Automated Ice-Water Classification using Dual Polarization SAR Imagery

by

Steven Leigh

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2013

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Mapping ice and open water in ocean bodies is important for numerous purposes including environmental analysis and ship navigation. The Canadian Ice Service (CIS) currently has several expert ice analysts manually generate ice maps on a daily basis. The CIS would like to augment their current process with an automated ice-water discrimination algorithm capable of operating on dual-pol synthetic aperture radar (SAR) images produced by RADARSAT-2. Automated methods can provide mappings in larger volumes, with more consistency, and in finer resolutions that are otherwise impractical to generate.

We have developed such an automated ice-water discrimination system called MAGIC. The algorithm first classifies the HV scene using the *glocal* method, a hierarchical region-based classification method. The glocal method incorporates spatial context information into the classification model using a modified watershed segmentation and a previously developed MRF classification algorithm called IRGS. Second, a pixel-based support vector machine (SVM) using a nonlinear RBF kernel classification is performed exploiting SAR grey-level co-occurrence matrix (GLCM) texture and backscatter features. Finally, the IRGS and SVM classification results are combined using the IRGS approach but with a modified energy function to accommodate the SVM pixel-based information.

The combined classifier was tested on 61 ground truthed dual-pol RADARSAT-2 scenes of the Beaufort Sea containing a variety of ice types and water patterns across melt, summer, and freeze-up periods. The average leave-one-out classification accuracy with respect to these ground truths is 95.8% and MAGIC attains an accuracy of 90% or above on 88% of the scenes. The MAGIC system is now under consideration by CIS for operational use.

Acknowledgements

I would like to thank many people for their assistance and support throughout the development of this thesis work. My supervisor Professor David Clausi provided valuable guidance and support in all aspects throughout my degree. Readers Alex Wong and Andrea Scott provided valuable feedback for this thesis. Lab mates Zhijie Wang and Fan Li contributed greatly to the theory and development of important components of this thesis work. External support was provided by CIS, specifically Don Isaacs who applied his ice analysis expertise to generate ground truth ice charts with assistance from Lesley Gamble from MDA. The 61 scene SAR data set was provided by MDA without which none of this research would be possible.

Contents

List of Tables

List of Figures

Nomenclature

CIS	Canadian Ice Service
DT	dynamic thresholding
EM	electromagnetic
GLCM	grey-level co-ocurrence matrix
HH	horizontal horizontal
HV	horizontal vertical
IRGS	iterative region growing using semantics
JPL	Jet Propulsion Laboratory
LEO	low Earth orbit
LOO	leave-one-out
MAGIC	map-guided ice classification
MAP	maximum a posteriori
MDA	MacDonald, Dettwiler and Associates
ML	maximum likelihood
MRF	Markov random field
PEF	performance evaluation framework
RBF	radial basis function
RGB	red green blue
SAR	synthetic aperture radar
SVM	support vector machine
WMO	World Meteorological Organization

A	effective antenna aperture area
C_S	SVM tuning weight
D	distance from sensor to target
E	total energy
G_t	transmitter gain
$K()$	SVM kernel function
L	length of SAR antenna
P_1	example satellite position 1
P_2	example satellite position 2
P_r	power received
P_t	power transmitted
R	set of all regions
R_a	SAR resolution in azimuth direction
R_{gr}	SAR resolution in range direction
$V_E()$	edge strength energy term
$V_G()$	Gaussian energy term
$V_S()$	SVM energy term
Δ_{ij}	total edge strength between regions i and j
α	look angle
$\bar{t}_i$	feature vector at pixel i
β	IRGS energy function tuning parameter
$\delta()$	Dirac delta function
γ	SVM kernel scaling paramter
μ_{x_i}	average of class x_i
$\sigma^2_{x_i}$	covariance of class x_i
σ_t	radar reflectance of target
τ	SAR pulse length

ξ	set of all pairwise cliques
$b_S()$	SVM decision boundary
c	speed of light
$g()$	edge penalty function
m_i	training label at pixel i
p_j	SVM classification label of pixel j
w_i	SVM weight for feature vector at pixel i
x_i	class label for region i
y_i	average of region i
Ω_i	set of all pixels in region i

Chapter 1

Introduction

Mapping ice extents in ocean bodies serves several important purposes including the facilitation of ship navigation, environmental science, and weather forecasting. Operational ice mapping of Canadian waters is performed by the Canadian Ice Service (CIS) on a daily basis. At CIS, ice analysts manually process RADARSAT-2 synthetic aperture radar (SAR) imagery to generate region-based ice concentration maps. CIS personnel process around 3500 to 4500 RADARSAT scenes per year in this manner [35].

CIS wants to augment their existing ice mapping process by incorporating an automated system to differentiate between ice and water in RADARSAT-2 scenes. Automated approaches have the potential to provide several valuable benefits to an ice classification process. These benefits include high volume throughput able to process hundreds of scenes per day, very high detail mapping down to the pixel level, and high consistency with no inter-operator bias.

A successful automated ice-water classification algorithm must overcome a number of challenges associated with SAR images of ice and water. SAR images are corrupted with significant multiplicative speckle noise due to the coherent nature of the imaging process [32]. The images are sensitive to the incidence angle resulting in statistical non-stationarities across scenes [30]. Ice types and water often have tremendous within class variability both within and across scenes and have highly nonlinear backscatter signatures. Automated methods must cope with the aforementioned challenges in computationally efficient ways due to the large sizes of SAR imagery, typically 10 000 by 10 000 pixels.

We have developed a binary ice-water classification system called MAGIC (MAp-Guided Ice Classification) [9] that overcomes the main limitations of SAR imagery and satisfies requirements stipulated by CIS. The MAGIC system is able to achieve state-of-the-art results by combining a *glocal* classification with a support vector machine (SVM) classification using the IRGS [44] framework (Chapter 4). The MAGIC system has been tested on 61 RADARSAT-2 dual polarization SAR scenes of the Beaufort Sea for classification accuracy, robustness, and computation time (Chapter 5).

This thesis consists of six chapters. The Background Chapter 2 discusses the basics of SAR systems, reviews numerous recently published ice classification sys-

tems, outlines the IRGS [44] algorithm, and discusses several SVM+MRF classification algorithms similar to the one presented in this thesis. The proposed ice classification system was developed using image data explained in Chapter 3. Information pertaining to the 61 scenes is presented as well as the process for generating detailed ground truth for model training and evaluation. The Ice Classification System Chapter 4 explains the theory and operation of the ice classification system developed for this thesis work. The chapter gives an overview of MAGIC, then explains the algorithm steps in the sequence in which it executes: glocal, then SVM, then SVM+MRF. Chapter 5 Testing and Results evaluates MAGIC on the 61 SAR scenes and gives some analysis of the performance on some specific scenes of interest. This thesis culminates with Chapter 6 which gives a brief summary and identifies possible future work.

Chapter 2

Background

2.1 SAR Basics

SAR is a specialized radar system capable of capturing two dimensional raster images of scenes using microwave bandwidths. Such imaging systems are commonly carried by aircraft (e.g. JPL AIRSAR), or launched into low Earth orbit (LEO) on satellites (e.g. RADARSAT-2).

A SAR sensor must move over a target, typically an approximately linear flight path or a set path in LEO, transmitting and receiving pulses at various locations. Two sample locations P_1 and P_2 are shown in Fig. 2.1 which presents key terminology common in SAR systems. Relative distances between a satellite, the Earth, and targets are measured using three axes. The nadir axis extends perpendicular from the Earth towards the satellite and measures height. The ground range axis extends perpendicularly from the nadir axis to the target. The slant range axis measures the distance from the satellite to the target. The angle formed between the slant range axis and a normal to the Earth's surface is called the incidence angle. Satellites heading north are said to be in ascending orbit and southward heading satellites are descending. Position and orientation of the imaging systems are accurately tracked using GPS and inertial systems which permit accurate georeferencing for the resulting images.

The operation of SAR systems is similar to a phased array radar system, but uses a single transmitter-receiver with variable position [32]. In this way SAR systems are able to achieve resolutions similar to phased arrays while using a relatively compact and portable imaging system. A SAR system transmits and receives electromagnetic (EM) pulses as it translates across various positions above a target. At each location the SAR system behaves similarly to a classical monostatic radar with an equation governing the EM behaviour of the form:

$$P_r = \frac{P_t G_t A}{(4\pi)^2 D^4} \sigma_t \tag{2.1}$$

where P_r is the power received, P_t is the power transmitted, G_t is the gain of the transmitter, A is the effective aperture area of the antenna, D is the distance to

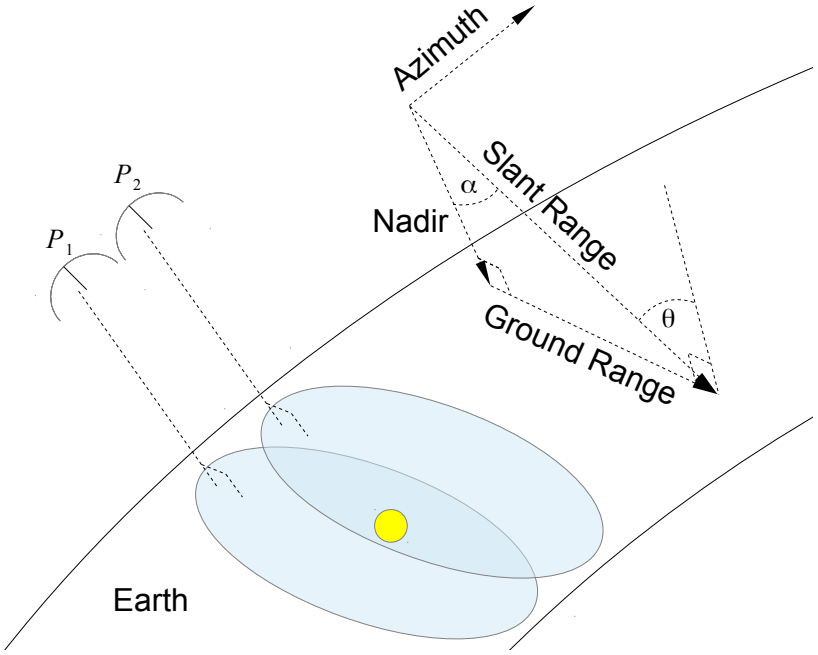

Figure 2.1: A single SAR satellite moving between two different locations P_1 and P_2. A sample sensed object is shown as a yellow circle on Earth. Blue ovals are overlapping radar footprints both of which capture a response from the sensed object. SAR satellites predominantly reside close to Earth in LEO, usually around 700km from the Earth's surface measured on the nadir axis. Slant range measures a distance from the satellite to an imaged object. Ground range measures a distance from nadir to an imaged object. The incidence angle is measured between the slant range and a normal to the Earth's surface shown as θ. The angle α is the look angle and is measured between the nadir axis and the slant range. The azimuth is parallel to the direction of travel of the sensor and is perpendicular to nadir, slant range, and ground range. This figure shows the satellite traveling north in an ascending orbit.

the target measured along the slant range, and σ_t is a measure of the reflectance of the target [37]. The resolutions achievable are one of the main benefits of SAR systems. The resolution of a SAR system in the ground range direction R_{gr} is intimately related to the length of each pulse transmitted:

$$R_{gr} = \frac{c\tau}{2\sin(\alpha)} \tag{2.2}$$

where c is the speed of light, τ is the pulse length, and α is the look angle. The resolution in the azimuth direction R_a is:

$$R_a = \frac{L}{2} \tag{2.3}$$

where L is the length of the SAR antenna [6]. These formulas show the surprising result that the resolution of SAR systems is independent of target distance and helps explain how modern spaceborne SAR systems are capable of resolutions as low as 1m by 1m. These formulas are only simplified approximations as significant research has contributed to further improving the resolution of SAR systems such as pulse chirp waveforms, beam steering, and modeling Doppler shifts.

Some SAR systems can control the polarization of the EM waves transmitted and can also measure the polarization of the returned signal. These modes are designated HH, HV, VH, and VV with the first letter indicating the horizontal or vertical polarization of the transmitted EM waves and the second letter indicating the polarization of the received wave. Transmitting and detecting the polarization of the EM waves is useful because different surface structures and materials react differently depending on the polarization. The polarizations used in this thesis work are HH and HV.

SAR systems have many desirable aspects making them highly effective for observing ice conditions. The microwave bandwidths used, between 1 to 10 GHz [6], can penetrate through weather allowing ice observation in almost any atmospheric conditions. Moreover, those bandwidths have been found to be sensitive to certain ice properties providing valuable information on ice thickness and type. SAR satellites are spaceborne systems, which if in polar orbits such as RADARSAT-1 and RADARSAT-2, are capable of imaging vast areas of the Arctic several times per day. SAR systems are active sensors emitting their own EM radiation and are not reliant on the sun meaning they can capture images any time of day. As mentioned previously SAR systems can capture images at fine resolutions even from space, such as the RADARSAT-2 satellite which achieves 1m by 1m resolution from LEO.

Though SAR systems have many benefits, they are accompanied by some unique challenges. Since SAR systems capture microwave bandwidths instead of optical bandwidths the resulting images display different information than the human visual system is used to interpreting. As a result the highly adept instinctual visual interpretation skills of humans are of less value for SAR scenes and humans must undergo training to read SAR scenes effectively. Backscatter statistics from SAR scenes are dependent on incidence angle resulting in statistical non-stationarities across a scene. These non-stationarities are most significant in co-pol images (HH

and VV) and can be compensated for algorithmically [22][30][18], or avoided using cross-pol (HV and VH) images as is done in this thesis work. SAR images are also corrupted by significant speckle noise due to the coherent nature of the image process. Speckle noise can be reduced using multilook sensing, though we mitigate the issue by using algorithms insensitive to speckle noise such as region based MRFs and window based GLCM features.

Many SAR satellites are currently in operation:

- RADARSAT-1 (Canada)
- RADARSAT-2 (Canada)
- Lacrosse (United States)
- ERS-2 (Europe)
- TerraSAR-X (Germany)
- TanDEM-X (Germany)
- SAR-Lupe 1-5 (Germany)
- COSMO-SkyMed (Italy)
- HJ-1C (China)
- METEOR-3M (Russia)
- RISAT-1 (India)
- RISAT-2 (India)
- TecSAR (Israel)

These satellites are used primarily for military reconnaissance and scientific observation. This thesis work utilizes images from the RADARSAT-2 satellite, though the algorithms developed should work on calibrated SAR images from other sources with minimal modifications. The details of the images captured from RADARSAT-2 are described in Chapter 3.

2.2 Previous Ice Classification Systems

A significant amount of work has been published in the automated SAR sea ice classification problem domain. The main challenge could be described as developing a model that captures domain specific expert knowledge for discriminating between ice and water using SAR backscatter characteristics. The most significant differences between previous work are the types of models used to achieve this goal. Many model types have been used including simple backscatter thresholding,

dynamic thresholding (DT), regression techniques, neural networks, maximum likelihood (ML) and Bayesian techniques, expert systems, and Markov random fields (MRF) among others. Some specific examples of recent work utilizing these models are briefly described next.

Work done by Fetterer et al. [14] distinguished between ice types using DT and incorporated spatial information using morphological closing. Haverkamp et al. [19] also used DT to identify homogenous regions of a single ice type. Region properties were then computed and processed by an expert system along with ancillary data to generate a final class prediction. Scheuchl et al. [33, 34] modeled the class SAR coherency matrices using Wishart distributions and discriminated classes using a Bayesian classifier. Linear and non-linear regression techniques were used by Lundhaug [28] to model the relationship between backscatter statistics and ancillary information with ice types and water. Gill [16] modeled amplitude and two additional SAR backscatter features using an empirical distribution function. Class discrimination was done using one of two distribution matching tests. Karvonen et al. [21, 23] used a Gaussian mixture model to capture class backscatter statistics and a neural network to incorporate spatial relationships. A system called ARKTOS, developed by Soh et al. [38], used watershed segmentation to incorporate spatial information. They also developed an extensive expert system to capture an expert human decision process based on a large set of region based features. Haarpaintner and Solbø [18] modeled simple backscatter statistics using three dimensional Gaussian distributions which were used as input to an ML classifier. Simple thresholds of backscatter intensities and some ancillary data were used by Geldsetzer [15]. Zakhvatkina et al. [45] analyzed a Bayesian classification of backscatter histograms. They also considered neural networks for capturing backscatter intensity and grey-level co-occurrence matrix (GLCM) texture information.

Our proposed algorithm builds upon IRGS (iterative region growing using semantics) [44]. IRGS models backscatter characteristics using Gaussian statistics and models spatial relationships using an MRF model. IRGS is an unsupervised classification algorithm that assigns arbitrary class labels to identified regions. Mapping the arbitrary class labels to ice-water labels is left as a manual process for human operators. One of the main contributions of our proposed algorithm is to remove the operator and automatically assign segmented regions ice-water labels. IRGS is explained in more detail in Section 2.3.

Ice classification systems can be differentiated by many important aspects in addition to the model type. Key characteristics of the thirteen aforementioned ice classification systems are summarized in Tables 2.1, 2.2, and 2.3, outlining algorithm details, features computed, and image details respectively.

Algorithm details are shown in Table 2.1 which identifies four key properties of the algorithms of each classification system. The *Pixel vs Region Based* column indicates whether the algorithm primarily operates on image pixels versus segmenting the image into regions and primarily operating on information pertaining to the regions. The *Level of User Interaction* identifies three categories: non-labeled, automated, and self training. Non-labeled means the algorithm only identifies distinct

Table 2.1: Previous Classification Systems: Algorithm Details. (NR: Not Reported)

Classification System	Pixel vs Region Based	Level of User Interaction	Computation Time [min]	Classification
Clausi et. al. [44][9]	region	non-labeled	1.65	ice type and water
Lundhaug [28]	pixel	automatic	NR	ice type and water
Karvonen [21]	pixel	non-labeled	"fast enough"	only ice
ARKTOS [38]	region	automatic	"near real-time"	ice type and water
Scheuchl [34]	pixel	non-labeled	NR	only ice
Scheuchl and Staples [33]	pixel	non-labeled	NR	ice type and water
Fetterer [14]	pixel	automatic	NR	only ice
Haverkamp [19]	region	automatic	5	ice type and water*
Geldsetzer [15]	pixel	automatic	NR	ice vs water
Haarpaintner [18]	pixel	automatic	NR	ice vs water
Gill [16]	pixel	non-labeled, self-training	NR	ice vs water
Zakhvatkina et al. [45]	pixel	automatic	NR	ice type and water*
Karvonen [23]	region	automatic	18	ice vs water
Proposed algorithm (MAGIC)	pixel and region	automatic	28	ice vs water

Table 2.2: Previous Classification Systems: Features

Classification System	Intensity Based	Texture	Shape	Ancillary	Other
Clausi et. al. [44][9]	I	None	None	None	None
Lundhaug [28]	I	S	None	wind speed, wind direction, air temp, location, time of year	local incidence angle
Karvonen [21]	I	None	None	None	None
ARKTOS [38]	I	S	Yes	historical ice charts, time of year, location	None
Scheuchl [34]	I, R	None	None	None	coherency matrix
Scheuchl and Staples [33]	I, R	None	None	None	coherency matrix
Fetterer [14]	I	None	None	None	None
Haverkamp [19]	I	None	Yes	historical ice charts, time of year, location	None
Geldsetzer [15]	I	None	None	None	None
Haarpaintner [18]	I	S	None	None	None
Gill [16]	I	S	None	None	None
Zakhvatkina et al. [45]	I	G, S	None	None	None
Karvonen [23]	I	S	Yes	None	None
Proposed algorithm (MAGIC)	I	G	None	None	None

Table 2.3: Previous Classification Systems: Image Details. (NR: Not Reported)

Classification System	Sensor	Band	Pol	Image Size [pixels]	Resolution [m]	# of Scenes	Location	Season
Clausi et. al. [44][9]	RADARSAT-1	C	HH	1209 by 865	NR	2	NR	NR
Lundhaug [28]	ERS 1, ERS 2	C	VV	1000 by 1000	100	105	Pechora Sea, Kara Sea	all
Karvonen [21]	RADARSAT-1	C	HH	5000 by 5000	100	9	Baltic Sea	freeze-up, winter
ARKTOS [38]	RADARSAT-1	C	HH	NR	NR	54	Beaufort Sea, Bering Sea, Chukchi Sea	all
Scheuchl [34]	JPL AIRSAR	C, P, L	full	NR	NR	1	Beaufort Sea	winter
Scheuchl and Staples [33]	JPL SIR-C	C	full	NR	100	1	Gulf of Saint Lawrence	melt
Fetterer [14]	ERS-1	C	VV	1000 by 1000	100	86	Central Arctic, Beaufort Sea, East Siberian Sea,	NR
Haverkamp [19]	ERS-1	C	VV	1024 by 1024	NR	90	Arctic Basin, Beaufort Sea	all
Geldsetzer [15]	RADARSAT-2	C	HH, HV	430 pixels total	12.5 and 50	8	Lakes in Canadian National Parks	all
Haarpaintner [18]	RADARSAT-1, Envisat ASAR	C	HH	NR	NR	8	Barents Sea	winter
Gill [16]	RADARSAT-1	C	HH	1250 by 1250	400	2	Disko Bay, Scoresby Sound	summer
Zakhvatkina et al. [45]	Envisat ASAR	C	HH	NR	NR	32	High Arctic Sea	winter
Karvonen [23]	RADARSAT-1	C	HH	5000 by 5000	100	52	Baltic Sea	winter
Proposed algorithm (MAGIC)	RADARSAT-2	C	HH, HV	2500 by 2500	200	21	Beaufort Sea	freeze-up, melt, summer

regions on the scene, but a human is required to assign each region to ice or water. Automated means the algorithm can perform a full ice-water classification without human interaction. Self-training means a human must first identify prototypical classes within the image and the algorithm will classify the remainder of the image into the identified classes. *Computation Time* is the reported runtime of the algorithm and is important for an operational implementation. The computation time varies depending on computer power and test image size, but is reported here for reference purposes. The *Classification* column shows what each classification system is designed to distinguish. Some only identify different ice types, some distinguish between different ice types as well as water, and some discriminate between only ice and water (those marked with * can not distinguish between water and certain types of ice).

Table 2.2 summarizes the types of features used in each classification system. Intensity information alone is usually insufficient to obtain satisfactory classifications in large part due to incidence angle variation. To overcome this, most algorithms use additional features as input to the classifiers. The *Intensity Based* column of the table shows that all algorithms use raw intensity as a feature (indicated by I) and a couple algorithms use simple ratios of polarizations (indicated by R). The *Texture* column shows that about half of the algorithms use texture information as features with some using statistical texture definitions (indicated by S) and some using GLCM methods (indicated by G). Three algorithms use information extracted from image region shapes as shown in the *Shape* column. The *Ancillary* column shows any supplementary information sources used such as image location or time of year. Information used for classification that does not fit into the previous columns is identified in the *Other* column.

Previous classification systems were designed and tested using distinct datasets with a variety of different image properties as summarized in Table 2.3. The imaging *Sensors*, *Bands*, and *Polarizations* used are shown in their respective columns. The size and resolution of images actually processed by the algorithms (after any downsampling in a pre-processing step) are shown in the *Image Size* and *Resolution* columns respectively. The number of scenes analyzed are identified in the *# of Scenes* column. The geographic locations of the scenes were all in various Arctic seas indicated in the *Location* column. The *Season* denotes likely types of ice formations within an image, with particularly difficult ice conditions arising in the melt and freeze-up seasons.

Tables 2.1, 2.2, and 2.3 facilitate the comparison of classification systems against CIS operational requirements. Notable CIS requirements are as follows:

- *Level of User Interaction*: automatic
- *Classification*: differentiate between ice and water
- *Polarization*: HH and/or HV (requires RADARSAT-2)
- *Seasons* tested: freeze-up, melt

Three classification systems satisfy all of these CIS requirements: ARKTOS [38], Geldsetzer [15], and MAGIC (proposed). The ARKTOS [38] system has been previously considered by CIS but is no longer under consideration. The Geldsetzer [15] algorithm was designed to work on freshwater lakes which have different backscatter signatures than sea ice. This algorithm is also based on intensity thresholding which is much less robust when applied to larger ocean bodies as will be done by CIS. Of those listed, our proposed algorithm is the only system that satisfies all requirements for operational implementation by CIS.

2.3 IRGS

The ice classification system presented in this thesis builds upon a previously developed classification algorithm called IRGS (iterative region growing using semantics) [44]. IRGS is a general purpose image classification algorithm. The classification that IRGS performs is unsupervised and as a result only arbitrary class labels are assigned. IRGS is based on an MRF framework that operates on many smaller regions of an image. IRGS is capable of processing large scenes common in remote sensing in large part due to the region based approach.

The first step of the IRGS process is to identify the initial regions, which is done by performing a watershed segmentation [42]. IRGS then operates on the fewer identified regions instead of individual pixels which can reduce the problem complexity by an order of magnitude. IRGS models region properties using Gaussian statistics computing the mean and covariance of all pixels in a region.

IRGS then proceeds with an iterative process, repeatedly alternating between a classification step and a merging step. The classification step involves fitting classes to sets of regions using a Gaussian mixture model which is optimized using expectation maximization [12]. Because the algorithm is unsupervised, the assigned classes have arbitrary labels. The merging step then merges connected regions of the same class if doing so reduces a specified energy function of the form:

$$E = \sum_{i \in R} V_G(x_i) + \beta \sum_{\langle i,j \rangle \xi} V_E(x_i, x_j) \tag{2.4}$$

where R is the set of all regions, x_i are region labels, β is a tuning parameter, ξ is the set of all cliques, and the $V_G()$ and $V_E()$ terms are explained next [44, eq. (3)-(6), (20)]. The $V_G()$ term of the left sum accounts for Gaussian statistics as follows:

$$V_G(x_i) = \frac{1}{2} \ln\left(2\pi\sigma_{x_i}^2\right) + \frac{\left(y_i - \mu_{x_i}\right)^2}{2\sigma_{x_i}^2} \tag{2.5}$$

where y_i are region averages, μ_{x_i} are the class averages, and $\sigma_{x_i}^2$ are the class covariances. The inclusion of this term encourages merges that result in small class covariances. The $V_E()$ term of the right sum incorporates edge strength information as:

$$V_E(x_i, x_j) = [1 - \delta(x_i, x_j)]\, g(\nabla_{ij}) \tag{2.6}$$

where $\delta()$ is the Dirac delta function, $g()$ is an edge penalty function, and ∇_{ij} is the total edge strength between two cliques. The $V_E()$ term encourages merging similarly labeled regions separated by weak edges. The influence of the two sums in eq. 2.4 is balanced using the weight β weight set to 3 which was previously found to perform well for ice scenes. Merges performed with this energy function are done using a simulated annealing [24] approach to avoid local minima. This energy function exists within a convenient MRF framework modeling the Gaussian statistics using unary potentials and edge strength relationships using first order clique potentials [25].

Since IRGS is unsupervised, the final step of mapping the arbitrary class labels to ice-water labels is left as a manual process for human operators. One of the main contributions of our proposed algorithm is to eliminate this manual process and automatically assign ice-water labels to the segmented regions. This is achieved by combining the MRF based IRGS algorithm with a supervised SVM model, something that has never been done on SAR data. Some previous SVM+MRF algorithms developed for other applications are discussed next.

2.4 Previous SVM+MRF Algorithms

SVM and MRF classifiers have previously been combined to achieve competitive results on image classification tasks. Tarabalka et al. [40] developed an SVM+MRF classifier for airborne hyperspectral imaging and tested on vegetation and urban classification environments. Bovolo and Bruzzone [3] tested the classification accuracy of an SVM+MRF classifier on eight ERS-1 SAR scenes identifying four classes: urban, forest, water, and fields. Farag et al. [13] developed an SVM+MRF model from an MAP framework and evaluated the classifier on a variety of images including synthetic, multispectral, and hyperspectral modalities. They were interested in agricultural and urban environments. Lui et al. [27] analyzed their SVM+MRF model using two airborne scenes capturing RGB and infrared bands. They identified three classes: bare land, dead trees, and forest. Finally, Zhang et al. [46] combined a probabilistic pixel-wise SVM model with an MRF Ising model [20] to classify synthetic as well as real hyperspectral scenes.

The SVM+MRF classifier proposed in this thesis improves upon the aforementioned algorithms in a number of ways. We use a region based MRF approach instead of a pixel-based approach used in other work enabling processing of much larger images at 2500 by 2500 pixels, vs 700 by 700 pixels for the next largest scene in an SVM+MRF publication. The availability of hyperspectral data, as is the case with [46][13][40], likely improves class separability. Unlike other methods, we only use a single imaging band and extracted texture features to improve class separability for the SVM model. Finally, the previous SVM+MRF models were developed to classify terrestrial phenomenon including agricultural and urban scenes. To the best of our knowledge, ours is the first SVM+MRF model used to

distinguish between sea ice and water.

Chapter 3

Data

3.1 Images

The image dataset we developed for testing consists of 61 scenes from the C-band SAR satellite RADARSAT-2. Each scene was captured in the ScanSAR Wide beam mode which provides images in HH and HV polarizations for each scene. The images have nominal pixel dimensions of 50m and image dimensions of around 10 000 by 10 000 pixels. This is the largest standard scene size available from RADARSAT-2 at 500km in both range and azimuth and is the most useful beam mode for operational mapping of the vast expanses of ice and water in Arctic seas. The incidence angle of these images ranges from 20 to 49 degrees and the scenes were captured from both ascending and descending satellite passes [5, 29]. The images were captured in the year 2010 at various locations over the Beaufort Sea displayed in Fig 3.1. The 61 scenes were taken from April through December inclusive of the more challenging freeze-up and melt seasons where many different ice types and ocean water conditions are present.

An example scene from our dataset is shown in Fig. 3.2. This scene was captured on October 27, 2010. It is a fairly complex scene of the north Alaskan coast containing many ice types, water and land. Both HH and HV polarizations are shown and the unique information captured by the HV pol is apparent, specifically its invariance to incidence angle effects in contrast to the HH pol.

We downsample the original images from 10 000 by 10 000 pixels to 2500 by 2500 pixels by performing 4 by 4 block averaging. This reduction in pixel count by a factor of 16 greatly increases the processing throughput and allows our system to compute a classification result in well under the CIS operational requirement of 1 hour. The primary tradeoff of downsampling is a coarser classification result, though our 2500 by 2500 pixel results with 200m pixel resolution still meet CIS needs and are far more detailed than could be expected from a manual human classification.

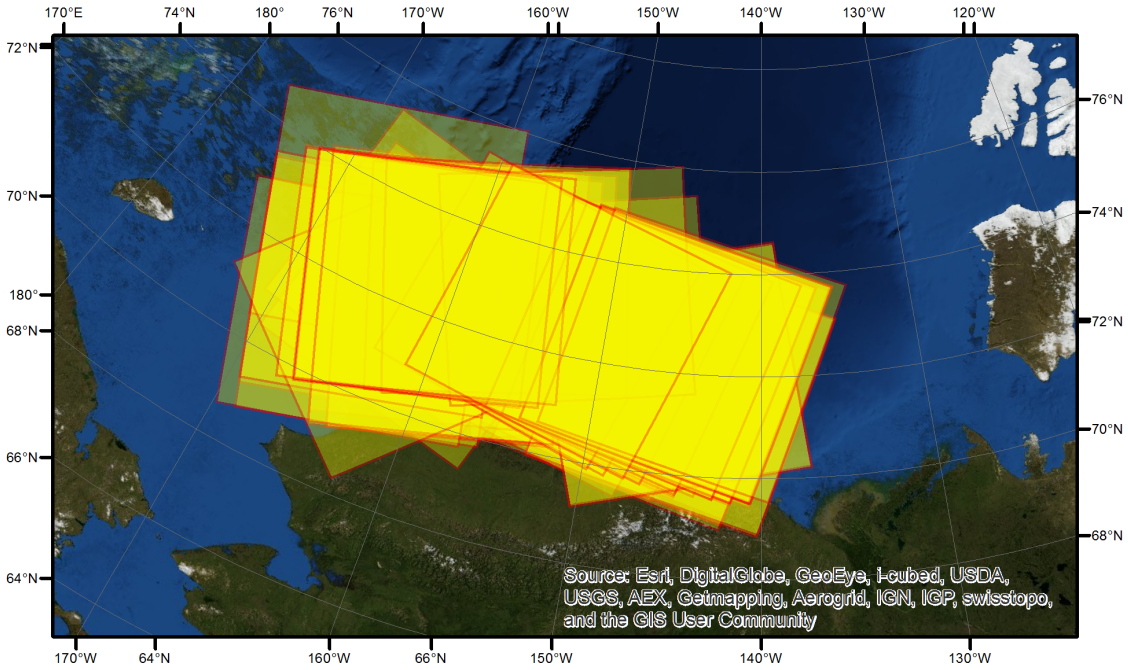

Figure 3.1: Locations of the 61 RADARSAT-2 scenes in the Beaufort Sea. Each scene covers an area of about 250 000km^2. All but one scene image a portion of the north Alaskan coast.

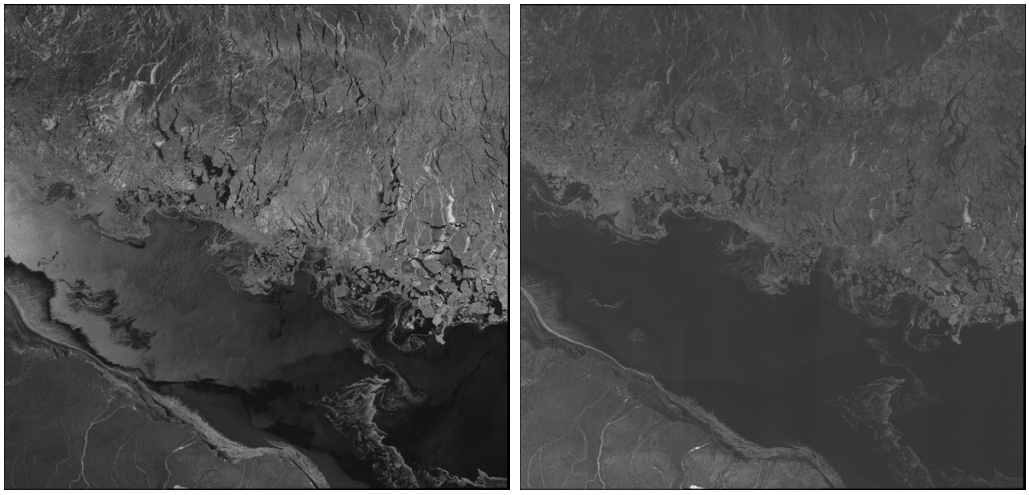

Figure 3.2: RADARSAT-2 scene captured October 27, 2010 (Scene ID: 20101027). Left: HH pol. Right: HV pol. Bottom left of scene is north Alaskan coast. Top third of scene is ice. Water in scene is most easily identified by dark regions in HV image. A very strong incidence angle effect is apparent in the HH image with water at far left appearing very bright and water at far right appearing very dark. The backscatter of ice is also affected across the HH image, but is affected differently than water: brighter at right and darker at left. The HV pol backscatter signatures are largely unaffected by incidence angle variations. This scene was reduced to 625 by 625 pixels and contrast enhanced for display purposes.

