, where meteorites have been found, the ice thickness has been computed by interpolation of radio-echo sounding observations, where the used interpolation method accounts for typical anisotropic patterns [Morlighem et al., 2020b]. The estimated error in these areas grows further away from the flight lines, with a value of 20m per km [Morlighem et al., 2020b]. As 50% of the ice sheet is further than 5 km from any measurement, the errors can become relatively large, with values of several 100 m to 1000 m [Morlighem et al., 2020a,b]. Also, this indicates that there is a spatial component to the error, which can influence the classification, as the training data is spatially not equally distributed.

Thin ice is presumably associated with BIAs in general, as the distribution of ice thickness at meteorite finding locations does not differ significantly from the distribution of thickness at the finding locations (Figure 4.5).

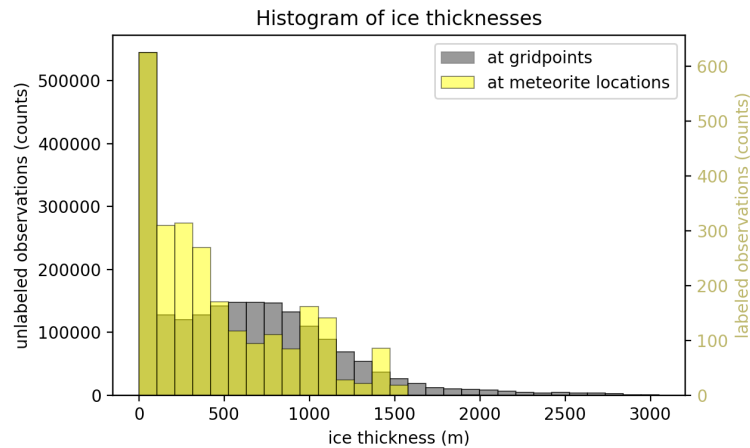


Figure 4.5: Histogram of estimated ice thicknesses of unlabeled observations (left axis) and of labeled observations (right axis).

4.4. Feature: Surface slope

The direction of the ice flow is directly related to the surfac slope [Cuffey and Paterson, 2010]. Consequently, the low ice flow rates at MSSs, expressed directly with the surface velocities (Section 4.2), can also be expressed indirectly by the surface slope, a quantity that like the ice thickness (Section 4.3) describes the geometry of the ice sheet. Surface undulations on a small scale do not govern the flow [Cuffey and Paterson, 2010]. Therefore the surface slope over larger distances is of interest. In the literature, the average surface slope is described as relatively gentle at most MSSs [Cassidy et al., 1992]. However, descriptions of the surface slope are inconclusive. For example, the surface slope of BIAs as in MSS model setting 3: open BIA (Figure 2.1) is described as steep, as katabatic winds are accelerated by the increasing surface slope, causing increased snowdrift transport and thus exposed blue ice [Bintanja, 1999; Takahashi et al., 1992].

The dataset used to calculate the surface slope is the Reference Elevation Model of Antarctica (REMA), at a 200 meter resolution [Howat et al., 2019]. Outcropping rocks are masked out, as the interest lies in the surface slope of the ice. For this, a classification of outcrops is used, provided by the Antarctic Digital Database [Burton-Johnson et al., 2016]. This dataset is discussed in more detail in Section 4.6. The values of the digital elevation model (DEM) at outcrop locations are replaced by NaN's. The remaining data represents ice surface elevation data. This data is filtered using an averaging filter with a circular shaped footprint. The circular shape is chosen in order to eliminate sensitivity to the orientation of the axes. The diameter of the footprint is 15 km, which corresponds to eight times the mean ice thickness at buffered BIAs¹.

After masking and smoothing the DEM, the maximum surface slope is calculated as:

$$\text{slope} = \sqrt{\text{slope}_x^2 + \text{slope}_y^2}. \quad (4.1)$$

The surface slopes in x- and y-direction are approximated using the central difference:

$$\text{slope}_{x_i} = \frac{\delta \text{DEM}}{\delta x} \approx \frac{\text{DEM}_{x_{i+1}} - \text{DEM}_{x_{i-1}}}{2h}, \quad (4.2)$$

where h is the distance between two data points, 200 meters in this case. The relative error of the high resolution (10 m) DEM is in the order of decimeters. By applying averaging filters this error results in almost negligible uncertainties in the estimation of the surface slope (in the order of 10^{-4} m/m). Errors in the masking of outcropping rocks (Section 4.6) can result in biased estimates of the ice surface slopes. However, these biases are (partly) eliminated, as the averaged elevation of neighboring pixels that are subtracted in the central difference (Equation (4.2)) are both influenced by a misclassified outcrop.

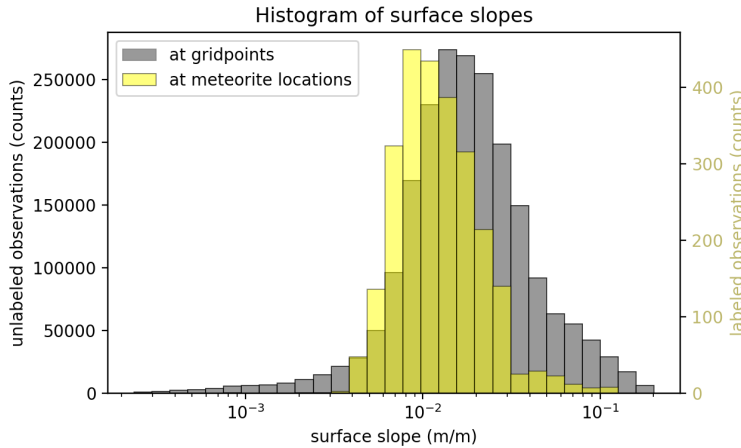


Figure 4.6: Histogram of average surface slopes of unlabeled observations (left axis) and of labeled observations (right axis).

Surface slopes at meteorite finding locations are slightly less steep than the surface slopes at the grid locations (Figure 4.6). The observations are log-transformed in order to reduce the skewness of the distribution.

¹Considering spatial scales corresponding to several ice thicknesses is a common approach in glaciology to account for the effect of non-local (longitudinal) stresses [e.g. Cuffey and Paterson, 2010; Mougintot et al., 2019a; Zekollari et al., 2019].

4.5. Feature: Yield stress

The surface slope is correlated to the ice thickness (Figure 4.7), as the product $H \times \alpha$ is constant under the assumption of perfect plasticity of the ice [Cuffey and Paterson, 2010]. This relation holds for the grid values over all buffered BIAs. However, for the meteorite finding locations the relation does not hold, as the ice is relatively thin given the value for the surface slope, or the slope is relatively flat given the value for the ice thickness (Figure 4.7). This abnormality is not described in the literature. To further explore this, a feature called 'yield stress' is defined.




Figure 4.7: Two dimensional histogram of ice thickness versus surface slope.

Gravitational forces acting on the ice sheet form a driving stress on the ice. The driving stress is balanced by resisting forces, which are the sum of the basal drag, wall drag and longitudinal drag. The magnitude of the driving stress is influenced by the ice thickness and the surface slope. Consequently, the magnitudes of the balancing resisting forces are also controlled by the ice thickness and surface slope. The magnitude of the resisting forces depends on the material properties of the ice and the bed, as well as on the flow velocity. The ice flow is thus governed by the ice thickness and surface slope, in combination with the properties of the ice and the bed. Under the assumption of perfect plastic deformation, the relation between the ice thickness (H), the surface slope (α), the stresses acting on the ice, and the material properties of the ice can be expressed as:

$$H = \frac{1}{f'} \frac{\tau_0}{\rho_i g \alpha}. \quad (4.3)$$

The material properties are expressed with the yield stress (τ_0), and the density of ice (ρ_i). The stresses acting on the ice are expressed with the gravitational acceleration (g), and f' , which denotes the relation between the basal drag τ_b and the driving stress τ_d ,

$$\tau_b = f' \tau_d. \quad (4.4)$$

Assuming constant values of 1 for f' , 917 kg/m^3 for ρ , and 9.81 m/s^2 for g , the yield stress can be calculated as:

$$\tau_0 = \rho g H \alpha. \quad (4.5)$$

As this feature is related to the ice flow, the ice thickness is averaged over a horizontal scale of several times the ice thickness, as to eliminate the irrelevant influence of small scale undulations in ice thickness (see also Section 4.4). The ice thickness, H , is filtered using an averaging filter with a circular shaped footprint. The circular shape is chosen in order to eliminate sensitivity to the orientation of the axes. The diameter of the footprint is set to 14.5 km^2 , which corresponds to approximately eight times the mean ice thickness at buffered BIAs. The surface slope α is the average surface slope (Section 4.4).




Figure 4.8: Histogram of estimated yield stress of unlabeled observations (left axis) and of labeled observations (right axis).

The typical value for τ_0 reported in the literature is 100 kPa [e.g. Cuffey and Paterson, 2010], and this is confirmed in the observations (Figure 4.8). The yield stress at meteorite finding locations is lower than at the grid cell locations. Despite the averaging filter applied to the ice thickness, the data may be subject to substantial biases influencing the estimates of the yield stress. Extra caution is required, as the feature is log-transformed in order to reduce the skewness of the distribution. This amplifies uncertainties/biases in the lower values of the yield stress. The correlation between the yield stress and the features 'ice thickness' and 'surface slope' is not very strong, despite the fact that the yield stress is calculated from the average ice thickness multiplied with the average surface slope. The correlation coefficient with the ice thickness (Section 4.3) is 0.53, and the correlation coefficient with the average surface slope (Section 4.4) is -0.26. This indicates that the feature 'yield stress' may contain additional information for the classification.

4.6. Feature: Distance to outcrops

Many MSSs are located close to bedrock outcrops (Figure 2.1). To investigate this relation, a feature expressing the distance to the closest outcrop is defined. To compute this feature, a classification of exposed bedrock is used [ADD, 2019; Burton-Johnson et al., 2016]. Buffers of increasing length (exponentially increasing from 400 m to 1500 m in 750 steps) are drawn around these polygons and rasterized with a spatial resolution of 500 m. A value of the summation of the previous buffers and half of the new buffer is assigned to the rasterized buffer. From these rasterized points the values for the observations are calculated. As the distance data is discrete, the number of unique values of the data is approximately 90,000. The large number of duplicates forms a problem in the cross-validation of the bandwidth of the Parzen density estimation, as simply the smallest bandwidth will result in the maximum log-likelihood (Section 3.2). To solve for this, random Gaussian noise with a standard deviation of 0.1 km is added to the observations. After that, the observations are log-transformed in order to reduce the skewness of the distribution. Although the used method to compute the distance to the nearest outcrop is computationally efficient, it is not as exact as calculating the shortest distance to exposed bedrock for every grid cell.

²This diameter is slightly smaller than the 15 km diameter used for averaging the surface elevation (Section 4.4). This is because the kernel of the averaging filter has to be symmetric in order to avoid biases. Consequently, only values equal to an odd number times the resolution can be used as diameter. As the resolution of the surface elevation data and the ice thickness data are not equal (Table 4.1), the diameters of the averaging filters are slightly different.

A more severe limitation of the obtained data lies in the sensitivity to the quality of the classification of exposed bedrock. Isolated exposed bedrock that has been missed in this classification, or an isolated ice- or snow patch misclassified as outcrop, can result in large biases in the observations. The accuracy of the bedrock classification is estimated pixelwise to be 74% for the area North of 82°40'S, and 39% for the area South of 82°40'S, where no Landsat 8 imagery is available [Burton-Johnson et al., 2016]. The (manual!) classification for areas South of 82°40'S has a low estimated accuracy, as the georeferencing is poor, shaded snow is misclassified as rock, and areas of exposed rock are overestimated. Of these three problems, only the misclassification of shaded snow as rock can result in large biases in the distance to exposed bedrock. The classification for areas North of 82°40'S often misses small outcrops of less than 60 m, also outcrops resulting pixels with similar spectral properties as clouds and snow are missed, and glacial debris over glaciers is misclassified as outcrops. Although the overall estimated accuracy of that data is higher, all three problems can result in large biases in the distance to exposed bedrock.




Figure 4.9: Histogram of estimated distances to outcrops of unlabeled observations (left axis) and of labeled observations (right axis).

Despite possible biases, the vast majority of observations is estimated to be within 100 km of exposed bedrock (Figure 4.9). There is no convincing difference between the distributions of labeled and unlabeled observations. That implies that the vicinity of exposed bedrock is related to blue ice areas in general, and not to meteorite concentrations.

4.7. Feature: Surface temperature

The temperature gives an indication of whether a site may contain meteorites or not, as high surface temperatures are unfavorable for meteorite concentrations (Section 2.1.3). The multi-annual mean surface temperature at the observation locations is estimated using a 8-day average surface temperature product derived from MODIS observations [Wan et al., 2015]. In retrieving this 8-day average surface temperature, observations on days where clouds were present are masked out using band 26 of the MODIS instrument for cirrus detection [Wan, 1999]. If the cloud masking has a consistent spatial character and/or a correlation to the surface temperature, this could lead to biased estimates for the multi-annual mean surface temperature. For example, in areas where there are persistent cloud covers in winter, the annual mean surface temperature will be overestimated. This effect, however, is assumed to be negligible, as 90% of the observations contain over 788 8-day average surface temperature observations (of a total of 872 epochs) to compute the multi-annual mean. To reduce the influence of outliers and/or the absence of data the multi-year mean temperature is calculated over a long time period, ranging from 01-01-2001 to 01-01-2020.

The observations confirm that the temperatures at meteorite finding locations are low (Figure 4.10). Harvey et al. [2001] state that MSSs are to be found above 1500 meter, as the temperatures are low there. The annual mean temperature of observations around 1500 meter is estimated from the data as -30 °C. The observations are thus in agreement with the values indicated by Harvey et al. [2001]. The histogram (Figure 4.10) shows clearly that a substantial part of the unlabeled observations can not act as MSS, as the temperatures there are too high.

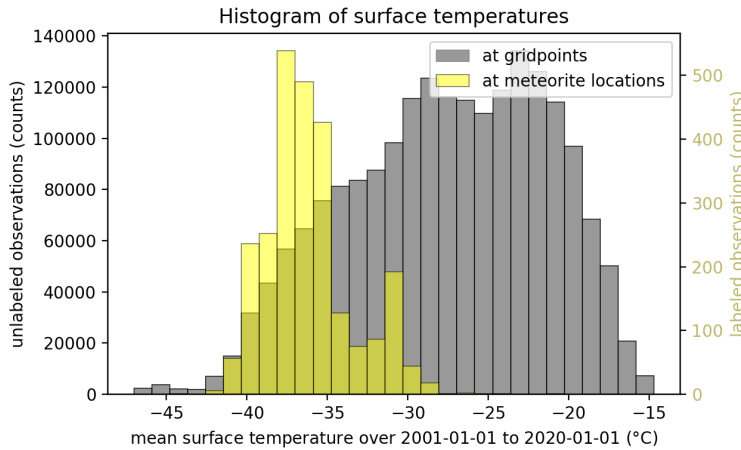


Figure 4.10: Histogram of estimated mean surface temperatures of unlabeled observations (left axis) and of labeled observations (right axis).

Before using the observations for the classification, Gaussian noise with a standard deviation of 0.04 °C is added to the observations, as the data contains many duplicates. The duplicates exist because the 1000 m resolution data has been interpolated using the nearest neighbor method to extract the values at grid locations and meteorite finding locations. Duplicates form a problem in the cross-validation of the bandwidth of the Parzen density estimation, as simply the smallest bandwidth will result in the maximum log-likelihood (Section 4.1). For the added Gaussian noise, a standard deviation of 0.04 °C is used, a conservative value based on the indicated accuracy of the individual measurements and the fact that averaging improves this accuracy.

The temperature is strongly correlated to the surface elevation (correlation coefficient of -0.86). However, the two quantities cannot be used interchangeably, as there is a relevant regional factor. In the Transantarctic mountains, where many meteorites are found, temperatures are relatively low considering the elevation. To find similar temperatures in other areas you would have to go to higher elevations. As a consequence, replacing the temperature, or combining it with elevation data, would result in biased training data, as the training data is not equally distributed over the Antarctic continent. Only the temperature itself is the best indication for the presence of meteorites.

4.8. Feature: Change of ice thickness along the flow line

MSSs often show upstream thinning of the ice along the flow line (Figure 2.1). The change of ice thickness along the flow line is estimated by calculating the difference between a filtered ice thickness at the location of the observation and a point 20 km upstream, (Figure 4.11). The ice thickness data (Section 4.3) is averaged




Figure 4.11: Illustration of the calculation of the change of ice thickness along the flow line. The arrow indicates the direction of the flow line, estimated using the maximum slope of the average filtered surface elevation data, for which the 15 km circular kernel is displayed in pink. The ice thickness at the two observation locations, displayed as black dots, is also filtered with an averaging filter with circular kernel of 14.5 km, displayed with diagonal shading. From the difference between the two average thicknesses, the change of ice thickness along the flow line is computed.

using a circular shaped footprint of diameter 14.5 km (Section 4.5). As the ice flow direction is parallel to the steepest average surface slope [Cuffey and Paterson, 2010], the location of the point 20 km upstream is calculated using the direction of the average surface slope (Section 4.4). The approach is a simplification by assuming locally straight flow lines. Despite the reduction of measurement uncertainties by using averaging filters, the estimates are subject to possible biases (Section 4.3, Section 4.4).

The approximation of the change of ice thickness over larger scales shows that in general ice is thinning along the flow line at blue ice areas, and that at MSSs this thinning is more pronounced (Figure 4.12).

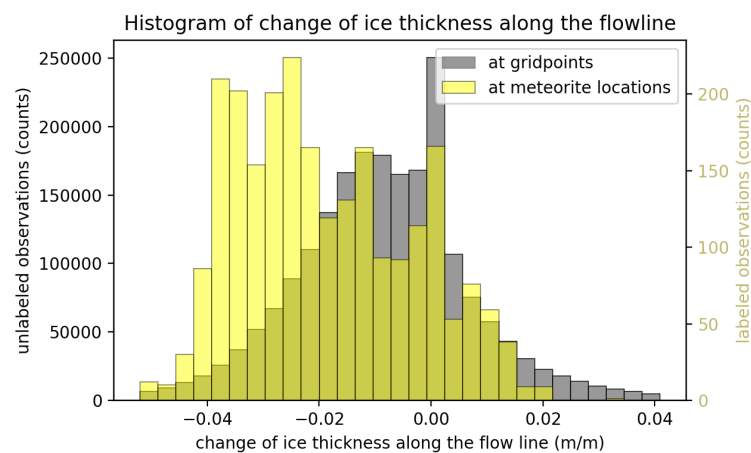


Figure 4.12: Histogram of estimated yield stresses of unlabeled observations (left axis) and of labeled observations (right axis).