3.2 Ground Truth

For training and validation purposes we created a detailed ground truth with ice-water labels for each scene. Vector-based ground truthing with associated World Meteorological Organization (WMO) egg codes [43] based on the SAR images was performed by an expert ice analyst for the earliest 21 scenes in the data set. A sample vector-based ground truth for the October 27, 2010 scene is shown in Fig. 3.3. Such vector-based ground truths, while having good ice condition predictions at larger scales, generally have a poor level of detail because of coarsely outlined ice-water boundaries and lack of identification of small floes and leads. Using coarse training data to develop an automated classification method with an aim to output highly detailed results is not feasible. To overcome this we generated our own highly detailed ground truth using a previously developed *performance evaluation framework* (PEF) [31].

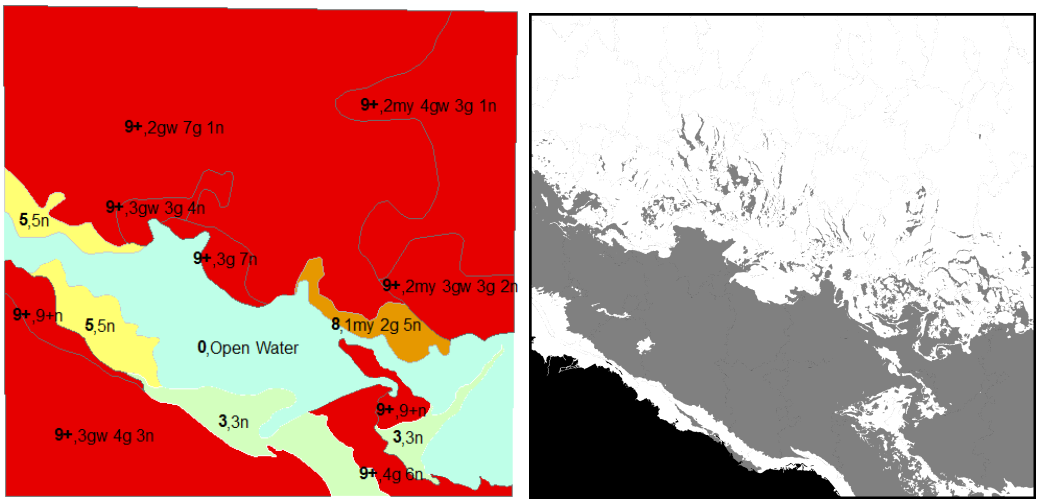

Figure 3.3: Ground truths of October 27, 2010 scene (Scene ID: 20101027). Left: vector-based ice map generated by expert ice analyst. Regions of similar ice conditions are outlined and ice conditions are identified by concentration labellings. An example decoding of the top right polygon labeled "9+, 2my 4gw 3g 1n" means this polygon contains over 90% ice, of which about 20% is multi-year ice, 40% grey-white ice, 30% grey ice, and 10% is new ice. Right: PEF detailed ground truth with ice-water labels. White is water, grey is ice, and black indicates no label due to image artifacts, ice-water boundaries, or land. Details such as ice leads, small floes, and complex ice-water boundaries are missing in the vector-based mapping, but are well outlined in the PEF ground truth.

The PEF is a framework capable of efficiently combining prior spatial contextual information with human interpretation to generate high quality ice-water estimates. The first step of the system is to divide the HV image into 144 evenly spaced square regions called *autopolygons*. Then, for each region, an IRGS classification is performed identifying anywhere from three to six classes with arbitrary class labels. The number of classes to identify is selected by the user and is related to

the complexity of the scene. More classes ensures regions are less likely to contain both ice and water providing a more detailed result, while fewer classes means less time spent by the user labeling regions. The user is then prompted one by one to provide an ice or water label for each class. For the case of six classes and 144 autopolygons that amounts to 864 user assigned labels. In this way the PEF is able assign ice-water labels to all 6 250 000 pixels of a scene using practical amounts of manual work. In practice this labeling process took less than 15 minutes for simple scenes and up to one hour for complex scenes. The PEF was run on 60 scenes to provide a highly detailed classification with binary ice-water labels. An example of the resulting highly detailed ground truth for the October 27, 2010 scene is shown in Fig. 3.3. Note that the leads, wispy ice, and small floes appear to be well outlined.

Chapter 4

Ice Classification System

4.1 Overview of Algorithm

Our ice classification system, MAGIC, consists of components as shown in the flowchart of Fig. 4.1. The inputs to the system are the HH and HV pols of a scene, the scene landmask which identifies each pixel as either land or not land, a previously trained SVM model, and a list of features to extract. The MAGIC algorithm begins with two initial independent classifications: a "glocal" IRGS classification shown as the left track in Fig. 4.1, and a pixel-based SVM classification shown as the right track. The glocal IRGS classification captures spatial contextual information from the SAR scene and identifies homogenous regions using a hierarchical approach. The glocal classification is unsupervised and therefore does not assign ice-water labels, but instead identifies regions using arbitrary class labels. The details of this glocal classification method are described in Section 4.2. Texture information of the SAR scene is exploited by the SVM classifier which is described in Section 4.3. The SVM classifier uses a previously trained SVM model and is able to provide ice-water labels. The glocal and SVM classification results are then combined using the IRGS approach with a modified energy function to balance the contributions of the contextual and texture based information. This is described further in Section 4.4. The final result of the algorithm, as shown at the end of Fig. 4.1, is a region image of the input scene with ice-water labels.

4.2 Unsupervised Classification with Arbitrary Class Labels: Glocal Approach

Large SAR scenes, on the order of 250 000km square, are necessary for mapping the large expanses of the Arctic. One of the main challenges that such large scene sizes introduce is statistical non-stationarities across the image i.e., the statistics for a particular class vary across the scene. The non-stationarities are a result of incidence angle effects, subtle sensor and processor imperfections, as well as intra-scene ice and water variations. However, these non-stationarities generally have low

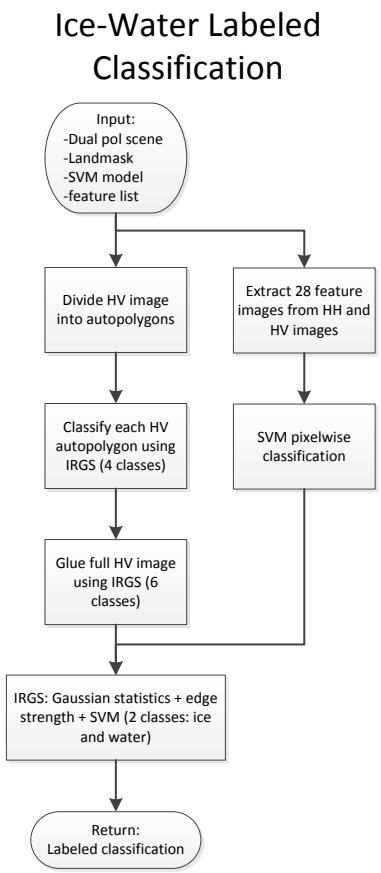

Figure 4.1: MAGIC system flowchart. Inputs are full size HH and HV images of a scene, a landmask, a previously trained SVM model, and a list of features to extract. The left track is glocal IRGS classification (Subsection 4.2). The right track is pixel-based SVM classification (Subsection 4.3). Classifications are merged using a modified IRGS approach (Subsection 4.4). The final result is a region image with ice-water labels.

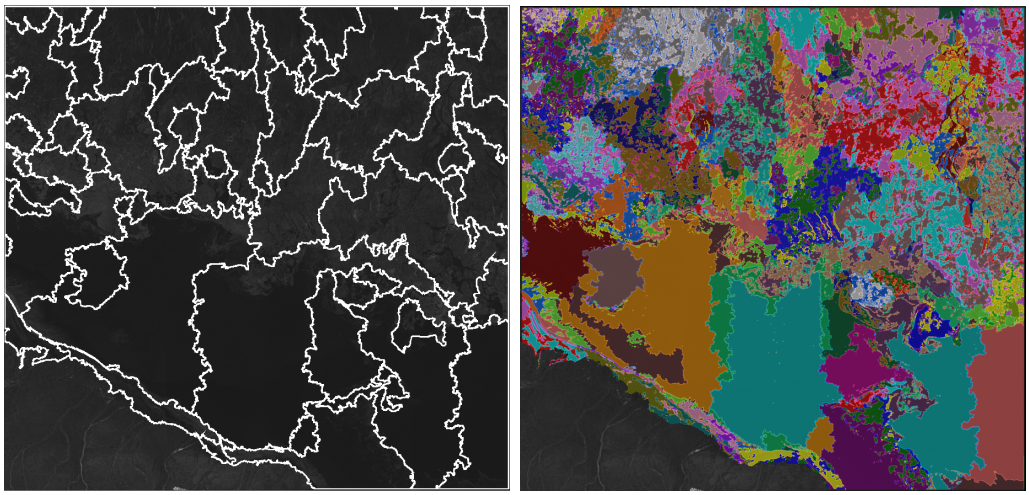

Figure 4.2: Local autopolygon classification of October 27, 2010 scene (Scene ID: 20101027). Left: autopolygon segmentation boundaries in white. Right: IRGS classification of each autopolygon. Up to four classes are distinguished in each autopolygon. Land in bottom left is ignored using a provided landmask.

spatial frequencies, only posing a problem at large scales and distances, typically greater than 50km. This observation motivates the use of the hierarchical classification method called "glocal" which combines large scale *global* and high detail *local* information in two steps. The local step divides the image into separate autopolygons and captures region-based information at a fine scale where class statistics can be considered stationary. This local step is explained next in Subsection 4.2.1. The global step then glues regions across autopolygons forming an image wide classification detailed in Subsection 4.2.2. Using this hierarchical approach the glocal method is able to provide robust, unsupervised, full scene classifications with arbitrary class labels.

4.2.1 Local Autopolygon Classification

The local autopolygon classification works on sub-regions, or autopolygons, of the HV image. The autopolygons are created using a watershed segmentation [42] with pruned minima. To select the pruned minima the image is divided into 144 square regions organized in a 12 by 12 grid. Within each square region the pixel with the lowest response from a gradient filter is selected as the pruned watershed seed. Autopolygons created in this way tend to be at roughly the same scale, have boundaries that follow natural image structures, and contain stationary class statistics. An example of an autopolygon segmentation is presented in Fig. 4.2.

For each autopolygon an IRGS classification is performed using only the HV polarization. IRGS is set to identify four classes in each autopolygon resulting in an over-segmentation ensuring all regions are homogenous and contain only ice or only water. Fig. 4.2 shows this classification for the October 27 scene.

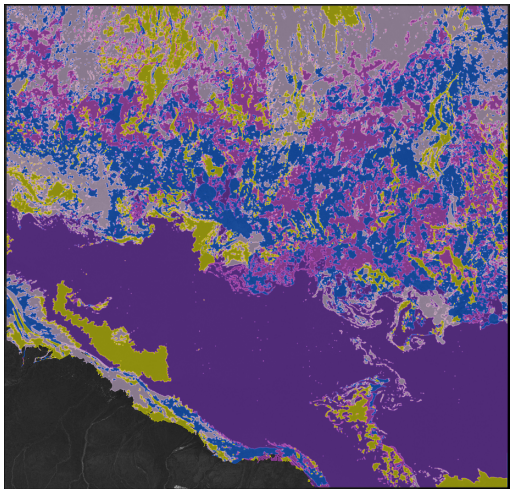

Figure 4.3: Glocal IRGS classification of October 27, 2010 scene (Scene ID: 20101027). The dark purple class segments out the water very well even across the large incidence angle variation. All other classes are various types of ice. Land in bottom left is ignored using a provided landmask.

4.2.2 Global IRGS Gluing

The global IRGS gluing step operates on regions identified during the local autopolygon classification and completes the full scene classification with arbitrary class labels. The autopolygon boundaries are first converted into region boundaries. IRGS then merges those regions across the entire HV scene identifying six distinct classes. An example of the final full scene glocal classification is shown in Fig. 4.3.

Experimentation was performed to determine preferred parameters for the glocal algorithm. The parameters mentioned previously (12 by 12 grid and 4 classes for local, and 6 classes for global) were found to perform well across all 21 scenes and algorithm performance was insensitive to minor parameter variations.

4.3 Classification with Ice-Water Labels: SVM Model

The result of the glocal classification is an identification of homogeneous regions with arbitrary class labels. To provide the required mapping to ice-water labels a soft-margin SVM [11] classifier is used. We implemented an SVM classifier using the popular LIBSVM [7] C/C++ library.

An SVM works by computing a linear decision boundary in a high dimensional space using the subset of labeled training samples near the decision boundary (called the support vectors). The SVM decision boundary equation is

$$b_S(x) = \sum_{\forall i} m_i w_i K(\bar{t}_i, \bar{t}) \tag{4.1}$$

where m_i are binary class labels, w_i are learned weights, $K()$ is a kernel function, $\bar{t}_i$

are the support vectors and $\bar{t}$ is the sample to be classified. Nonlinear kernels can be used to model nonlinear decision boundaries. The most common kernel function, and the one used in this thesis, is a radial basis function (RBF) kernel of the form

$$K(\bar{t}_i, \bar{t}) = \exp\left(-\gamma \left|\bar{t}_i - \bar{t}\right|^2\right) \tag{4.2}$$

where γ is a scaling hyperparameter.

SVM classifiers have numerous beneficial properties making them suitable for our application. Their ability to model nonlinear decision boundaries is essential for differentiating between the nonlinear SAR signatures of ice and water. SVM models trained with large amounts of data consume relatively little memory making them ideal for our large GIS datasets. Evaluation of test points in an SVM classifier is computationally efficient which is necessary for classifying the many points in a large RADARSAT-2 image. SVM classifiers have a strong theoretical foundation and often provide state-of-the-art classification accuracy.

Our SVM model was trained using a leave-one-out (LOO) method to gauge the likely operational performance. Training points were always *left out* for the scene being tested, leaving only the sample points from the remaining 60 scenes for training. This avoids testing on training data, gives realistic classification accuracies, and mimics an operational environment. The SAR scenes were found to contain significant amounts of redundant information making only a small subset of the image data useful for SVM training. Only 400 random sample points from each scene were necessary for training as more samples showed no improvement in classification performance. The training time for an LOO SVM model was under 5 minutes.

SVM models with RBF kernels, such as ours, generally require tuning of the complexity vs precision hyperparameter C and the RBF scale hyperparameter γ. A multiscale grid search [7] was performed to find satisfactory values for these hyperparameters of $C = 1$, $\gamma = 1$.

4.3.1 Feature Search

The choice of features as inputs to the SVM model can have a significant performance impact. Generally, using more features allows for better class separability, but too many features increases computation time as well as increasing the risk of the "curse of dimensionality" [1] with a reduction in generalizability. A feature search was used to determine a good set of SVM features and is shown in Fig 4.4. The feature search only considered 21 scenes from the full 61 scene dataset to reduce computation time. The subset of 21 scenes is presented in Table 4.1 and contains a variety of ice conditions including winter, melt, summer, and freeze-up seasons. Even with the significant scene count reduction, the execution time for the search was about four days. Execution time of the feature search is less critical because the search is performed offline and only needs to be done once. However, shortening the computation time contributed to an accelerated algorithm development process while still maintaining adequate operational performance.

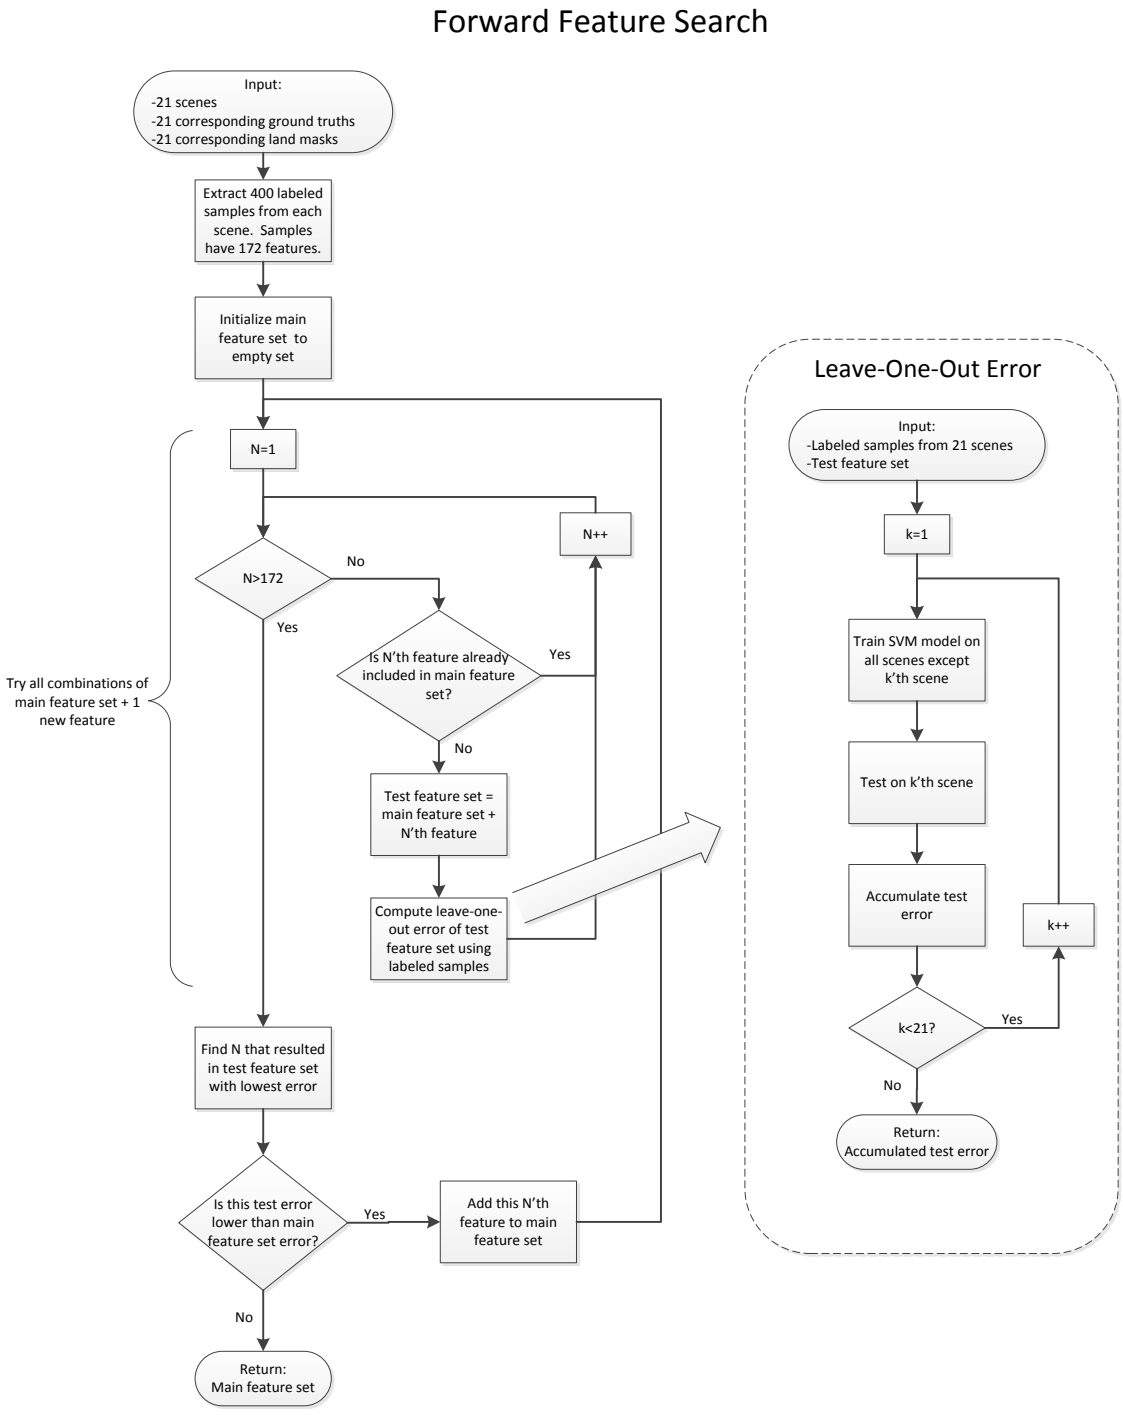

Figure 4.4: Feature search for SVM model. This is a forward feature search through 172 features minimizing leave-one-out error.

Table 4.1: Scenes used for feature search. The scenes encompass many ice conditions including winter, melt, summer, and freeze-up seasons.

Scene ID
20100418
20100426
20100510
20100524
20100605
20100623
20100629
20100712
20100721
20100730
20100807
20100816
20100907
20100909
20101003
20101021
20101027
20101114
20101120
20101206
20101214

The type of search used was a forward feature search [17] minimizing the 21 scene average LOO error. The feature search was used to select a subset of the 172 features in an initial feature set. All features were computed on both HH and HV images. The initial feature set included simple pixel intensities, local averages and maximum intensities in 5 by 5 and 25 by 25 pixel windows. Many GLCM features were included as well because they were found to proficiently differentiate between ice types and water [8]. The most significant concern when using GLCM features is the large amount of computation time required, however this issue is adequately mitigated by using a previously developed accelerated GLCM algorithm [10]. The GLCM measures used were:

- ASM: applied second moment
- CON: contrast
- COR: correlation
- DIS: dissimilarity
- ENT: entropy
- HOM: homogeneity
- INV: inverse moment
- MU: mean
- STD: standard deviation

A number of different GLCM window and isotropic displacements were computed for each measure:

Window Size [pixels]	Displacement [pixels]
5 by 5	1
11 by 11	1
25 by 25	1
25 by 25	5
51 by 51	5
51 by 51	10
51 by 51	20
101 by 101	10
101 by 101	20

The forward feature search terminated with 28 selected features shown in Table 4.2 in the order in which they were selected. Of the 28 features selected only four operated on the HV pol. A visual analysis of HH and HV pols suggests that the HV pol contains more valuable intensity information and the HH pol provides more textural information. Since it is well known that simple intensity cannot adequately

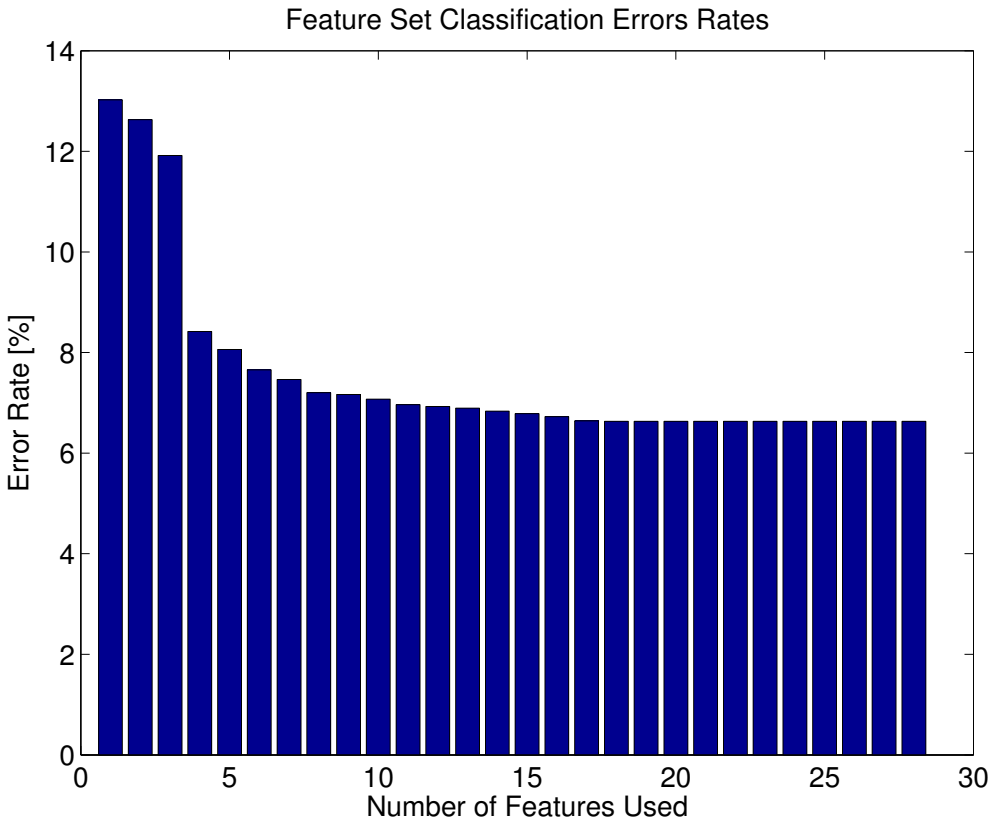

Figure 4.5: Classification error rates as more features are added to the feature set. The first featured added to the feature set is a GLCM mean of the HV pol. The first feature, when used alone, achieves an error rate of 13.02%. The second feature selected in the feature search is a GLCM correlation of the HH pol, and when used in conjunction with the first feature achieves an error rate of 12.63%. When all 28 features are used together the error rate is 6.63%. The first few features reduce the error rate significantly. At a feature set size of eight adding more features only improves the error rate by tenths of a percent.

distinguish ice from water it is not surprising that many more HH texture features than HV features were selected. Intensities are still important however, as six of the top ten features compute either intensity or some sort of averaging in a window. The error rate as each successive feature is added to the set is also shown in the table as well as plotted in Fig. 4.5. The error rate is decreased most significantly when the feature set is small. As the feature set grows beyond eight features adding additional features reduced the error rate by less than a tenth of a percent.

An example pixel-wise LOO SVM classification using the 28 selected features for the October 27 scene is displayed in Fig. 4.6. The SVM result is able to correctly assign ice-water labels to most pixels. Pixels that are near ice-water boundaries are more likely misclassified. The labels assigned by the SVM also have some salt-and-pepper-like noise due to a lack of spatial context. These limitations are overcome by combining the glocal classification described next.

Table 4.2: List of 28 selected features in order of selection by forward feature search. Error rate as more features are successively added to the feature set is shown in rightmost column and plotted in Fig. 4.5.

#	Pol	Feature	Error Rate [%]
1	HV	GLCM MU 25 by 25 step 5	13.02
2	HH	GLCM COR 51 by 51 step 5	12.63
3	HH	GLCM MU 25 by 25 step 1	11.92
4	HH	GLCM DIS 51 by 51 step 20	8.42
5	HH	GLCM ASM 101 by 101 step 5	8.06
6	HH	Intensity	7.65
7	HV	Average 25 by 25 window	7.46
8	HH	Average 5 by 5 window	7.20
9	HH	GLCM DIS 51 by 51 step 5	7.17
10	HH	GLCM MU 101 by 101 step 20	7.07
11	HH	GLCM MU 25 by 25 step 5	6.96
12	HH	GLCM ASM 51 by 51 step 5	6.93
13	HH	GLCM ASM 101 by 101 step 20	6.89
14	HH	GLCM MU 5 by 5 step 1	6.83
15	HV	GLCM COR 25 by 25 step 5	6.79
16	HV	GLCM COR 5 by 5 step 1	6.73
17	HH	Average 25 by 25 window	6.64
18	HH	GLCM STD 101 by 101 step 20	6.63
19	HH	GLCM CON 101 by 101 step 20	6.63
20	HH	GLCM CON 101 by 101 step 5	6.63
21	HH	GLCM ASM 11 by 11 step 1	6.63
22	HH	GLCM CON 11 by 11 step 1	6.63
23	HH	GLCM CON 25 by 25 step 1	6.63
24	HH	GLCM CON 25 by 25 step 5	6.63
25	HH	GLCM CON 5 by 5 step 1	6.63
26	HH	GLCM CON 51 by 51 step 10	6.63
27	HH	GLCM STD 11 by 11 step 1	6.63
28	HH	GLCM CON 51 by 51 step 20	6.63

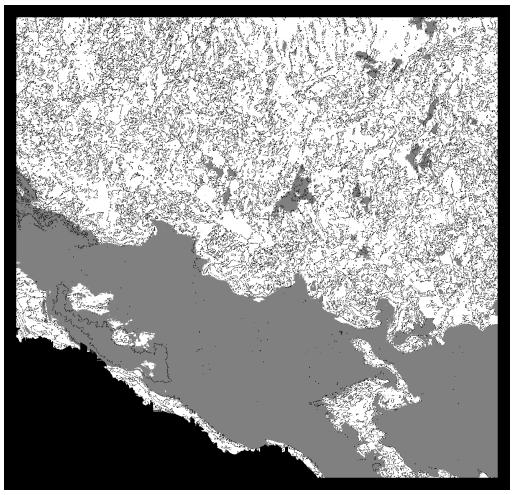

Figure 4.6: Pixel-wise LOO SVM classification with ice-water labels of October 27, 2010 scene (Scene ID: 20101027). White: water. Grey: ice. Black: unclassified. Pixels on or near land, near image borders, or on image region boundaries are not classified by the SVM. Though the pixel-wise classification provides valuable ice-water labels, it lacks spatial context information as is evident by the noisy result and poorly outlined ice-water boundaries.

4.4 Combination of Classifiers: Glocal + SVM

The glocal classification explained in Section 4.2 uses spatial context information to provide a robust classification with arbitrary class labels giving well identified ice-water boundaries. The SVM classification explained in Section 4.3 is able to provide the required ice-water labels using backscatter and texture information. We combine these two complimentary algorithms to generate a final binary ice-water classification.

Different combinations of classifiers have previously been used with success on SAR data [41]. Combinations of similar SVM+MRF classifiers have also been developed, but for different types of satellite data [40, 3, 13, 27, 46]. These SVM+MRF classifiers are discussed in Section 2.4. The glocal and SVM classifications are uniquely and inherently combined using the IRGS framework. The regions for IRGS are initialized from the glocal classification result. The ice-water labels for the regions are initialized using the SVM classification result. The overall IRGS energy function displayed in eq. 2.4 is modified to incorporate the additional SVM information as follows

$$E = \sum_{i \in R} [V_G(x_i) + V_S(x_i)] + \sum_{\langle i,j \rangle \xi} V_E(x_i, x_j) \tag{4.3}$$

where the energy term $V_G()$ accounts for the Gaussian statistics of all regions R. The energy term $V_E()$ scores the MRF clique potentials as edge strength between connected regions or cliques ξ. The x terms denote current region labels which are optimized to find the minimal overall energy. The new SVM energy term $V_S()$

accounts for the SVM label information on a per-region basis. The SVM energy term is calculated for each region as:

$$V_S(x_i) = C_S \sum_{j \in \Omega_i} \left[1 - \delta\left(x_i, p_j\right)\right] \tag{4.4}$$

p_j	SVM classification label of pixel j (0 = water , 1 = ice)
Ω_i	set of pixels in region i
x_i	current label of region i (0 = water , 1 = ice)
C_S	tuning weight
$\delta()$	Dirac delta function

The Dirac delta function of the energy term evaluates to 1 whenever both arguments are equal and 0 otherwise. The summation in the energy function counts all the pixels in the region who's labels p_j are not the same as the currently assigned region label x_i. High energy results when many pixel labels conflict with the region label, and when strong agreement exists between pixel labels and the region label the energy is low. The inclusion of the energy term $V_S()$ has the effect of biasing IRGS towards the labellings generated by the SVM classifier while still maintaining some influence from the Gaussian and edge strength terms. The influence of $V_S()$ on the overall IRGS energy function is controlled by the tuning weight C_s which is selected using empirical testing.

This final IRGS classification combining glocal and SVM was run with two classes to identify only ice and water. Results of this final classification are shown next.

Chapter 5

Testing and Results

5.1 Overall Results

The combined classifier was tested on all 61 scenes using an LOO testing scheme and the results were compared against our detailed ground truths. The resulting confusion matrix [39] entries for each scene are shown in Table 5.1. Some image pixels were not included in these statistics such as areas of land, areas lacking sensor data near image borders, and pixels on segmentation region boundaries. The classification accuracies for the data set are also shown in Table 5.1. The MAGIC system attained an overall average classification accuracy of 95.8%. The lowest classification accuracy occurred on the July 26, 2010 scene (Scene ID: 20100726) at 69.04% (shown later in Fig 5.3). Seven scenes had perfect classifications (Scene IDs: 20100418, 20100426, 201001120, 20101206, 20101214, 20110430, 20111121). These scenes were mainly comprised of a single class with no ice-water mixing leading to easier classifications. Of all the ice pixels in the data set 97.2% were classified correctly as ice and of all the water pixels 93.8% were classified correctly as water. A histogram of the classification accuracies, presented in Fig. 5.1, shows that the distribution of classification accuracies is heavily weighted at high percentages and has a tail towards lower percentages. Only two scenes had classification accuracies below 80%, and only six were below 90%. Two scenes from the data set have no classification statistics (Scene IDs: 20100816, 20111013) because the scenes are very difficult for even humans to interpret leaving no ground truth for an algorithm to be compared against. The scene 20100816 is specifically analyzed in Section 5.2.

Table 5.1: Classification confusion matrix tallies for all 61 scenes. Columns 2-5: true class/predicted class. Farthest right column is accuracy percentage.

Scene ID	ice/ice	water/water	ice/water	water/ice	%
20100418	6260105	0	0	0	100.00
20100426	6353751	0	0	0	100.00
20100510	6131265	258162	24376	15106	99.39
20100524	5555925	647726	22129	82727	98.34

Scene ID	ice/ice	water/water	ice/water	water/ice	%
20100605	5646913	0	0	46827	99.18
20100623	3926655	2072476	69151	31825	98.34
20100629	4511802	534951	144894	229422	93.10
20100704	3890930	1604946	944	26204	99.51
20100712	4715843	898954	246397	247416	91.92
20100719	3397073	1645727	117285	491383	89.23
20100721	3135627	2146542	111224	340462	92.12
20100726	2564477	1280370	931758	792224	69.04
20100730	2885209	2808890	244373	318584	91.00
20100806	1982669	3174060	307511	46676	93.57
20100807	1498534	4085618	186988	382525	90.75
20100816	NA	NA	NA	NA	NA
20100822	962768	4467757	77512	58071	97.56
20100829	1102687	4458510	494809	44251	91.16
20100907	0	5800642	0	12712	99.78
20100909	666491	5125252	167344	288135	92.71
20100927	0	5166842	0	672338	88.49
20100929	264275	5701840	45751	13567	99.02
20101001	0	5443729	0	28914	99.47
20101003	1410256	4568017	91416	115732	96.65
20101006	1528345	4379252	40218	9224	99.17
20101008	108497	5356617	6283	22027	99.48
20101013	3104941	2521180	448397	26970	92.21
20101014	0	5822952	0	109162	98.16
20101017	3210965	1346141	70382	850740	83.19
20101021	1252400	5085562	116999	76317	97.04
20101025	1504910	3737731	119519	106236	95.87
20101027	3640716	1938652	117699	194531	94.70
20101030	3568515	1221466	392024	77167	91.08
20101114	4867454	1003105	56585	236893	95.24
20101120	6245693	0	0	0	100.00
20101206	5631313	0	0	0	100.00
20101214	5880914	0	0	0	100.00
20110404	6351314	0	13194	0	99.79
20110418	6413314	0	33466	0	99.48
20110430	6220384	0	0	0	100.00
20110514	6237895	0	128923	0	97.98
20110530	5853051	189706	6722	19053	99.58
20110613	5987320	92663	17464	805	99.70
20110619	5659465	632243	10689	28440	99.38
20110626	5654142	704944	31282	50218	98.73
20110627	5681030	345927	87001	5164	98.49

Scene ID	ice/ice	water/water	ice/water	water/ice	%
20110709	3707524	2632010	5173	38946	99.31
20110710	4499308	1452921	28528	102142	97.85
20110720	3664242	2724413	11299	49069	99.06
20110725	3125155	3128457	101637	16752	98.14
20110811	2671216	3553044	16564	36237	99.16
20110817	1584890	4566107	40185	219004	95.96
20111005	1832395	4417450	141367	44843	97.11
20111006	10232	6299997	168	5136	99.92
20111013	NA	NA	NA	NA	NA
20111015	2431940	3616340	38857	102465	97.72
20111029	4613307	1745991	51660	1749	99.17
20111105	5661791	3905	27978	781181	87.50
20111106	3028317	2982835	3053	369510	94.16
20111113	2623149	2314882	325625	1027086	78.50
20111121	6133323	0	0	0	100.00
Total:	207052622	135707504	5772803	8892168	95.81

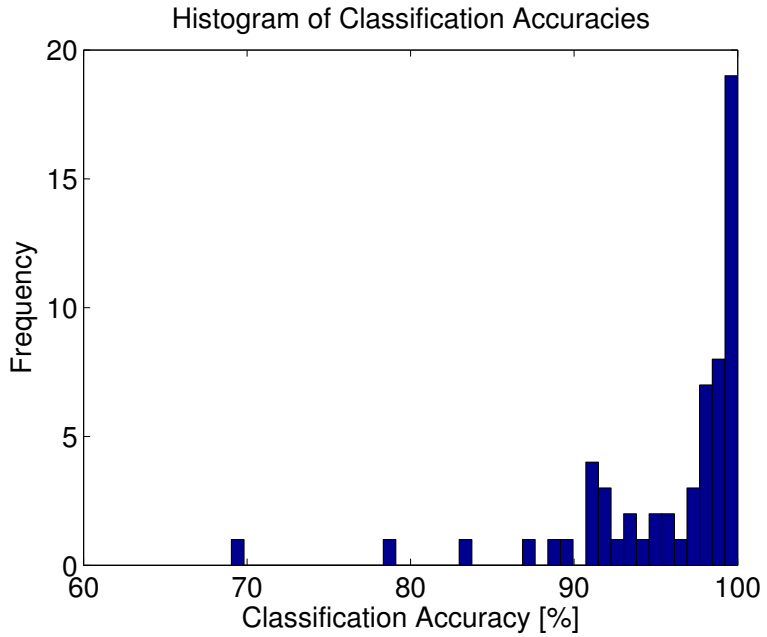

Figure 5.1: A histogram of classification accuracies for 60 scenes. All but two scenes are above 80% accurate and all but six are above 90% accurate.

The total execution time for the MAGIC system to classify a 2500 by 2500 pixel scene is under 30 minutes which is well within the CIS requirement of 1 hour. The processing times are scene dependent, but a typical breakdown is as follows. The local autopolygon classification takes about 2 minutes, IRGS glueing takes 10 seconds, the GLCM feature extraction takes about 15 minutes, and the final

combined IRGS+SVM step takes around 10 minutes. About 2GB of memory are required throughout the classification process.

5.2 Specific Scene Analysis

The final classification result for the October 27 scene, which was used previously in this paper in Figs 3.2, 3.3, 4.2, 4.3, and 4.6 to demonstrate all steps of the ice classification system, is displayed in Fig 5.2. This scene has a variety of ice types and water and is one of the more challenging scenes in the dataset. The MAGIC system achieves a classification accuracy of 94.70% on this scene. Some of the erroneous labellings are caused by the thin ribbons of wispy ice and small amounts of grease ice.

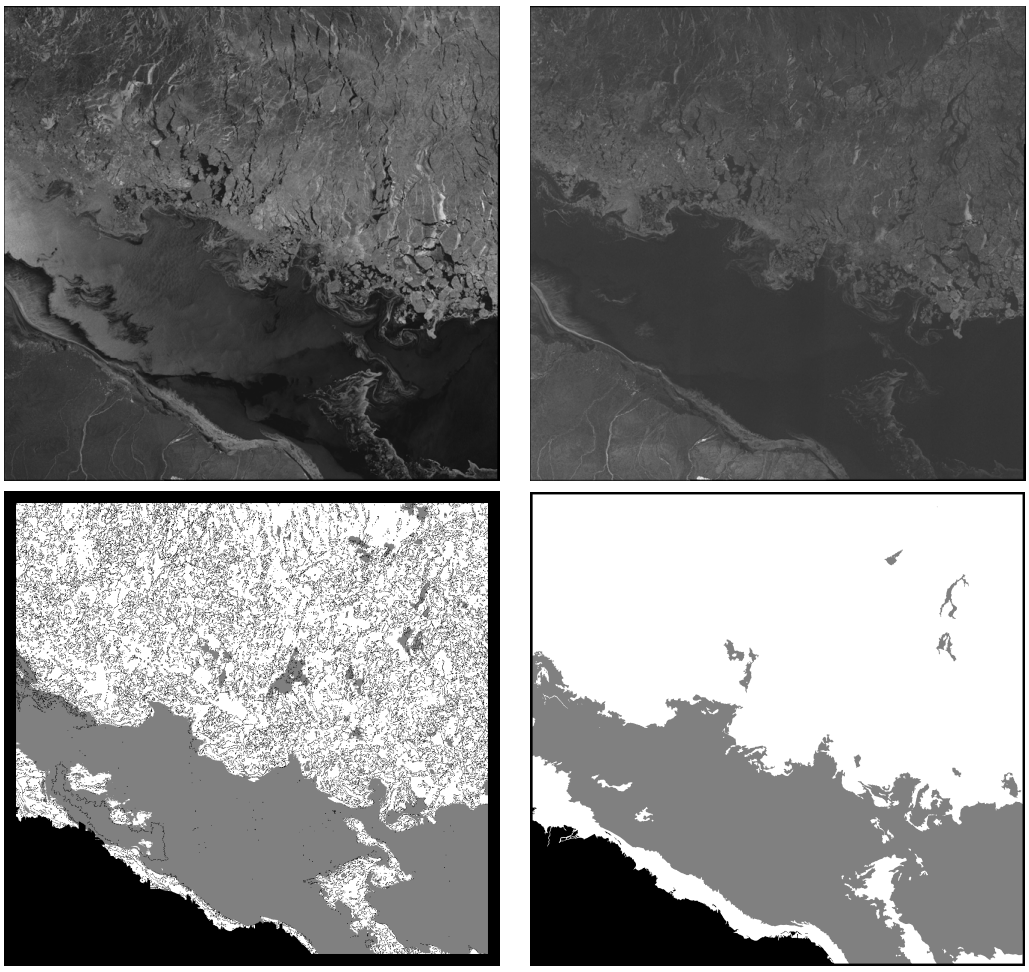

Figure 5.2: Final classification of October 27, 2010 scene (Scene ID: 20101027). Top left: HH pol. Top right: HV pol. Bottom left: SVM pixel-wise result. Bottom right: final result. In this scene land is at the bottom left, ice consists of the top third, and water is the dark area in the HV image. Classification accuracy: 94.70%. Refer to Appendix A for ground truth.

Of all 61 scenes MAGIC achieves the lowest accuracy of 69.04% on the July 26, 2010 scene shown in Fig 5.3. This is a complex scene with wispy ice and countless small floes. Since the pixel resolution for the downsampled images processed by MAGIC is only 200m by 200m there is a strong possibility of imaging subpixel ice floes. The significant ice and water mixing of this scene exacerbates this issue with many ambiguous pixels containing both ice and water leading to classification inaccuracies.

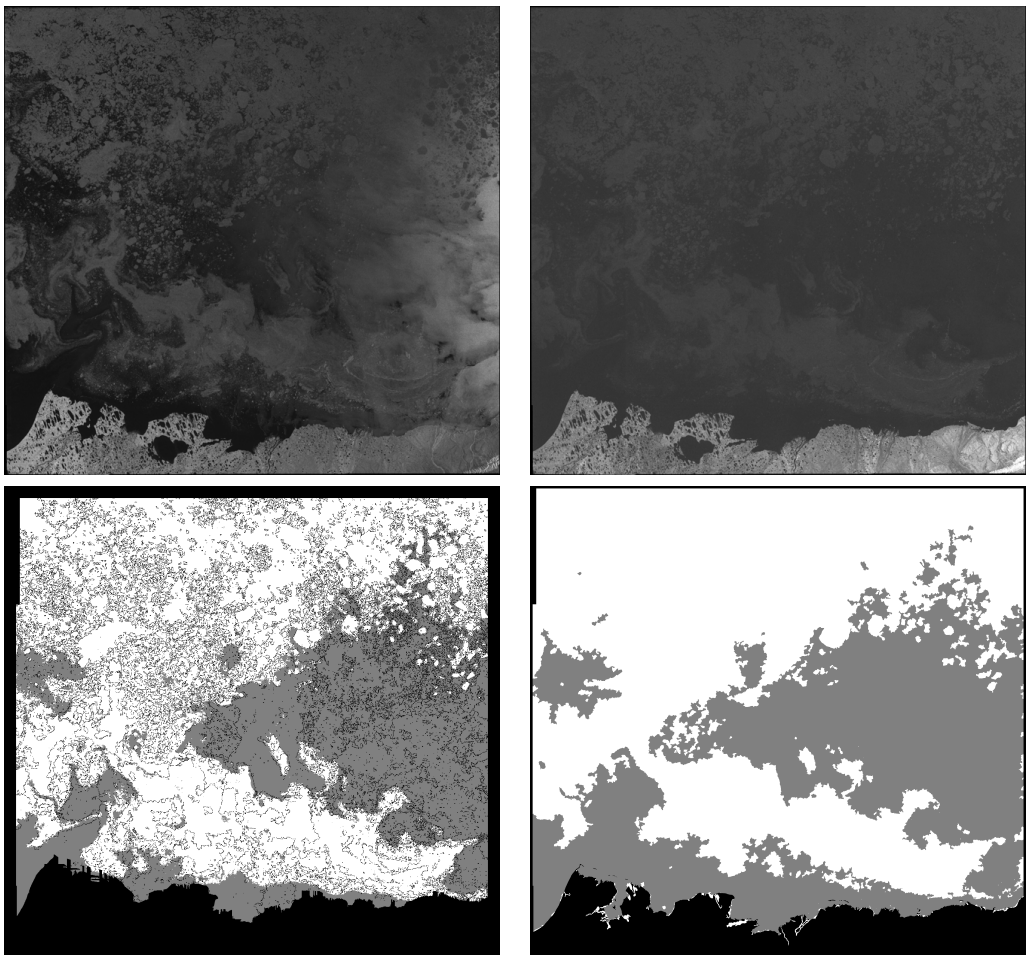

Figure 5.3: Final classification of July 26, 2010 scene (Scene ID: 20100726). Top left: HH pol. Top right: HV pol. Bottom left: SVM pixel-wise result. Bottom right: final result. Land is present in the bottom right corner. The top left corner contains high concentration ice. The remainder of the scene consists of water with differing concentrations of ice. Classification accuracy: 69.04%. Refer to Appendix A for ground truth.

A ground truth was never generated for the August 18, 2010 scene presented in Fig 5.4. The upper third of this scene has weather conditions that make discrimination of ice from water very difficult. As a result humans had low confidence in the ground truth generated for this scene. Leaving this scene out of the training process helped ensure the quality of the training samples. A classification result

can still be generated for this scene, though the classification statistics cannot be computed without a ground truth to compare to. The scene captured October 13, 2011 (Scene ID: 20111013) was omitted from accuracy assessments for similar reasons.

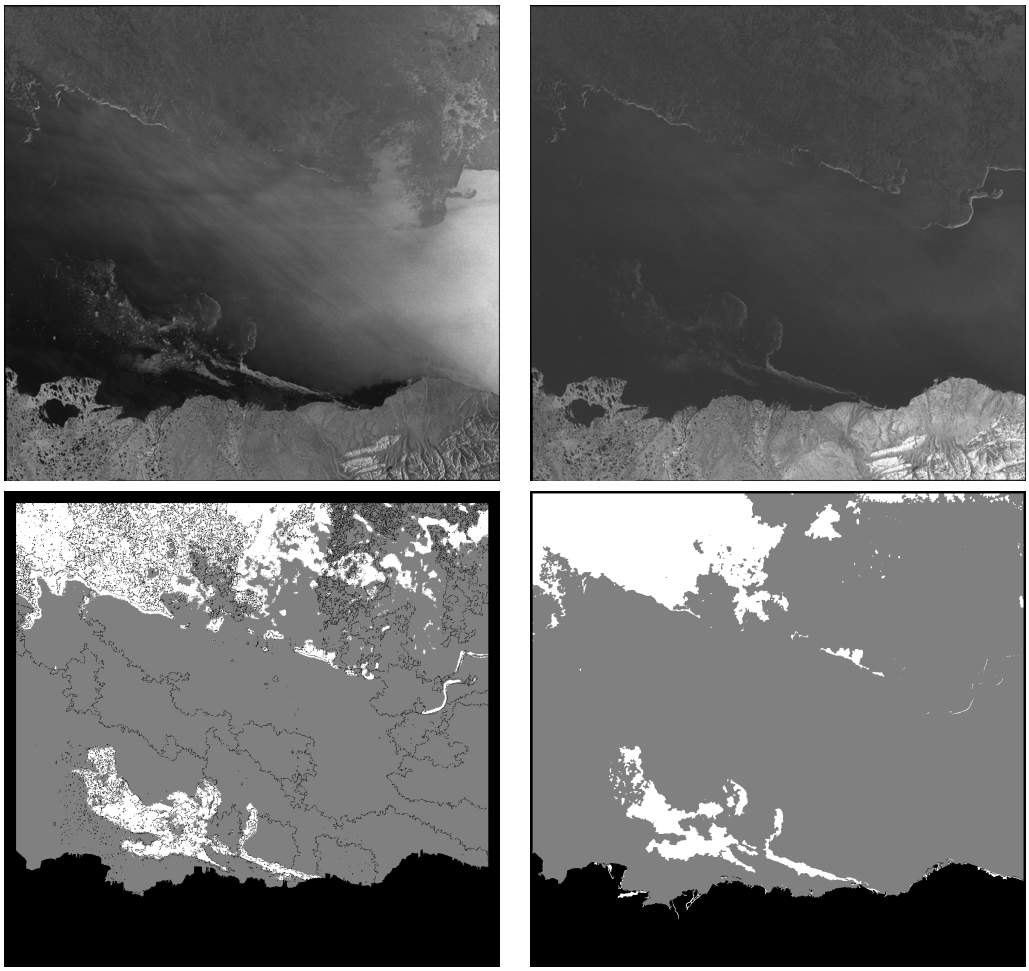

Figure 5.4: Final classification of August 18, 2010 scene (Scene ID: 20100816). Top left: HH pol. Top right: HV pol. Bottom left: SVM pixel-wise result. Bottom right: final result. Land is present along the bottom. Above the land is a thorough mixture of ice and water. This is the only scene of the data set with no ground truth. This is because the complexity of the scene, attributed to ice-water, mixing made mapping for humans very difficult. As a result no classification accuracy was determined for this scene. Refer to Appendix A for ground truth.

The scene of October 21, presented in Fig 5.5, has a number of challenging aspects. Some minor sections of grease ice exist in the upper left of the scene which has markedly different backscatter characteristics than most other types of ice and appears very similar to calm water. These minor sections in this scene are misclassified as water. Another challenging aspect is the noticeable banding artifacts most prevalent in the HV image appearing as about four vertical bands

contributing false structure and misleading statistics to the image. The MAGIC system effectively mitigates this issue with no noticeable effect on the final result. A third, and most challenging aspect of this scene is the highly textured nature of the water, likely a result of significant wind roughening. Typically ice has a much stronger texture response than water leading to confusion when water becomes highly textured. Again the MAGIC system is robust to these phenomenon and achieves a high classification accuracy for this scene of 97.04%. The mutual benefit from the combination of the glocal and SVM classifications is also evident in this scene. Along the right edge the SVM classifier incorrectly identifies sizable sections of ice as water and many speckles of water as ice. Most of these classification errors are eliminated by the combined classifier because they are not significant enough to affect the overall average label of the regions identified in the glocal step.

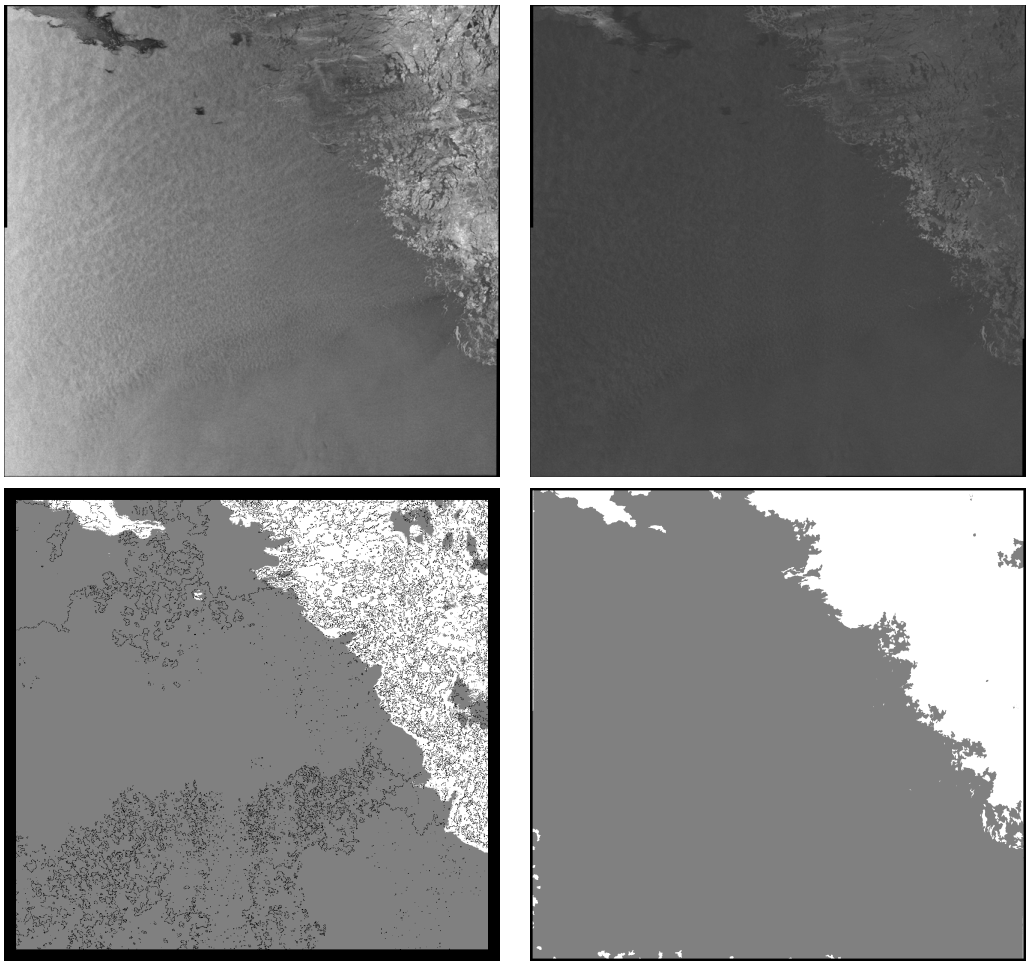

Figure 5.5: Final classification of October 21, 2010 scene (Scene ID: 20101021). Top left: HH pol. Top right: HV pol. Bottom left: SVM pixel-wise result. Bottom right: final result. There is no land in this image. The ice covers the top right and a small section at the top left of the scene. The remainder is wind roughened water. Classification accuracy: 97.04%. Refer to Appendix A for ground truth.

5.3 Identified Limitations

The full data set testing revealed two limitations of the current system. One limitation is the misclassification of grease ice as water shown in Fig 5.5. Grease ice has a very dark backscatter profile and a weak texture response, both of which are properties more commonly found in water. This similarity with water leads to poor class separability and corresponding classification errors. Though robust distinction of all types of ice from water is desirable, identification of grease ice is less of an issue than most other ice types. Grease ice is one of the first stages of new ice growth and is a non-issue for navigational purposes because ships can easily cut through its very thin and soupy consistency [4].

Subpixel ice floes are another challenge for any classification system. The 200m by 200m pixel size makes it highly possible for a single pixel to contain both small sub-resolution ice floes and water. A binary ice or water label is insufficient in such cases and can result in misleading over or under estimation of ice concentrations. The simplest solution to this problem would be to eliminate image downsampling and use original image sizes with a resolution of 50m by 50m, though this would increase computation time and still not address floes smaller than 50m. Another approach for dealing with subpixel floes could be to estimate the ice concentrations of each pixel. However, it remains to be seen whether issues attributed to subpixel floes are significant enough to necessitate further development.

The region based approach of the MAGIC system permits simple correction of some classification errors. A large error attributed to more ambiguous ice signatures on the November 6, 2011 scene is shown in Fig. 5.6. A large section of water above the land is misclassified as ice. This large section is contained within a single IRGS region. By manually flipping this large region along with two smaller regions (three mouse clicks) the classification accuracy for the scene can be increased from 94.16% to 99.44% in less then a minute of manual processing time.

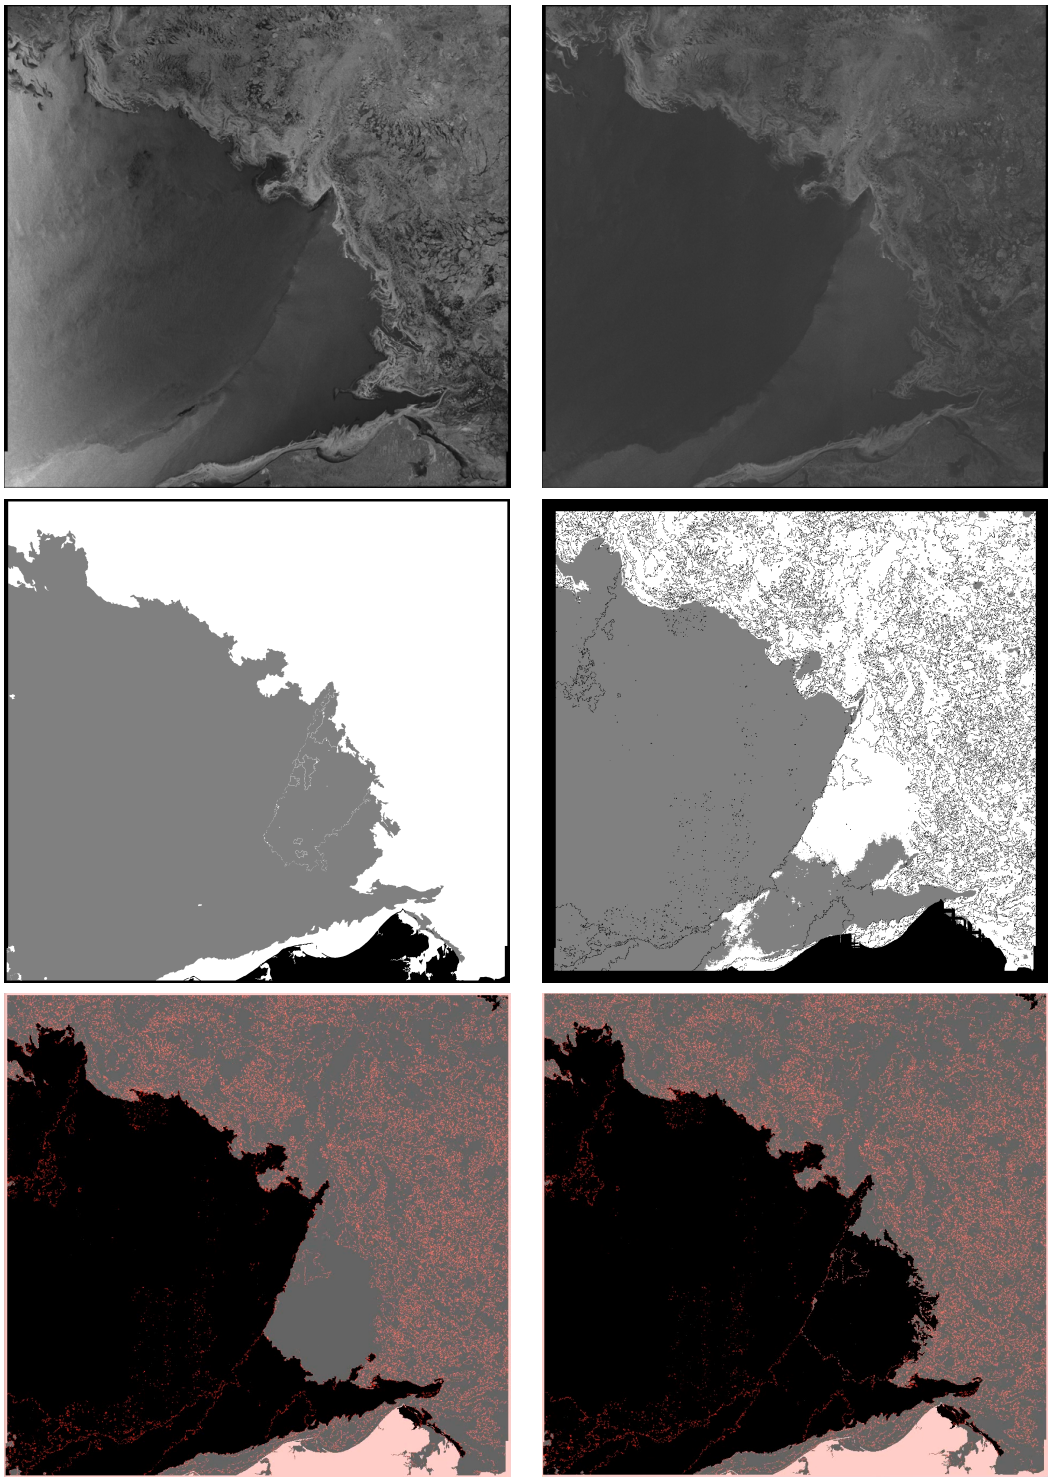

Figure 5.6: Polygon correction. Scene ID: 20111106. Top left: HH. Top right: HV. Mid left: ground truth. Mid right: pixel-wise classification. Bottom left: initial classification result. Bottom right: corrected classification result. Classification accuracy before: 94.16%. Classification accuracy after: 99.44%. Manual processing time <1min. (note: colour codes are not consistent. For mid row: white=ice, grey=water, black=land. For bottom row: white=land, grey=ice, black=water). Refer to Appendix A for ground truth.

Chapter 6

Conclusion

6.1 Summary

Mapping ice locations in ocean bodies is crucial for several valuable activities including maritime navigation and environmental science. CIS has traditionally generated ice maps using a manual process involving expert ice analysts' interpretation of satellite imagery. CIS has expressed a desire for an automated system able to process more scenes, in higher detail, and with less bias to augment the manual process. The MAGIC classification system presented in this thesis aimed to satisfy this need.

The MAGIC system consists of three main steps: glocal, SVM, and SVM+MRF. The glocal step hierarchically incorporates spatial context information into the classification process using the watershed algorithm at a local scale and the IRGS algorithm at the global scale. The SVM works on a pixel-wise basis and computes true ice-water labels primarily using GLCM texture features as inputs. The information from the glocal and SVM steps is then combined in the SVM+MRF step by again using the IRGS algorithm, but with a modified energy function. The final result is a region image with true ice-water labels.

The MAGIC system was able to overcome several key challenges inherent in ice-water classification of SAR imagery including very large scene sizes, speckle noise, non-stationary statistics, and nonlinear class statistics. The MAGIC system is capable of processing large 2500 by 2500 pixel scenes in less then 30 minutes by using an efficient region based MRF approach and a fast SVM classifier. The segmentation steps of the MAGIC system are able to greatly reduce the influence of multiplicative speckle noise by modeling statistical properties of regions of pixels instead of focusing on individual pixel intensities. The MAGIC system operates on dual pol imagery incorporating the HH pol which contains valuable textural information and the HV pol which is significantly less affected by any incidence angle effects. The highly nonlinear class statistics are modeled using GLCM texture features and an SVM classifier with a nonlinear RBF kernel.

Using these methods MAGIC is able to overcome significant challenges and achieve classification accuracies adequate for operational performance. MAGIC

was tested in an LOO manner on a 61 scene data set and attained an overall classification accuracy of 95.8%. MACIC achieved a classification accuracy of 90% or greater 88% of the time. Ice pixels were classified correctly 97.2% of the time, and water pixels were classified with 93.8% accuracy. The MAGIC system is now under consideration by CIS for operational implementation.

6.2 Future Work

Though MAGIC has solved key challenges and achieved adequate classification accuracies there are still many areas for improvement which are listed below.

Fast features. The feature search, explained in Subsection 4.3.1, made no consideration for feature computation time. If two features were equally good at discriminating ice and water the feature search would select between the two randomly, even if one feature took significantly longer to compute. The final feature set of 28 features used about 15 minutes of the total 25 minutes it took to process a scene with the MAGIC system. The feature search could be modified to optimize a discrimination efficacy per computation time metric. Another option could be to consider adding features that are naturally fast to compute to the feature search set (e.g. Haar-like features [26]). Faster feature computation could allow for faster generation of classification results, or for more detailed classifications in the same amount of time. In the later case higher detailed classification would help mitigate the sub-pixel floe issue identified in Section 5.3.

Classification confidence. Currently the MAGIC system only outputs binary values indicating the predicted ice-water class. Providing an estimate of the confidence of those predictions would likely also be of value. Ice analysts could better judge when to trust a computer classification or when to override it with a manual classification. The confidence information would also be valuable when incorporating the generated ice maps into other scientific models. Providing confidence estimates would likely require only minimal additional work as the SVM model developed in this work has the capability to provide probabilistic outputs.

Ancillary features. Expert ice analysts typically use several information sources in addition to SAR imagery to improve the accuracy of their predictions. Other sources that could be explored include incorporating a previous day's classifications, local weather predictions such as temperature and wind speed, and other imaging modes such as optical or hyperspectral bands. Some information that exists within the current data set that was not utilized includes incidence angle at each pixel, lat/lon of each pixel, and distance from land.

Active and on-line learning. Generating accurate ground truths is a difficult and time consuming task. Active learning theory attempts to maximizing the accuracy of a classifier using a minimal amount of purposefully selected training samples

[36]. Active learning could greatly reduce the amount of time needed to generate ground truth, especially for SAR scenes that contain significant amounts of redundant data. Moreover, since far fewer samples need to be labeled, the training data could be of higher quality because the user has more time to dedicate to labeling each sample. Another method that has potential to improve the performance of a remote sensing classification system is on-line learning [2]. Expert ice analysts would likely review each automated classification and could fix any errors using something like the polygon correction algorithm shown in Section 5.3. The labeling information from the corrected polygon could then be used to updated the classifier so similar mistakes are less likely in the future.

Ice typing. The goal of this thesis work was to discriminate between ice and water, though a natural extension would be also to discriminate between the various ice types. Significant research has already been published on automated ice typing methods, some of which is reviewed in Section 2.2. An ice typing algorithm could avoid the difficult problem of first separating the ice from the water by using the MAGIC system as a preprocessing step.

New regions. In this thesis work the Beaufort Sea was the only area studied. Automated ice-water discrimination algorithms have applications in numerous locations containing mixtures of water and ice. Many different areas have been previously studied as shown in Table 2.3. The presented algorithm would likely work to some degree on new areas, though different ice formations and weather patterns could cause confusion. It would be of interest to know the performance of this algorithm on new regions, and whether region specific parameters and features could improve performance.