4.9. Comparison to the literature

The summary statistics of the eight features are summarized in Table 4.2. In general, the values of the described features are in good agreement with the qualitative descriptions in the literature.

- The radar backscatter does confirm that MSSs are exposing blue ice, while absent of melt.
- The surface velocities do confirm that ice flow rates are low at MSSs, with observed values up to 12.5 m/year.
- Both the ice thickness and the distance to outcrops seem to be features that can be related to blue ice areas in general and do not clearly show to be a distinguishing feature of MSSs.
- Average surface slopes at MSSs are similar to the average surface slopes at BIAs.
- The slopes are relatively flat given the ice thickness, a relation that is not described directly in literature. This phenomenon is expressed in the feature 'yield stress'.
- The surface temperatures do confirm that the temperatures are low at MSSs, with multi-annual mean observed values of -36 °C on average.
- The change of ice thickness along the flow line indicates that at most BIAs, the ice is constant or thinning along the flow line, and this thinning is stronger at MSSs.

Besides the eight features described here, based on the literature, other features such as the ice flux, the divergence curl of the ice flow, and the snow grain size have been investigated. They are not considered further for different reasons. The flux strongly correlates with the surface velocity (correlation coefficient of 0.86) and consequently no new information is added by this feature. The divergence and the curl describe the compression of flow, and stress and strain rates. Both quantities are sensitive to the uncertainties of the surface velocities, as derivatives of the velocities are considered. Not only the magnitude of the velocity is needed in calculating the divergence and the curl, but also the direction of the velocity. After visual inspection of the velocity data, the direction of the surface velocity did not appear to follow the direction of the maximum

Feature	Set of observations	Mean	Median	Std Dev	Range
Radar backscatter (-) (Section 4.1)	Unlabeled	183.47	184.65	33.61	1.6 - 255.0
	Labeled	132.65	130.43	27.85	42.25 - 232.15
Surface velocity (m/year) (Section 4.2)	Unlabeled	64.7	10.9	168.64	0.0 - 2561.9
	Labeled	1.49	1.13	1.08	0.05 - 12.47
Ice thickness (m) (Section 4.3)	Unlabeled	587.79	522.78	534.59	0.0 - 4342.35
	Labeled	452.96	323.83	411.71	0.0 - 1931.97
Surface slope (cm/m) (Section 4.4)	Unlabeled	2.48	1.73	2.51	0.0 - 57.36
	Labeled	1.45	1.16	1.11	0.08 - 11.48
Yield stress (kPa) (Section 4.5)	Unlabeled	80.81	79.12	60.77	0.0 - 1898.68
	Labeled	46.9	45.63	27.47	0.23 - 152.17
Distance to outcrops (km) (Section 4.6)	Unlabeled	26.68	8.22	50.52	0.0 - 577.82
	Labeled	16.22	4.11	23.0	0.0 - 86.44
Surface temperature (°C) (Section 4.7)	Unlabeled	-27.6	-27.19	6.26	-48.92 - -7.37
	Labeled	-36.13	-36.59	2.73	-44.4 - -20.28
Change of ice thickness along the flow line (cm/m) (Section 4.8)	Unlabeled	-0.83	-0.82	1.71	-12.45 - 16.97
	Labeled	-2.15	-2.29	1.88	-8.52 - 3.45

Table 4.2: Summary statistics (mean, median, standard deviation (std dev), and range) of the eight features. The set of unlabeled observations contains both MSSs and non-MSSs, whereas the set of labeled observations contains only MSSs. Units are indicated in the first column. Corresponding histograms are shown in the corresponding Sections.

surface slope in some areas. This, in combination with the given precision of the data, makes estimates of the divergence and curl unreliable. Lastly, the snow grain size has been extracted from an Antarctic-wide mosaic of snow grain sizes, created with MODIS data collected between 1 November 2008 and 28 February 2009 [Haran et al., 2014; Scambos et al., 2007]. The same product with data collected in the austral summer of 2003-2004 has been used by Hui et al. [2014] to classify blue ice areas South of 82.5°S. This classification of blue ice areas has been used to define the locations of unlabeled observations (Section 3.1.3). The snow grain size is estimated by calculating the normalized difference of MODIS band 1 (~647 nm) and MODIS band 2 (~857 nm) [Scambos et al., 2007]. This value correlates to the grain size, as snow reflectivity decreases in the infrared band when snow grain sizes are larger [Scambos et al., 2007]. The translation from normalized difference ratio to snow grain size is made using a look-up-table [Scambos et al., 2007]. Coarse snow grain sizes are related exposed blue ice as the surface can be seen as very large snow grains. Scambos et al. [2007] mentions a threshold of 400 μm to estimate the blue ice area in a region, and this value has been applied by Hui et al. [2014]. However, this threshold is not clearly reflected in the extracted values at meteorite finding locations and the grid cells over the blue ice areas. After visual inspection of the data, sharp transitions of snow grain sizes are found at blue ice areas, for which a physical explanation is not obvious. Also, the link between the estimated snow grain size and the physical phenomenon of the concentration mechanism is only addressed in a fieldwork report that mentions that the lack of meteorites on the Far Northern ice patches in the Allan Hills area can be explained by the fact that the ice has a firm-like surface [Huss et al., 1988]. The capacity of the snow grain size to distinguish snow covered areas, is already covered by the feature 'radar backscatter'. Hence, there is no added value in evaluating the influence of the snow grain size on the classification.

In conclusion, the eight selected features indicate that positive observations (MSSs) differ from unlabeled observations (MSSs and non-MSSs). The influence of individual features on the classification is investigated in Chapter 5, in which the classification is presented.

5

Results

The features defined in Chapter 4 are combined in a classification algorithm (Section 3.2) to classify unlabeled observations as MSS or non-MSS. Before performing the final classification to obtain a meteorite hotspot map (Figure 5.1), the influence of the input data on the classification has been analyzed. First, the influence of the individual features on the classification is investigated (Section 5.1), based on which several features are removed for further analyses. Subsequently, a principal component analysis is performed on the remaining features and differences between the positive observations and the unlabeled observations are investigated (Section 5.2). The influence of the number of principal components used as input for the classifier and the influence of the definition of the positive observations is evaluated (Section 5.3). Finally, the classification is optimized according to the findings and a meteorite hotspot map is created (Section 5.4), which is evaluated with independent test data.




Figure 5.1: Antarctic meteorite hotspot map. A high resolution overview map, as well as local maps are provided in Appendix E

5.1. Feature selection

The influence of the individual features on the classification is investigated by comparing cross-validated ROC curves (Section 3.3) obtained by classifications with different (sub)sets of features as input data. There are 255 possible combinations of the eight features (Chapter 4). Performing 255 classifications in an exhaustive feature selection method is costly and optimizing the classification in following this method may lead to overfitting to the validation data [Guyon and Elisseeff, 2003]. Therefore, the influence of the features is investigated with a sequential forward selection (SFS) method, and a sequential backward selection method (SBS).

In the SFS method (Section 5.1.1), the classification is performed recursively with an increasing amount of features. The first classification is based on a single feature, and the best predicting feature is selected. After that, the feature that improves the classification based on a single feature the most, is added. This is repeated until none of the features improves the classification. A drawback of this method is that possible correlations between features that contribute to the separability of the classes can be missed, as only one additional feature is evaluated per iteration.

In the SBS method (Section 5.1.2), the classification is performed recursively with a decreasing amount of features. The classifiers are trained with all but one feature, and the results are compared to a classifier that is trained using all features. This procedure is repeated until none of the features affects the performance of the classification negatively. As long as the dimension of the input data is larger than five, the dimensionality is reduced to five, using a principal component analysis (Section 3.2). This implies that for some of the analyses, not all variance of the data is considered. However, this does not influence the results to a large degree, as in the worst case 86.9% of the variance of the data is considered (Figure D.5).

The addition or removal of features is based on a visual inspection of the ROC curves. Despite the fact that a quantitative evaluation of the ROC curves with a performance measure would provide a less objective selection procedure, the selection of the performance measure itself introduces objectivity in the results. Namely, the selected performance measure optimizes only on a single aspect of the ROC curve. For example, when the performance measure is the area under the curve, maximizing the area under the curve when curves cross, does not necessarily result in the best ROC curve. A visual evaluation of the ROC curves allows a simultaneous comparison of the area under the curve, the smoothness, and the steepness of the curves. This approach is feasible as there are only eight features. Based on the visual inspection, a feature is selected or rejected. The motivation for selection or rejection are briefly described (Section 5.1.1, Section 5.1.2). All ROC curves of intermediate steps are provided in Appendix D.1.

To evaluate the influence of the features on the classifications, the results of the SFS and SBS are used in combination with a qualitative summary of the features (Section 5.1.3). In Figure 5.2 the features are placed in a space spanned up by the connection to the physical phenomenon and the measurement uncertainties and/or biases. This qualitative summary (Figure 5.2) must be seen as a sketch, as placing the features objectively in this space is not practicable.

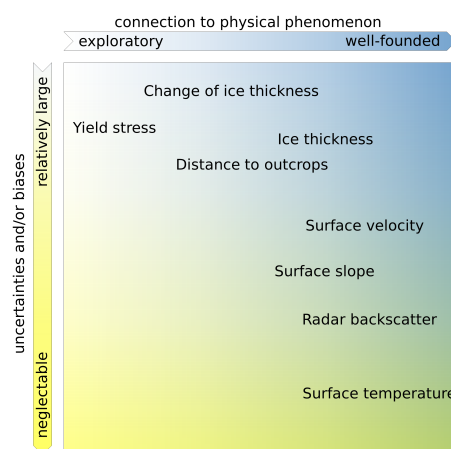


Figure 5.2: Qualitative summary of features, illustrating observation uncertainties and the link to the physical phenomenon of MSSs.

5.1.1. Sequential forward selection (SFS)

As a starting point for the sequential forward selection, the ROC curves obtained by performing a classification based on a single feature are evaluated. Not all individual features can discriminate equally well between MSSs and non-MSSs, as the ROC curves of the classifications using a single feature are not aligned (Figure 5.3). Features with similar distributions of the labeled and unlabeled observations, such as the ice thickness, the surface slope and the distance to outcrops, as well as features with a similar range, such as the radar backscatter, the yield stress, and the change of ice thickness along the flow line, perform less well than the surface velocity and surface temperature features (Figure 5.3).

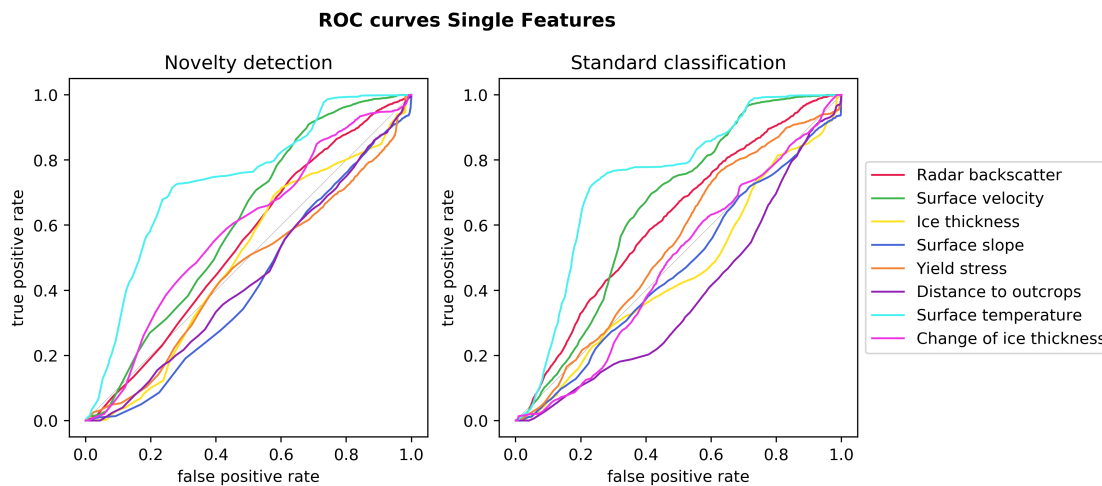


Figure 5.3: ROC curves when training the classifiers with a single feature. The eight curves are labeled with the name of the single feature.

The classifications based on the surface temperature result in the best separability between the classes (Figure 5.4, **iteration 0**). Adding surface velocity to the training data does improve the separability for almost all thresholds (Figure 5.4, **iteration 1**).




Figure 5.4: ROC curves when training the classifiers with an increasing amount of features. The three curves are labeled with the number of the iteration and the additional feature in that iteration.

A strongly irregular ROC curve indicates that the distributions of the training data contain multiple peaks in the area where they overlap. These peaks can be related to observations of individual sites. A smoother curve indicates that the sensitivity of the classifier to individual sites is reduced. Adding the radar backscatter to the classification as a third feature, results in a smoother curve, despite the fact that it does not improve

the separability for all thresholds (Figure 5.4, **iteration 2**). Adding any additional feature to the classification based on the surface velocity, the surface temperature, and the radar backscatter, does not improve the ROC curve.

5.1.2. Sequential backward selection (SBS)

As a starting point for the sequential backward selection, the ROC curve obtained through a classification with eight features (Figure 5.5, **iteration 0**) is compared to the eight ROC curves obtained with seven features, where one of the eight features is left out alternately. From this, it appears that the change of ice thickness has a negative influence on the classifications (Figure 5.5, **iteration 1**).

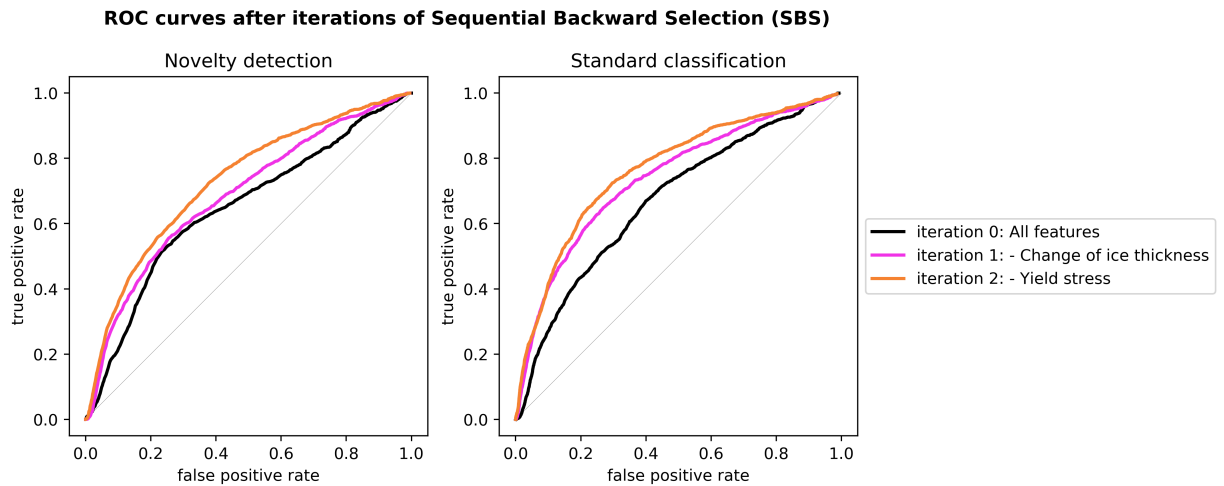


Figure 5.5: ROC curves when training the classifiers with a decreasing number of features. The three curves are labeled with the number of the iteration and the discarded feature in that iteration.

In the second iteration of the SBS the ROC curve obtained through a classification with seven features is compared to seven ROC curves obtained with six features, where one of the seven features is left out alternately. From this, it appears that the yield stress has a negative influence on the classifications (Figure 5.5, **iteration 2**), and consequently this feature is rejected. Repeating the same procedure with six features shows that all remaining features contribute to the standard classification.

5.1.3. Comparison between SFS and SBS

The SFS method and the SBS method do converge, as there are no features rejected by the SBS method that have been selected in the SFS method (Figure 5.6). However, the resulting subsets of features are not the

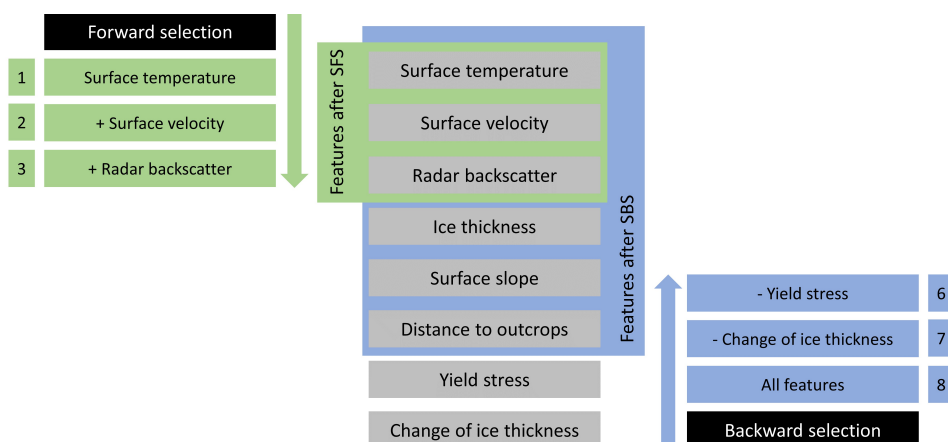


Figure 5.6: Subsets of features selected by the SFS and SBS method. The number of features in each iteration is indicated. The methods converge, but do not result in the same subset of features.

same, which (i) might be an indication that the classifiers do not generalize well, (ii) may be due to the fact that the influence of features are not re-evaluated in the SFS and SBS method, or (iii) may directly relate to the test data. The straightforward conclusion that features selected in both methods are beneficial to the separability of the classes, whereas features rejected in both methods do negatively influence the separability of the classes, is therefore not airtight. However, the rejection by the SBS of the 'change of ice thickness along the flow line', and the 'yield stress' is supported with respect to the relation to the physical phenomenon and the sensitivity to measurement uncertainties and/or biases (Figure 5.2). Therefore, these two features are removed for the further analyses.