References

[1] Christopher M Bishop. *Pattern Recognition and Machine Learning*. New York: springer, vol. 1 edition, 2006. 23

[2] Avrim Blum. On-line algorithms in machine learning, 1998. 42

[3] Francesca Bovolo and Lorenzo Bruzzone. A context-sensitive technique based on support vector machines for image classification. *Pattern Recognition and Mach. Intell.*, pages 260–265, 2005. 13, 29

[4] Canadian Ice Service - Environment Canada. MANICE: Manual of standard procedures for observing and reporting ice conditions, 2005. 38

[5] Canadian Space Agency (CSA). RADARSAT-2, 2011. 15

[6] Yee Kit Chan and Voon Chet Koo. An introduction to synthetic aperture radar (SAR). *Progress in Electromagnetics Research*, B(2):27–60, 2008. 5

[7] Chih Chang and Chih Lin. LIBSVM: A library for support vector machines. *ACM Trans. on Intell. Syst. and Technol.*, 2(3):27:1–27:27, 2011. 22, 23

[8] D. A. Clausi. Comparison and fusion of co-occurrence, Gabor and MRF texture features for classification of SAR sea-ice imagery. *Atmosphere-Ocean*, 39(3):183–194, 2001. 26

[9] D A Clausi, A K Qin, M S Chowdhury, P Yu, and P Maillard. MAGIC: MAp-Guided Ice Classification system. *Canadian J. of Remote Sens.*, 36(1):S13—-S25, December 2010. 1, 8, 9, 10

[10] D. A. Clausi and Y. Zhao. Grey level co-occurrence integrated algorithm (GLCIA): a superior computational method to rapidly determine co-occurrence probability texture features. *Comput. & Geosci.*, 29(7):837–850, August 2003. 26

[11] Cortes Corinna and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995. 22

[12] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39(1):1–38, 1977. 12

[13] Aly A Farag, Refaat M Mohamed, and Ayman El-baz. A unified framework for MAP estimation in remote sensing image segmentation. *IEEE Trans. Geosci. Remote Sens.*, 43(7):1617–1634, 2005. 13, 29

[14] F. Fetterer, C. Bertoia, and J. P. Ye. Multi-year ice concentration from RADARSAT. *IEEE Trans. Geosci. Remote Sens.*, 1:402–404, 1997. 7, 8, 9, 10

[15] Torsten Geldsetzer. Mapping and monitoring lake ice using SAR satellites. Technical Report March, CCRS, 2010. 7, 8, 9, 10, 12

[16] Rashpal S Gill. SAR surface cover classification using distribution matching. Technical report, Danish Meteorological Institute, 2002. 7, 8, 9, 10

[17] Isabelle Guyon and Andre Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003. 26

[18] Jörg Haarpaintner and Stian Solbø. Automatic ice-ocean discrimination in SAR imagery. Technical report, Norut IT-report, 2007. 6, 7, 8, 9, 10

[19] D. Haverkamp, L. K. Soh, and C. Tsatsoulis. A comprehensive, automated approach to determining sea ice thickness from SAR data. *IEEE Trans. Geosci. Remote Sens.*, 33(1):46–57, 1995. 7, 8, 9, 10

[20] Ernst Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925. 13

[21] J. A. Karvonen. Baltic Sea ice SAR segmentation and classification using modified pulse-coupled neural networks. *IEEE Trans. Geosci. Remote Sens.*, 42(7):1566–1574, 2004. 7, 8, 9, 10

[22] Juha Karvonen, Markku Similä, and Marko Mäkynen. An iterative incidence angle normalization algorithm for sea ice SAR images. *IEEE Int. Geosci. Remote Sens. Symp.*, 3:1524–1527, 2002. 6

[23] Juha Karvonen, Markku Similä, and Marko Mäkynen. Open water detection from Baltic Sea ice RADARSAT-1 SAR imagery. *IEEE Geosci. Remote Sens. Lett.*, 2(3):275–279, 2005. 7, 8, 9, 10

[24] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. 13

[25] Stan Z. Li. *Markov random field modeling in image analysis*. Springer, 2009. 13

[26] Rainer Lienhard and Jochen Maydt. An extended set of Haar-like features for rapid object detection. *IEEE Int. Conf. on Image Processing*, 1:900–903, 2002. 41

[27] Desheng Liu, Maggi Kelly, and Peng Gong. A spatial-temporal approach to monitoring forest disease spread using multi-temporal high spatial resolution imagery. *Remote Sens. of Environment*, 101(2):167–180, March 2006. 13, 29

[28] Maria Lundhaug. ERS SAR studies of sea ice signatures in the Pechora Sea and Kara Sea region. *Canadian J. of Remote Sens.*, 28(2):114–127, April 2002. 7, 8, 9, 10

[29] MacDonald Dettwiler and Associates (MDA). RADARSAT-2 – product details. Technical Report 1, 2011. 15

[30] Marko P Mäkynen, A Terhikki Manninen, Markku H Similä, Juha A Karvonen, and Martti T Hallikainen. Incidence angle dependence of the statistical properties of C-band HH-polarization backscattering signatures of the Baltic Sea ice. *IEEE Trans. Geosci. Remote Sens.*, 40(12):2593–2605, 2002. 1, 6

[31] Shuhratchon Ochilov and David A Clausi. Operational SAR sea-ice image classification. *IEEE Trans. Geosci. Remote Sens.*, 50(11):4397 – 4408, 2012. 17

[32] Chris Oliver and Shaun Quegan. *Understanding Synthetic Aperture Radar*. Artech House Boston London, 1998. 1, 3

[33] B. Scheuchl, R. Caves, I. Cumming, and G. Staples. Automated sea ice classification using spaceborne polarimetric SAR data. *IEEE Int. Geosci. Remote Sens. Symp.*, 7:3117–3119, 2001. 7, 8, 9, 10

[34] B Scheuchl, I Hajnsek, and I Cumming. Sea ice classification using multifrequency polarimetric SAR data. Technical report, 2002. 7, 8, 9, 10

[35] Bernd Scheuchl, Dean Flett, Ron Caves, and Ian Cumming. Potential of RADARSAT-2 data for operational sea ice monitoring. *Canadian J. of Remote Sens.*, 30(3):448–461, 2004. 1

[36] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin, Madison, 2010. 42

[37] Merrill I Skolnik. Introduction to Radar. *Radar Handbook*, 1962. 5

[38] L. K. Soh, C. Tsatsoulis, D. Gineris, and C. Bertoia. ARKTOS: An intelligent system for SAR sea ice image classification. *IEEE Trans. Geosci. Remote Sens.*, 42(1):229–248, 2004. 7, 8, 9, 10, 12

[39] Stephen V. Stehman. Selecting and interpreting measures of thematic classification accuracy. *Remote Sens. of Environment*, 62(1):77–89, October 1997. 31

[40] Yuliya Tarabalka, Mathieu Fauvel, Jocelyn Chanussot, and Jon Atli Banadiktsson. SVM- and MRF-based method for accurate classification of hyperspectral images. *IEEE Geosci. Remote Sens. Lett.*, 7(4):736–740, 2010. 13, 29

[41] R. Touzi, W M Boerner, J S Lee, and E. Lueneburg. A review of polarimetry in the context of synthetic aperture radar: Concepts and information extraction. *Canadian J. of Remote Sens.*, 30(3):380–407, June 2004. 29

[42] Luc Vincent and Pierre Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(6):583–598, 1991. 12, 21

[43] World Meteorological Organization (WMO). World Meteorological Organization. 17

[44] Q. Yu and D. A. Clausi. IRGS: Image segmentation using edge penalties and region growing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(12):2126–39, December 2008. 1, 2, 7, 8, 9, 10, 12

[45] Natalia Yu. Zakhvatkina, Vitaly Yu. Alexandrov, Ola M. Johannessen, Stein Sandven, and Ivan Ye. Frolov. Classification of sea ice types in ENVISAT synthetic aperture radar images. *IEEE Trans. Geosci. Remote Sens.*, 51(5):2587–2600, May 2013. 7, 8, 9, 10

[46] Bing Zhang, Shanshan Li, Xiuping Jia, Lianru Gao, and Man Peng. Adaptive Markov random field approach for classification of hyperspectral imagery. *IEEE Geosci. Remote Sens. Lett.*, 8(5):973–977, 2011. 13, 29

Appendix A

Full Data Set Classification Result Images

This appendix contains the final classification results from the MAGIC system for the full 61 scene data set.

Data set display format:

Scene ID

HH - Reduced to 625 by 625 pixels for display purposes.	**HV** - Reduced to 625 by 625 pixels and contrast enhanced for display purposes.
landmask - black: land, white: other	**groundtruth** - black: unclassified, grey: water, white: ice
pixelwise SVM - black: unclassified, grey: water, white: ice	**final result** - black: unclassified, grey: water, white: ice

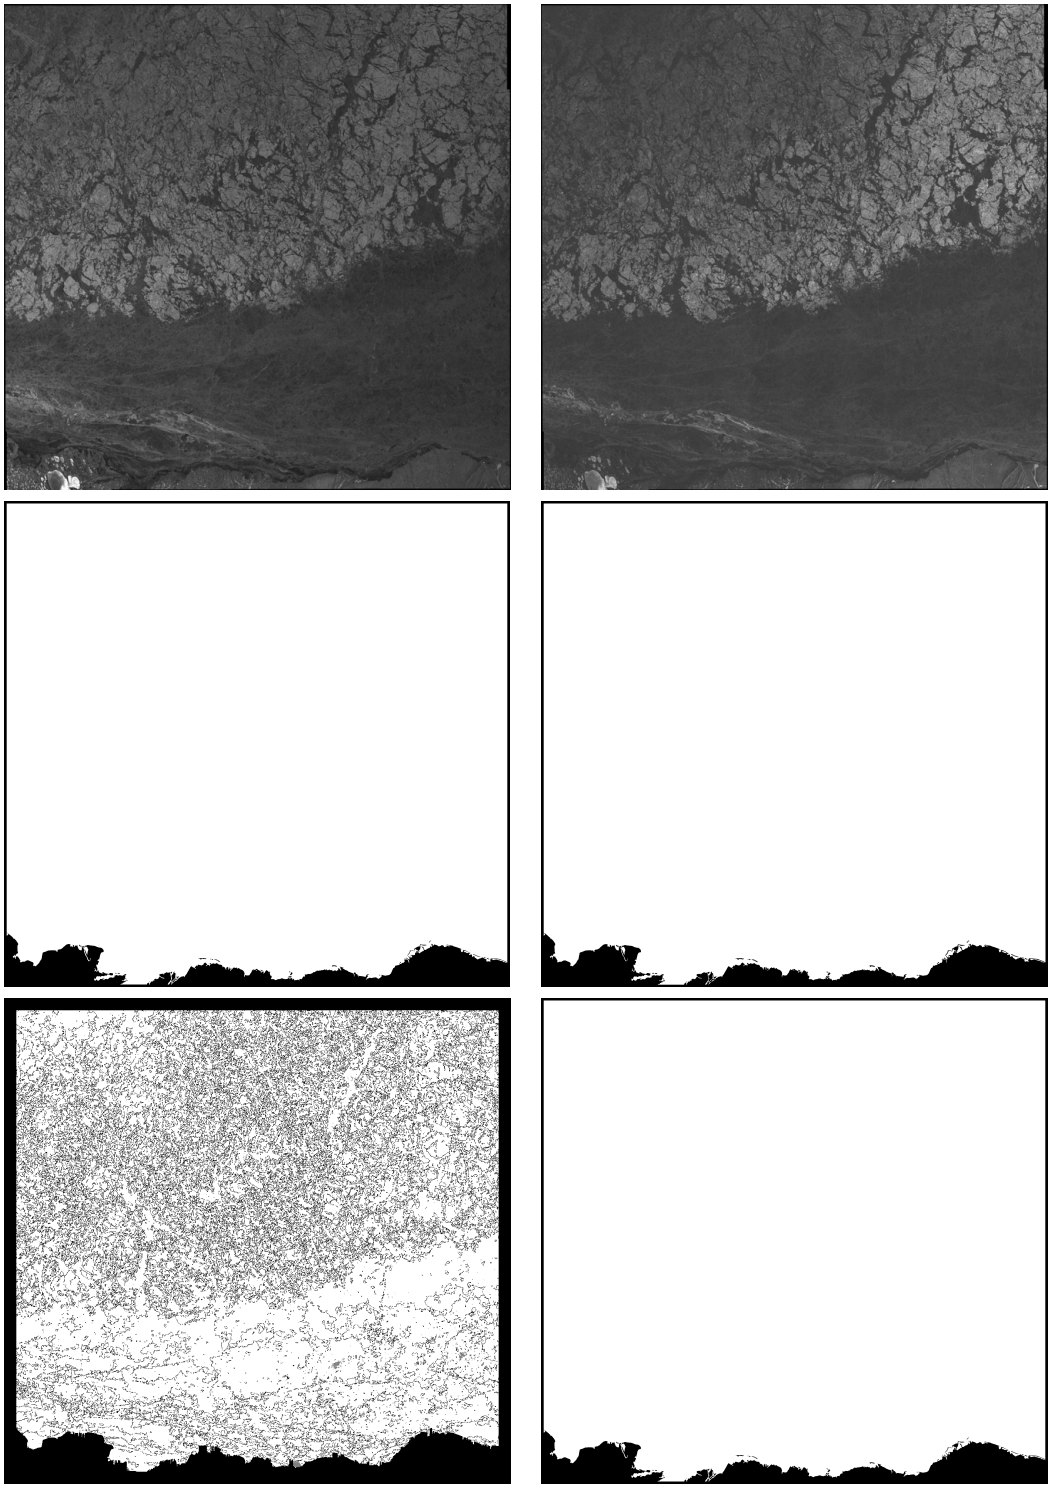

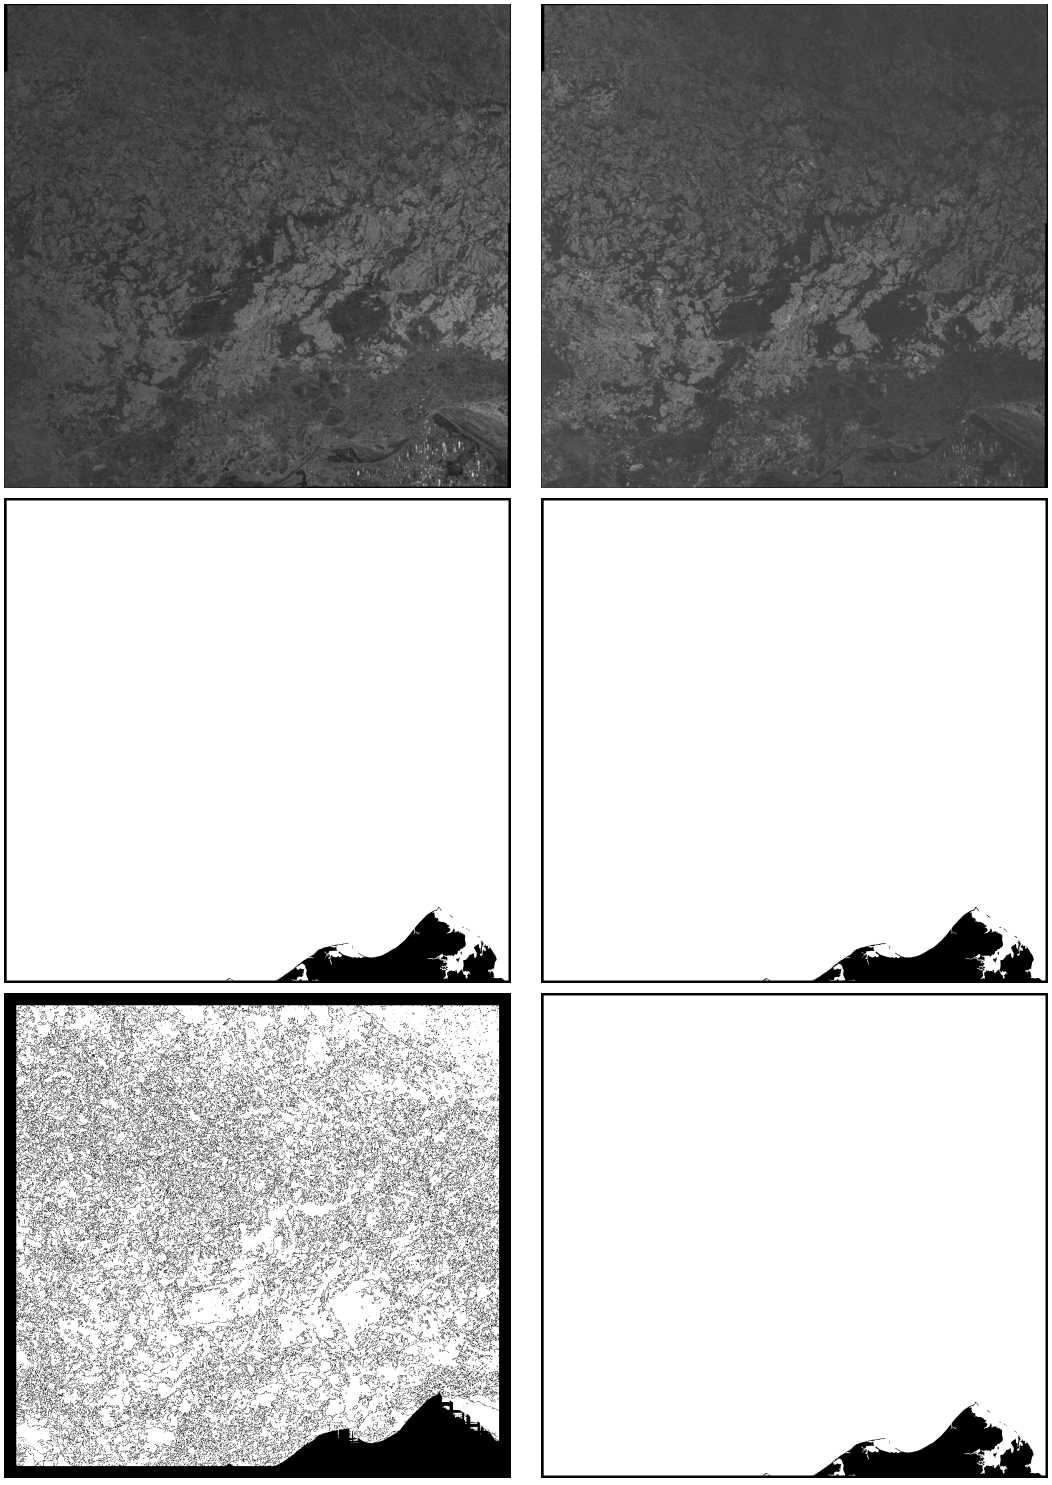

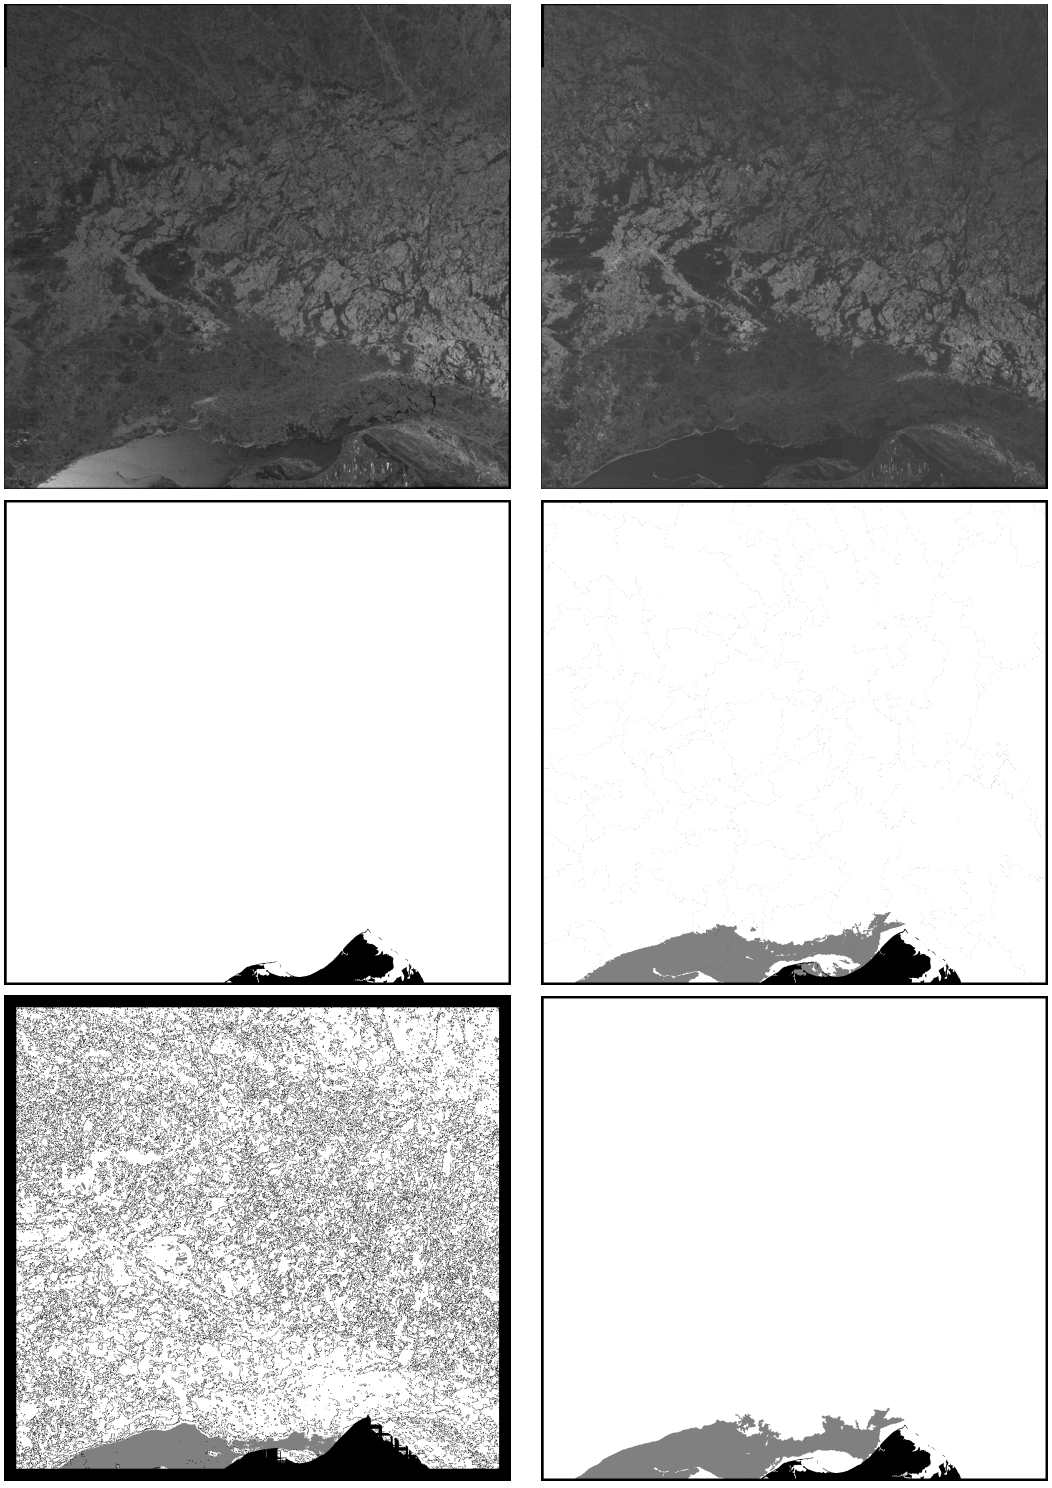

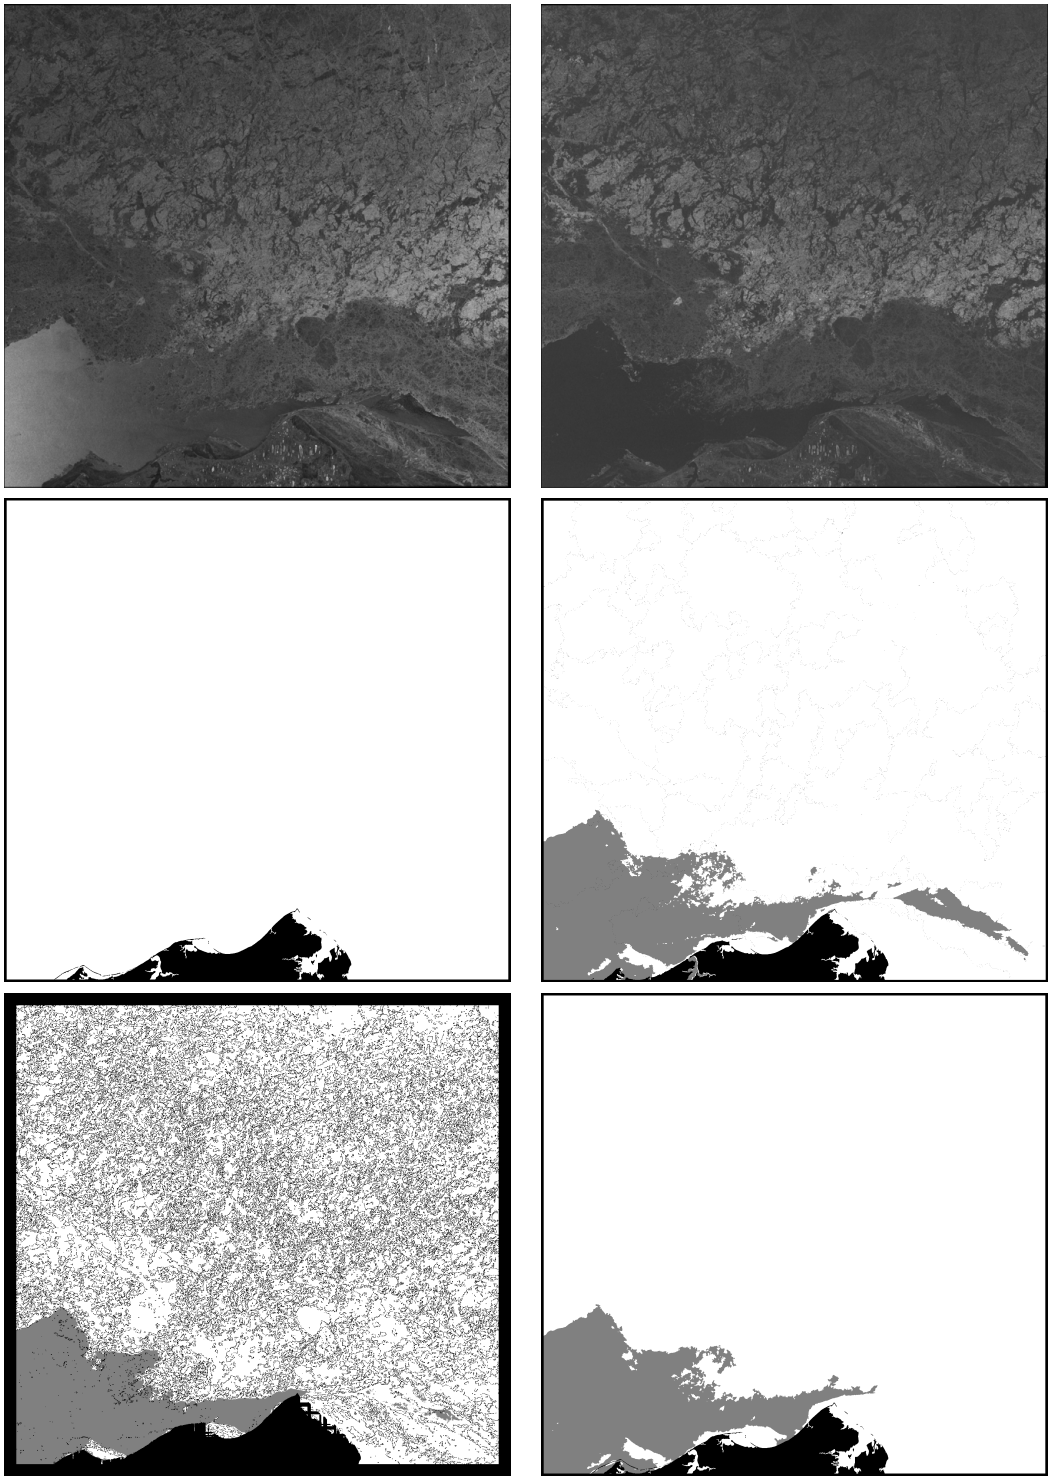

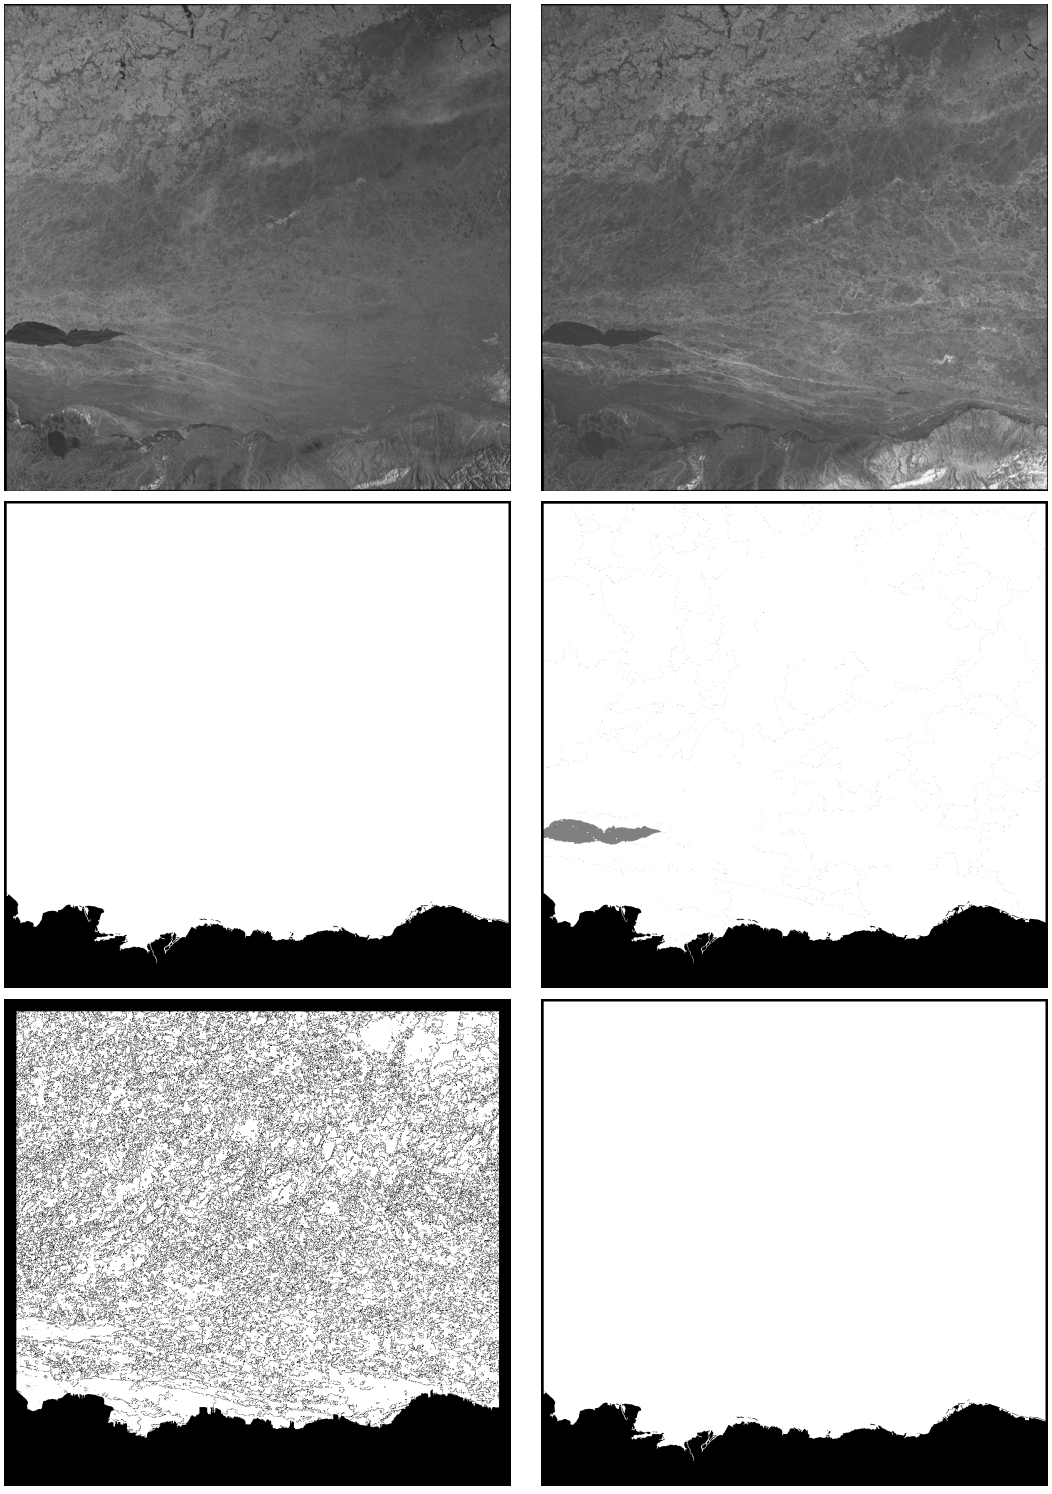

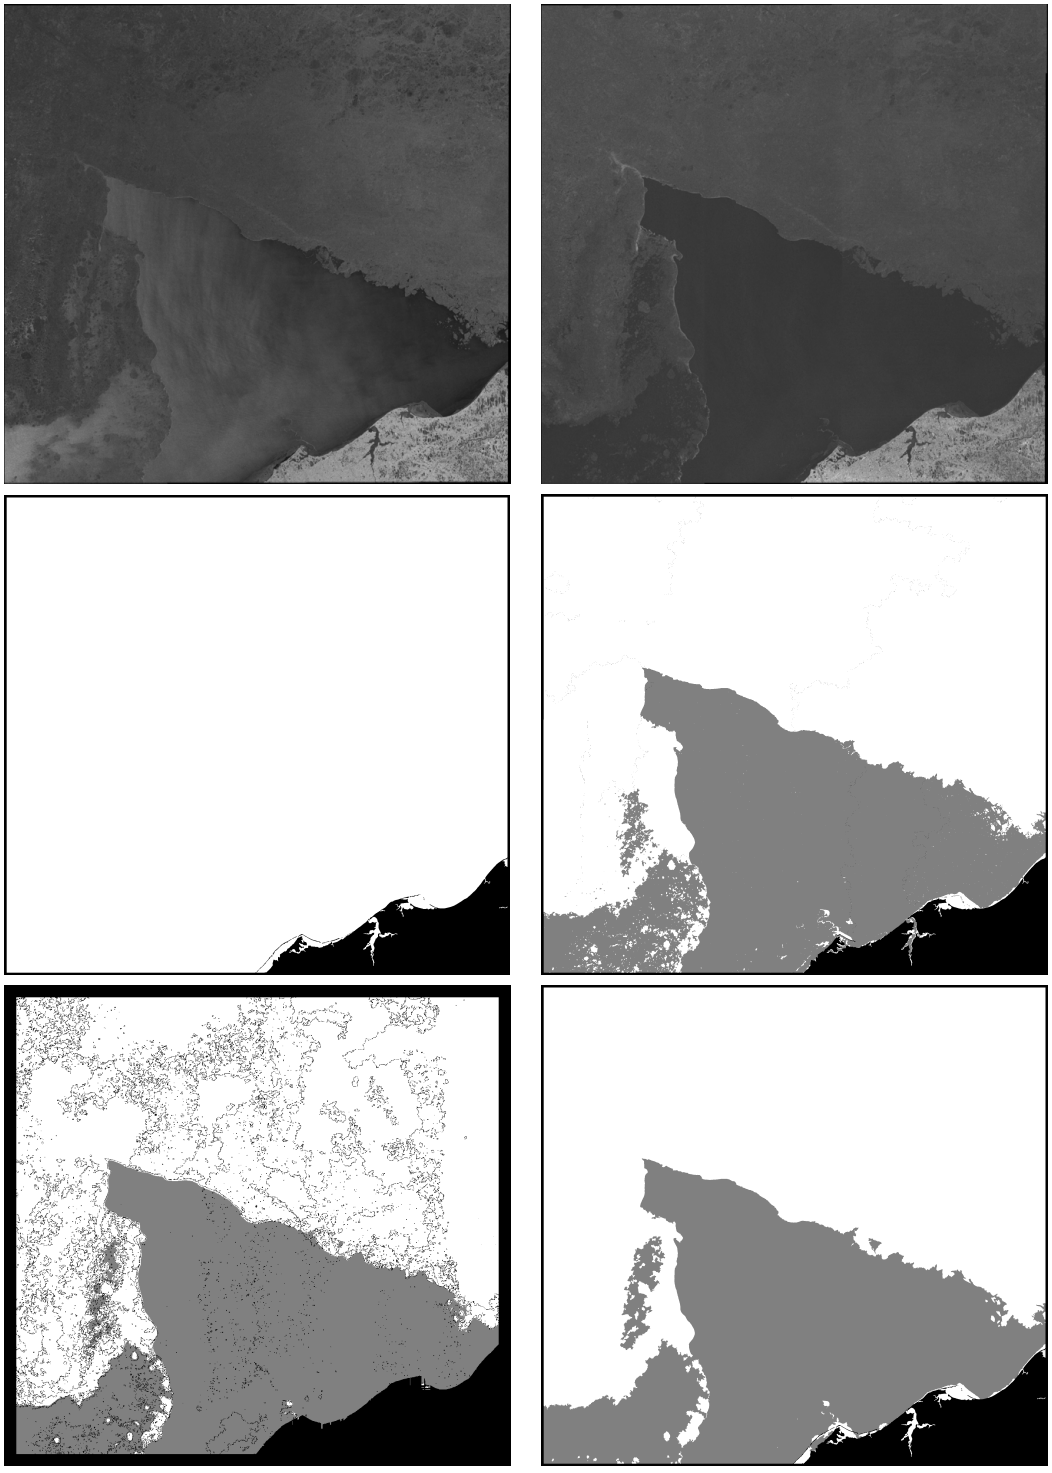

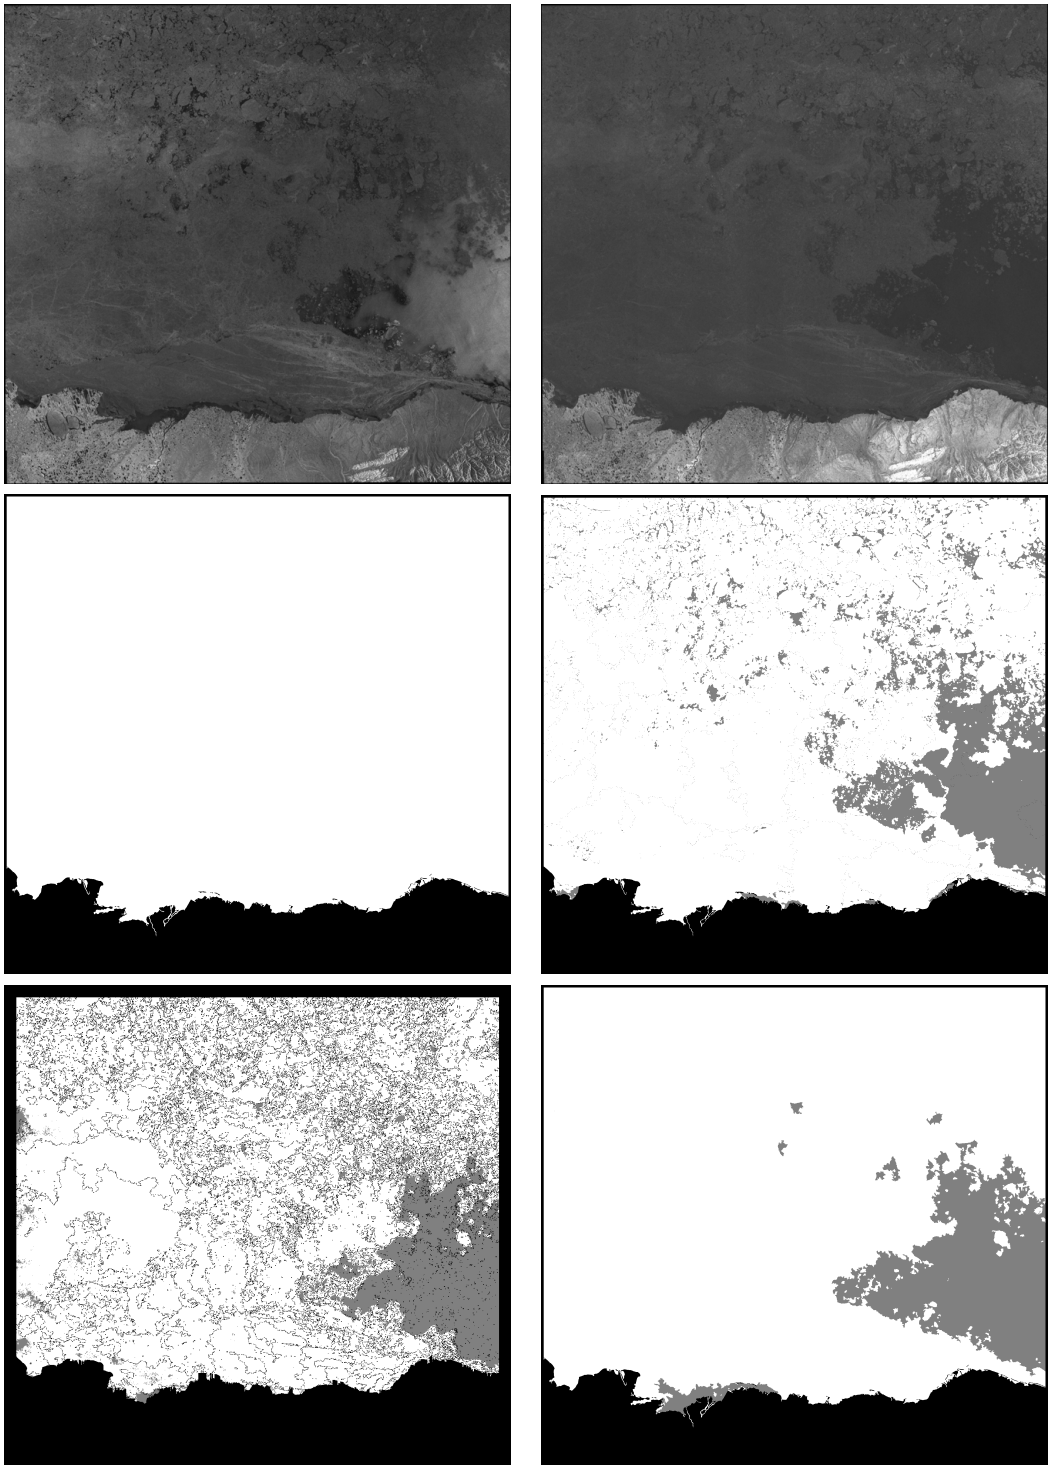

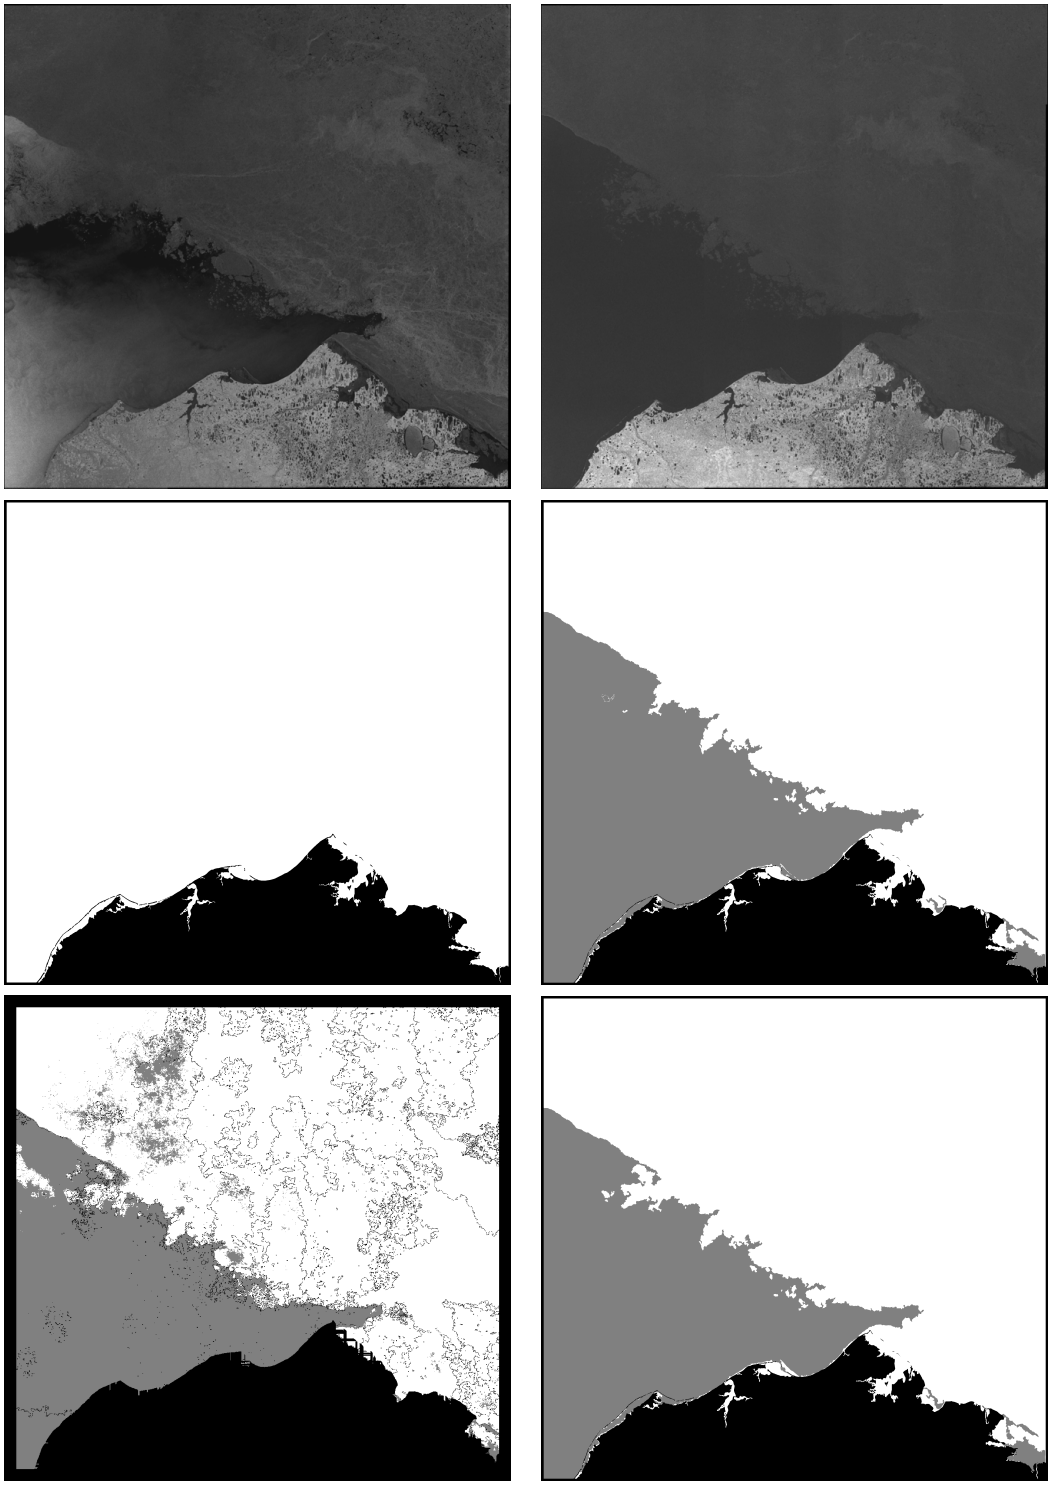

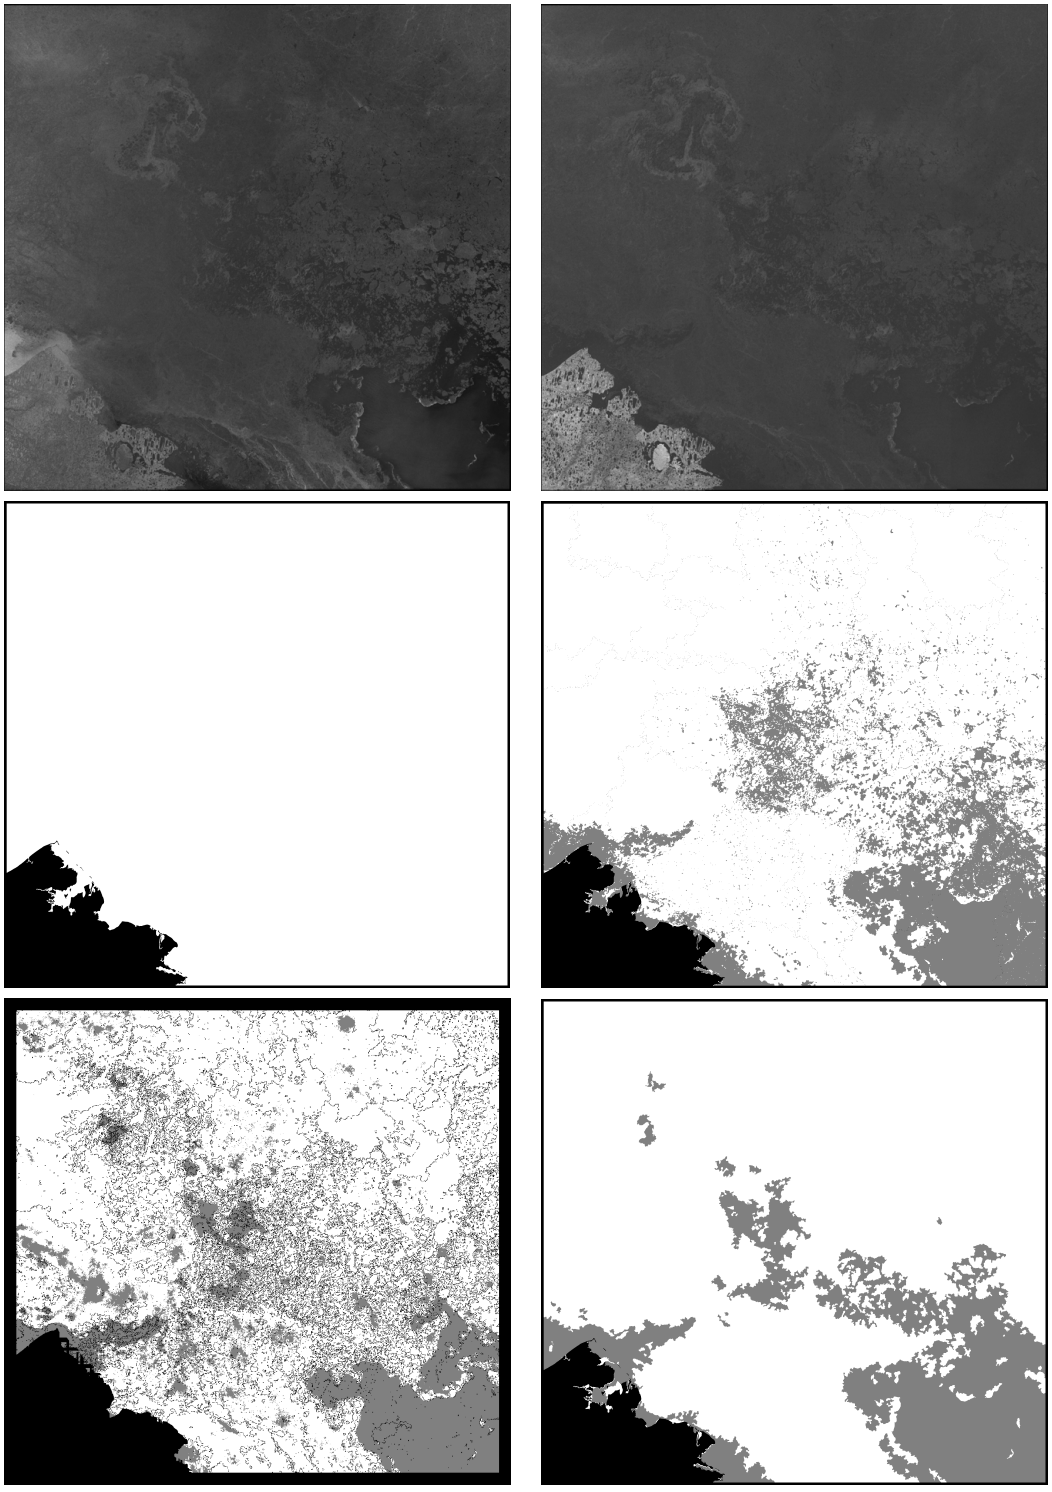