ROC curves after Sequential Forward Selection (SFS) and Sequential Backward Selection (SBS)

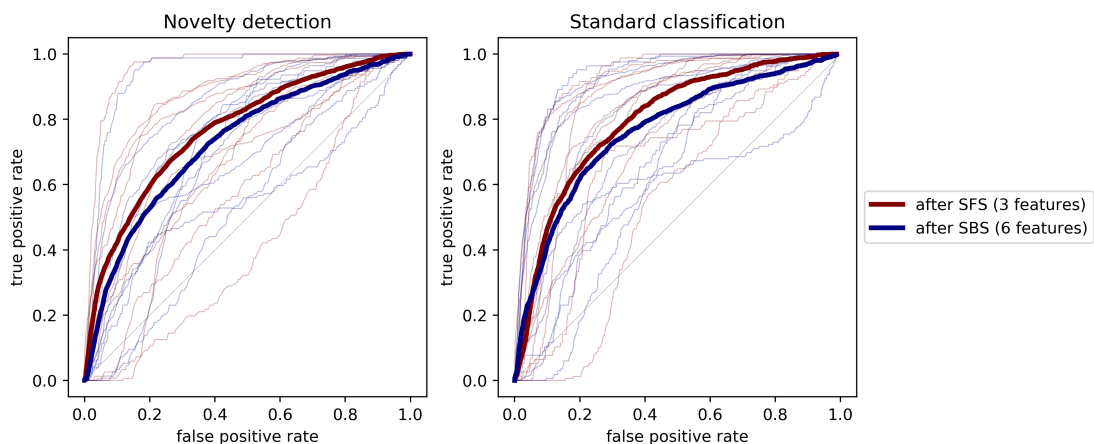


Figure 5.7: ROC curves after the sequential forward selection (SFS) of features and the sequential backward selection (SBS) of features. ROC curves of individual sites (thin lines) from which the weighted average ROC curve (bold line) is computed are plotted for comparison.

Despite the fact that the influence of the remaining features on the classification is not indisputably positive, further analyses (Section 5.2, Section 5.3) are performed using the set of six features for two reasons. (i) Firstly, the individual ROC curves obtained in the cross-validation using the six features after the SBS show a narrower spread for lower values of the false positive rate as compared to the three features after the SFS (Figure 5.7). This indicates that the classifiers with six features are more stable. This stability is favored over a slightly better average performance of the classifiers, as it allows a detection of all different field sites. (ii) Secondly, the negative test data is not representative for all locations that are not a MSS (Section 3.1.2). Therefore, subgroups in the unlabeled observations might be misclassified as MSS, but not detected. This problem is less likely to occur when more descriptive features are used.

5.2. Principal component analysis (PCA)

With a principal component analysis (PCA), the original features are combined to create new variables that describe the data. The new variables, referred to as principal components (PCs), are linear combinations of the original data and are uncorrelated. The goal of a PCA is to find a subset of variables that describe the data, in order to reduce the number of dimensions of the data [Webb, 2002]. A reduction of the dimensionality of the data is desirable, as the dimensionality of the data, together with the number of observations, imposes a limitation to the accuracy of the estimated density distributions (a problem often referred to as 'the curse of dimensionality') [Silverman, 1986]. In order to obtain an accurate density estimation with 2,554 observations (i.e. the number of positive observations), the dimensionality of the data is limited to five [Silverman, 1986]. The PCs are ranked based on how much of the variation of the original data they describe (Figure 5.8).

Together, all PCs account for all variation of the original features (100%). Each principal component derived from the six features is a linear combination of these six features (Table 5.1). The interpretation of the individual PCs is difficult, as there is no obvious link between the PCs and the physical phenomenon. Nevertheless, a PCA as unsupervised learning technique does allow to see groupings within the data. In this case, a two dimensional histogram plot of the first two components does shows that there are three groups visible in

the unlabeled data (Figure 5.9). Also, it is clear that the values of the PCs of the unlabeled data differ from the values of the PCs of the positive labeled training data. Two dimensional histogram plots for the other PCs are provided in Appendix D.2.

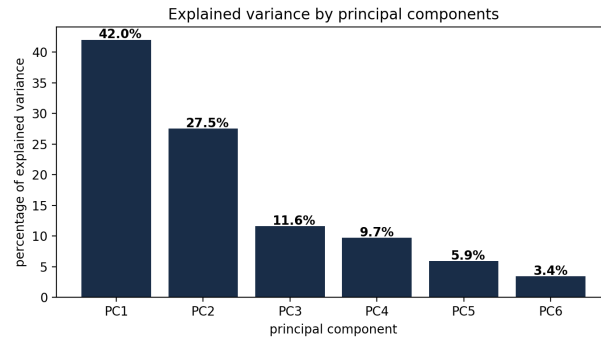


Figure 5.8: Explained variance by the principal components.

PC1	PC2	PC3	PC4	PC5	PC6
0.55 Dist. to outcrops	-0.62 Surface temp.	-0.83 Radar backscatter	-0.63 Surface velocity	-0.76 Dist. to outcrops	0.78 Ice thickness
0.55 Ice thickness	-0.49 Radar backscatter	0.51 Surface temp.	0.53 Surface temp.	-0.53 Surface slope	0.35 Surface slope
-0.44 Surface slope	-0.43 Surface velocity	0.23 Surface velocity	-0.51 Surface slope	0.36 Surface velocity	-0.32 Dist. to outcrops
0.38 Surface velocity	-0.37 Surface slope	0.09 Dist. to outcrops	-0.19 Ice thickness	0.06 Ice thickness	-0.31 Surface velocity
0.21 Radar backscatter	0.24 Ice thickness	-0.02 Surface slope	0.18 Radar backscatter	0.04 Radar backscatter	0.28 Surface temp.
0.08 Surface temp.	0.03 Dist. to outcrops	-0.02 Ice thickness	0.07 Dist. to outcrops	0.02 Surface temp.	0.02 Radar backscatter

Table 5.1: Linear combinations of original features from which the principal components are computed.

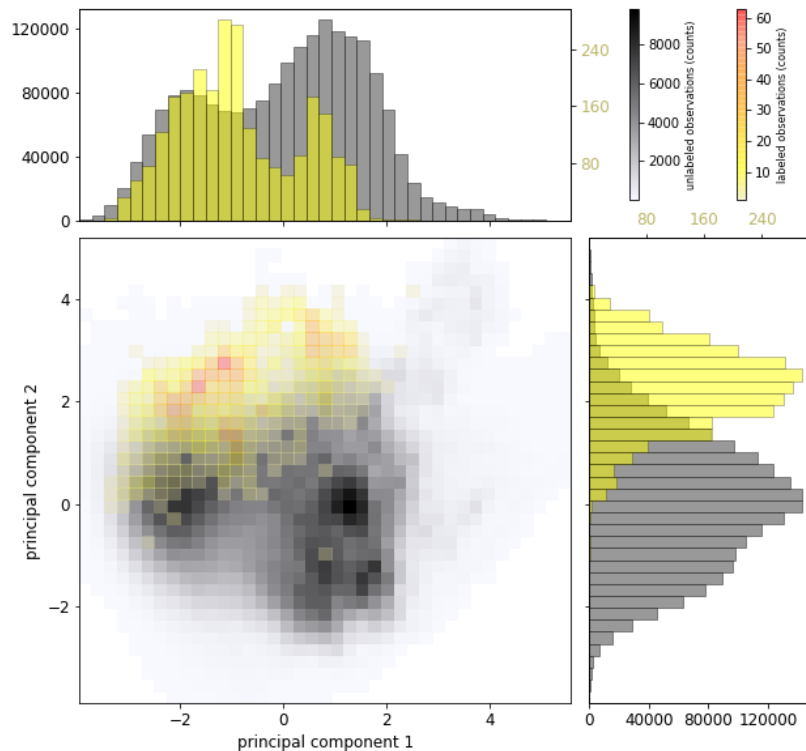


Figure 5.9: Two dimensional histogram of the first principal component and the second principal component, clearly indicating that the unlabeled data differs from the positive labeled data.

The number of PCs used as input to the classifications influences the results (Figure 5.10). The best ROC curve is obtained by using all five PCs. The best ROC curve is slightly better when considering the standard classification as compared to the novelty detection (Figure 5.10).




Figure 5.10: ROC curves when training the classifiers with a different number of principal components. Subsets of PCs are always defined so that the selected PCs explain the maximum variance (Figure 5.8). The ROC curves of the left panel are indicated in the right panel using dashed lines, and the other way around.

5.3. Sensitivity analysis

To evaluate the sensitivity of the classification to the definition of the positive observations, the individual ROC curves obtained in the cross-validation are used. The 10 different sets of positive validation data corresponding to 10 different field sites (Section 3.3), result in 10 different ROC curves. These are not all equally well, indicating that some sites are easier to classify than others (Figure 5.11).




Figure 5.11: ROC curves when training the classifier with the first five PCs obtained in a PCA on six features (Section 5.1). Alternately, the observations at a single field site (as listed in the legend) are kept apart from the training data and used as positive validation data (Section 3.3). The ROC curve labeled with 'other classification' displays the weighted average ROC curve of the right panel in the left panel, and the other way around.

5.3.1. Meteorite density

One of the reasons for the varying performance of the different sites can be related to the definition of the positive observations. As there is no weight applied to the number of meteorites found in a single grid cell,

grid cells with multiple meteorite finds contribute equally to the classification as grid cells with a single meteorite find. In other words, the spatial density of meteorite finds is not taken into account, and isolated finds are not ignored. To investigate the sensitivity of the algorithm to this spatial density of meteorites, the algorithm is trained using only grid cells containing more than one meteorite find. This reduces the number of positive observations to 1,359. As to allow easy comparison, the validation data contains all grid cells and is thus identical to the validation data used in Figure 5.11. The resulting ROC curves indicate that on average the classification does perform almost exactly the same (Figure 5.12). However, for individual field sites there are some differences, for example EET performs worse in the novelty detection as well as in the standard classification, whereas LEW performs slightly better. As the differences are not very large, it can be concluded that the classifier is not sensitive to isolated and spatially widespread meteorite finds.




Figure 5.12: ROC curves when training the classifier with the almost the same training data as in Figure 5.11, with as difference that positive observations with a single meteorite in a grid cell are not used for training, but only for validation. The curves labeled with 'reference' are the weighted averages as displayed in Figure 5.11.

5.3.2. Meteorite finding locations

The re-projection of the meteorite finds on the 450 meter spaced grid cells reduces the original number of 12,906 finding locations to 2,554 (Section 3.1.1). The influence of this reduction of the number of positive ob-




Figure 5.13: ROC curves when training the classifier with the exact meteorite finding locations as positive observations. The positive validation data also consists of the exact meteorite finding locations. The curves labeled with 'reference' are the weighted averages as displayed in Figure 5.11.

servations is evaluated by performing a classification using the meteorite finding locations directly as positive observations. This alternative observation definition also accounts for the density of meteorite finds. When a lot of meteorites are found on locations close to each other, the observations will be similar because of two reasons. (i) Firstly, the geospatial data from which the features values are extracted does not show any sharp transitions. (ii) Secondly, when the distances of meteorite finds are smaller than the spatial resolution of the data (Table 4.1), the only differences in feature values are caused by the fact that the data is linearly interpolated. This concern is justified, as on average a grid cell of 450 meter by 450 meter contains approximately 5 meteorites ($\approx 12,906/2,554$), i.e. the average meteorite density is 25 meteorites per square kilometer. Consequently, the training data will contain a lot of correlated samples. Performing the classification using the 12,906 meteorite finding locations directly, reveals that some sites benefit from this alternative observation definition (e.g. QUE, MAC), whereas other sites perform worse (e.g. LEW, GRV) (Figure 5.13). On average the positive observation definition using grid cells (Figure 5.11) performs better and is more stable.

5.4. Classification

In order to perform the classification, decisions have to be made on the features considered, the number of principal components used, the definition of positive observations, the type of classification, and the selected threshold (for the novelty detection) or cost (for the standard classification). The final set of features contains the 'radar backscatter', 'surface velocity', 'ice thickness', 'surface slope', 'distance to outcrops' and 'surface temperature' (Section 5.1). The number of principal components is five (Section 5.2). The definition of positive observations is as described in Section 3.1.1 (Section 5.3). And the type of classification is the standard classification, as the ROC curve is slightly better than the novelty detection, and the individual ROC curves obtained in the cross-validation show a narrower spread (Figure 5.11). The selection of the cost is discussed in Section 5.4.1, the visualization of the classification in Section 5.4.2 and the quantitative evaluation of the classification in Section 5.4.3.

5.4.1. Selection of the threshold

In selecting the cost, in other words, selecting an operating point on the ROC curve, a trade-off has to be made between the false positive rate and the true positive rate. This trade-off can be optimized using the metrics precision and recall (equal to the true positive rate) (Table 5.2), which are computed using the number of true positives, the number of false positives, and the number of false negatives. Depending on the use of the meteorite hotspot map, it is more important to have a high precision or to have a high recall. If the meteorite hotspot map is used for reconnaissance missions, a high precision is required. This is because it is very expensive to carry out a reconnaissance mission, and therefore it is more important that a classified MSS is classified correctly, than that all possible MSSs detected. On the contrary, if the meteorite hotspot map is used for research purposes, a high recall is required. Carrying out a comprehensive research on MSSs, requires to include all MSSs, despite the fact that some MSSs are classified incorrectly.




Figure 5.14: Metrics obtained by the number of true positives, false positives and false negatives. The F1 score represents the harmonic mean between precision and recall. The maximum value of the F1 score is indicated with an x.

To find a threshold (i.e. cost) that satisfies both requirements, the F1 score is used, which represents the harmonic mean between precision and recall (Table 5.2). The precision, recall and F1 score are calcu-

lated using the true positive and false positive rates corresponding to the different values of the cost (i.e. the operating points on the ROC curve of Figure 5.11, standard classification). To obtain the number of true positives and the number of false positives, the true positive rate is multiplied with the number of positive observations (2,554) and the false positive rate is multiplied with the number of negative observations (8,870), respectively. Using the number of true positive observations, the number of false negative observations is calculated ($FN = 2554 - TP$). The obtained result shows that for an increasing cost the recall is increasing, whereas the precision is decreasing (Figure 5.14). In order to select a value for the cost, the F1 score is maximized (Figure 5.14). This results in a value of 72.65 for the cost, which relates to a value of 0.63 for the true positive rate and a value of 0.20 for the false positive rate. The corresponding performance metrics obtained by the validation data are listed in Table 5.3.

Metric	Equation
Precision	$PREC = \frac{TP}{TP+FP}$
Recall	$REC = \frac{TP}{TP+FN}$
F1 Score	$F1 = \frac{2TP}{2TP+FP+FN}$
Accuracy	$ACC = \frac{TP+TN}{TP+TN+FP+FN}$

Table 5.2: Equations of classification performance metrics, where TP denotes the number of true positives, FP the number of false positives, FN the number of false negatives, and TN the number of true negatives.

Metric	Value (fraction)
Precision	0.47
Recall	0.63
F1	0.54
Accuracy	0.76

Table 5.3: Performance metrics estimated using the validation data.

5.4.2. Visualization

As to prevent costly reconnaissance missions to incorrectly classified MSSs, the estimated *a posteriori* probabilities are visualized (Figure 5.15).

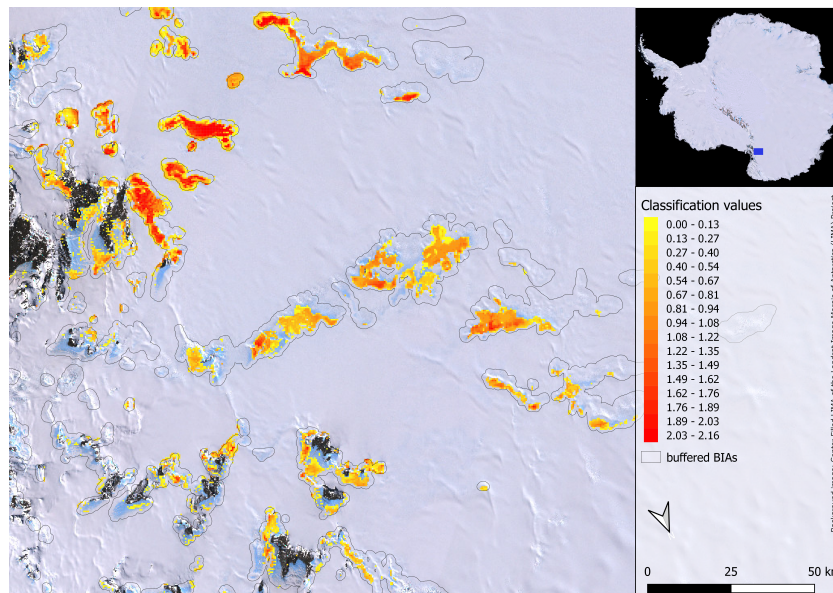


Figure 5.15: Snapshot of the meteorite hotspot map in the area of the Allan Hills and Elephant Moraine. The classification values correspond to the classification values of Figure 5.16. The area over which the classification is performed is indicated with black outlines, labeled as 'buffered BIAs' (Section 6.1).

To visualize these estimated a posteriori probabilities, an arbitrarily defined 'classification value' is used, that can be interpreted as an indication of how likely it is to find meteorites at a MSS-classified observation (i.e. the precision of the classification). Low values of the classification value (yellow) correspond to a lower precision, whereas high values of the classification value (red) correspond to a high precision. This leaves the interpretation of the obtained classification to the user.

The classification value is defined as:

$$\text{classification value} = \log(p(s = 1|x) \times \lambda_{\text{selected}}) - 0.3, \quad (5.1)$$

where $(p(s = 1|x) \times \lambda_{\text{selected}})$ equals the estimated *a posteriori* probability multiplied with the selected value for the cost, $\lambda_{\text{selected}}$ (Section 3.2.4).

The obtained classification values range from 0 to 2.2. The values are related to the metrics precision, recall, F1 score, and accuracy (Table 5.2). This relation is displayed in Figure 5.16, where the values for the estimated metrics are obtained by using the validation data. The estimated value for the metric corresponding to a cost equal to $\lambda_{\text{selected}} / (p(s = 1|x) \times \lambda_{\text{selected}} \times 2)$ is looked up in Figure 5.14. The distribution of classification values is also visualized, with color scales corresponding to the color scales of the provided maps (Appendix E).