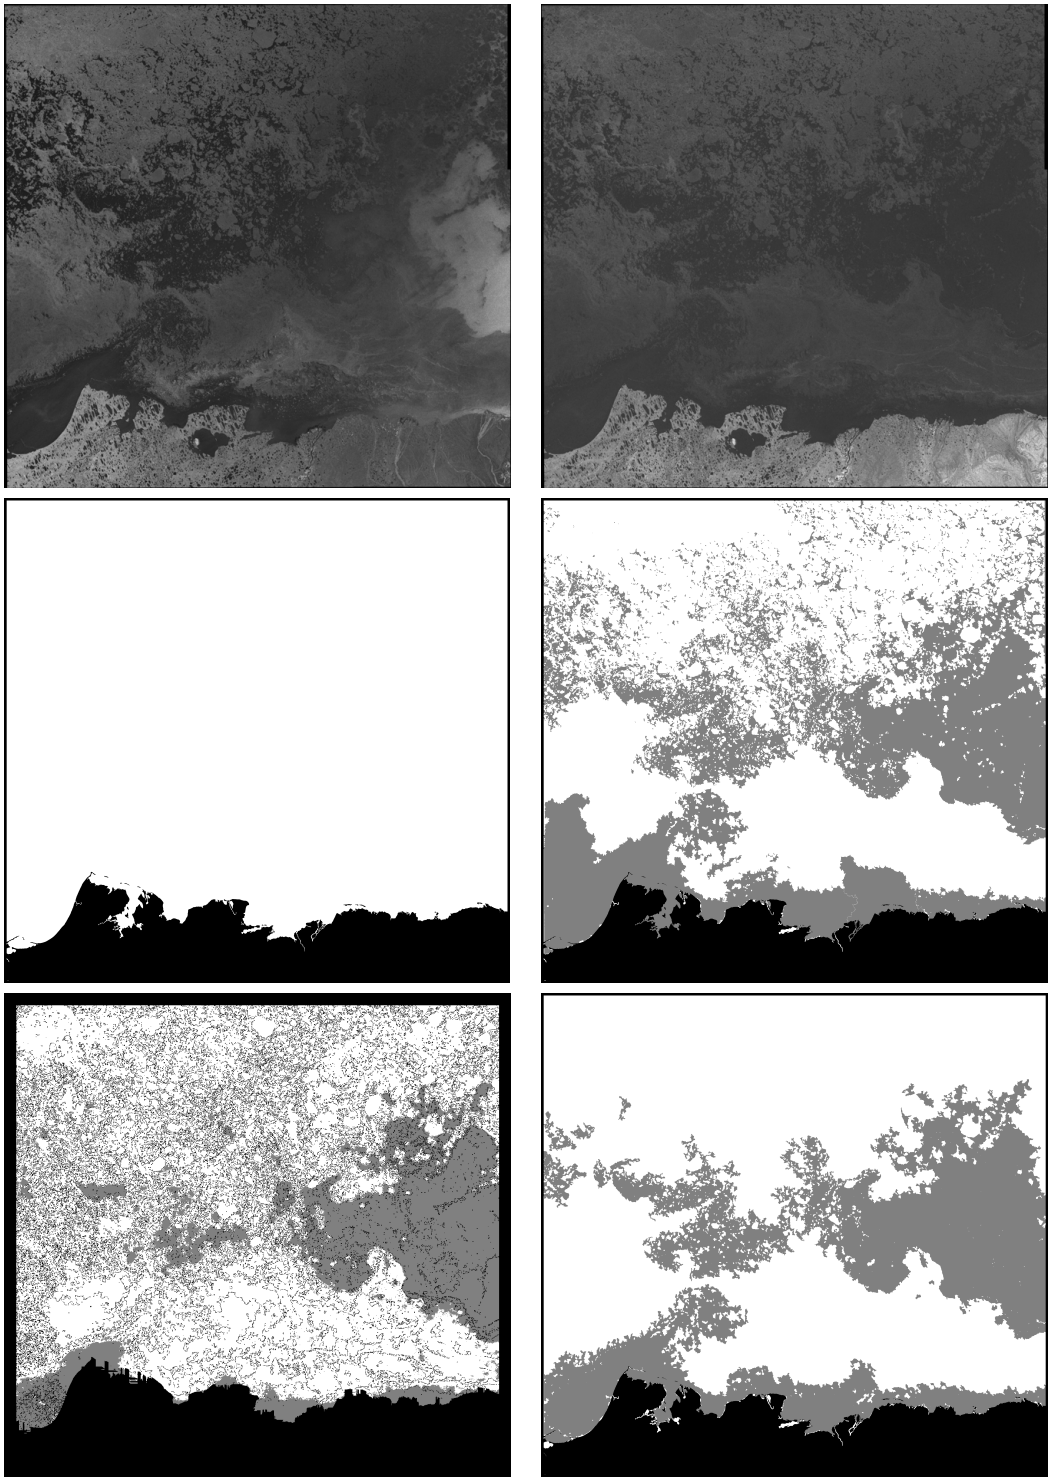

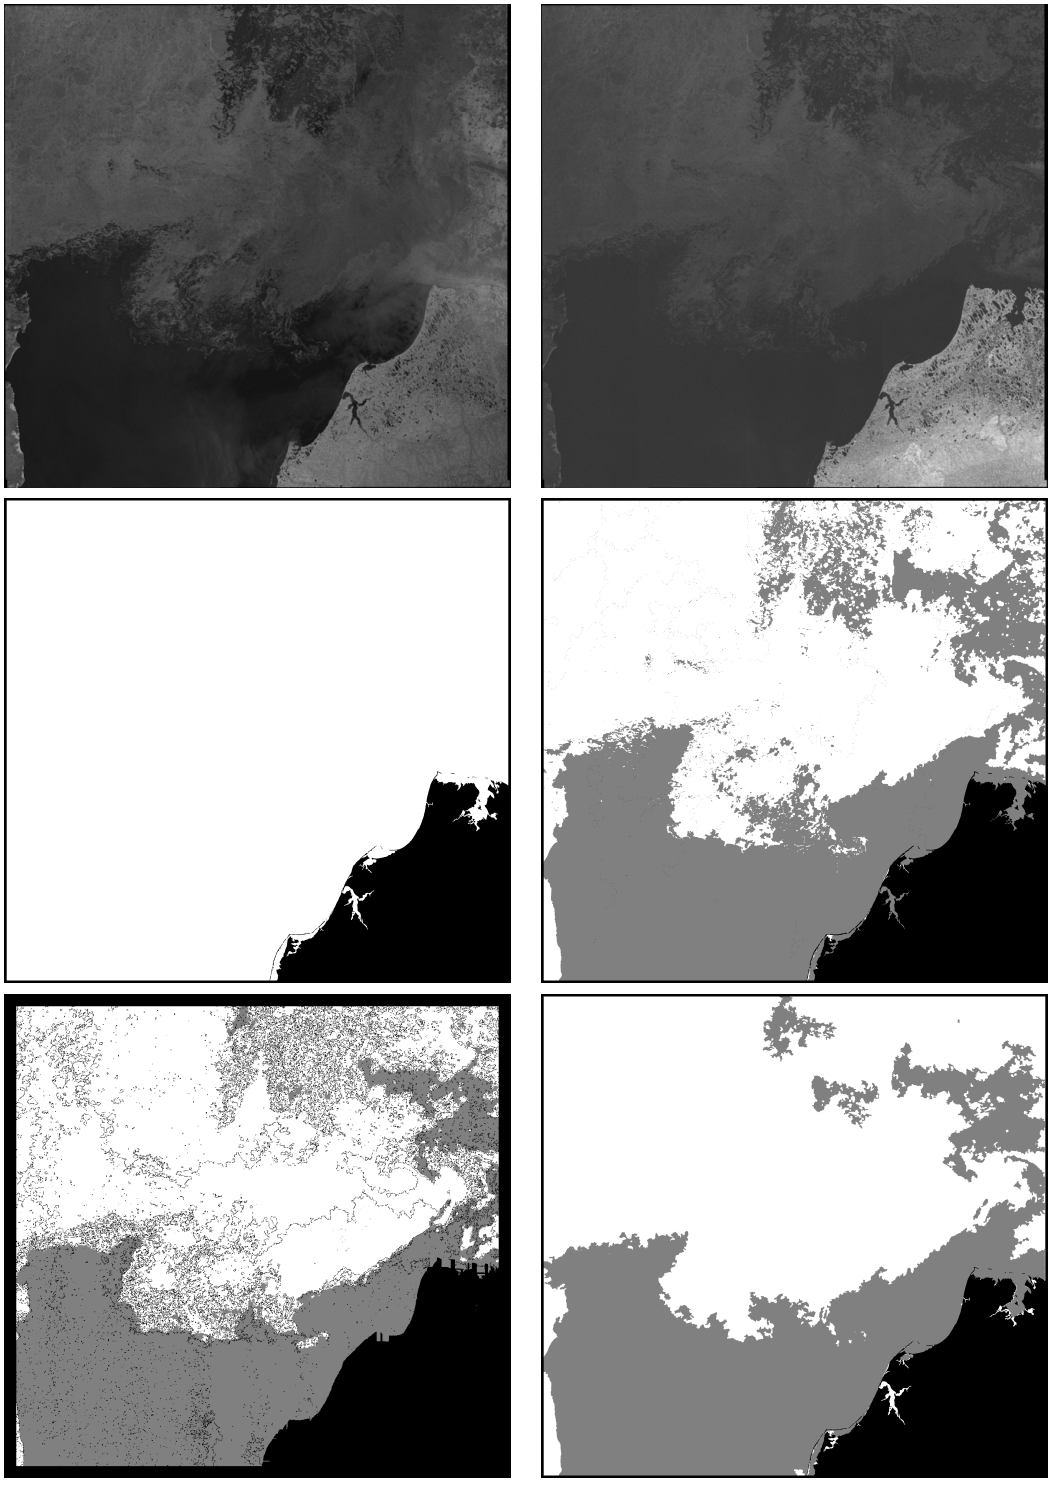

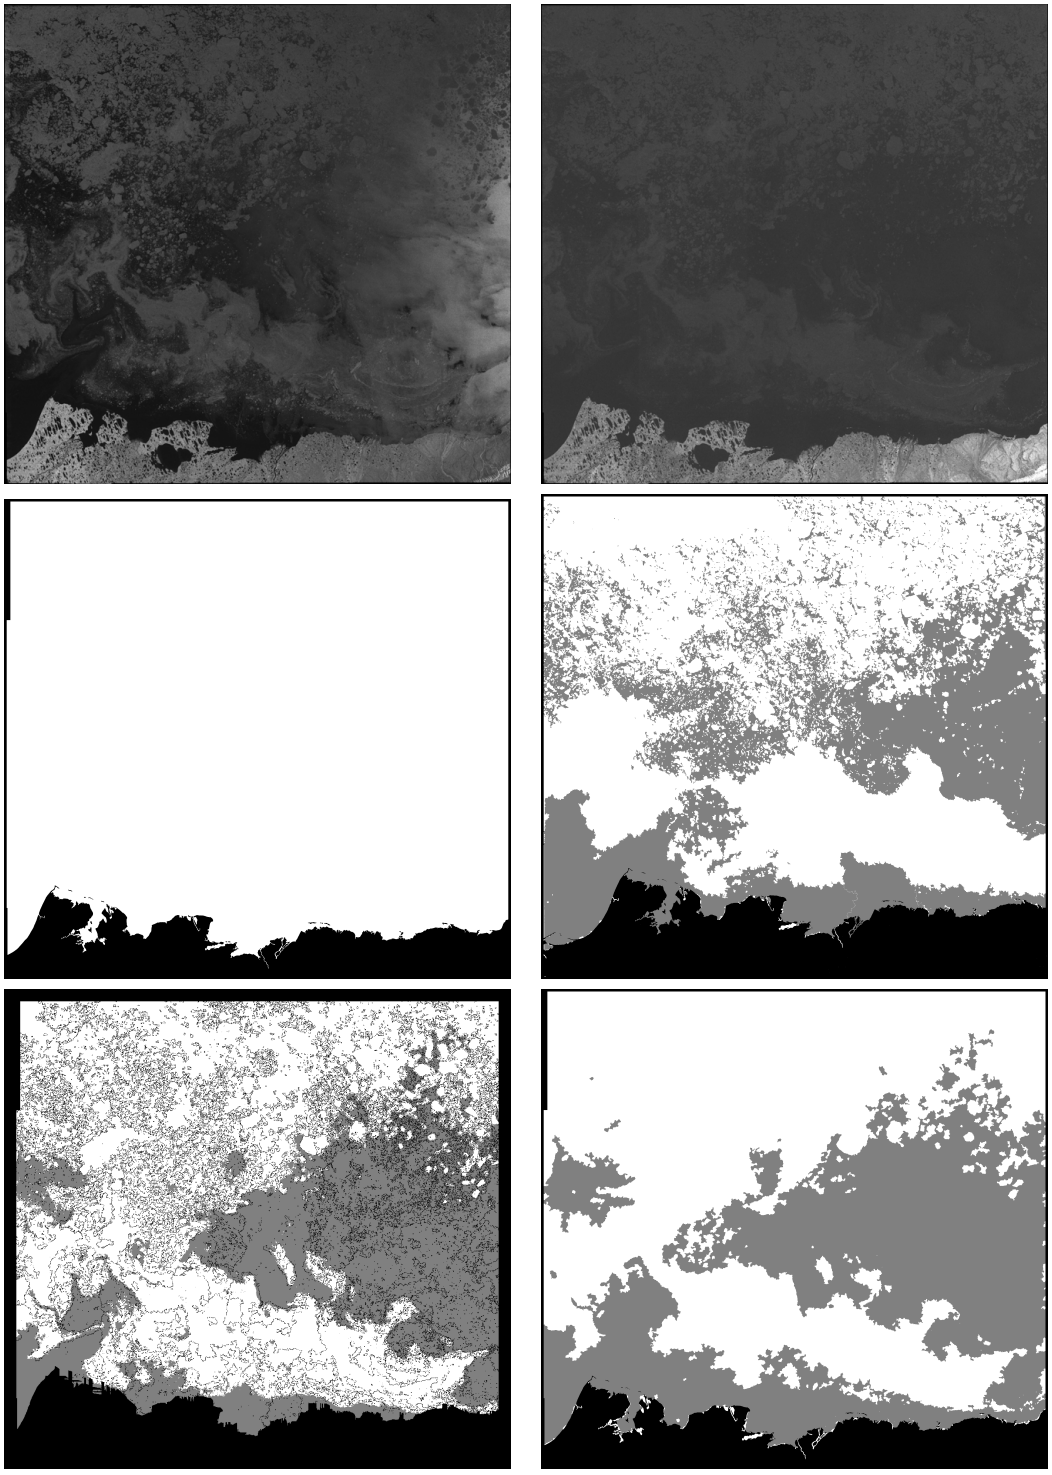

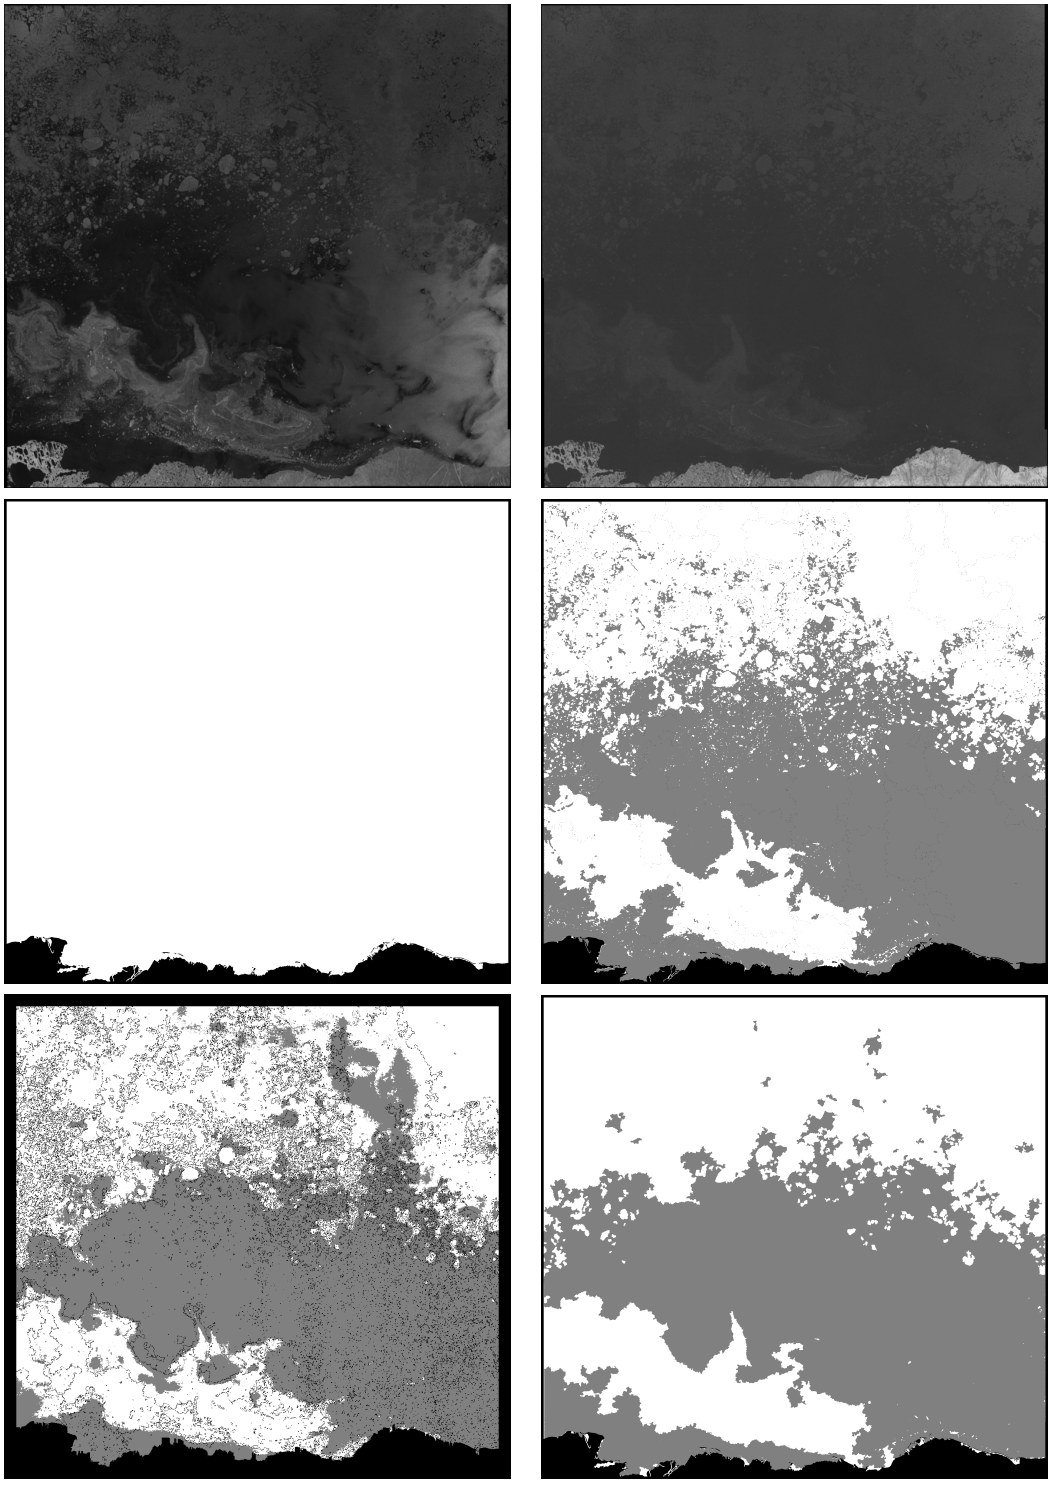

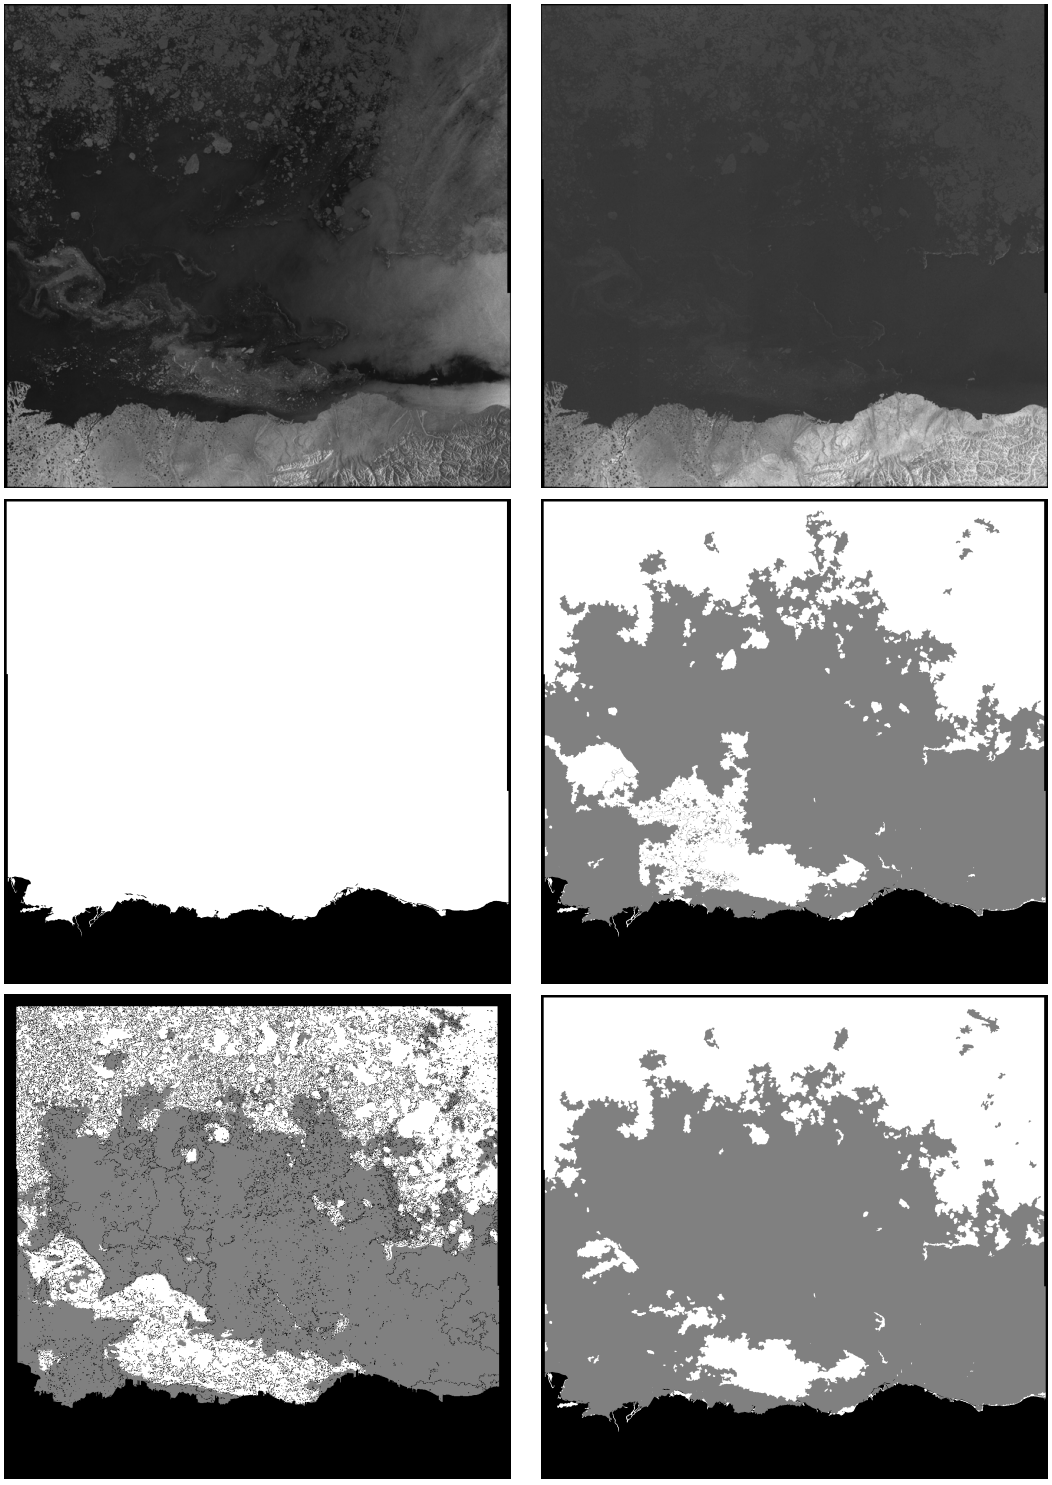

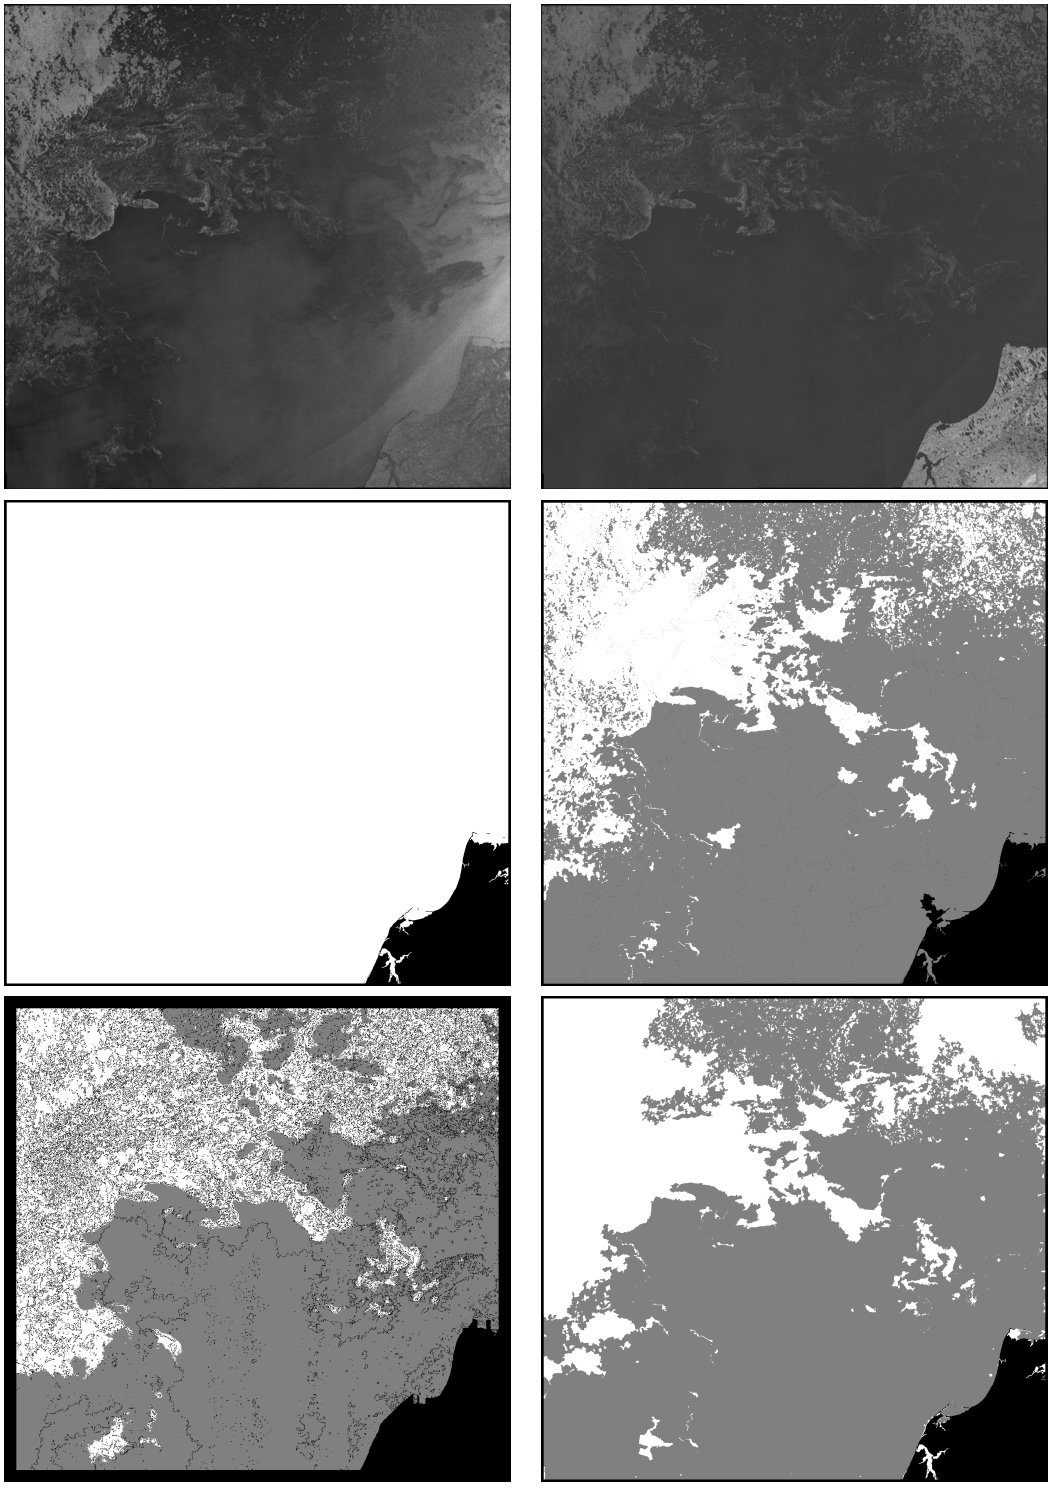

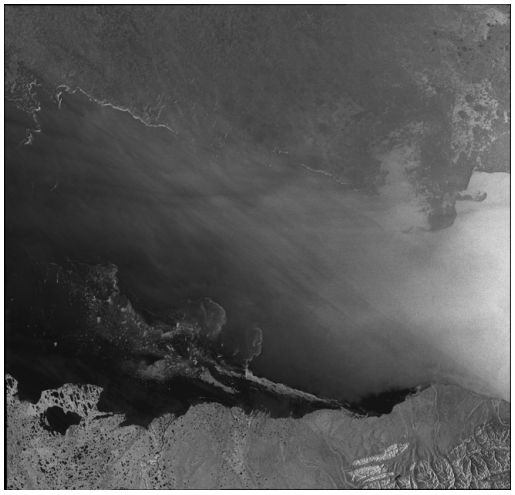

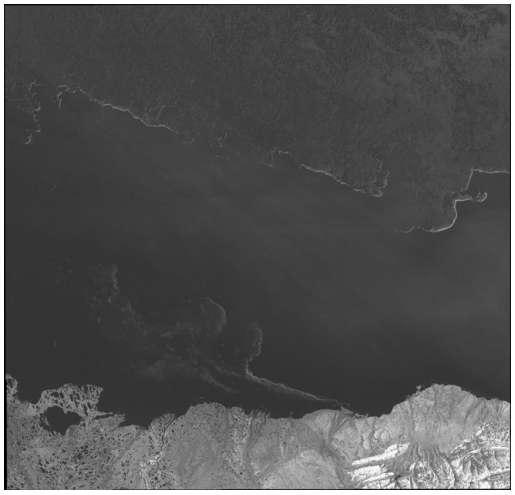

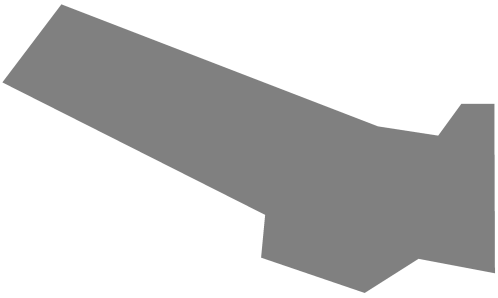

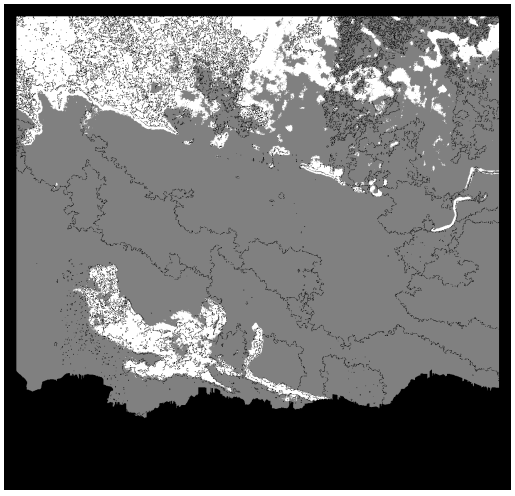

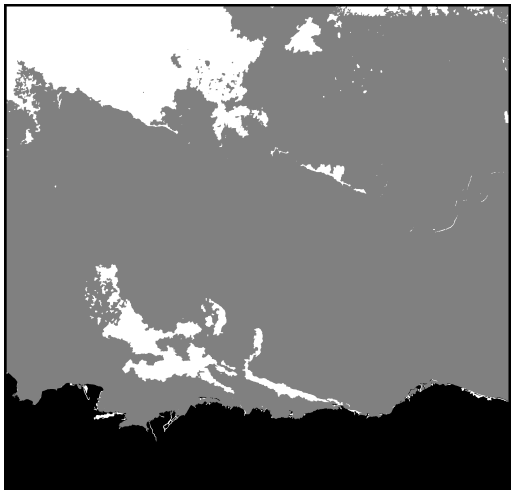

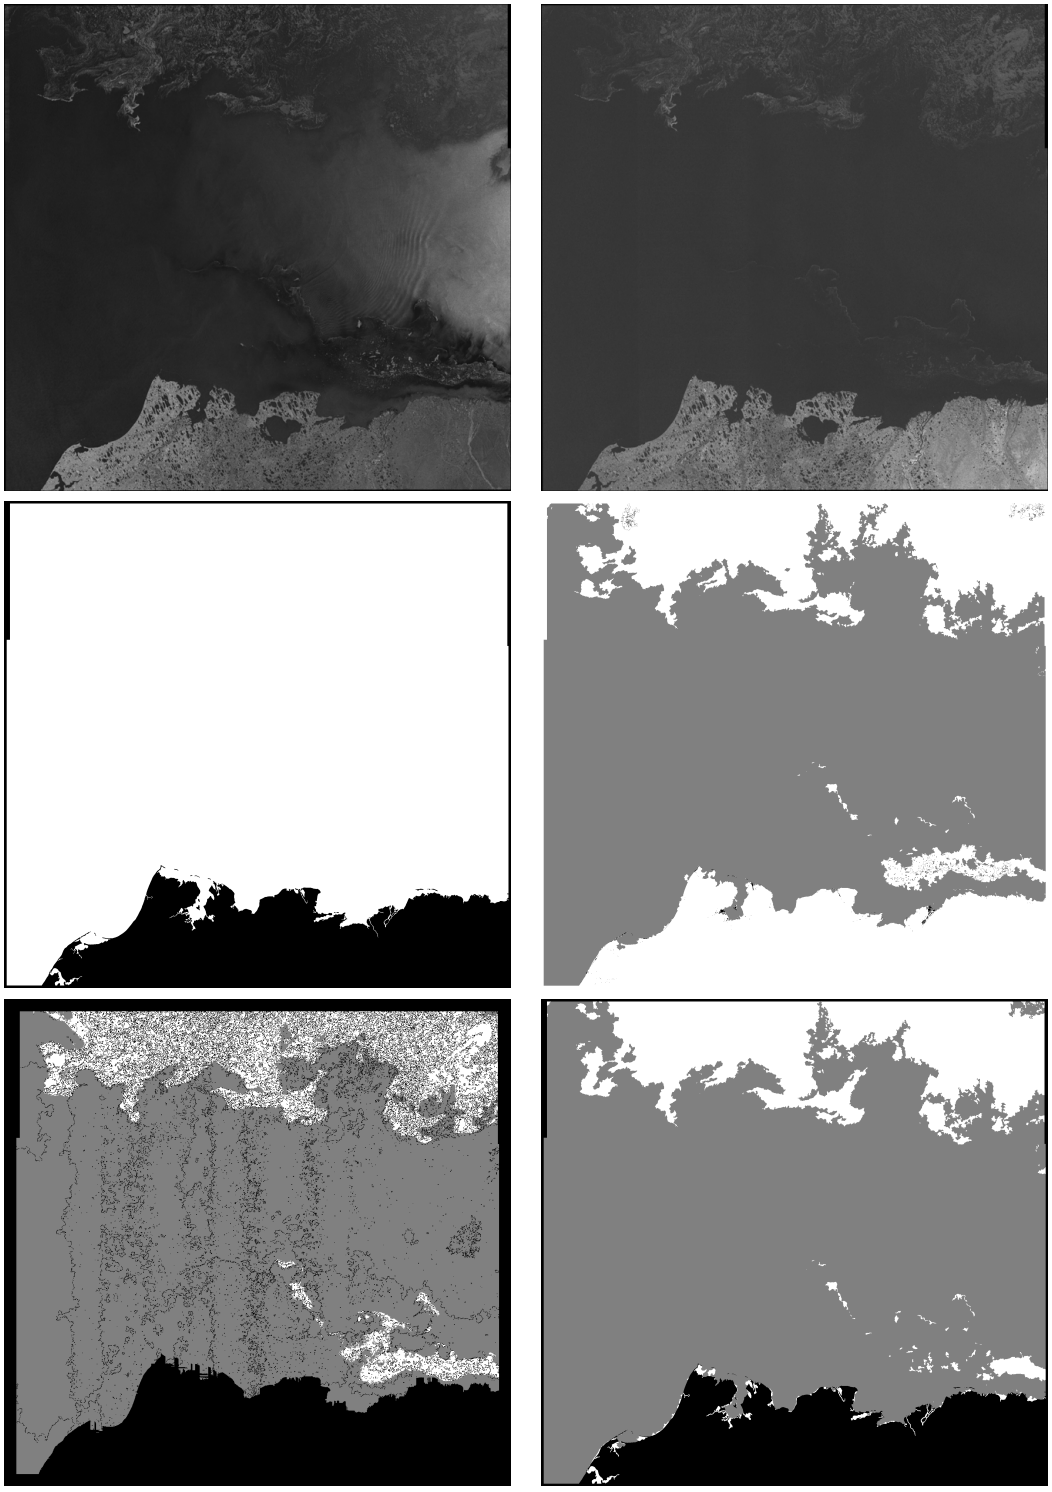

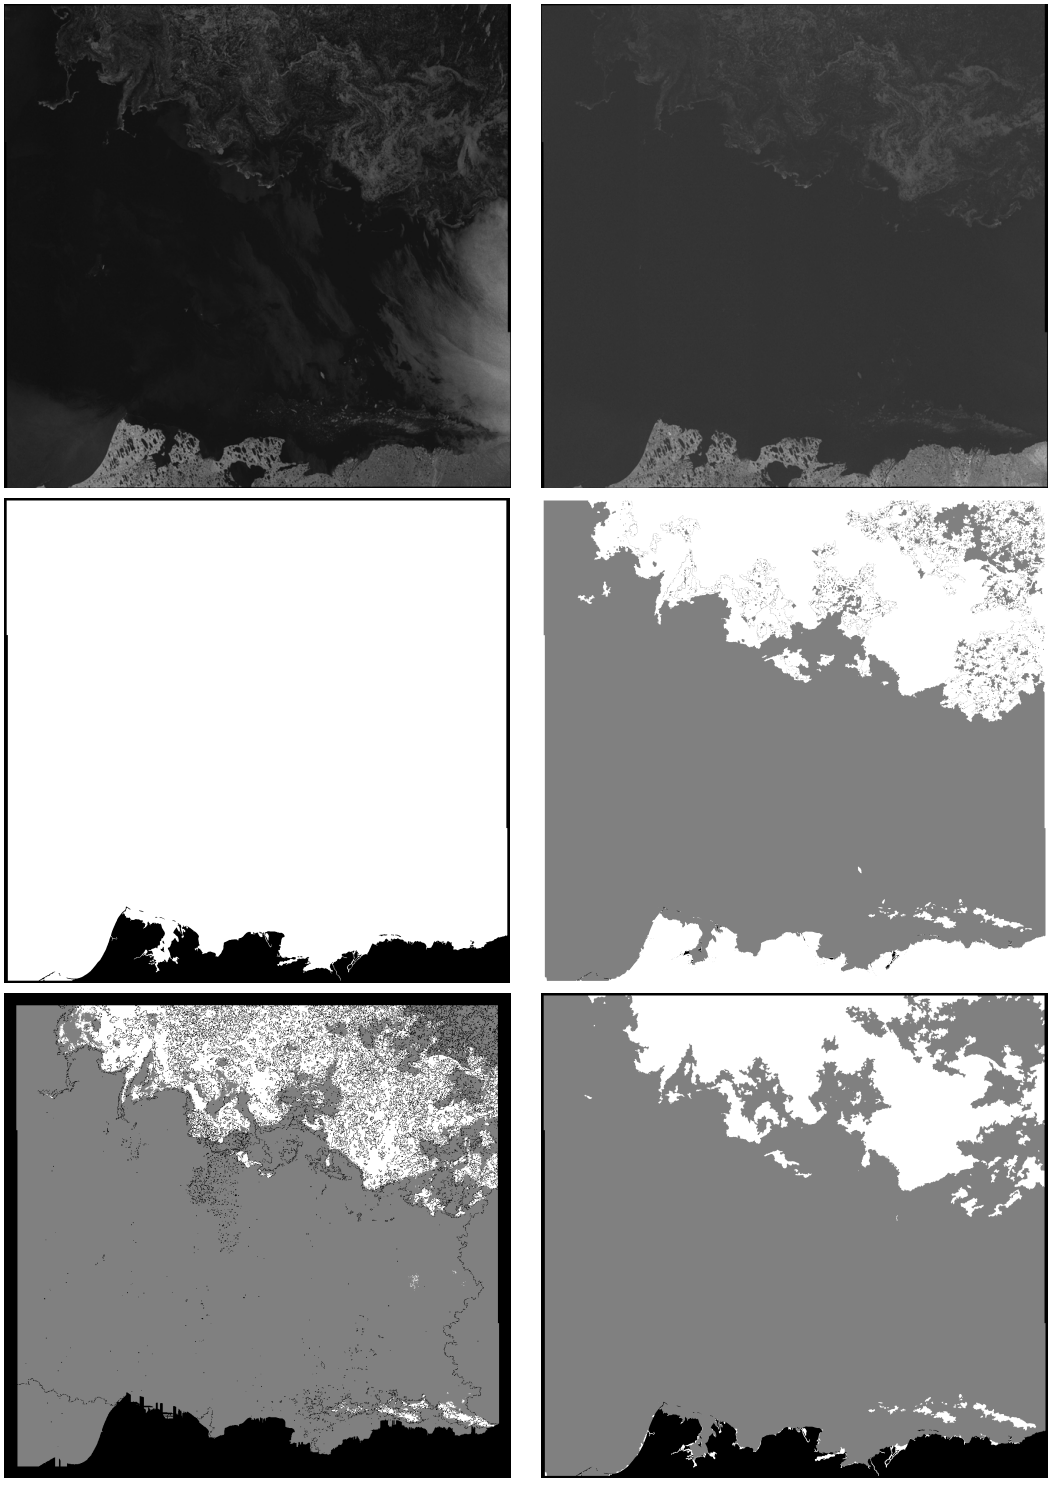

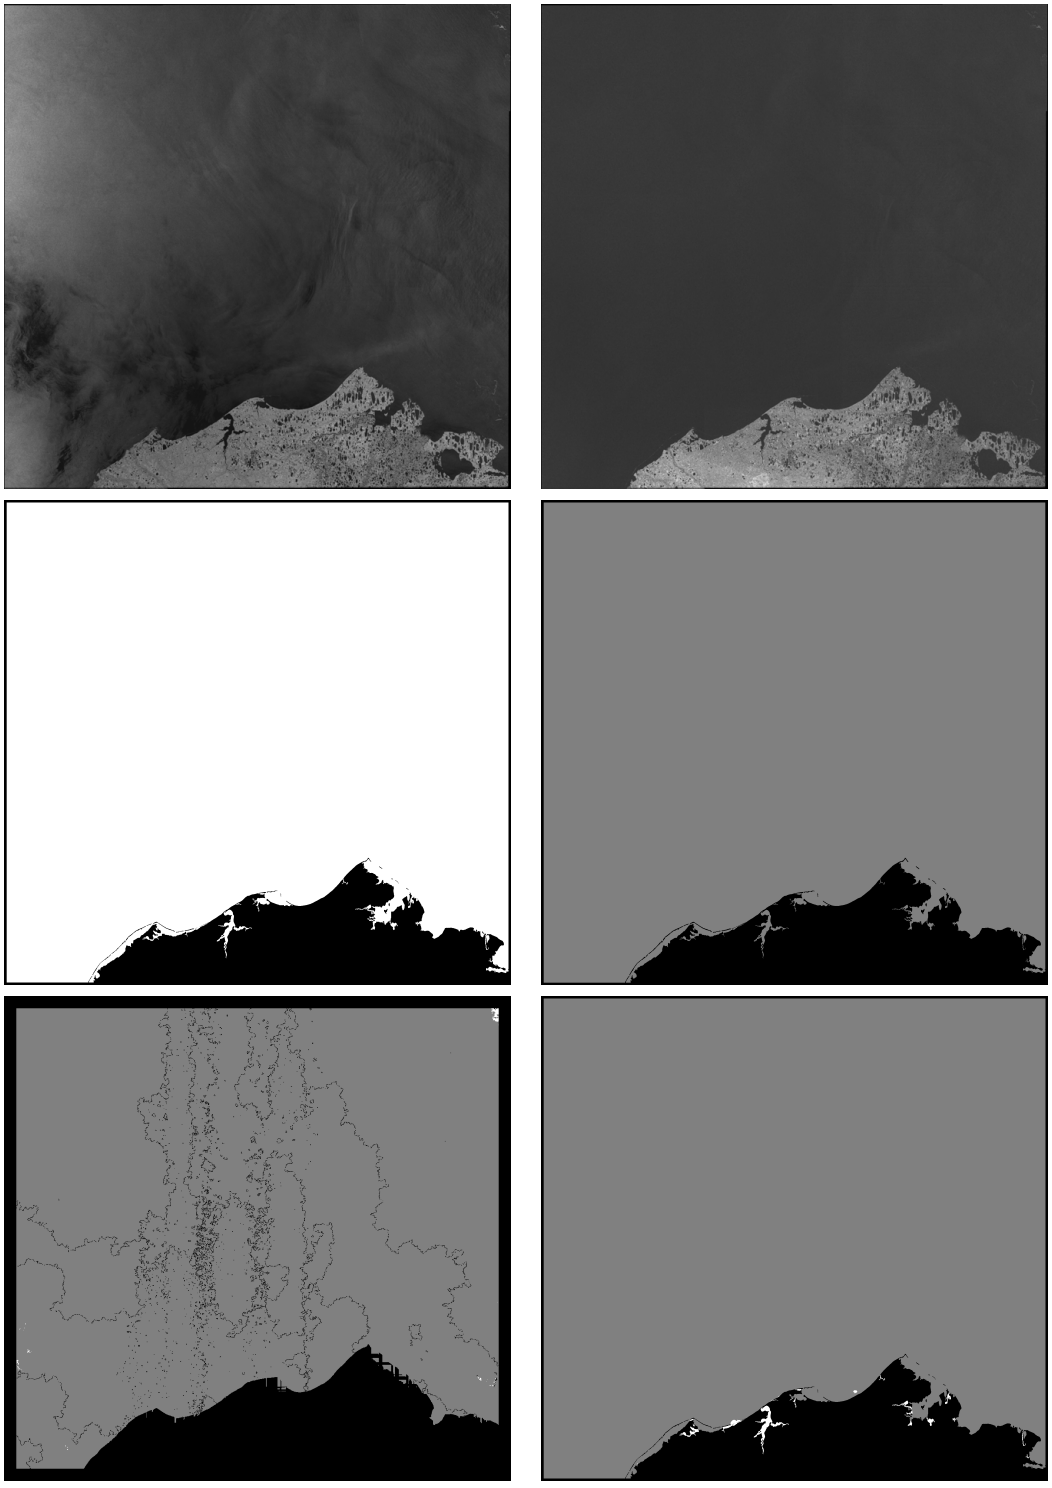

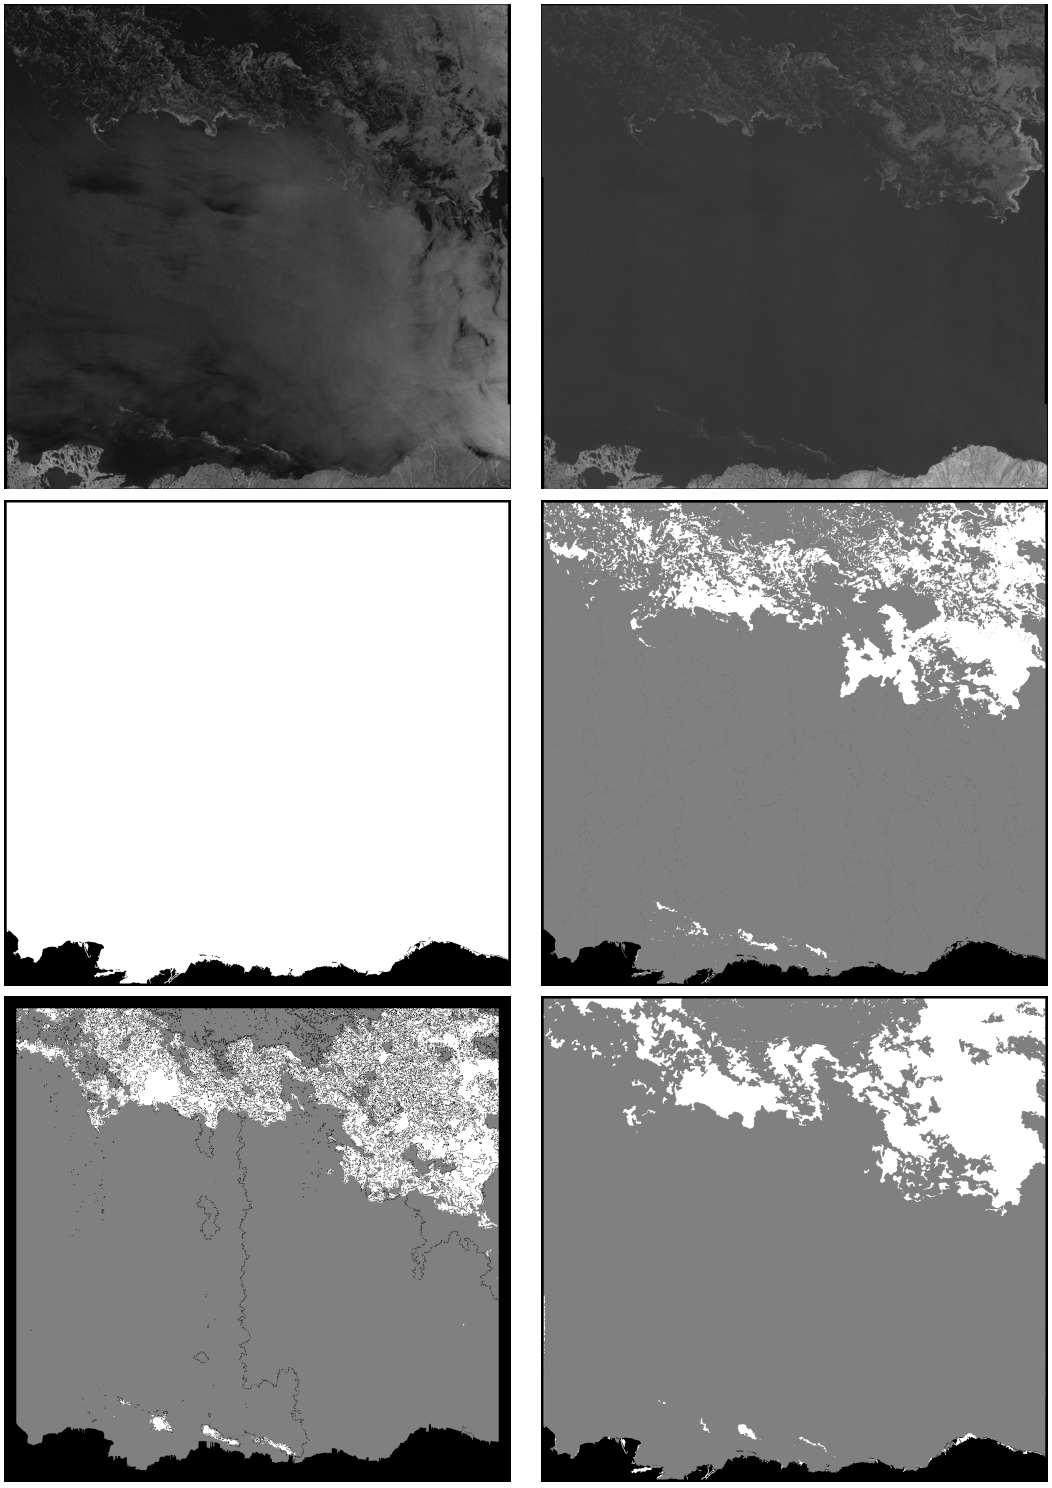

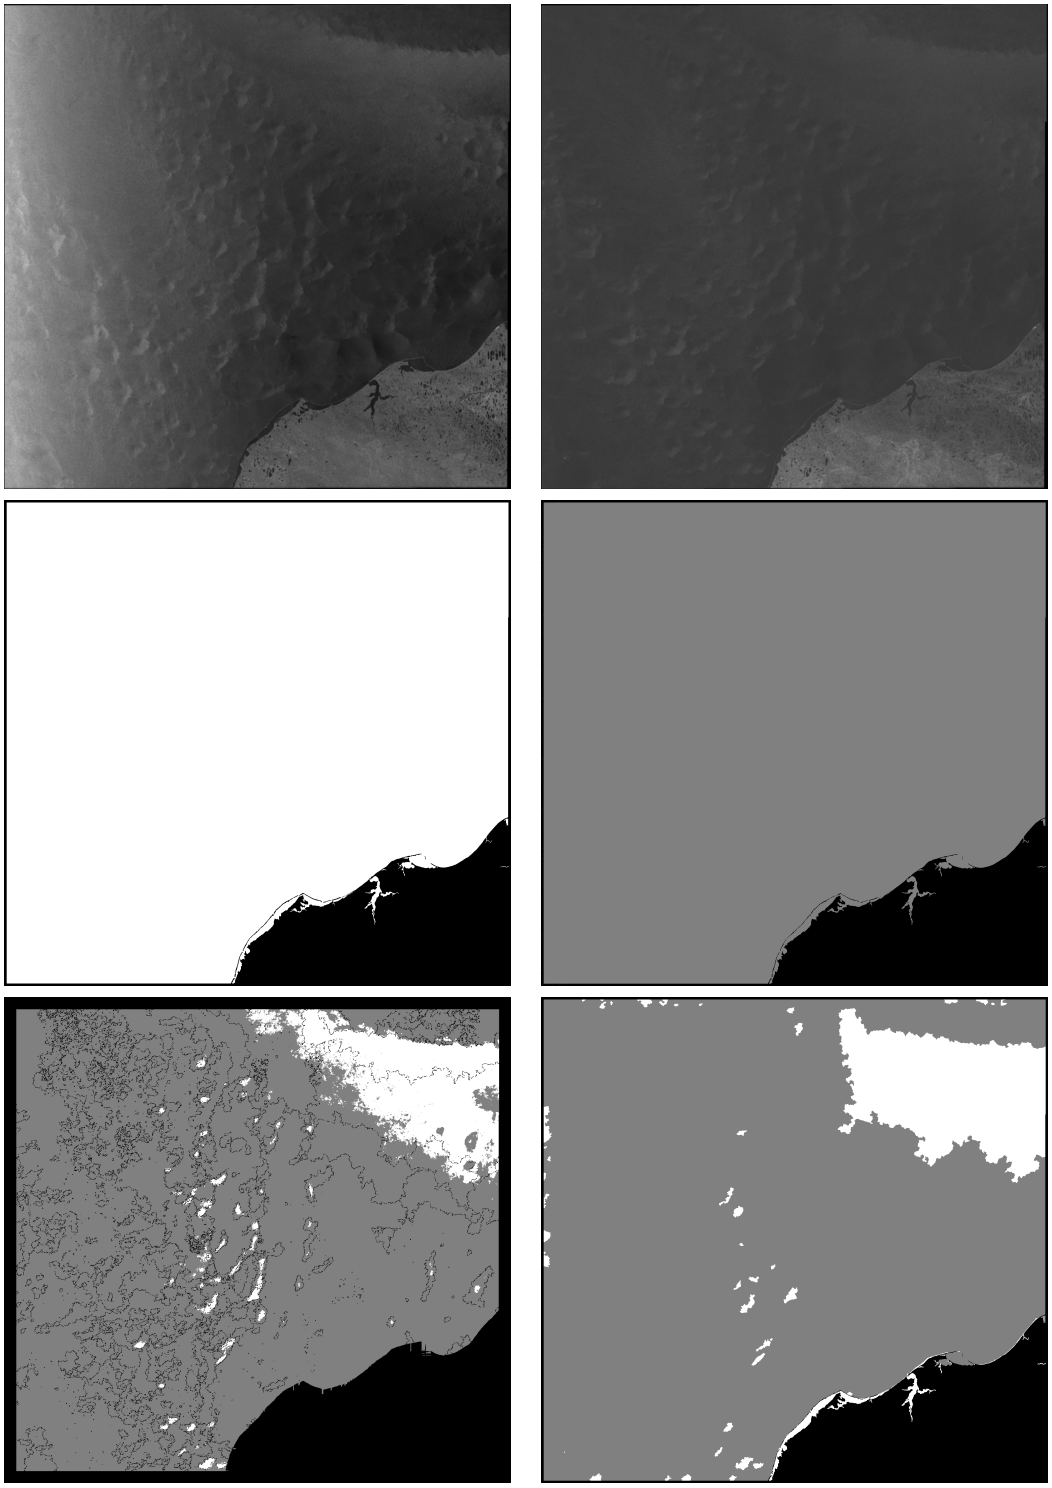

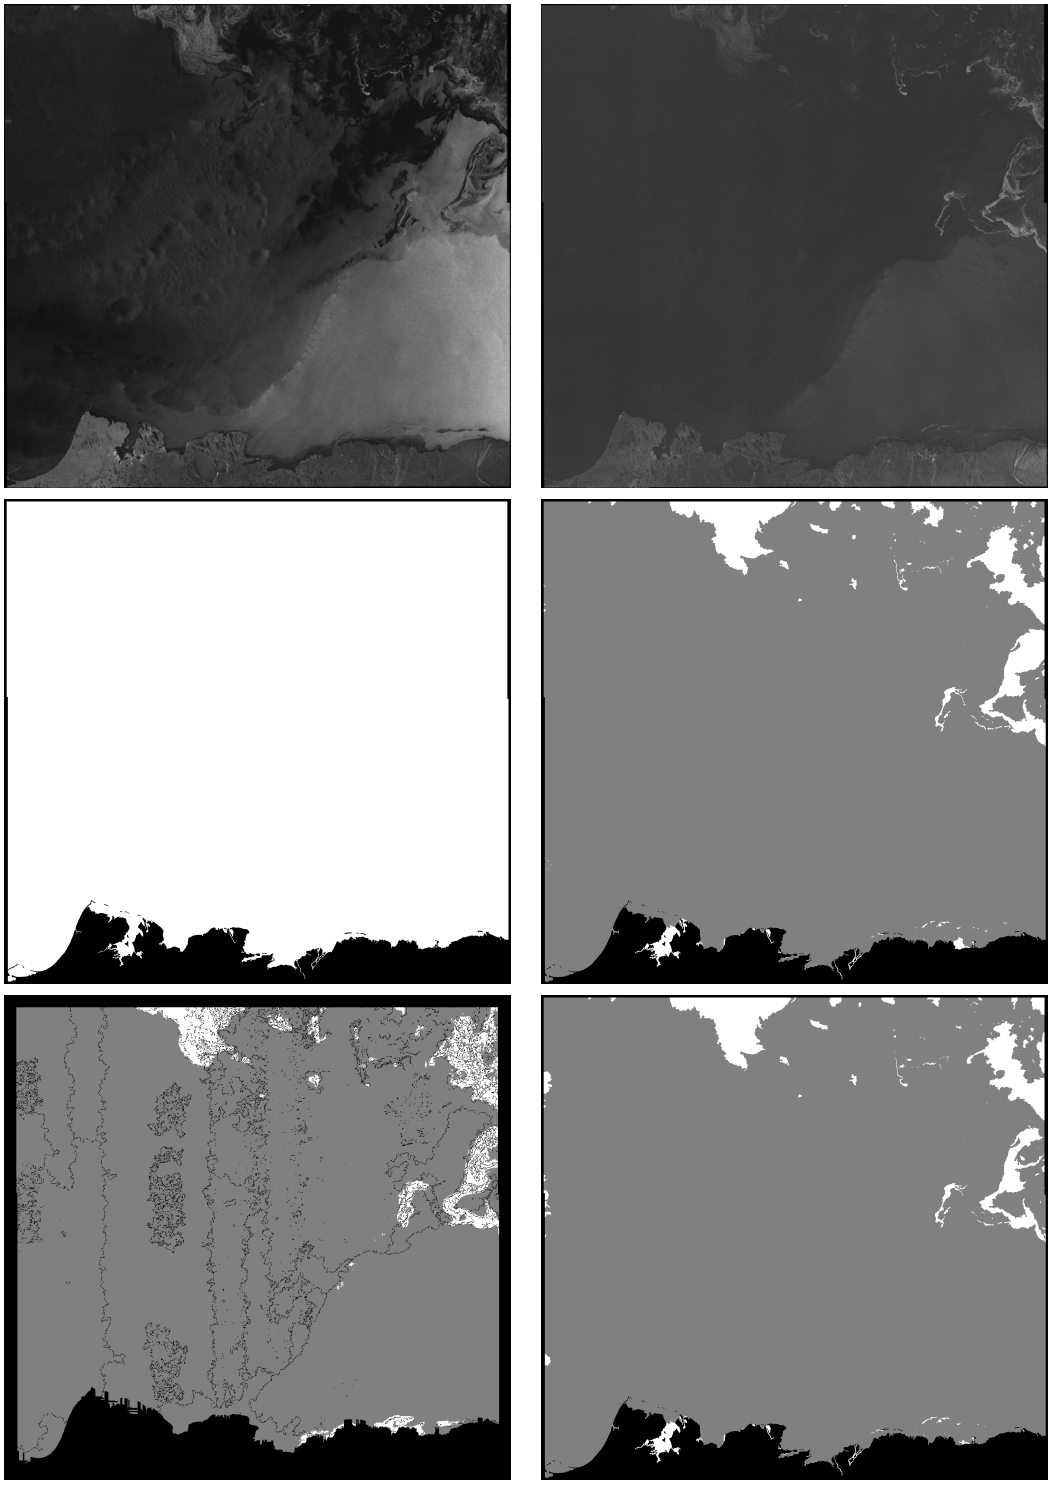

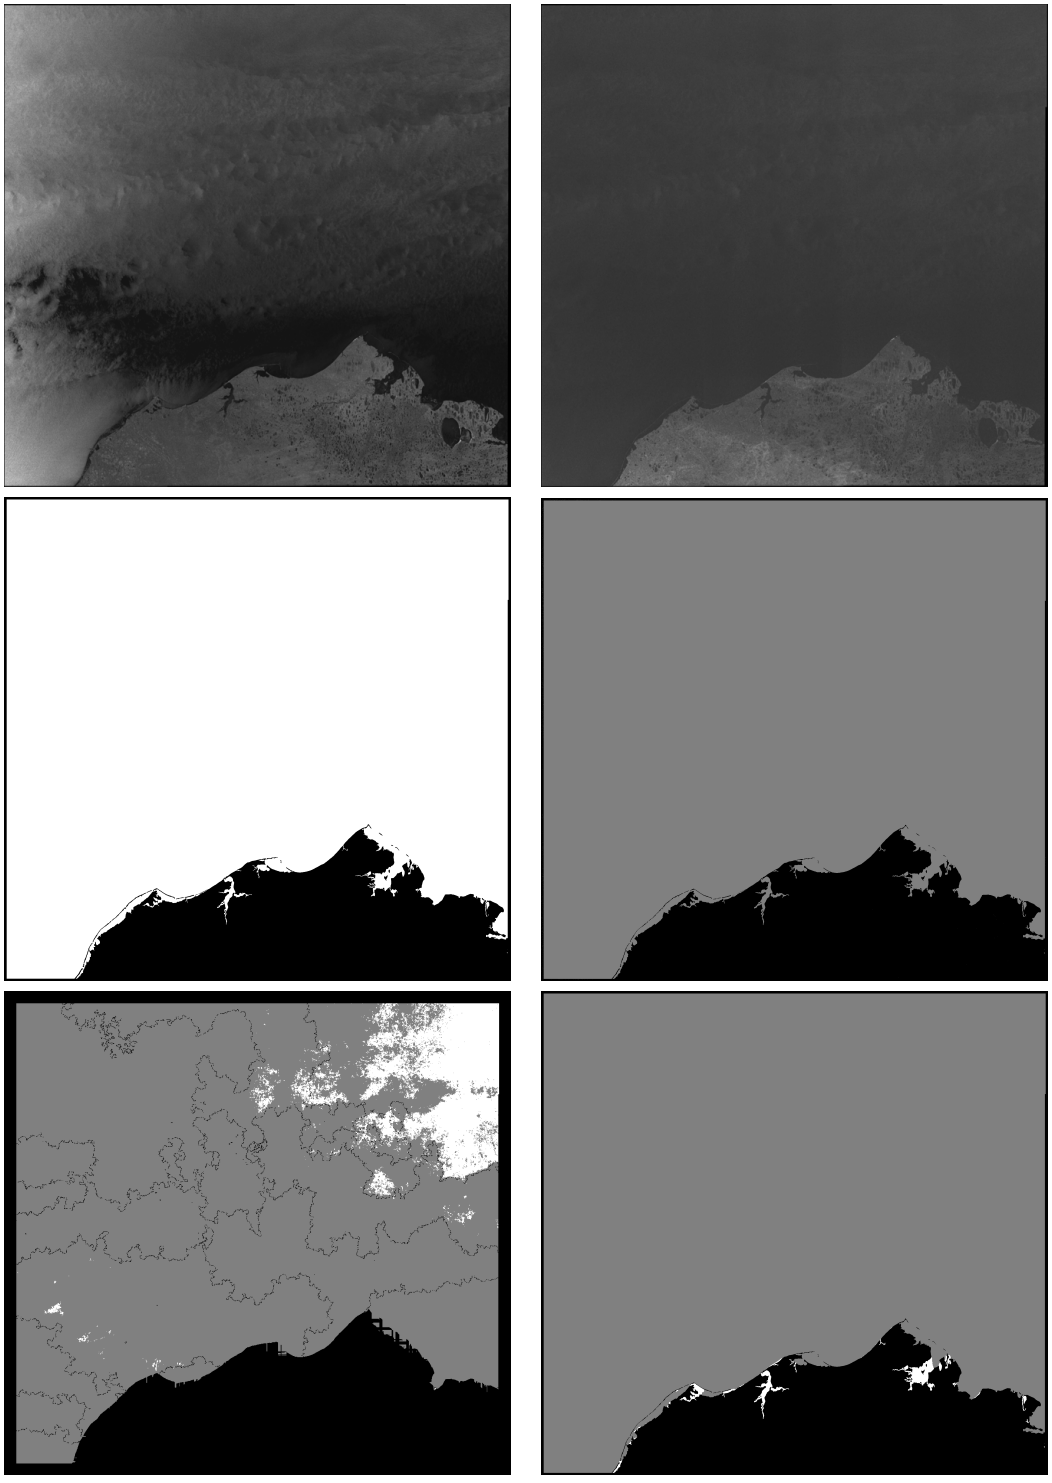

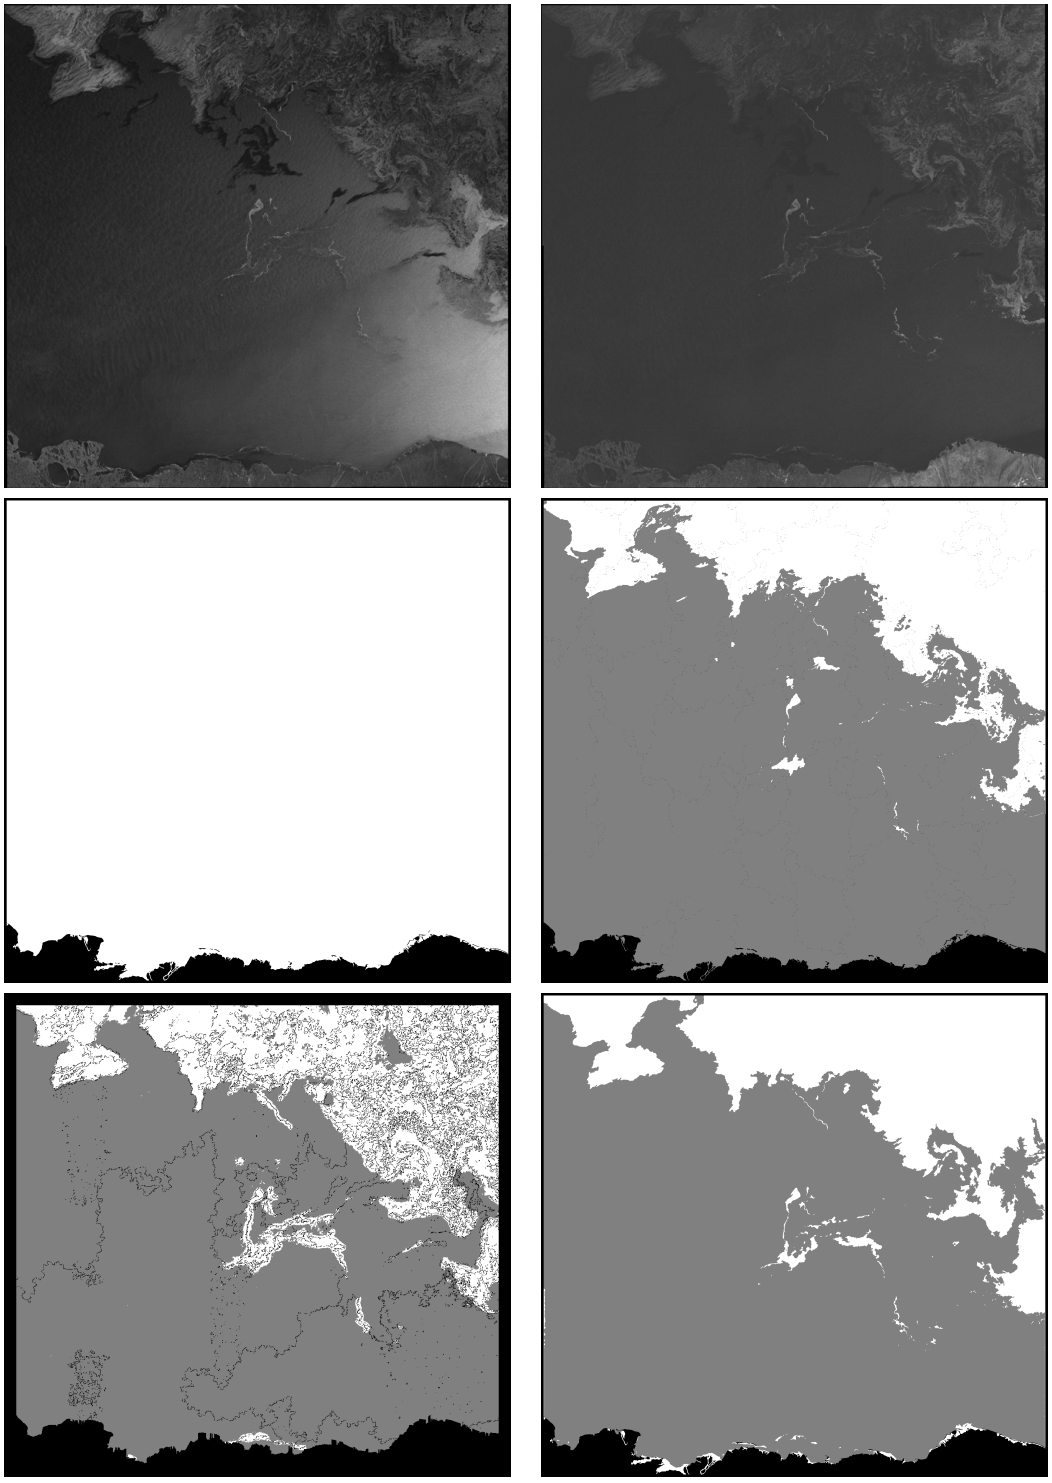