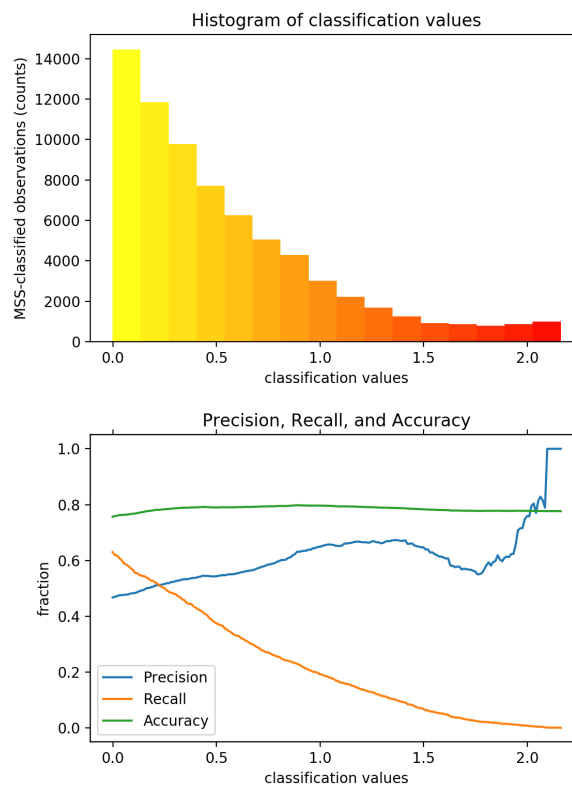


Figure 5.16: Histogram of classification values (upper panel) and related performance metrics (lower panel). The color scheme of the upper panel corresponds to the classification values displayed on the maps in Appendix E.

It is remarkable that the precision is not a monotonically increasing function. This indicates that for increasing costs the number of true positives of the classified validation data stagnates, whereas the number of false positives of the classified validation data increases. Despite the fact that this is counter-intuitive, it merely indicates that we are dealing with real data. For the interpretation of the classification values, the general trend of the precision should be considered. This trend is opposed to the trend of the recall. The estimated accuracy remains rather stagnant, as it reflects the fraction of correctly classified observations, which is the summation of the number of true positives and the number of true negatives. When the number

of true negatives increases, the number of true positives decreases, consequently, the accuracy does remains rather stable.

5.4.3. Evaluation

The quality of classification is assessed using independent test data (Appendix C). By evaluating the performance of the classifier on the independent test data, a confusion matrix is obtained, displaying the number of true positives, true negatives, false positives and false negatives (Table 5.4). From these values the performance metrics are computed (Table 5.5).

	Observed MSS	Observed non-MSS	Total
Classified MSS	191	286	477
Classified non-MSS	662	3498	4160
Total	853	3784	4637

Table 5.4: Confusion matrix obtained by using the independent test data.

Metric	Value (fraction)
Precision	0.40
Recall	0.22
F1	0.29
Accuracy	0.80

Table 5.5: Performance metrics estimated using the independent test data.

The estimated performance metrics 'precision' and 'accuracy' do correspond well to the performance metrics estimated using the validation data (Table 5.3). However, the recall is much lower, and as the F1 score is the harmonic mean between precision and recall (Table 5.2), the F1 score is also much lower. This indicates that the classifier has a low sensitivity, in other words, for the classifier it is difficult to pick up the positive observations from the test data. Reasons for the low sensitivity remain speculative, as they can be related to the sets of observations used for training, validation, and testing, and to the selected features used in the classification.

The total number of classified MSSs is 72,027. This number can be put in perspective by performing some back-of-the-envelope calculations to obtain a very rough estimate of what percentage of all exposed meteorites we found until now:

- The total number of observations classified as MSS is approximately 72,000.
- With an estimated precision of 0.40, this corresponds to $72,000 \times 0.40 \approx 28,800$ true positives.
- On average a positive grid cell contains $12.906/2554 \approx 5$ meteorites.
- This means there are a total of 144,000 meteorites exposed at MSSs in Antarctica.
- Up until today we found almost 44,000 meteorites in Antarctica.
- This corresponds to $44,000/144,000 \approx 31\%$ of all meteorites exposed at MSSs.

With an estimated accuracy of 80%, the obtained meteorite hotspot map is a first step towards the recovery of the remaining meteorites exposed at MSSs. The obtained results are discussed and put in context in Chapter 6.

6

Discussion

The discussion is centered around the four research questions introduced in Chapter 1.

6.1. Representativeness of locations of labeled and unlabeled observations

What are representative locations of labeled and unlabeled observations?

In the creation of the meteorite hotspot map, locations that needed to be classified were selected. The obtained meteorite hotspot map shows good accordance with the validation data (Figure 6.1), indicating that the used locations of labeled and unlabeled observations represent MSSs and potential MSSs well. Despite this, the assumptions made in this discretization of the real world do influence the obtained classification. Therefore, it is essential to evaluate how the locations represent MSSs and potential MSSs, and to discuss their limitations.

The locations of the observations are regularly spaced points at the buffered BIAs. The points are aligned with the surface velocity observations, resulting in a resolution of 450 meter. As the sizes of BIAs are typically a few tens to a few hundreds of square kilometers, this resolution is sufficient for obtaining locations that are not contaminated by neighboring snowfields and exposed bedrock. The resolution even allows to observe possible intra-BIA variations in the likeliness to find meteorites. However, these patterns need to be interpreted with restraint, as some of the data are simply obtained by interpolating coarser resolution data (Table 4.1).

As the obtained meteorite hotspot map shows good accordance with the validation data, it indicates that settings beneficial for meteorite concentrations can be captured on a resolution of 450 m. This confirms that the observed **expression** of the processes influencing the meteorite concentration are of a local (500-5000 meter) character. The scale of the processes themselves, on the other hand, is described as ranging from very local (e.g. intra-MSS wind patterns) to rather regional (e.g. regional ice flow). The high resolution of the Antarctic-wide state-of-the-art datasets allowed a relatively dense spacing of observations. However, data of higher resolution is available (such as a 8 meter resolution DEM [Howat et al., 2019]), allowing to investigate these intra-MSS meteorite distributions. Working with this high resolution data may provide new insights. Another tool to be applied for specific cases at high resolution is ice flow modeling, which will allow to investigate the regional processes resulting in a meteorite concentration.

The definition of the observations implies that all locations that are not a BIA are not considered in the classification. This reduction of the amount of locations allows computations on a local machine, and it is well-supported by the literature, as MSSs are described as a subset of BIAs (Chapter 2). However, to exclude non-BIAs, an existing dataset of BIA outlines is used (Section 3.1.3). The quality of this dataset does impose severe limitations to the results obtained in this research. Two measures are taken to mitigate these limitations. (i) Firstly, a buffer of 1 km is drawn around the outlines of the BIAs. The importance of the buffer is illustrated by the fact that 2070 (approximately 16%) of the meteorite finds are located within this buffer. (ii) Secondly, these buffered outlines are visualized in the meteorite hotspot map, as to indicate which areas are

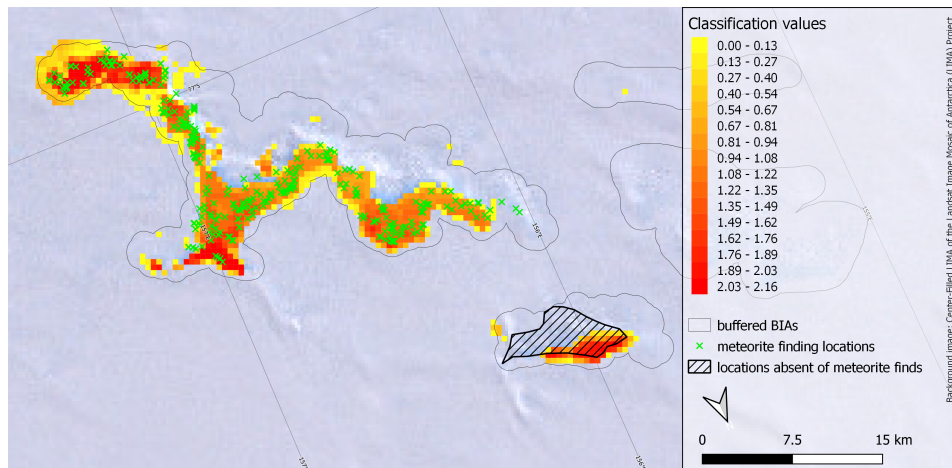


Figure 6.1: Snapshot of the meteorite hotspot map, displaying the Far Western Icefield in the area of the Allan Hills (77°S, 157°E). The positive and the negative observations used for validation are displayed. The obtained classification shows good accordance with the validation data, except in the Northern part of the area absent of meteorites.

not considered, despite the fact that they might be BIAs.

Positive observations are defined as the center of a grid cell containing one or multiple meteorites. Despite the fact that this definition greatly reduces the number of observations (from ca. 13,000 to ca. 2,500), a sensitivity analysis indicated that this definition results in a better performance of the classification as compared to using the meteorite finding locations directly as training data (Section 5.3.2). This can be related to the fact that the often densely spaced meteorite finding locations result in highly correlated observations. The negative effect on the classification of correlated observations is mitigated by re-projecting the meteorite finds onto the grid cells. However, this results in a skewed representation of the meteorite finding locations, as a grid cell containing multiple meteorites does contribute equally to the classification as a grid cell containing a single meteorite. This effect is mitigated by applying weights in the calculation of the average ROC curve from the individual cross-validated curves (Section 3.3), which is used for the validation of the classification. A sensitivity analysis where the classifier is trained with only grid cells containing more than one meteorite find does confirm that isolated or sparsely spaced meteorite finds do not have a negative influence on the classification (Section 5.3.1). Future work lies in the collection of more positive observations, as only a third of the nearly 44,000 meteorites has been used for the classification, as no location data was available for the other meteorites.

Finally, a fundamental assumption in the creation of the meteorite hotspot map lies in the fact that the mechanism is considered to be a stagnant system (Section 2.3). One of the reasons for misclassifications may reside in this assumption, but as the results are promising, it is suggested that the temporal factor is of minor importance.

6.2. Quantitative features versus qualitative descriptions in the literature

Which quantitative features do represent the characteristics of MSSs that are to date described qualitatively in literature? And do the feature-values of the labeled observations confirm the qualitative descriptions in the literature?

The meteorite hotspot map is created using six features: radar backscatter, surface velocity, ice thickness, surface slope, distance to outcrops, and surface temperature. The values of these features do confirm the qualitative descriptions in the literature (Section 4.9). Additional features considered in this thesis are the 'yield stress' and 'change of ice thickness along the flow line'. These are not used for the meteorite hotspot map, as a sequential backward selection showed that these features negatively influence the classification. This is because these features are either subject to substantial biases and/or uncertainties, or not representative for the physical phenomenon. The other features capture the properties of MSSs well enough to see differences between the positive observations and the unlabeled observations (Section 5.2), i.e. to see in the

observations that MSSs are a subset of BIAs.

Some aspects influencing the meteorite concentration described in the literature are not well represented in the six features. There is no feature that can be linked directly to the high sublimation rates, and there is no feature that can be linked directly to the exposure to wind. Expressing these phenomena directly or indirectly using observations extracted from geospatial datasets obtained by remote sensing is a future challenge that could improve the understanding of the concentration mechanism.

As the six features do confirm the current understanding of the concentration mechanism as described by experts (i.e. no unknown feature is revealed), the meteorite hotspot map should overlap with a manual classification by experts. However, performing the task of comparing six features over the entire Antarctic continent manually is extremely time consuming and has never been performed. By performing this task using big data and machine learning, an encompassing meteorite hotspot map is created. This encompassing result allows comparisons between MSSs, answering questions such as why no meteorites have been found at a certain BIA, or why a certain MSS is not classified. This qualitative evaluation of the obtained results will lead to a better understanding of the phenomenon, and eventually to a better meteorite hotspot map, resulting in a significantly improved advantage over an expert judgment.

6.3. Classification with positive and unlabeled data

To what degree is a classification feasible given the lack of negative observations and the expected diversity of unlabeled observations?

It is possible to perform a classification using only positive and unlabeled data. However, if the obtained classification can not be validated or evaluated with respect to the false positive rate, the result rather subjective. Therefore, negative observations are defined by investigating fieldwork reports. These negative observations are not used in training the classifier, as they are unreliable and not representative for all non-MSSs. Training the classifier with positive and unlabeled data, and tuning the classifier with the negative observations, results in a classification of MSSs with a value of 0.8 for the accuracy. This indicates that it is feasible to perform a MSS-classification with positive and unlabeled observations.

As the influence of the set of unlabeled observations on the classification is not by definition positive, the obtained ROC curves in the validation of the classification are always compared to a novelty detection, where only the positive observations are used. From these comparisons, it can be concluded that the unlabeled observations used for training the classifier do improve the classification, as the obtained ROC curves in the validation of the classifications are better in the standard classification (positive and unlabeled data) as compared to the novelty detection (only positive data) (Figure 5.11).

The unlabeled observations do represent potential MSSs. As an Antarctic-wide product is created, unlabeled observations can be located at any place on the continent. This implies a very diverse set of observations, covering e.g. accumulating snow fields, exposed bedrock, and melt-induced BIAs. However, the diversity and number of the unlabeled observations is greatly reduced by discarding all non-BIAs, a decision well-founded by the literature. Nevertheless, unlabeled observations are still diverse, given the fact the BIAs are buffered with a 1 km buffer to mitigate uncertainties in the defined outlines. In the two-dimensional principal components histogram groupings related to this diversity are visible (Figure 5.9). As the unlabeled observations is diverse, and the negative validation data is not representative for all negative observations (as it only includes sites that have been visited because of their potential to bear meteorites), (unnoticed) biases can be introduced in the classification. This is illustrated by an example. A classification using only the feature 'surface velocity', will classify exposed bedrock as MSS, as the surface velocities of exposed bedrock are low, just like the surface velocities at MSSs. This misclassification will not influence the estimation of the false positive rate of the classification, as the negative test data consists of only blue ice areas (i.e. no exposed bedrock). Therefore, a bias in the classification remains unnoticed. This indicates that all relevant features of MSSs need to be represented, and the results need to be carefully examined, not only by metrics, but also qualitatively. With respect to the qualitative evaluation, some visual inspections have been performed, but in-depth analyses of case-study areas are required to assess the quality of the obtained meteorite hotspot map.

The degree of feasibility of a classification of observations at BIAs as MSS or non-MSS does depend on the expectations of the intended users, i.e. the requirements of the meteorite hotspot map. As the encompassing meteorite hotspot map is entirely novel, the obtained results are already a great advancement for the planning of meteorite recovery missions. However, the results can still be refined significantly. Namely, theoretically it is feasible to achieve the same separability between the MSSs and non-MSSs using positive and unlabeled data, as can be obtained in a regular binary classification problem using positive and negative data. In practice this is a challenge, because (i) the set of features does not fully capture all components of the concentration mechanism, (ii) the observations contain measurement uncertainties and/or biases influencing the classification, (iii) and the labels of the MSSs are not 100% accurate. In conclusion, the classification will always be subject to errors, but obtaining a useful result is feasible.

6.4. Use of the meteorite hotspot map

In what way is the obtained classification to be used in meteorite recovery missions?

Meteorite recovery missions can be split up into two parts: reconnaissance missions and actual recovery missions. Despite the fact that the meteorite hotspot map does show intra-BIA variations in the likeliness to find meteorites, these patterns need to be interpreted with restraint, as these variations can be simply caused by uncertainties and/or resampling of the data used for the classification. Therefore, the meteorite hotspot map is advised to be used only in reconnaissance missions, at the BIA level. Nevertheless, the results are presented at the intra-BIA level, as giving the results at the BIA-level would make the results very dependent on the outlines of the BIAs.

An insight in the intra-BIA variations would be very relevant for meteorite recovery missions, where the high costs of grid searches could be reduced greatly. To obtain reliable intra-BIA predictions, high resolution data needs to be used. Also, the processes responsible for redistributing the meteorites within a MSS should be considered, such as transport by wind and wind-scoops. As intra-BIA predictions would also be useful in reconnaissance missions by imposing restrictions on the area that needs to be considered, the creation of such a product is very useful.

Preliminary results of my thesis have been shared with scientists in the context of two meteorite recovery missions¹. The enthusiasm with which the results were received does indicate that there is a need in science to make the recovery of meteorites more efficient. These meteorites are after all the witnesses of geological processes on other planetary bodies that form our understanding of the origin and sustainability of life on Earth and on other planets.

¹Preliminary results are shared with Prof. Dr. T. Mikouchi in the context of a Korean mission, and with Prof. Dr. S. Goderis in the context of a Belgium mission.

Conclusion and recommendations

The first steps towards an Antarctic meteorite hotspot map are taken by performing a classification of observations at BIAs in Antarctica. In the classification, two classes are considered, MSSs and non-MSSs. Six features are used in the classification: radar backscatter, surface velocity, ice thickness, surface slope, distance to outcrops, and surface temperature. With an accuracy of 80 %, the classifier can determine whether an observation is a MSS or a non-MSS.

For the next steps towards an Antarctic meteorite hotspot map the following considerations can be taken into account:

- **Case studies.** To understand why the recall of the test data is much lower than the recall of the validation data (0.22 versus 0.63), case studies should be performed. This recommendation is related to the fact that specific sites perform much better in the cross-validation than other sites, i.e. some MSSs are easier to detect than others. The reason for this precarious quality could be unraveled by performing tests related to the influence of the different settings of the concentration model (Figure 2.1) and the influence of wind displacement.
- **Definition of additional features.** The estimated precision of the obtained classification is 40 %. To improve this precision, additional discriminating features can be explored. An exhaustive feature selection method can subsequently lead to new insights with respect to the concentration mechanism. Also, complementary feature selection based on the performance of a single MSS can contribute to the understanding of the various settings of the concentration mechanism (Figure 2.1).
- **More advanced machine learning.** The obtained results can be seen as a baseline, as the machine learning algorithm used to create the meteorite hotspot map is basic. An obvious improvement of the obtained results resides in the use of more advanced pattern recognition techniques that consider the spatial distribution of observations (versus a pixel wise classification). The consideration of the spatial correlations between the observations and their labels would make the results physically more profound.

The great advantage of the results obtained by the big data approach resides in the fact that the results are encompassing. These results allow a quantification of why a certain area does not contain any meteorites, whereas another area does contain many meteorites. As MSSs are a subset of BIAs, this also implies a first step towards the quantification of the properties of BIAs. The key to a better understanding of MSSs might lie in the understanding of the processes resulting in these BIAs, which are to date not well understood (i.e. why is blue ice exposed at a certain location?). With the current technological possibilities in machine learning and numerical modeling, and the available remote sensing observations of the remote Antarctic ice sheet, processes leading to the occurrence of blue ice areas, which can act as traps for meteorites, can be better understood. The Antarctic meteorite hotspot map is an embodiment of this statement.

Bibliography

- ADD. Rock Outcrop medium resolution v7.1. URL <https://www.add.scar.org/>. [Accessed on: 2019-11-22].
- Bintanja, R. (1999). On the glaciological, meteorological, and climatological significance of Antarctic blue ice areas. *Reviews of Geophysics*, 37(3):337–359.
- Brown, H. (1960). The density and mass distribution of meteoritic bodies in the neighborhood of the Earth's orbit. *Journal of Geophysical Research*, 65(6):1679–1683.
- Brown, I. C. and Scambos, T. A. (2004). Satellite monitoring of blue-ice extent near Byrd Glacier, Antarctica. *Annals of Glaciology*, 39:223–230.
- Burton-Johnson, A., Black, M., Peter, T. E., and Kaluza-Gilbert, J. (2016). An automated methodology for differentiating rock from snow, clouds and sea in Antarctica from Landsat 8 imagery: A new rock outcrop map and area estimation for the entire Antarctic continent. *Cryosphere*, 10(4):1665–1677.
- Campell, J. B., (1996). *Introduction to Remote Sensing*. Taylor & Francis Ltd, London, second edition.
- Cassidy, W. (1977). Antarctic search for meteorites. *Antarctic Journal of the United States*, XII(4):96–98.
- Cassidy, W. (1978). Antarctic search for meteorites during the 1977-78 field season. *Antarctic Journal of the United States*, XIII(4):39–40.
- Cassidy, W. (1980). Field Occurrences and Collecting Procedures. *Smithsonian Contributions to the Earth Sciences: Catalog of Antarctic Meteorites, 1977-1978*, 23:3–7.
- Cassidy, W. and Rancitelli, L. (1982). The Traverse to Reckling Peak, 1979-1980. *Smithsonian Contributions to the Earth Sciences: Catalog of Meteorites from Victoria Land, Antarctica, 1978-1980*, 24:9–11.
- Cassidy, W. and Schutt, J. (1985). Antarctic search for meteorites: Field program 1984-1985. *Antarctic Journal of the United States*, XIX(5):54.
- Cassidy, W., Harvey, R., Schutt, J., Delisle, G., and Yanai, K. (1992). The meteorite collection sites of Antarctica. *Meteoritics*, 27(5):490–525.
- Cassidy, W. A. (1979). Antarctic search for meteorites (ANSMET 1978-79). *Antarctic Journal of the United States*, XIV(5):41–42.
- Cassidy, W. A. (1984). The 1980-1981 Field Season. *Smithsonian Contributions to the Earth Sciences: Field and Laboratory Investigations of Meteorites from Victoria Land, Antarctica*, 26:5–8.
- Cassidy, W. A. (1989). Meteorite search at Lewis Cliff ice tongue: Systematic recovery program completed. *Antarctic Journal of the United States*, XXIV(5):44.
- Cassidy, W. A. and Annestad, J. O. (1981). Antarctic search for meteorites, 1980-1981. *Antarctic Journal of the United States*, XVI(5):61–62.
- Cassidy, W. A. and Schutt, J. W. (1991). Antarctic search for meteorites 1990-1991 field season. *Antarctic Journal of the United States*, XXVI(5):52–53.
- Choi, B.-g. and Kusakabe, M. (2007). Antarctic Meteorites Recovered from Thiel Mountains, West Antarctica by the First Korea Expedition for Antarctic Meteorites. *70th Annual Meteoritical Society Meeting*.
- Claeys, P., Mattielli, N., Debaille, V., and Goderis, S., (2015). Search of Antarctic meteorites: Belgian Activities "SAMBA", final report.


- Committee on Meteorite Nomenclature, (1980). Guidelines for meteorite nomenclature. URL <https://www.lpi.usra.edu/meteor/>.
- Coren, F., Delisle, G., and Sterzai, P. (2003). Ice dynamics of the Allan Hills meteorite concentration sites revealed by satellite aperture radar interferometry. *Meteoritics and Planetary Science*, 38(9):1319–1330.
- Corrigan, C. M., Welzenbach, L. C., Richter, K., McBride, K. M., McCoy, T. J., Harvey, R. P., and Satterwhite, C. E. (2014). *A Statistical Look at the U.S. Antarctic Meteorite Collection*, chapter 10, pages 173–187. American Geophysical Union (AGU).
- Corti, G., Zeoli, A., and Bonini, M. (2003). Ice-flow dynamics and meteorite collection in Antarctica. *Earth and Planetary Science Letters*, 215(3-4):371–378.
- Corti, G., Zeoli, A., Belmaggio, P., and Folco, L. (2008). Physical modeling of the influence of bedrock topography and ablation on ice flow and meteorite concentration in Antarctica. *Journal of Geophysical Research: Earth Surface*, 113(1).
- Cuffey, K. and Paterson, W., (2010). *The physics of glaciers*. Elsevier, fourth edition.
- Dehant, V., Debaille, V., Dobos, V., Gaillard, F., Gillmann, C., Goderis, S., Grenfell, J. L., Höning, D., Javaux, E. J., Karatekin, Ö., Morbidelli, A., Noack, L., Rauer, H., Scherf, M., Spohn, T., Tackley, P., Van Hoolst, T., and Wünnemann, K. (2019). Geoscience for Understanding Habitability in the Solar System and Beyond. *Space Science Reviews*, 215(6).
- Delisle, G., Frachi, I., Rossi, A., and Wieler, R. (1993). Meteorite finds by EUROMET near Frontier Mountain, North Victoria Land, Antarctica. *Meteoritics*, 28:126–129.
- Elkan, C. and Noto, K., (2008). Learning classifiers from only positive and unlabeled data. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD 08*, pages 213–220. ACM, New York.
- Evatt, G. W., Coughlan, M. J., Joy, K. H., Smedley, A. R., Connolly, P. J., and Abrahams, I. D. (2016). A potential hidden layer of meteorites below the ice surface of Antarctica. *Nature Communications*, 7:10679.
- Folco, L., Capra, A., Chiappini, M., Frezzotti, M., Mellini, M., and Tabacco, I. E. (2002). The Frontier Mountain meteorite trap (Antarctica). *Meteoritics and Planetary Science*, 37(2):209–228.
- Folco, L., Welten, K. C., Jull, A. J., Nishiizumi, K., and Zeoli, A. (2006). Meteorites constrain the age of Antarctic ice at the Frontier Mountain blue ice field (northern Victoria Land). *Earth and Planetary Science Letters*, 248(1-2):209–216.
- Grinsted, A., Moore, J., Spikes, V. B., and Sinisalo, A. (2003). Dating Antarctic blue ice areas using a novel ice flow model. *Geophysical Research Letters*, 30(19):1–5.
- Guyon, I. and Elisseeff, A. (2003). An Introduction to Variable and Feature Selection. *Journal of Machine Learning Research*, 3:1157–1182.
- Halliday, I. (2001). *The Present-day Flux of Meteorites to the Earth*, pages 305–318.
- Halliday, I. and Griffin, A. A. (1982). a Study of the Relative Rates of Meteorite Falls on the Earth'S Surface. *Meteoritics*, 17(1):31–46.
- Haran, T., Bohlander, J., Scambos, T., Painter, T., and Fahnestock, M. (2014). MODIS Mosaic of Antarctica 2008-2009 (MOA2009) Image Map, Version 1. *NSIDC: National Snow and Ice Data Center [Accessed on: 2019-10-03]*, (updated 2019).
- Harvey, R. (2003). The origin and significance of Antarctic meteorites. *Chemie der Erde*, 63(2):93–147.
- Harvey, R. and Schutt, J. (1992). Meteorite recovery and reconnaissance near Pecora Escarpment and surrounding regions. *Antarctic Journal of the United States*, XXVII(5):26–28.
- Harvey, R. and Schutt, J. (1993). Meteorite recovery and reconnaissance in the Allan Hills-David Glacier region, 1992-1993. *Antarctic Journal of the United States*, XXVIII(5):51–52.

- Harvey, R. and Schutt, J. (1997). Meteorite recovery and reconnaissance in the Allan Hills-David Glacier and Darwin Glacier regions, 1996-1997. *Antarctic Journal of the United States, U.S. Antarctic Program, 1996-1997*.
- Harvey, R. P., Meibom, A., and Haack, H. (2001). Meteorite stranding surfaces and the Greenland icesheet. *Meteoritics and Planetary Science*, 36(6):807–816.
- Harvey, R. P., Schutt, J., and Karner, J. (2015). *Fieldwork Methods of the U.S. Antarctic Search for Meteorites Program*, pages 23–41. American Geophysical Union and John Wiley & Sons, Inc., Washington.
- Howat, I. M., Porter, C., Smith, B. E., Noh, M. J., and Morin, P. (2019). The Reference Elevation Model of Antarctica. *Cryosphere*, 13(2):665–674.
- Hughes, D. W. (1981). Meteorite Falls and Finds: Some Statistics. *Meteoritics*, 16(3):269–281.
- Hui, F., Ci, T., Cheng, X., Scambos, T. A., Liu, Y., Zhang, Y., Chi, Z., Huang, H., Wang, X., Wang, F., Zhao, C., Jin, Z., and Wang, K. (2014). Mapping blue-ice areas in Antarctica using ETM+ and MODIS data. *Annals of Glaciology*, 55(66):129–137.
- Huss, G. R., Wagstaff, J., Wasilewski, P. J., and Thompson, C. (1988). Search for meteorites north and west of Elephant Moraine, Victoria Land, 1987-1988. *Antarctic Journal of the United States*, XXIII(5):47–49.
- Jezek, K. C. (1999). Glaciological properties of the Antarctic ice sheet from RADARSAT-1 synthetic aperture radar imagery. *Annals of Glaciology*, 29:286–290.
- Jezek, K. C., Curlander, J. C., Carsey, F., Wales, C., and Barry, R. G. (2013). RAMP AMM-1 SAR Image Mosaic of Antarctica, Version 2. *NSIDC: National Snow and Ice Data Center [Accessed on: 2019-10-15]*.
- Jull, A. J. T. (2001). *Terrestrial Ages of Meteorites*, pages 241–266. Springer US, Boston, MA.
- Lee, P., Cassidy, W., Apostolopoulos, D., Deans, M., Foessel, A., Krause, C., Parra, J., Pedersen, L., Schwer, K., and Whittaker, W. L. (1998). Search for meteorites in the Patriot Hills area, Ellsworth Mountains, West Antarctica. *Meteoritics and Planetary Science*, 33(4):A92.
- Lee, P., Cassidy, W., Apostolopoulos, D., Bassi, D., Bravo, L., Cifuentes, H., Deans, M., Foessel, A., Moorehead, S., Parris, M., Puebla, C., and Pedersen, L., (1999). Search for meteorites at Martin Hills and Pirrit Hills, Antarctica. In *Proceedings of Lunar and Planetary Science Conference XXX*.
- Li, W., Guo, Q., and Elkan, C. (2011). A positive and unlabeled learning algorithm for one-class classification of remote-sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 49(2):717–725.
- Lindstrom, M. and Score, R. (1994). *Populations, Pairing, and Rare Meteorites in the U. S. Antarctic Meteorite Collection*, pages 43–45. Lunar and Planetary Institute, Houston, lpi techni edition.
- Liu, H., Wang, L., and Jezek, K. C. (2006). Automated delineation of dry and melt snow zones in Antarctica using active and passive microwave observations from space. *IEEE Transactions on Geoscience and Remote Sensing*, 44(8):2152–2163.
- Martel, L. Searching Antarctic Ice for Meteorites. URL <http://www.psr.d.hawaii.edu/Feb02/meteoriteSearch.html>. [Accessed on: 2020-02-14].
- Marvin, U. B. (1982). The Field Season in Victoria Land, 1978-1979. *Smithsonian Contributions to the Earth Sciences: Catalog of Meteorites from Victoria Land, Antarctica, 1978-1980*, 24:3–8.
- Marvin, U. B. (2014). *The Origin and Early History of the U.S. Antarctic Search for Meteorites Program (ANS-MET)*, chapter 1, pages 1–22. American Geophysical Union (AGU).
- Meteoritical Society. Bulletin of classified and named meteorite samples. URL <https://www.lpi.usra.edu/meteor/>. [Accessed on: 2019-07-05].
- Meteoritical Society. Bulletin of classified and named meteorite samples. URL <https://www.lpi.usra.edu/meteor/>. [Accessed on: 2020-02-11].

- Miao, B., Xia, Z., Zhang, C., Ou, R., and Sun, Y. (2018). Progress of Antarctic meteorite survey and research in China. *Advances in Polar Science*, 29(2):61–78.
- Morlighem, M. (2019). MEaSUREs BedMachine Antarctica, Version 1. *NASA National Snow and Ice Data Center Distributed Active Archive Center* [Accessed on: 2019-12-20].
- Morlighem, M., Rignot, E., Binder, T., Blankenship, D., Drews, R., Eagles, G., Eisen, O., Ferraccioli, F., Forsberg, R., Fretwell, P., Goel, V., Greenbaum, J. S., Gudmundsson, H., Guo, J., Helm, V., Hofstede, C., Howat, I., Humbert, A., Jokat, W., Karlsson, N. B., Lee, W. S., Matsuoka, K., Millan, R., Mouginit, J., Paden, J., Pattyn, F., Roberts, J., Rosier, S., Ruppel, A., Seroussi, H., Smith, E. C., Steinhage, D., Sun, B., van den Broeke, M. R., van Ommen, T. D., van Wessem, M., and Young, D. A. (2020a). Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the Antarctic ice sheet. *Nature Geoscience*, 13:132–137.
- Morlighem, M., Rignot, E., Binder, T., Blankenship, D., Drews, R., Eagles, G., Eisen, O., Ferraccioli, F., Forsberg, R., Fretwell, P., Goel, V., Greenbaum, J. S., Gudmundsson, H., Guo, J., Helm, V., Hofstede, C., Howat, I., Humbert, A., Jokat, W., Karlsson, N. B., Lee, W. S., Matsuoka, K., Millan, R., Mouginit, J., Paden, J., Pattyn, F., Roberts, J., Rosier, S., Ruppel, A., Seroussi, H., Smith, E. C., Steinhage, D., Sun, B., van den Broeke, M. R., van Ommen, T. D., van Wessem, M., and Young, D. A. (2020b). Supplementary Information for Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the Antarctic ice sheet. *Nature Geoscience*, 13.
- Mouginit, J., Rignot, E., and Scheuchl, B. (2019a). Continent-wide, interferometric SAR phase, mapping of Antarctic ice velocity. *Geophysical Research Letters*, 46(16):9710–9718.
- Mouginit, J., Rignot, E., and Scheuchl, B. (2019b). MEaSUREs Phase-Based Antarctica Ice Velocity Map , Version 1. *NASA National Snow and Ice Data Center Distributed Active Archive Center* [Accessed on: 2019-08-07].
- Scambos, T. A., Haran, T. M., Fahnestock, M. A., Painter, T. H., and Bohlander, J. (2007). MODIS-based Mosaic of Antarctica (MOA) data sets: Continent-wide surface morphology and snow grain size. *Remote Sensing of Environment*, 111(2):242–257.
- Scherer, P., Schultz, L., Neupert, U., Knauer, M., Neumann, S., Leya, I., Michel, R., Mokoš, J., Lipschutz, M. E., Metzler, K., Suter, M., and Kubik, P. W. (1997). Allan Hills 88019: An Antarctic H-chondrite with a very long terrestrial age. *Meteoritics and Planetary Science*, 32(6):769–773.
- Schultz, L. and Annexstad, J. O. (1984). Ablation and Ice Movement at the Allan Hills Main Icefield between 1978 and 1981. *Smithsonian Contributions to the Earth Sciences: Field and Laboratory Investigations of Meteorites from Victoria Land, Antarctica*, 26:17–21.
- Schutt, J. (1982). Results of the antarctic search for meteorites, 1981-1982. *Antarctic Journal of the United States*, XVII(5):56–57.
- Schutt, J., Rancitelli, L. A., Krähenbühl, U., and Crane, R. (1983). Exploration for meteorite concentrations in the Thiel Mountains/Pecora Escarpment region, 1982-1983. *Antarctic Journal of the United States*, XVIII(5): 83–86.
- Schutt, J., Schultz, L., Zinner, E., and Zolensky, M. (1986). Search for meteorites in the Allan Hills region, 1985-1986. *Antarctic Journal of the United States*, XXI(5):82–83.
- Schutt, J., Fessler, B., and Cassidy, W. A., (1993). Antarctic Meteorite Location and Mapping Project, LPI Technical Report. 93-07, Lunar and Planetary Institute.
- Schutt, J. W. (1989). The Expedition to the Thiel Mountains and Pecora Escarpment, 1982-1983. *Smithsonian Contributions to the Earth Sciences: Field and Laboratory Investigations of Meteorites from Victoria Land and the Thiel Mountains Region, Antarctica, 1982-1983 and 1983-1984*, 28:10–15.
- Silverman, B., (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, London, first edition.
- Sinisalo, A. and Moore, J. C. (2010). Antarctic blue ice areas - Towards extracting palaeoclimate information. *Antarctic Science*, 22(2):99–115.

- Takahashi, S., Endoh, T., Azuma, N., and Meshida, S. (1992). Bare ice fields developed in the inland part of Antarctica. *Proc. NIPR Symp. Polar Meteorol. Glaciol.*, 5:128–139.
- Wan, Z. (1999). MODIS land-surface temperature algorithm theoretical basis document (LST ATBD): version 3.3.
- Wan, Z., Hook, S., and Hulley, G. (2015). MOD11A2 MODIS/Terra Land Surface Temperature/Emissivity 8-Day L3 Global 1km SIN Grid V006. *NASA EOSDIS Land Processes DAAC* [Accessed on: 2020-01-13].
- Webb, A. R. (2002). *Feature Selection and Extraction*, chapter 9, pages 305–360. John Wiley & Sons, Ltd, second edition.
- Welten, K., Alderliesten, C., Van Der Borg, K., Lindner, L., Loeken, T., and Schultz, L. (1997). Lewis Cliff 86360: An Antarctic L-chondrite with a terrestrial age of 2.35 million years. *Meteoritics and Planetary Science*, 32: 775–780.
- Wetherill, G. W. (1976). Where do the meteorites come from? A re-evaluation of the earth-crossing apollo objects as sources of chondritic meteorites. *Geochimica et Cosmochimica Acta*, 40(11):1297–1317.
- Winther, M. N. and Liston, G. E. (2001). Blue-ice areas in Antarctica derived from NOAA AVHRR satellite data. *Journal of Glaciology*, 47(157):325–334.
- Yanai, K., Kojima, H., and Naraoka, H. (1993). The Asuka-87 and Asuka-88 Collections of Antarctic meteorites: search, discoveries, initial processing, and preliminary identification and classification. *Proc. NIPR Symp. Antarct. Meteorites*, 6:137–147.
- Yanai, K., Shiraishi, K., and Kojima, H. (1994). The Asuka-90 meteorites collection from Antarctica: searching, initial processing and preliminary identification. *Proc. NIPR Symp. Antarct. Meteorites*, 7:1–8.
- Yoshida, M. (2010). Discovery of the Yamato Meteorites in 1969. *Polar Science*, 3(4):272–284.
- Yoshida, M., Ando, H., Omoto, K., Naruse, R., and Ageta, Y. (1971). Discovery of Meteorites near Yamato Mountains, East Antarctica. *Antarctic Rec.*, 39:62–65.
- Zekollari, H., Goderis, S., Debaille, V., van Ginneken, M., Gattacceca, J., Timothy Jull, A. J., Lenaerts, J. T., Yamaguchi, A., Huybrechts, P., and Claeys, P. (2019). Unravelling the high-altitude Nansen blue ice field meteorite trap (East Antarctica) and implications for regional palaeo-conditions. *Geochimica et Cosmochimica Acta*, 248:289–310.
- Zolensky, M., Bland, P., Brown, P., and Halliday, I. (2006). *Flux of extraterrestrial Materials*, pages 869–888.

Appendices



Meteoritical Bulletin dataset

The Meteoritical Society [a] provides a large online database of meteorites, providing authoritative information on meteorite names, together with basic information such as classification and place and year of discovery. For this study mainly the geographic information of the meteorites is relevant.

The Meteoritical Bulletin Database, consulted on 05/07/2019, provides geographic information on meteorites in Antarctica both in a "normal table" as well as in a "bulletin table" and the two datasets are not identical. Of the 43,988 entries in the "normal table" 28,941 entries contain geographical coordinates. Of the 31,709 entries with unique names in the "bulletin table" 4,474 entries contain geographical coordinates. Merging the "bulletin table" to the "normal table" adds 0 meteorites to the "normal table", resulting in a dataset of 43,988 meteorites of which 28,941 entries have geographic coordinates. This information is visualized on the maps in this Appendix. Comparing the geographic coordinates of the two tables reveals several small discrepancies due to rounding off and four larger discrepancies, presumably due to typos. These four entries have not been considered in further evaluation. Also disregarded are 222 entries with location "Yamato", as the latitude is indicated falsely as 0.0.

Locations with single meteorites

Of the 28,937 coordinates, 14,128 are unique and 12,906 are locations with a single meteorite. There are thus 1,222 locations where multiple meteorites have the exact same coordinates. Presumably, the actual coordinates of the meteorite finding place are unknown and therefore these are approximated for all meteorites in the region with a single coordinate pair. These locations contain on average 13 meteorites.

Isolated locations

Out of the 12,906 locations with a single meteorite, 38 have no other meteorite locations within a radius of 4 km. This means that it is either an outlier, a meteorite that has not been found on a meteorite stranding surface, or a meteorite that has been found on a meteorite stranding surface that has not yet been searched thoroughly.

Field site names

Besides the geographic coordinates, the tables of the Meteoritical Society [a] also provide information on finding location of the meteorite in the name and so called origin. Namely, the guidelines on nomenclature state that a meteorite name should convey the meteorites' fall or find location [Committee on Meteorite Nomenclature, 1980]. In the case of Antarctic meteorites the name seems to refer to the larger region of the find, indicated by Corrigan et al. [2014] as the 'field site'. The label 'Origin' in the table, indicated for 11,631 entries, seems to refer to the specific icefield of the find. Corrigan et al. [2014] mentions that in the beginning of Antarctic meteorite collection, some field sites that were glaciologically and geographically unconnected have been grouped together, while they would be named as distinct field sites nowadays due to the large amount of specimens collected in these areas. The location, fieldsite, and icefield information in the database of the Meteoritical Society [a] is thus neither complete nor consistent, but it is helpful for structuring and understanding the data as shown in Table A.1.

Pairing relationships

As described in Section 2.1.2, some meteorite specimens can be paired. When obviously fragmented proximal stones are found they are usually collected and labelled as one single meteorite [Corrigan et al., 2014]. As for the data of the Meteoritical Society [a], it is not clear whether the paired groups are listed as a single meteorite or as individual specimens. However, it is known that the minority of the possible pairing groups is actually identified, as the resources for classification are limited and the scientific interest in performing a thorough pairing campaign is low [Corrigan et al., 2014]. Therefore, it is assumed that the vast majority of the entries in the database of the Meteoritical Society [a] represent individual specimens and not classified pairing groups.

Table A.1: Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].

Field site (abbreviation)	Field site (full name)	Icefield	no. meteorites	no. single coordinates
n.i.	Adelie Land	not indicated	1	1
ALH	Allan Hills	Allan Hills main icefield	910	374
		Far North	1	0
		Far Western	318	312
		Manhaul Bay	2	0
		Middle Western	82	80
		Near Western Icefield	175	131
		not indicated	357	65
AMU	Amundsen Glacier	not indicated	1	0
DNG	Angelo Bluff	not indicated	4	0
A	Asuka	not indicated	3611	0
BTN	Bates Nunataks	Bates North	8	0
		not indicated	7	0
BEC	Beckett Nunatak	not indicated	2	0
B	Belgica	not indicated	37	0
BOW	BowdenNévé	not indicated	1	1
BUC	Buckley Island	not indicated	29	0
CMS	Cumulus Hills	not indicated	79	0
DAV	David Glacier	not indicated	10	1
DRP	Derrick Peak	not indicated	27	0
DOM	Dominion Range	Dominion Range Main	2	0
		Mt Ward Ice Tongue	27	0
		not indicated	2097	0
EET	Elephant Moraine	Blue Lagoon	4	0
		Elephant Moraine	74	72
		Main icefield	336	255
		Meteorite City	438	3
		Northern Ice Patch	223	219
		Shoodabin Icefield	5	0
		Texas Bowl	1184	1165
		not indicated	124	26
FIN	Finger Ridge	not indicated	9	0
FRO	Frontier Mountain	Frontier Mountain icefield	77	55
		not indicated	760	437
GDR	Gardner Ridge	not indicated	4	0
GEO	Geologists Range	not indicated	32	0
GRA	Graves Nunataks	98 Camp Patch	7	0
		Inuksuq NE	5	0
		Inuksuq SW	2	0
		Lower Central	165	0
		Lower NE	8	0
		Lower West Graves	22	0
		Upper West Graves	8	0
		not indicated	149	0
GRO	Grosvenor Mountains	A-1	5	0
		Inner Cecily	17	0
		Mt. Bumstead	15	0
		Outer Cecily	125	0
		South Raymond	3	0
		not indicated	269	0
GRV	Grove Mountains	Ice Field	991	175
		not indicated	2436	1357
ILD	Inland Forts	not indicated	1	1
JOH	Johannessen Nunataks	not indicated	20	15
KLE	Klein Glacier	not indicated	1	1
LAP	La Paz Icefield	Bridle Path Trail	1	0
		Canyonero Mesa (West Paradise Palms)	2	0
		Forbidden Cove	1	0
		Paradise Palms	3	0
		Pebble Beach	1	0
		Pinnacle Vista Estates	6	0
		Terrapin Landing Royale	4	0
		Whispering Bamboo Grove	3	0
		not indicated	1654	0
LAR	Larkman Nunatak	not indicated	1020	0

Continued on next page


Table A.1: Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].

Field site (abbreviation)	Field site (full name)	Icefield	no. meteorites	no. single coordinates
n.i.	Lazarev	not indicated	1	1
LEW	Lewis Cliff	Central Walcott Névé	3	3
		Lower Ice Tongue	839	839
		Meteorite Moraine	310	279
		Near Meteorite Moraine	2	0
		S. Lewis Cliff Moraine	4	4
		South Lewis Cliff	143	143
		South Walcott Icefield	20	20
		Upper Ice Tongue	528	523
		Upper Walcott Névé	8	8
		not indicated	103	1
LON	Lonewolf Nunataks	not indicated	10	0
MAC	Mac Alpine Hills	Bottom Lip	134	91
		Camp Ice	13	13
		Fuel Ice	1	1
		Harvaine Moraine	43	32
		Jacobs	10	9
		Jacobs, Schuttaine Moraine	6	6
		Mac Top Lip	1	1
		Overlook Hill of Law Glacier	1	1
		Quiche Moraine	104	47
		Schuttaine Moraine	3	1
		Supersonic Icefield	1	1
		That Moraine	40	23
		The Neck	25	15
		Top Mac	1	1
		Upper Macalpine	1	1
		lower Macalpine	36	36
		not indicated	640	263
		snowshow ice	1	1
MCY	Mac Kay Glacier	McKay Glacier Icefields	12	3
		not indicated	76	5
MET	Meteorite Hills	Lower Vee	4	0
		not indicated	1129	0
MIB	Miller Butte	Miller Butte, summit	1	1
		not indicated	4	3
MIL	Miller Range	Milan Ice Tongue	1	1
		not indicated	2963	2266
MBR	Mount Baldr	not indicated	2	2
CRA	Mount Cranfield	not indicated	1	0
CRE	Mount Crean	not indicated	1	0
DEW	Mount De Witt	Blue ice patch 40 km WSW of Mount DeWitt	1	1
		Mt. DeWitt	1	1
		not indicated	13	10
FLE	Mount Fleming	not indicated	2	1
HOW	Mount Howe	not indicated	5	0
PRA	Mount Pratt	not indicated	22	0
PRE	Mount Prestrud	Prestrud-Bjaaland Ice Tongue	14	0
		Upper Norway Glacier	3	0
WAL	Mount Walton	not indicated	2	2
n.i.	Mount Wegener	not indicated	1	1
WSG	Mount Wisting	Lower Wisting	2	0
		Upper Wisting	7	0
n.i.	Neptune Mountains	not indicated	1	1
ODE	Odell Glacier	not indicated	3	0
OTT	Outpost Nunatak	not indicated	1	1
PTT	Patriot Hills	Moraine	1	1
PAT	Patuxent Range	Brazitis Nunatak	30	30
		Lekander Nunatak	1	1
		Main icefield	20	20
		Mount Tolchin	1	1
PCA	Pecora Escarpment	not indicated	132	0
		Cliff Bowl	14	14
		Damschroder	16	16
		Halfway Icefield	1	1
		Kink Bowl	82	80
		Lulow Bowl	1	1
		Main	394	369
		Middle Eastern	1	0
		North Forty	62	3
		Northeast Stepp	30	28
		Southwest #1	1	0
		Upper Lulow	1	1
		not indicated	30	11
PGP	Purgatory Peak	not indicated	1	1
QUE	Queen Alexandra Range	Central Nunataks Icefield	27	26
		Foggy Bottom Icefield	73	64
		Foggy Bottom Moraine	491	276

Continued on next page

Table A.1: Overview of recovered Antarctic meteorites as listed by the Meteoritical Society [a].

Field site (abbreviation)	Field site (full name)	Icefield	no. meteorites	no. single coordinates
		Footrot Flats	349	182
		Goodwin Nunataks Icefields	968	760
		Gordon Valley	2	2
		Lower SW	3	3
		Mare Meteoriticus	1043	887
		Pwellam Icefield	7	7
		Round Bottom Moraine	48	20
		Scoraine Moraine	147	88
		Scoraine Moraine Icefield	7	7
		Tail's End Icefield	94	94
		not indicated	221	148
RKP	Reckling Peak	not indicated	146	132
ROB	Roberts Butte	not indicated	1	1
RBT	Roberts Massif	not indicated	229	0
SAN	Sandford Cliffs	Mt Wyman	5	0
		Sandford Cliffs	15	0
		not indicated	40	0
SCO	Scott Glacier	not indicated	42	0
STG	Steingarden Nunataks	not indicated	14	12
		snow-covered ice slope	1	1
STE	Stewart Hills	not indicated	1	1
SZA	Szabo Bluff	not indicated	44	0
TYR	Taylor Glacier	not indicated	3	1
TEN	Tentacle Ridge	not indicated	2	0
TIL	Thiel Mountains	Bermel Escarpment	4	4
		Davies Escarpment	5	0
		Moulton Escarpment	54	54
		Mt. Powell	1	1
		Mt. Walcott	2	2
		Mt. Wrather	2	2
		not indicated	22	22
WIS	Wisconsin Range	2250 North	7	7
		East Spear	9	9
		East Strickland	3	3
		Spear Nunatak	13	9
		Strickland	1	1
Y	Yamato	not indicated	13717	93




Background image: Center-Filled LIMA of the Landsat Image Mosaic of Antarctica (LIMA) Project

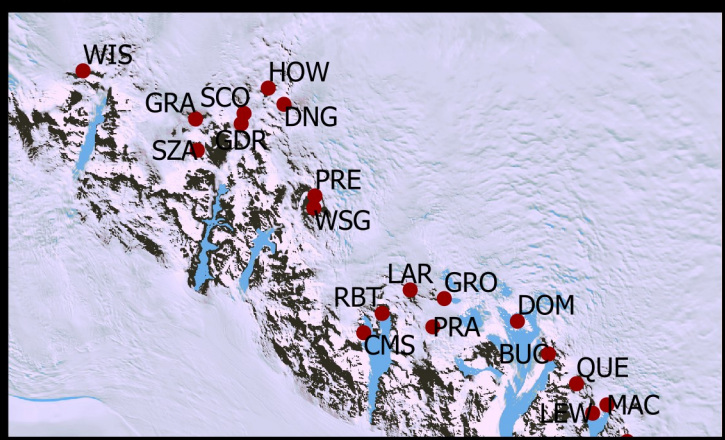
map A



map B




map C




0 75 150 km


map D




map E



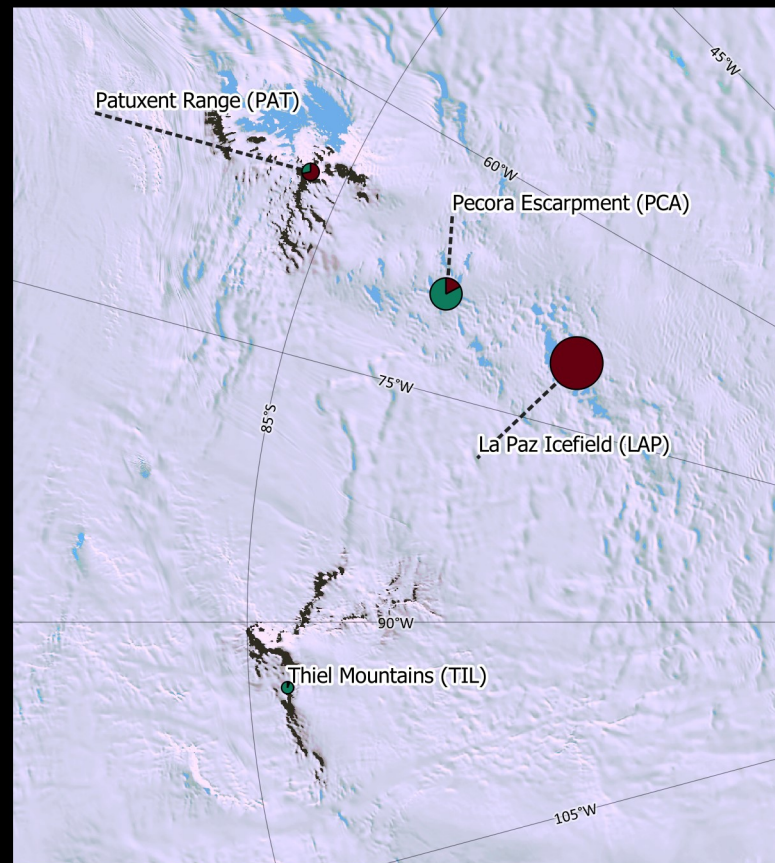
map F



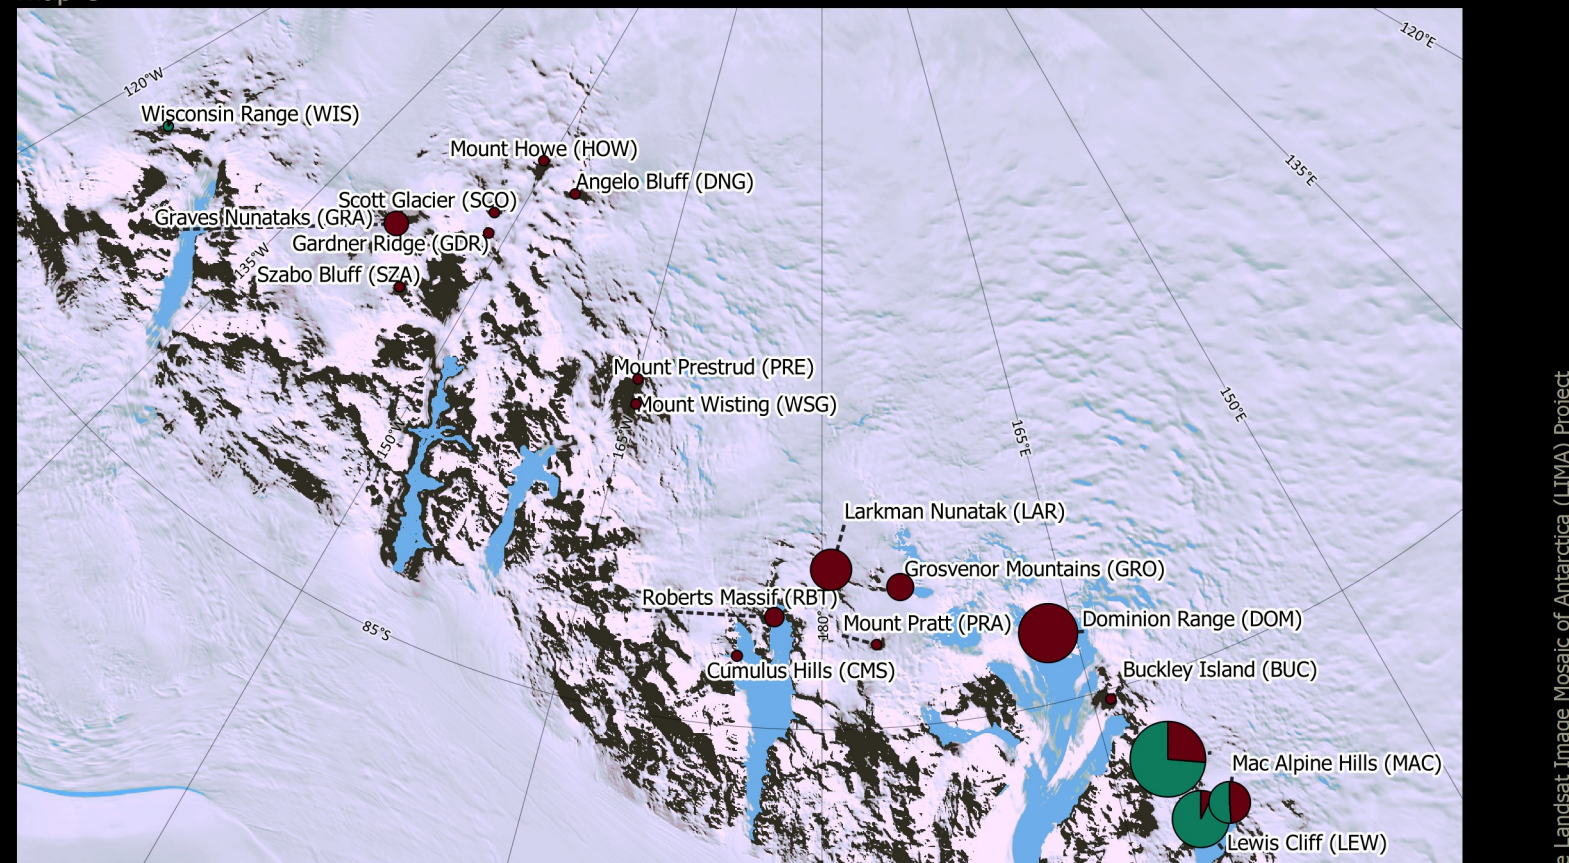
map A




map B



map C



Number of meteorites




■ meteorites with known coordinates
■ meteorites with unknown coordinates

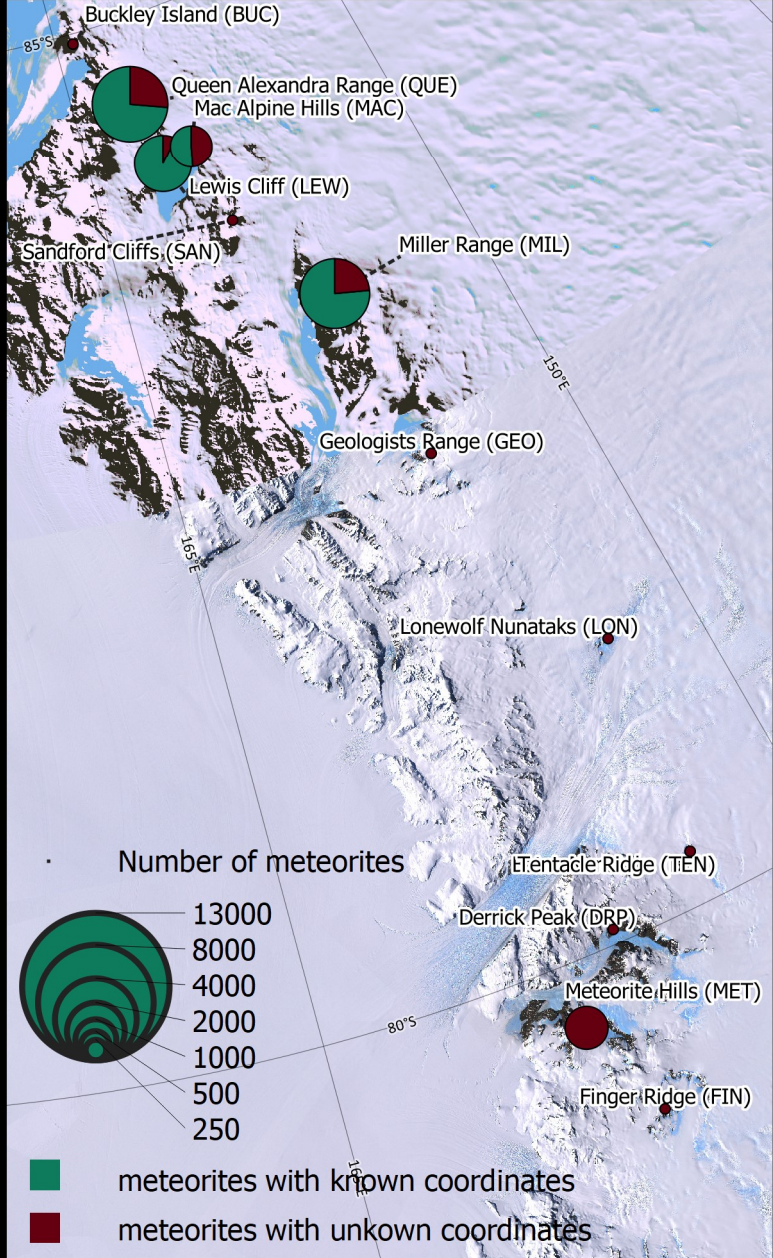
0 75 150 km




map D



map E



map F



Background image: Center-Filled LIMA of the Landsat Image Mosaic of Antarctica (LIMA) Project

B

Negative examples

As discussed in Section 3.1.2, the location of non-MSSs as mentioned in literature is often not clear. Also information on the intensity of the search at these locations is not indicated. It is thus not very sure that these locations do not contain any meteorite, as the search might have been not extensive enough or the circumstances can have been unfavorable. For example, a temporary snow cover or sastrugi can have hidden specimens [Cassidy et al., 1992; Zekollari et al., 2019]. The presented table (Table B.1) does contain supplementary information on sites with a lack of meteorite concentration that have been located on a map using the description of the location and used for the validation of the classification of MSSs as presented in this thesis. The majority of the information is found in the online available fieldwork reports of the Antarctic Search for Meteorites (ANSMET) program.

Table B.1: Overview of negative examples as extracted from the literature.

Name	Literature	Description of the location	Information on search
Allan Hills Northwestern icefield	Cassidy et al. [1992], Schutt et al. [1986]	Small and separate blue ice area approximately 16 kilometers north of the far western icefield in Allan Hills	"Thorough reconnaissance search" in 1985-86 ANSMET season, only one meteorite find
Battlements Nunatak icefields (2 parts)	Cassidy and Schutt [1985]	Large exposed icefields around Battlements Nunatak	Reconnaissance search in 1984-85 ANSMET season
Bessinger Nunatak icefield	Harvey and Schutt [1992]	near Bessinger Nunatak	Extensive searching in 1991-92 ANSMET season, only two meteorites found
Boomerang Ranges icefield	Marvin [1982], Cassidy et al. [1992]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season
Brimstone Peak icefield	Cassidy [1984], Coren et al. [2003], Cassidy and Annexstad [1981]	Blue ice located downstream of the outcrop Brimstone Peak	Visited by snowmobiles in the 1980-81 ANSMET season
Butcher Ridge icefield	Marvin [1982], Cassidy [1979]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season
Carapace Nunatak icefield	Marvin [2014], Cassidy [1977]	Carapace Nunatak has its own small icefield	Helicopter-assisted survey of some smaller patches near Carapace Nunatak in 1976-77 season, visited in 1981-82 ANSMET season, many geodes at base of Carapace Nunatak, no meteorites found
Colbert Hills icefield	Cassidy et al. [1992]	Above and below submerged parts of Colbert Hills, Figure 21 [Cassidy et al., 1992, p. 506]	Only one meteorite find below ice ramp, supports view that main body Walcott Névé is not a productive area
David Glacier icefield	Harvey and Schutt [1993]	Along north facing escarpments there are areas of blue ice, not clear if all BIAs around David Glacier Icefields in Figure [Harvey and Schutt, 1993] are meant	Five days of rigorous searching in ANSMET season 1992-93
Emlen Peaks icefields (6 parts)	Schutt [1982]	Bare ice patches at Emlen Peaks, downstream (as 'there were no such exposed icefields upstream' in the areas visited)	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Far Northern ice patches (2 parts)	Huss et al. [1988]	Map [Huss et al., 1988, p. 47]	Side by side traverse on snowmobiles in 1987-88 ANSMET season
Finger Ridge icefield	Marvin [1982], Cassidy [1979]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season
Gallipoli Heights ice patches (2 parts)	Schutt [1982]	Ice patches in Gallipoli Heights	Snowmobile traverse in 1981-82 ANSMET season

Continued on next page

Table B.1: Overview of negative examples as extracted from the literature.

Name	Literature	Description of the location	Information on search
Gordon Valley icefield	Cassidy et al. [1992]	Some blue icefields along edge of Queen Alexandra Range in Walcott Névé, Figure 21 [Cassidy et al., 1992, p. 506]	Only two specimens encountered, supports view that main body Walcott Névé is not a productive area
Griffin Nunatak icefield	Cassidy [1984], Coren et al. [2003], Cassidy and Annexstad [1981]	Blue ice located downstream of the outcrop Griffin Nunatak Cassidy [1984], icefield near Ambalada Peak Coren et al. [2003]	Visited by snowmobiles in the 1980-81 ANS-MET season
Ice patch between Mount Baldr and Mount Fleming	Cassidy [1980], Cassidy et al. [1992], Cassidy [1977]	Wright Upper Glacier receives small amount of ice from area between Mount Baldr and Mount Fleming	Two isolated finds when flying over by helicopter, searched this area and Wright Upper Glacier for six weeks during 1976-77 ANS-MET season
Jarina Nunatak to Trinity Nunatak icefields (4 parts)	Schutt et al. [1986]	Blue ice areas stretching from Jarina Nunatak to Trinity Nunatak	Searched in 1985-86 ANSMET season with no succes
Largest icefield along Davies Escarpment	Schutt et al. [1983], citetSchutt1989	Figure 3 [Schutt et al., 1983, p. 85], Figure 3-5 [p. 13]Schutt1989	Extensive search in 1982-83 ANSMET season, six meteorite specimens found in northern part of icefield, 'search of this icefield indicates that large concentrations of meteorites do not exist on it'
Lekander Nunatak icefield	Cassidy et al. [1992]	Figure 27 [Cassidy et al., 1992, p. 512]	Few isolated meteorite finds
Lewis Nunatak icefields (3 parts)	Schutt et al. [1983], Schutt [1989]	Figure 3 [Schutt et al., 1983, p. 85], Figure 3-5 [p. 13]Schutt1989	Searched by snowmobiles in 1982-83 ANS-MET season, only one meteorite found below ice cliffs
Lonely One Nunatak icefield	Schutt [1982]	Bare ice patch at Lonely One Nunatak, downstream (as 'there were no such exposed icefields upstream' in the areas visited)	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Lonewolf Nunataks icefield	Marvin [1982], Cassidy [1979]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season
Manhaul Bay icefield	Cassidy et al. [1992], Cassidy [1978]	Ice patch between arms of Allan Hills	Two isolated meteorite finds, one in 1977-78 ANSMET season, when foot searches were conducted
Monument Nunataks icefields (6 parts)	Schutt [1982]	Bare ice patches at Monument Nunataks, downstream (as 'there were no such exposed icefields upstream' in the areas visited)	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Mount Crean icefield	Cassidy [1977]	Large area of blue ice at Mount Crean	Helicopter-assisted survey in 1676-77 ANS-MET season
Mount Dewitt icefield	Cassidy [1977]	Extensive plateau area of blue ice around Mount Dewitt	Searched in 1676-77 ANSMET season
Mount Howe icefield	Cassidy [1989], Cassidy et al. [1992]	Icefield of ca. 18 km ² at west-facing foot of Mount Howe	Extensive week-long search in 1988-89 ANS-MET season by two team members, only four specimens found
Mount Tolchin icefield	Cassidy et al. [1992]	Figure 27 [Cassidy et al., 1992, p. 512]	Few isolated meteorite finds
Onlooker Nunatak icefield	Schutt [1982]	Bare ice patche at Onlooker Nunatak	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Outback Nunataks icefields (16 parts)	Schutt [1982]	Bare ice patches at Outback Nunataks, downstream (as 'there were no such exposed icefields upstream' in the areas visited)	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Outpost Nunatak icefield	Cassidy [1984], Cassidy et al. [1992], Marvin [2014], Cassidy and Annexstad [1981]	Blue ice located downstream of the outcrop Outpost Nunatak	Visited by snowmobiles in the 1980-81 ANS-MET season, found one isolated meteorite specimen
Patuxent Main icefield	Harvey and Schutt [1992], Harvey et al. [2001], Cassidy et al. [1992]	Expansive blue ice area (ca. 300 km ²) east of Anderson Hills [Harvey and Schutt, 1992]/Patuxent Mountains [Cassidy et al., 1992; Harvey et al., 2001], Figure 1 [Harvey and Schutt, 1992, p. 27], Figure 27 [Cassidy et al., 1992, p. 512]	Five days exploring in 1991-92 ANSMET season, only 22 meteorites recovered
Renerie Rocks icefield	Schutt [1982]	Bare ice patch at Renerie Rocks, downstream (as 'there were no such exposed icefields upstream' in the areas visited)	Reconnaissance by helicopter with occasional ground checks in 1981-82 ANSMET season
Scharffenbergbotnen (2 parts)	Grinsted et al. [2003]	Heimefrontfjella, Figure 4 [Grinsted et al., 2003, p. 3]	Not indicated
Shooda Bin icefield	Harvey and Schutt [1997]	Bowl-shaped large icefield, Figure 2, [Harvey and Schutt, 1997]	Reconnaissance trip in 1996-97 ANSMET season, only five specimens in large area
South and west of Reckling Peak	Schutt et al. [1986]	Extensive areas south and west of the peak	Traversed in 1985-86 ANSMET season
Tent Rock icefield	Cassidy [1984], Marvin [2014], Cassidy and Annexstad [1981]	Blue ice located downstream of the outcrop Tent Rock	Visited by snowmobiles in the 1980-81 ANS-MET season
Turnstile Ridges icefield	Marvin [1982], Cassidy [1979]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season

Continued on next page

Table B.1: Overview of negative examples as extracted from the literature.

Name	Literature	Description of the location	Information on search
Unnamed Nunatak icefield	Harvey and Schutt [1992]	Small local blue ice patches at an unnamed nunatak (84°51'S 68°40'W)	Few hours exploring in 1991-92 ANSMET season
Upstream icefield of Chastain Peak	Schutt et al. [1983], Schutt [1989]	Figure 4 [Schutt et al., 1983, p. 85], Figure 3-6 [p. 14]Schutt1989	Searched by snowmobiles in 1982-83 ANSMET season
Westhaven Nunatak icefield	Marvin [1982], Cassidy [1979]	Map [Marvin, 1982, p. 5]	Careful search by helicopter in the 1978-79 ANSMET season
Wright Upper Glacier	Cassidy [1980], Cassidy et al. [1992], Cassidy [1977]	Debris at the bottom end of Wright Upper Glacier	Searched some days in the terminal moraine (also with a mine detector), and also at the surface of the glacier in 1976-77 ANSMET season

C

Test data

The evaluation of the classification is performed using independent test data (Section 5.4.3). This data has been collected after the classification has been optimized using the validation data (Section 5.1). Therefore, this test data is presented separately from the validation data.

C.1. Negative test data

The negative test data is gathered by studying additional field work reports on meteorite reconnaissance and/or recovery missions not in the context of the Antarctic Search for Meteorites (ANSMET) program. The same limitations apply to this data as to the data presented in Appendix B: (i) The location of non-MSSs as mentioned in literature is often not clear. (ii) Information on the intensity of the search at these locations is often not indicated. It is thus not very sure that these locations are absent of meteorites, as the search might have been not extensive enough or the circumstances can have been unfavorable, for example a temporary snow cover or sastrugi can have hidden specimens [Cassidy et al., 1992; Zekollari et al., 2019].

The presented table (Table C.1) is in the same format as Table B.1. The table presents supplementary information on the sites, which have been located using imagery provided by Google Earth Pro (v. 7.3.2.5776). After translating the location information into polygons, observations are labeled as negative according to the same procedure as described in Section 3.1.2.

Table C.1: Overview of additional negative examples used for testing.

Name	Literature	Description of the location	Information on search
A 140-A 180	Yanai et al. [1993]	Large area of bare ice south of A 140-A 180, 50 km long E-W and 5 km wide, Fig 2. [Yanai et al., 1993, p. 139]	Searched by six members of the Asuka party in mid-October, as fourth exploration of the 29th Japanese Antarctic Research Expedition (JARE-29), 1987-89
above Morris Cliff	Lee et al. [1998]	Blue ice fields above Morris Cliff	Assessed by aircraft and searched on foot in the 1997-1998 field in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Between Mt. Simmons and Mt. Geissel	Lee et al. [1998]	Blue ice fields between Mt. Simmons and Mt. Geissel	Assessed by aircraft and searched on foot in the 1997-1998 field in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Independence Hills	Lee et al. [1998]	Blue ice fields near Independence Hills	Traversed and searched in the 1997-1998 field season by snowmobile and on foot in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Marble Hills	Lee et al. [1998]	Blue ice fields near Marble Hills	Traversed and searched in the 1997-1998 field season by snowmobile and on foot in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Martin Hills	Choi and Kusakabe [2007]; Lee et al. [1999]	Fig 1. [Lee et al., 1999, p. 1], blue ice fields at Martin Hills	Foot search in November 1998 in the context of the Robotic Antarctic Meteorite Search (RAMS) Program, searched in January 2007 in the context of the first Korea Expedition for Antarctic meteorites (KOREAMET)

Continued on next page

Table C.1: Overview of additional negative examples used for testing.

Name	Literature	Description of the location	Information on search
Minaret Bowl	Lee et al. [1998]	Blue ice fields near Minaret Peak	Traversed and searched in the 1997-1998 field season by snowmobile and on foot in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Morris Cliff	Lee et al. [1998]	Blue ice fields near Morris Cliff	Traversed and searched in the 1997-1998 field season by snowmobile and on foot in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Mount Bamse	Yanai et al. [1993]	Bare icefield south of Mt. Bamse and Mt. Nils Larsen, Fig 2. [Yanai et al., 1993, p. 139]	Searched by six members of the Asuka party in mid-October, as fourth exploration of the 29th Japanese Antarctic Research Expedition (JARE-29), 1987-89
Mount Nils Larsen	Yanai et al. [1993]	Bare icefield south of Mt. Bamse and Mt. Nils Larsen, Fig 2. [Yanai et al., 1993, p. 139]	Searched by six members of the Asuka party in mid-October, as fourth exploration of the 29th Japanese Antarctic Research Expedition (JARE-29), 1987-89
Patriot Hills	Lee et al. [1998, 1999]	Blue ice fields near Patriot Hills	Traversed and searched in the 1997-1998 field season by snowmobile and on foot and visited again in November 1998 in the context of the Robotic Antarctic Meteorite Search (RAMS) Program
Pirrit Hills	Choi and Kusakabe [2007]; Lee et al. [1999]	Fig 4. [Lee et al., 1999, p. 1], blue ice fields at Pirrit Hills	Foot search in November 1998 in the context of the Robotic Antarctic Meteorite Search (RAMS) Program, searched in January 2007 in the context of the first Korea Expedition for Antarctic meteorites (KOREAMET)
Sequence Hills	Delisle et al. [1993]	Similar to Frontier Mountain ice field, valleys open to the NE.	Searched during EUROMET 1990/91 season, valley floors drowned in large meltwater lakes.

C.2. Positive test data

The positive test data is gathered by collecting maps showing the spatial distribution of meteorite finds provided in field work reports and the literature. These maps are georeferenced using the software QGIS. After that, a shapefile with the coordinates of individual meteorite finds is created. These finds are rasterized according to the same procedure as described in Section 3.1.1. Errors of the finding locations introduced by georeferencing might be substantial. However, these errors do not influence the obtained data to a large extent, given that the resolution of the grid to which the finding locations are projected is 450 m (Section 3.1.1). The maps used to collect the positive test data are listed in Table C.2. The positive test data is complemented with a small set of recently published location data by the Meteoritical Society [b], also indicated in Table C.2.

Name	Literature	Map
Asuka	Meteoritical Society [b]	3 locations of meteorite finds, published in July 2019
Dominion range	Meteoritical Society [b]	23 locations of meteorite finds, published in August 2019
Meteorite Hills	Martel [2002]	Last figure
Mount Balchen	Yanai et al. [1994]	Fig. 5, p. 5
Nansen (Asuka)	Claeys et al. [2015]	Fig. 12, p. 19
Yamato	Cassidy et al. [1992]	Fig. 30, p. 513

Table C.2: Overview of additional positive examples used for testing



Additional Figures

D.1. ROC curves sequential forward selection (SFS) and sequential backward selection (SBS)

D.1.1. Sequential forward selection (SFS)

In the SFS method (Section 5.1.1), the classification is performed recursively with an increasing amount of features. The first classification (iteration 0) is based on a single feature, and the best predicting feature is selected. After that, the feature that improves the classification the most is selected. This is repeated until none of the features improves the classification. The ROC curves of each iteration are shown here.

Iteration 0




Figure D.1: ROC curves when training the classifiers with a single feature. The eight curves are labeled with the name of the single feature. Figure is identical to Figure 5.3.

Iteration 1




Figure D.2: ROC curves when training the classifiers with two features: the surface temperature and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on a single feature, the surface temperature, and is identical to the light blue curve in Figure D.1.

Iteration 2




Figure D.3: ROC curves when training the classifiers with three features: the surface temperature, the surface velocity, and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on two features, surface temperature and surface velocity, and is identical to the green curve in Figure D.2.


Iteration 3

Figure D.4: ROC curves when training the classifiers with four features: the surface temperature, the surface velocity, the radar backscatter, and an additional feature. The curves are labeled with the feature that has been added. The black curve shows the classification based on three features, surface temperature, surface velocity, and radar backscatter, and is identical to the red curve in Figure D.3.

D.1.2. Sequential backward selection (SBS)

In the SBS method (Section 5.1.2), the classification is performed recursively with a decreasing amount of features. The classifiers are trained with all but one feature, and the results are compared to a classifier that is trained using all features. This procedure is repeated until none of the features affects the performance of the classification negatively. As long as the dimension of the input data is larger than five, the dimensionality is reduced to five, using a principal component analysis (Section 3.2). This implies that for some of the analyses, not all variance of the data is considered. However, this does not influence the results to a large degree, as in the worst case 86.9% of the variance of the data is considered (Figure D.5).


Iteration 0

Figure D.5: ROC curves when training the classifiers with seven features, all features but one. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on all features. The first five principal components used for the classifications with seven features explain between 87.6 and 93.1% of the variance of the data. The first five principal components used for the classifications with all eight features explain 86.9% of the variance of the data.

Iteration 1




Figure D.6: ROC curves when training the classifiers with six features, all features but the change of ice thickness and another feature. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on seven features, all features but the change of ice thickness, and is identical to the pink curve in Figure D.5. The first five principal components used for the classifications with six features explain between 94.6 and 97.0% of the variance of the data. The first five principal components used for the classifications with seven features explain 92.8% of the variance of the data.

Iteration 2




Figure D.7: ROC curves when training the classifiers with five features, all features but the change of ice thickness, the yield stress, and another feature. The curves are labeled with the name of the feature that has been left out. The black curve shows the classification based on six features, all features but the change of ice thickness and the yield stress, and is identical to the orange curve in Figure D.6. The first five principal components used for the classifications with five features explain 100% of the variance of the data, as there are only five features considered. The first five principal components used for the classifications with six features explain 96.6% of the variance of the data.

D.2. Two dimensional histograms of principal components (PCs)

A principal component analysis (PCA) does allow to see groupings within the data. The two dimensional histogram plots of the different combinations of principal components, show that there are up to three groups visible in the unlabeled data (Figure D.8). Also, it is clear that the values of the PCs of the unlabeled data differ from the values of the PCs of the positive labeled training data.

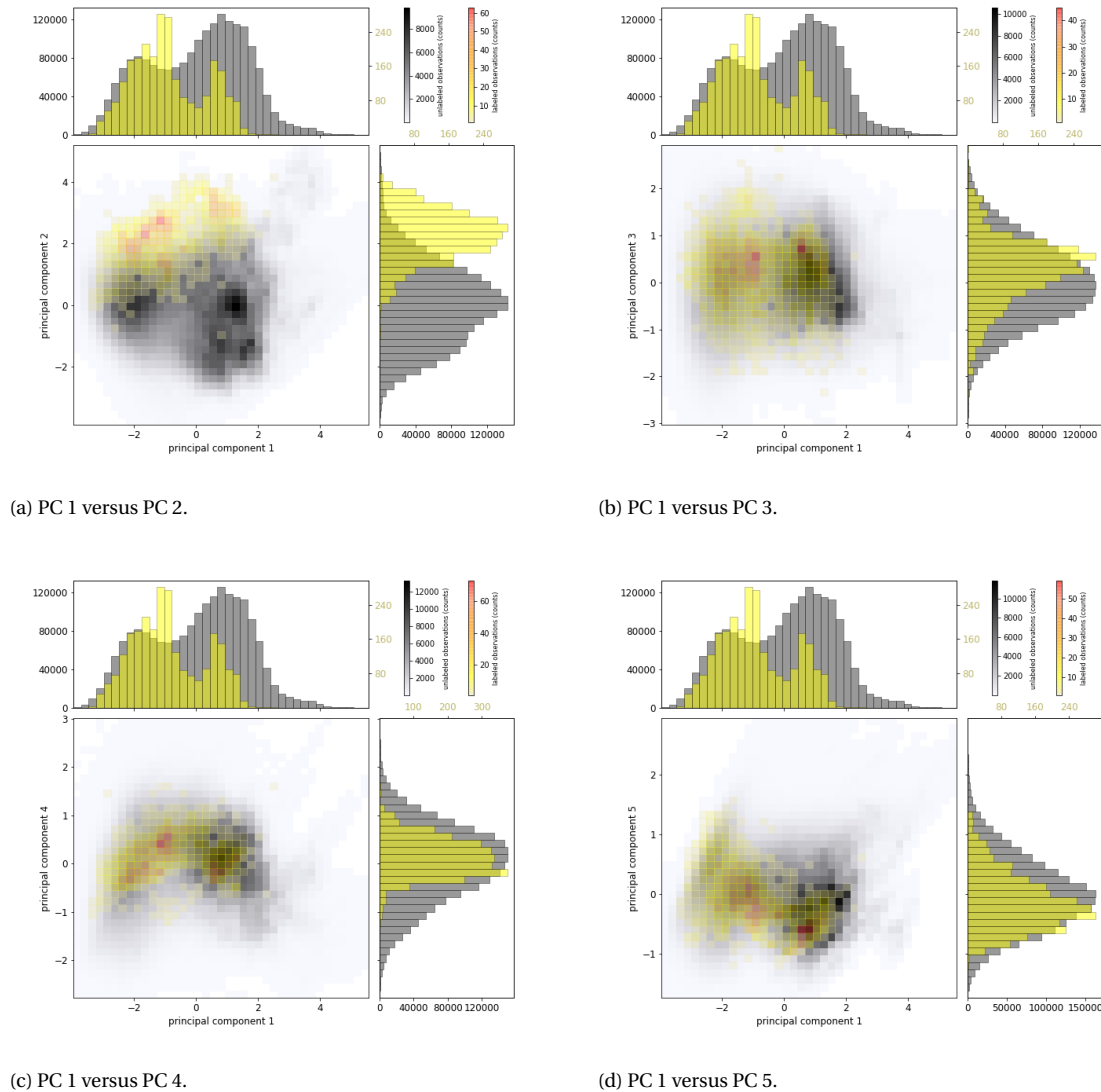





Figure D.8: Two dimensional histograms of the principal components (continued on next page).




(e) PC 1 versus PC 6.




(f) PC 2 versus PC 3.




(g) PC 2 versus PC 4.



(h) PC 2 versus PC 5.




(i) PC 2 versus PC 6.




(j) PC 3 versus PC 4.


Figure D.8: Two dimensional histograms of the principal components (continued on next page).




(k) PC 3 versus PC 5.




(l) PC 3 versus PC 6.



(m) PC 4 versus PC 5.




(n) PC 4 versus PC 6.




(o) PC 5 versus PC 6.

Figure D.8: Two dimensional histograms of the principal components.




Meteorite hotspot map

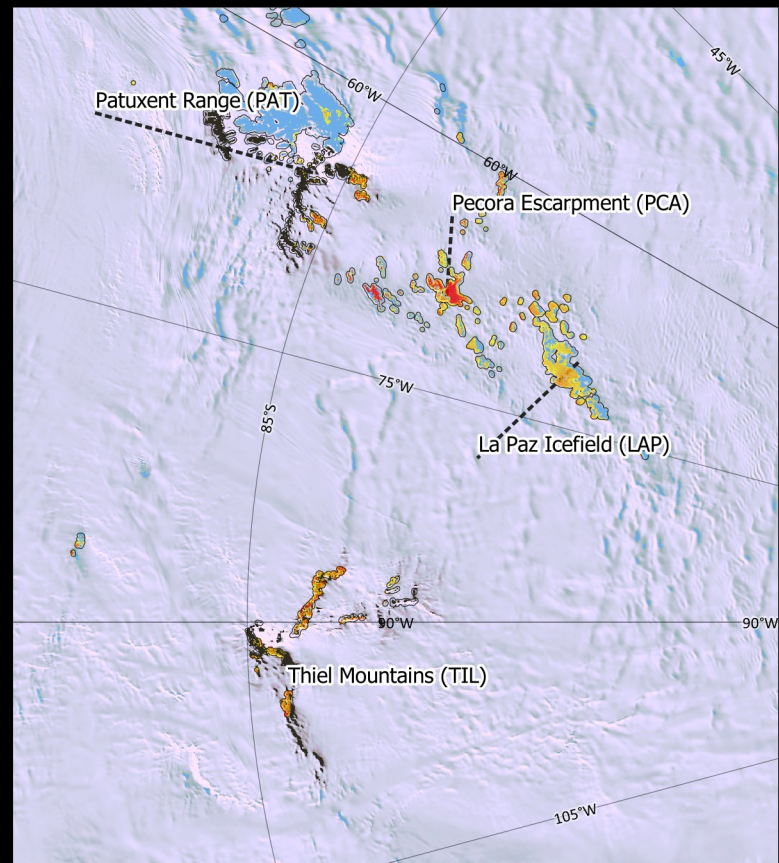
The presented meteorite hotspot map is obtained by performing a classification of observations at (buffered) blue ice areas in Antarctica. These areas are indicated on the provided maps (buffered BIAs). Observations that are classified as meteorite stranding surface (MSS) are displayed on the maps. Their corresponding 'classification value' (Section 5.4.2) can be interpreted as an indication of how likely it is to find meteorites at a MSS-classified observation (i.e. the precision of the classification). This leaves the interpretation of the obtained results to the user. The overview map does allow to zoom up to 6400 %, a level on which individual grid cells can be distinguished. An overview of the locations of the sub maps is provided in Appendix A.



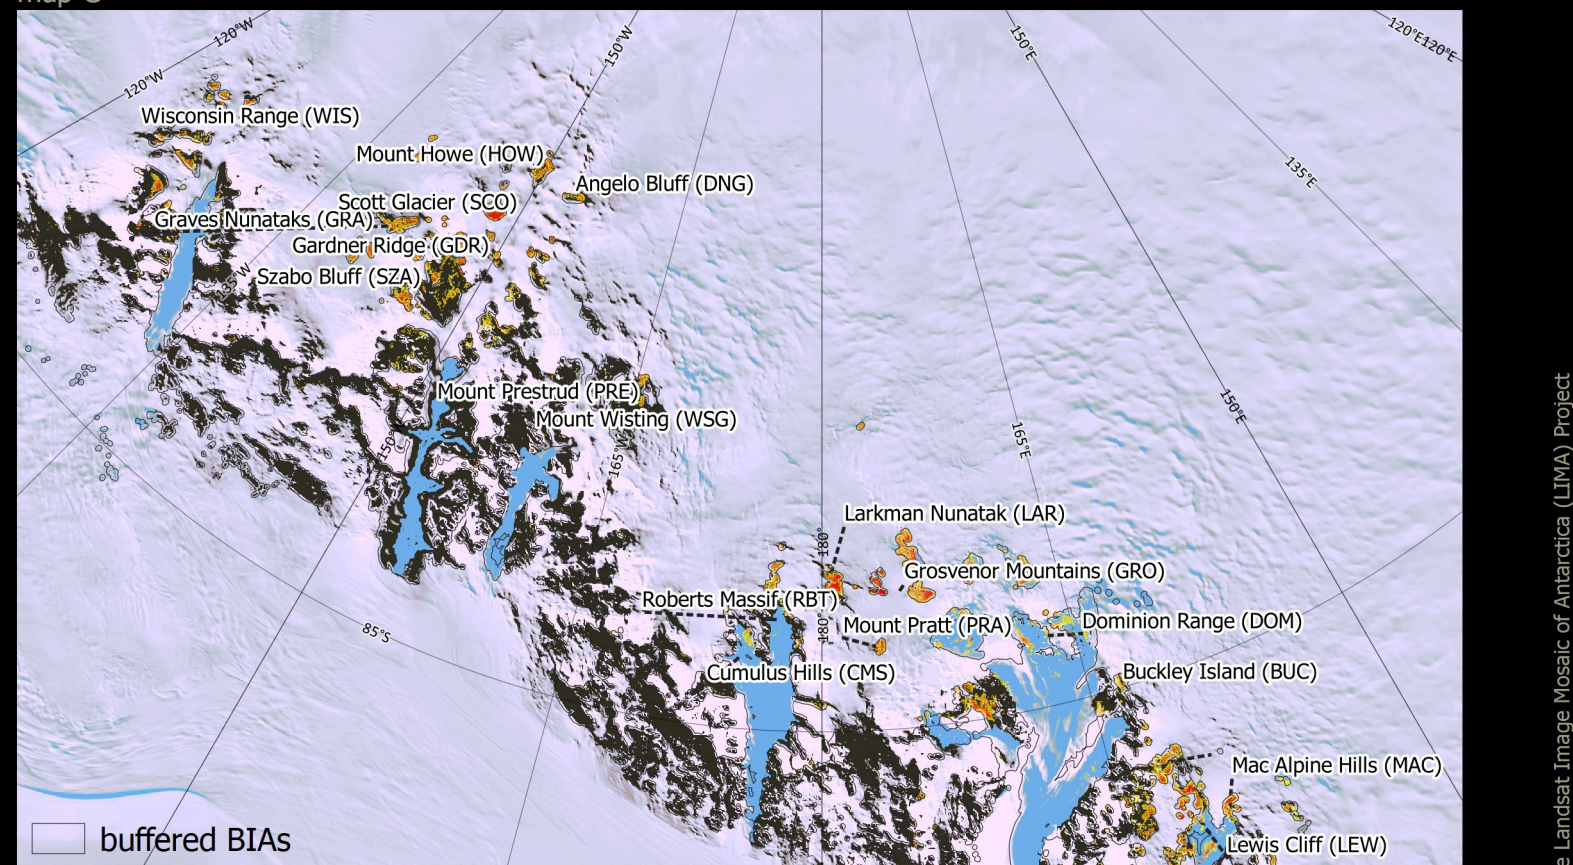
map A



map B



map C



0.00 - 0.13
 0.13 - 0.27
 0.27 - 0.40
 0.40 - 0.54
 0.54 - 0.67
 0.67 - 0.81
 0.81 - 0.94
 0.94 - 1.08
 1.08 - 1.22
 1.22 - 1.35
 1.35 - 1.49
 1.49 - 1.62
 1.62 - 1.76
 1.76 - 1.89
 1.89 - 2.03
 2.03 - 2.16



Classification values

0 75 150 km