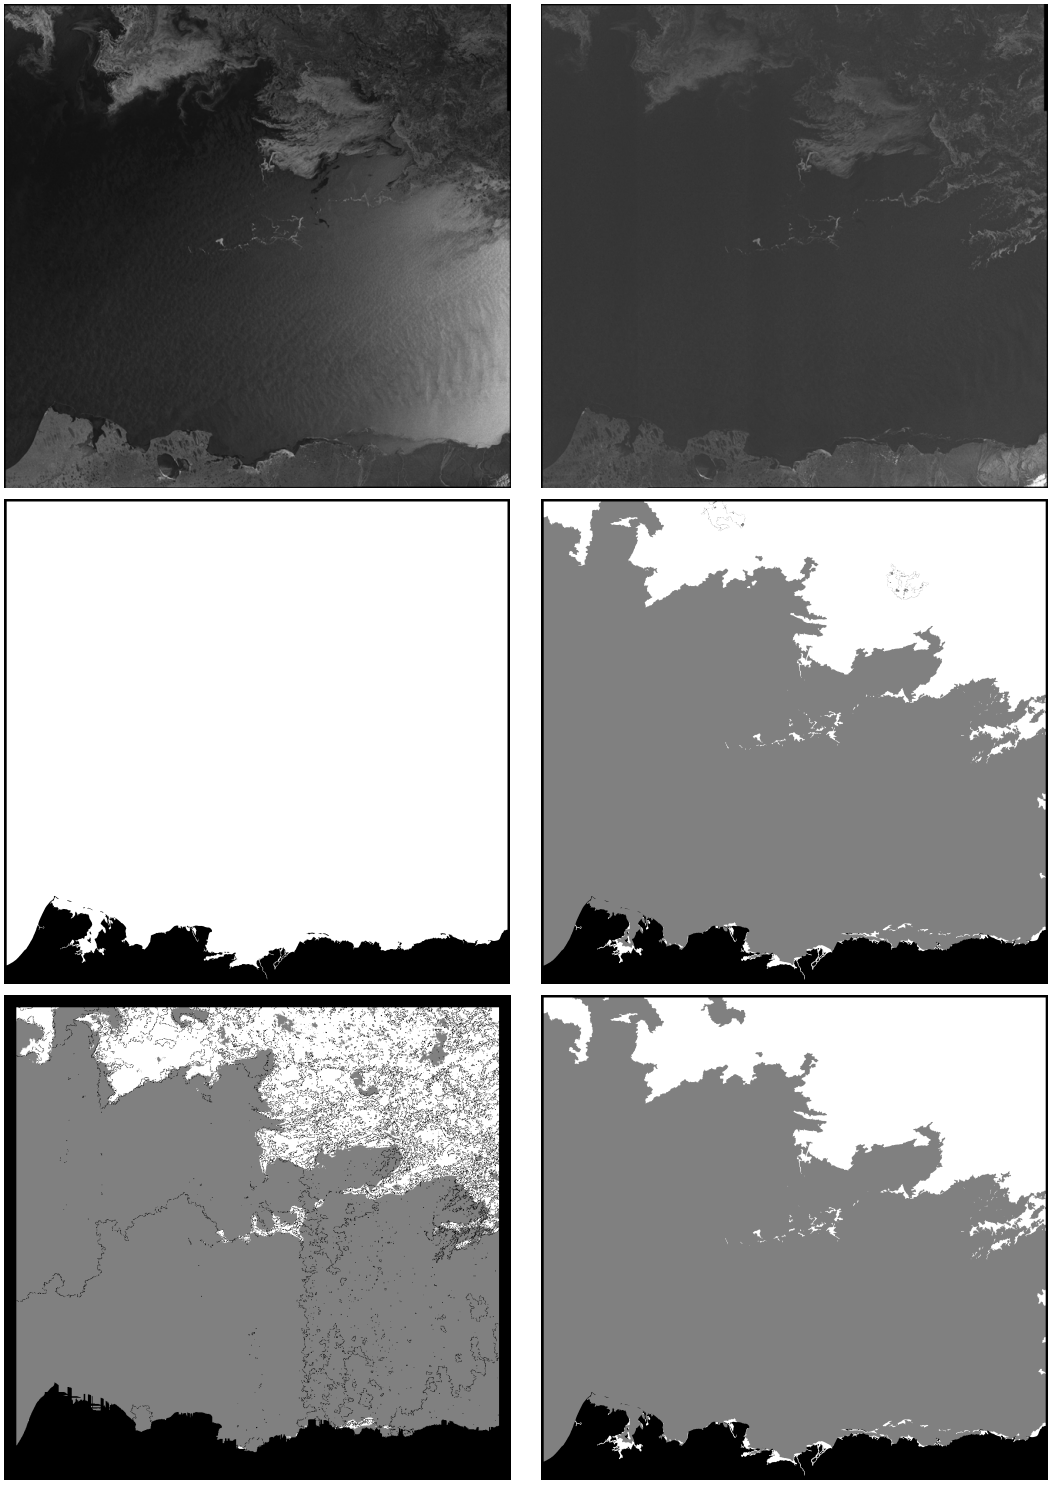

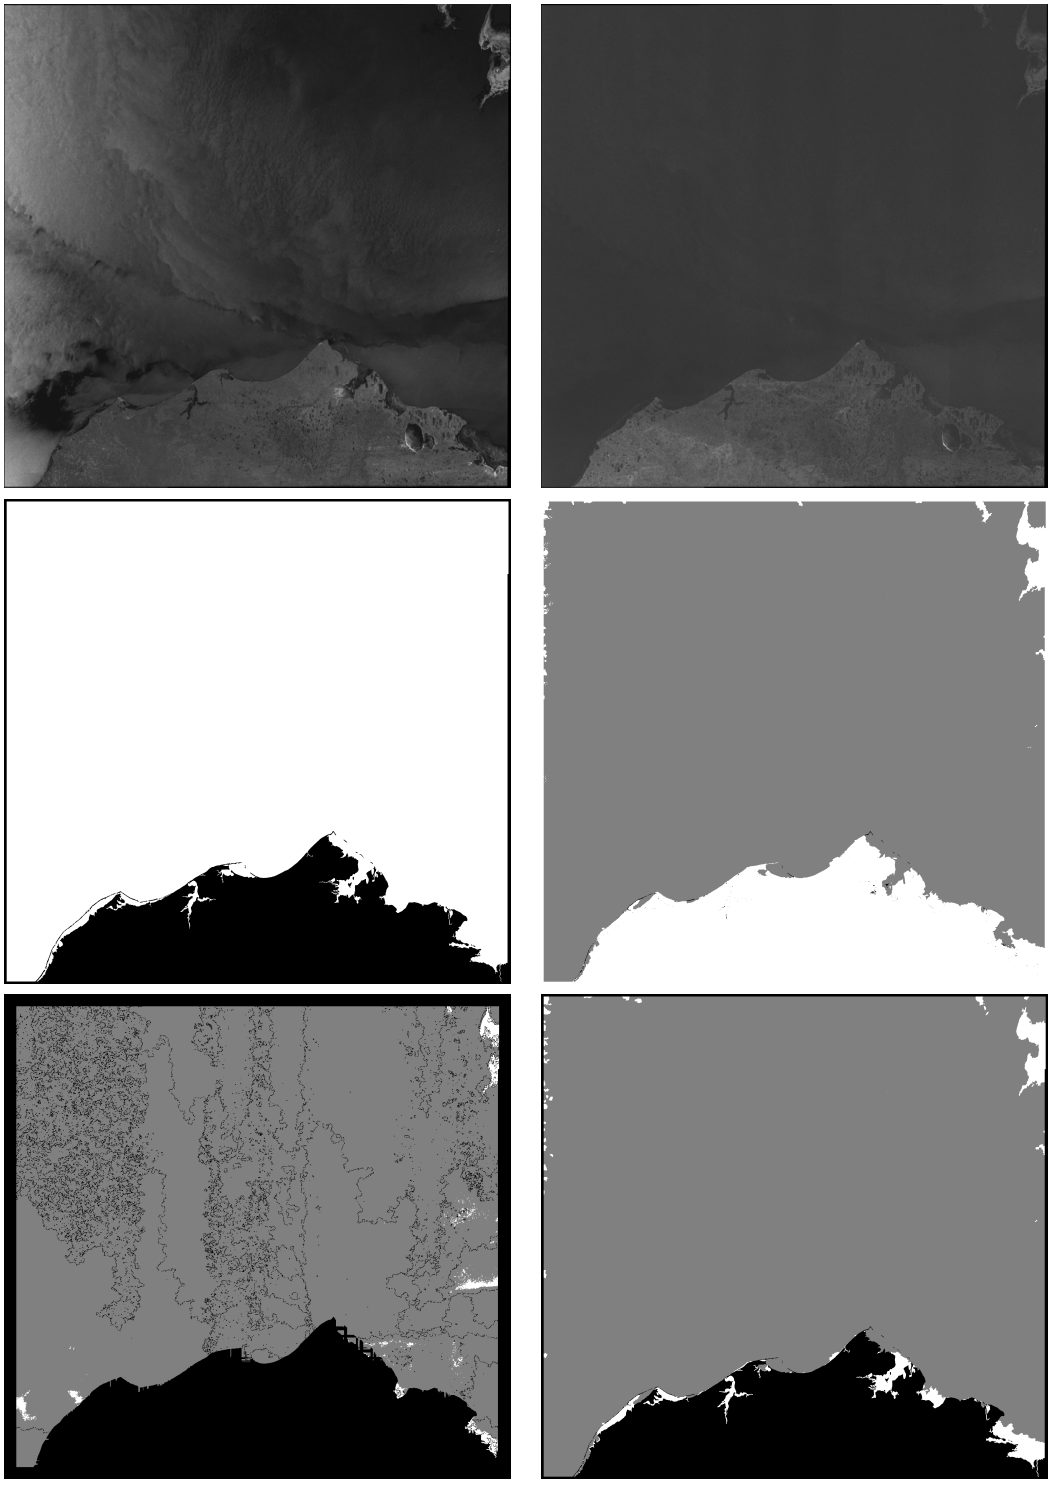

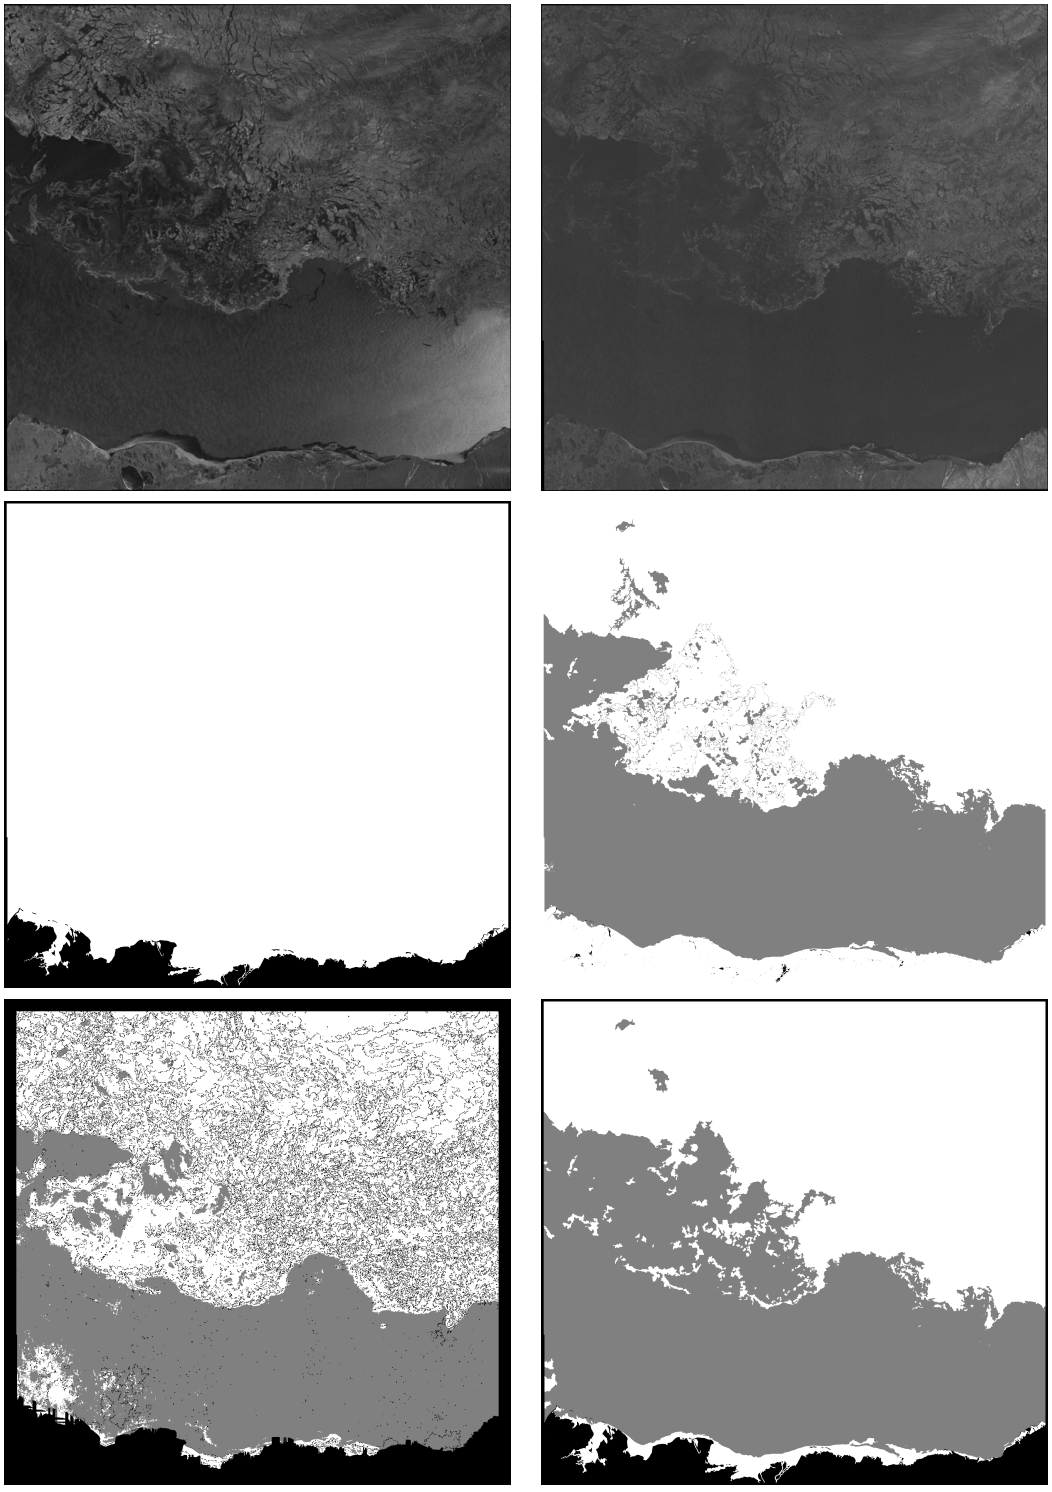

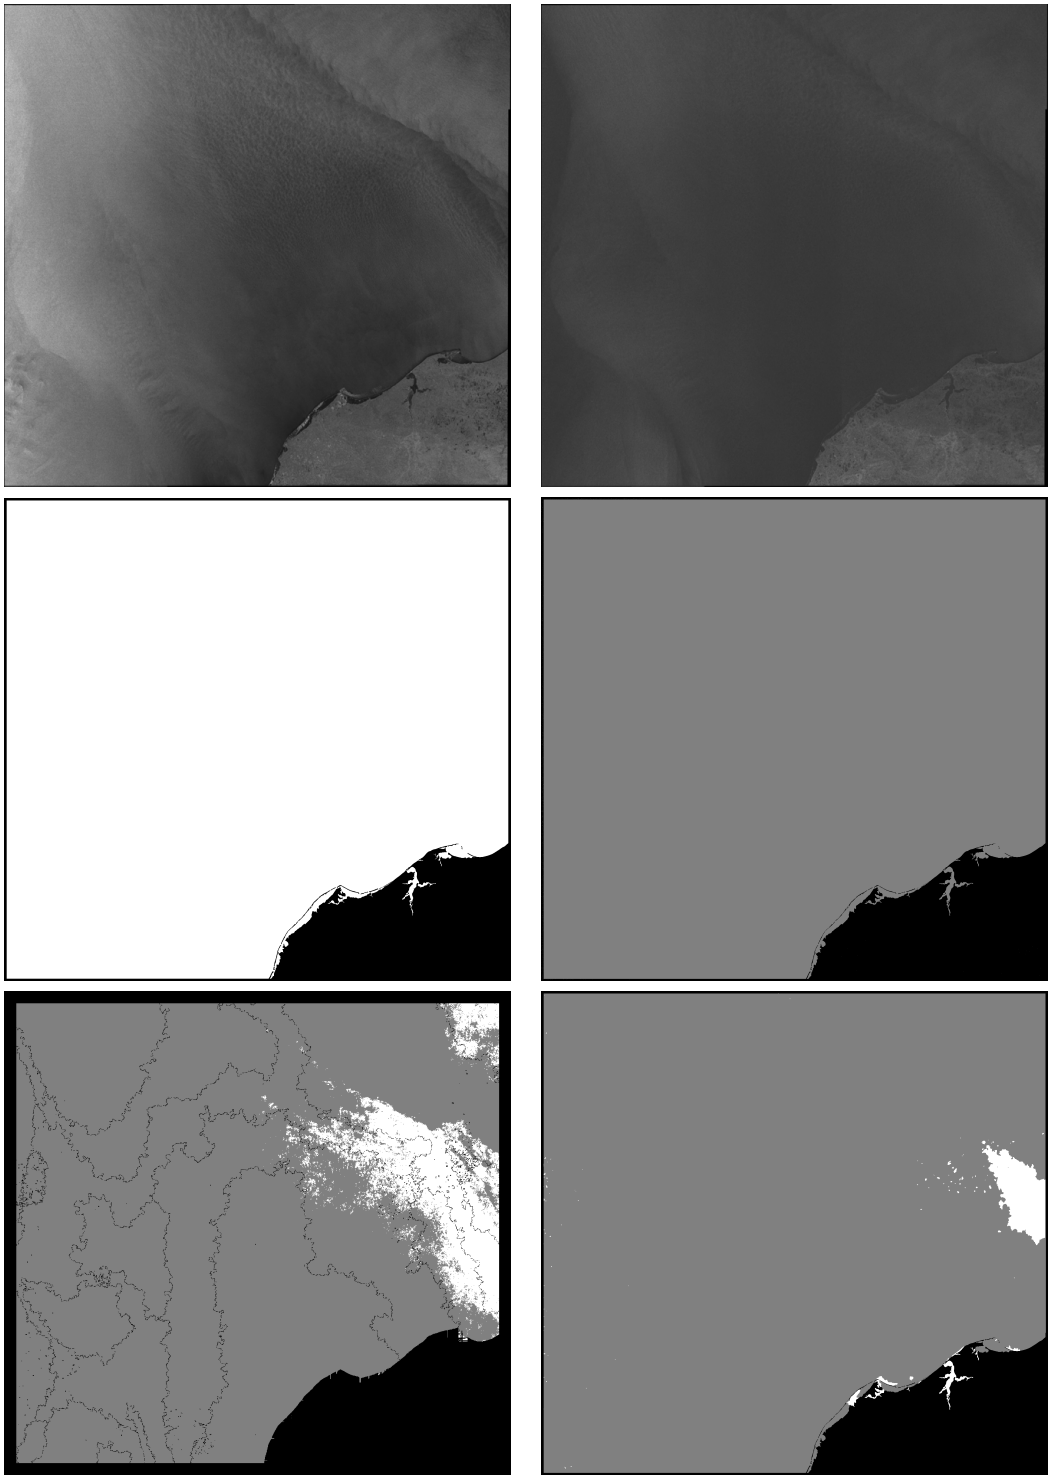

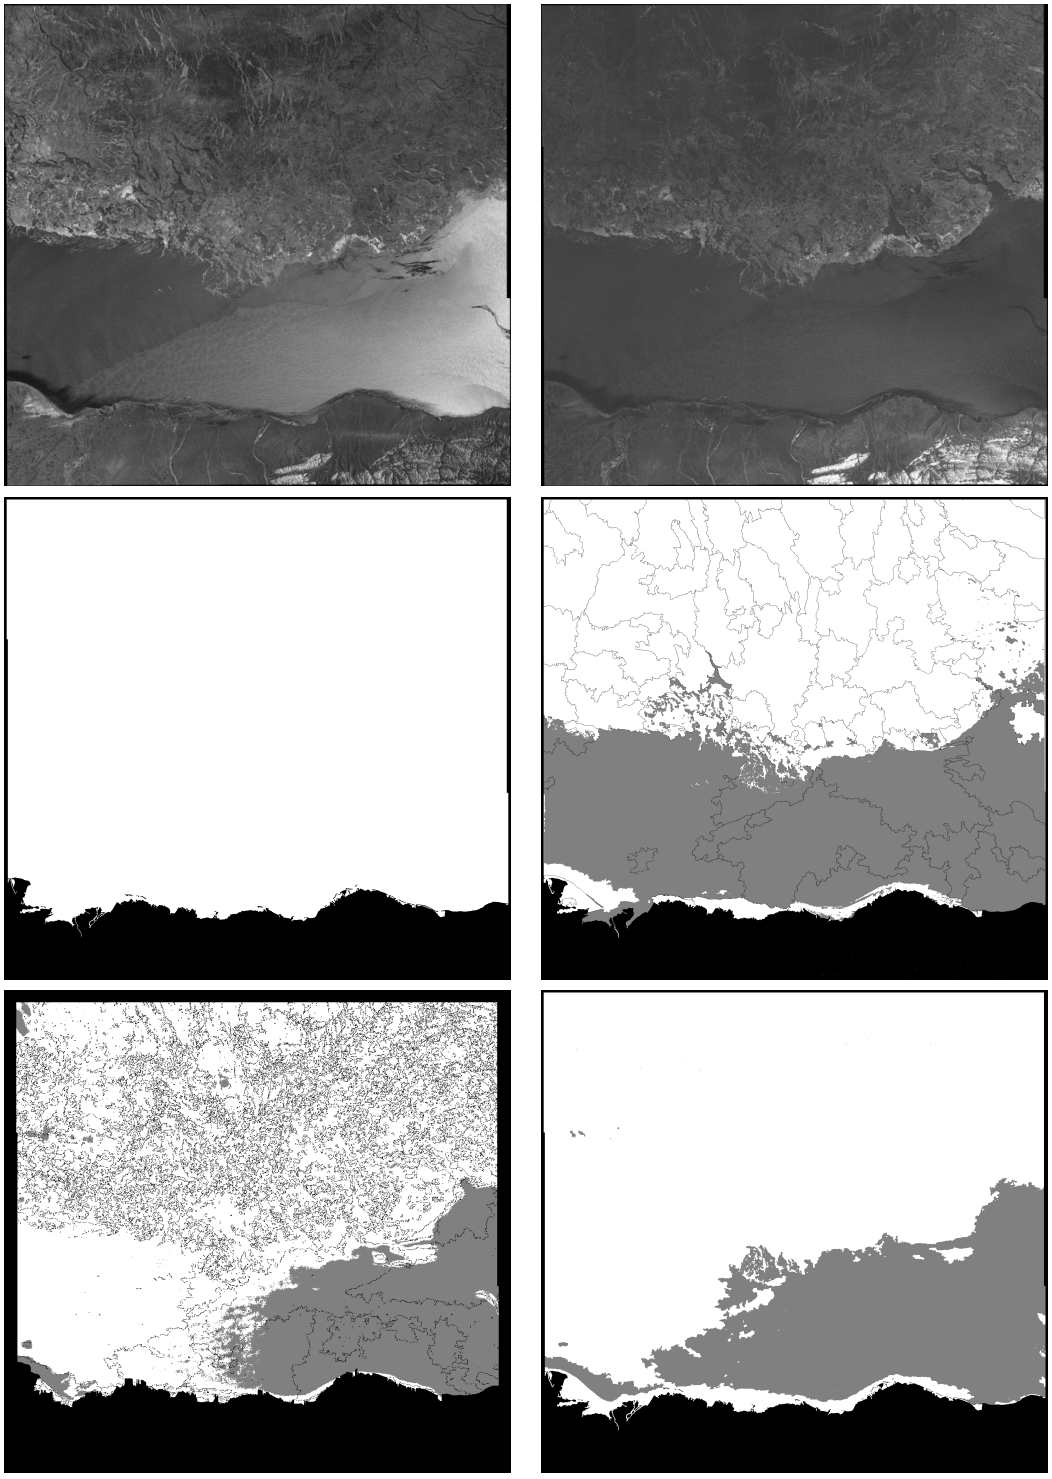

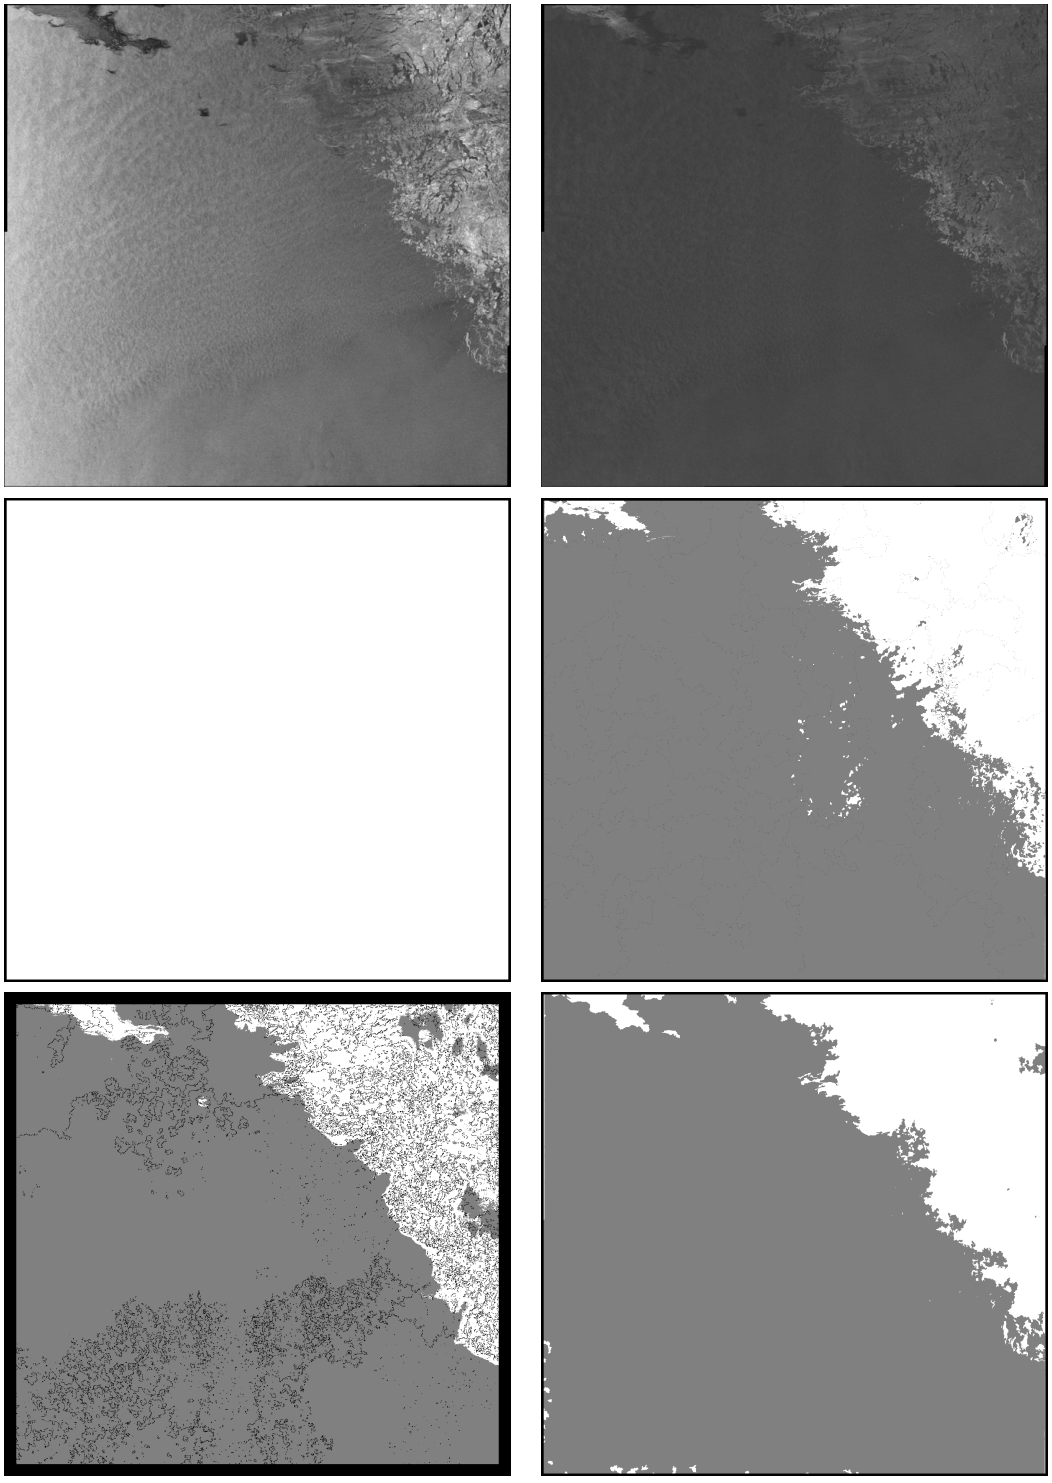

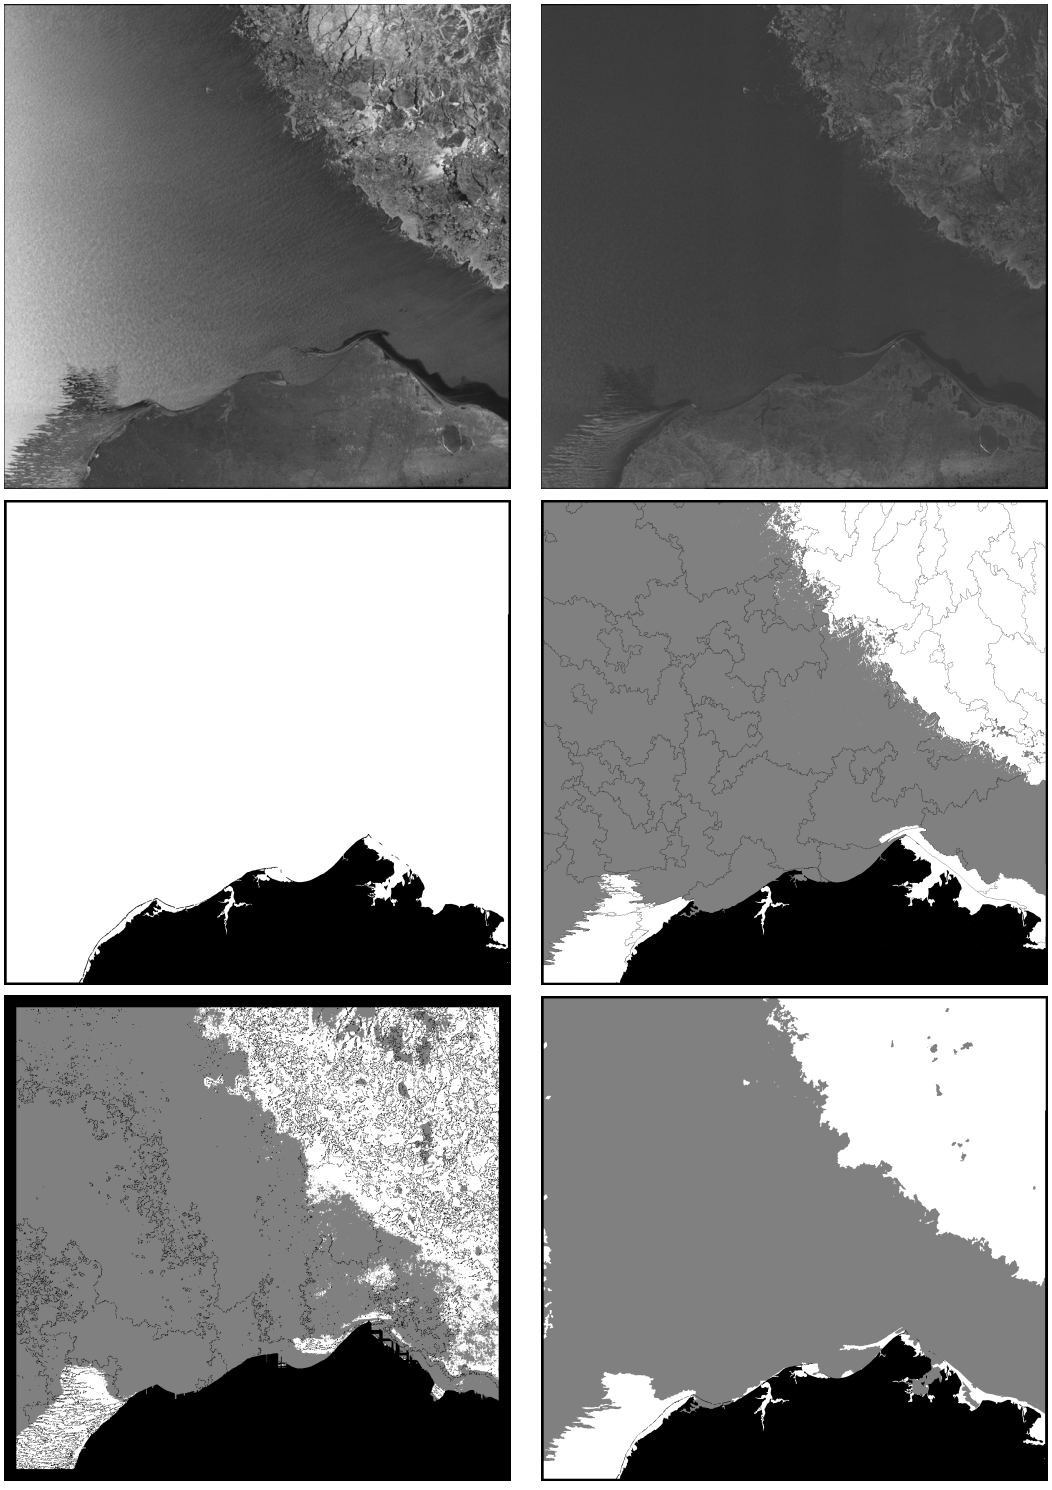

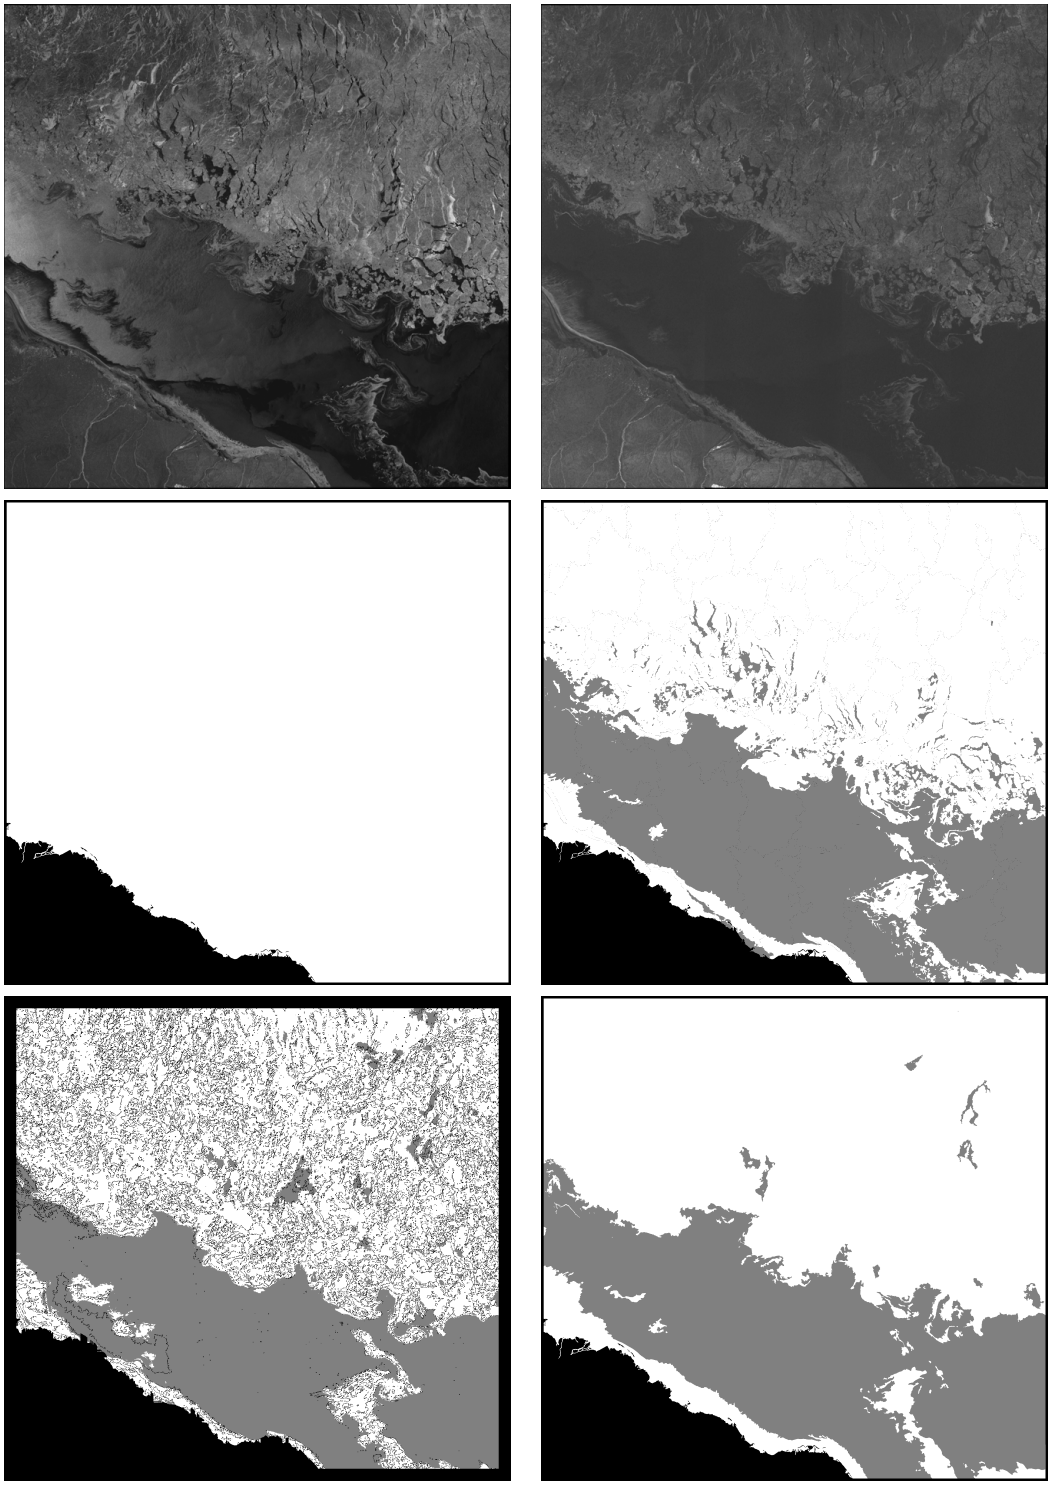

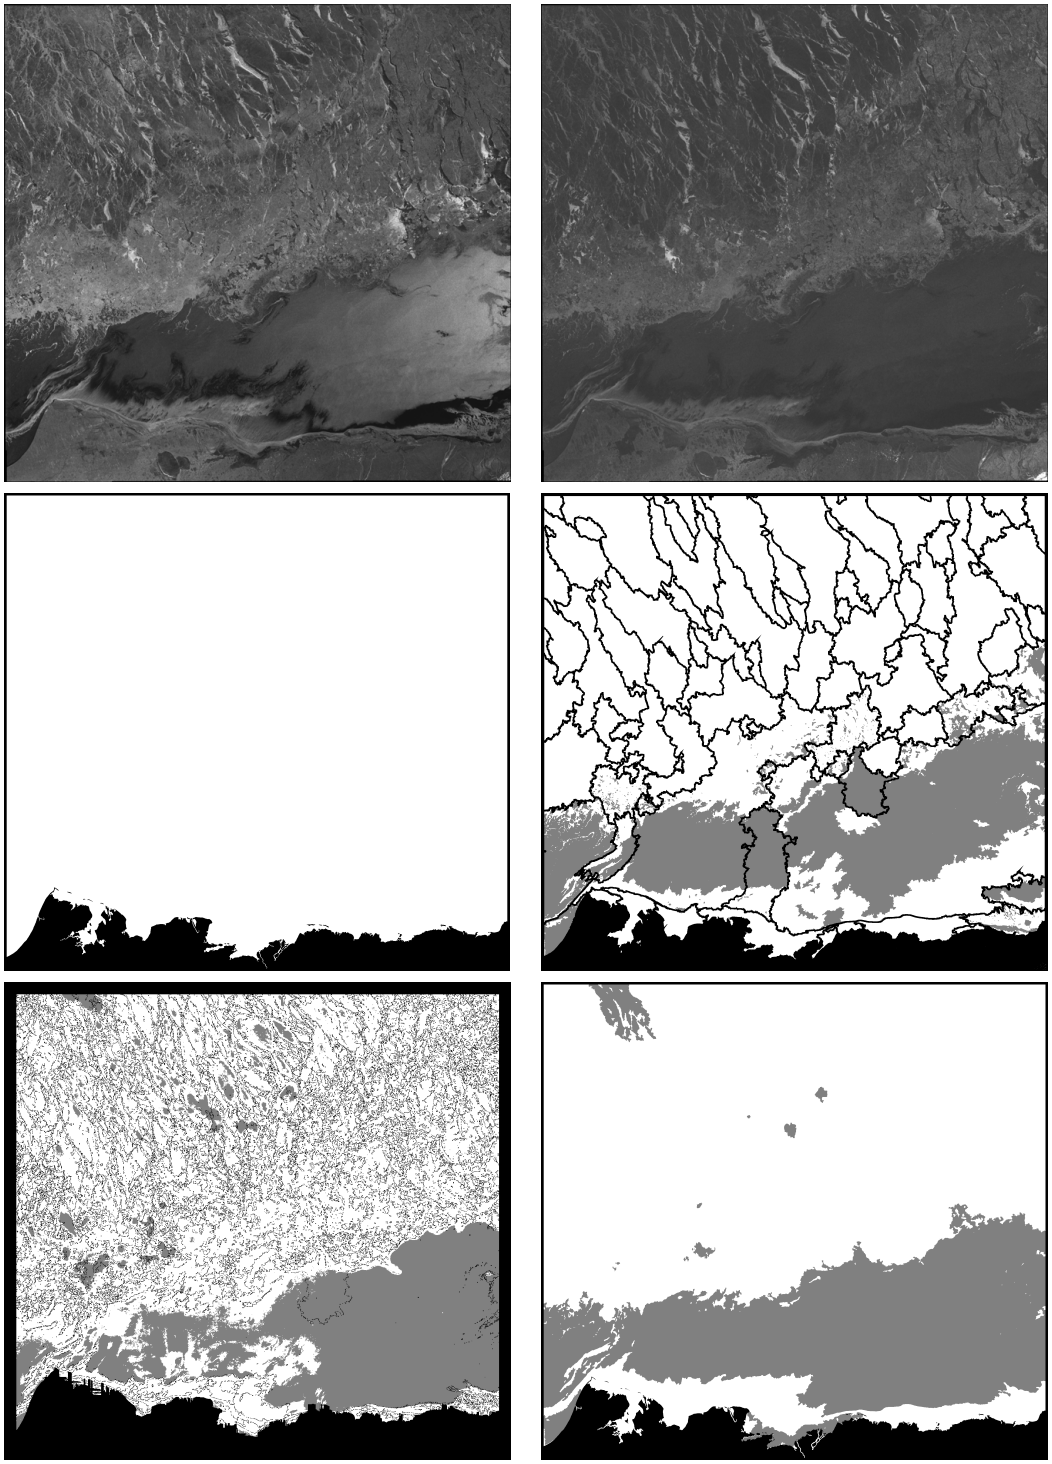

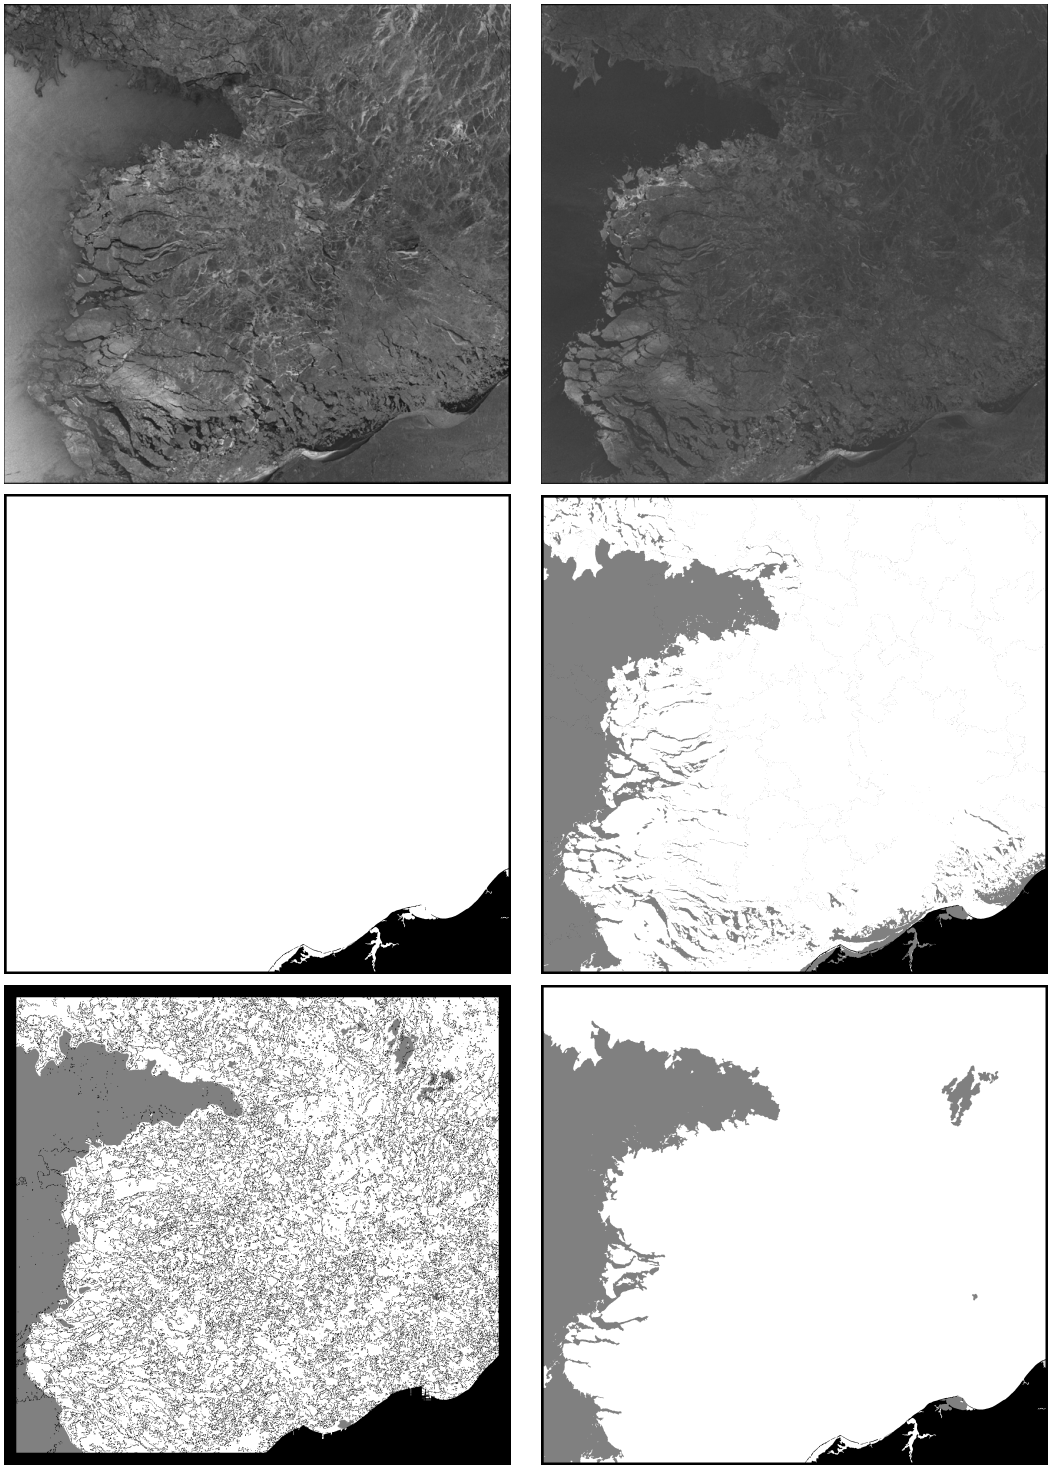

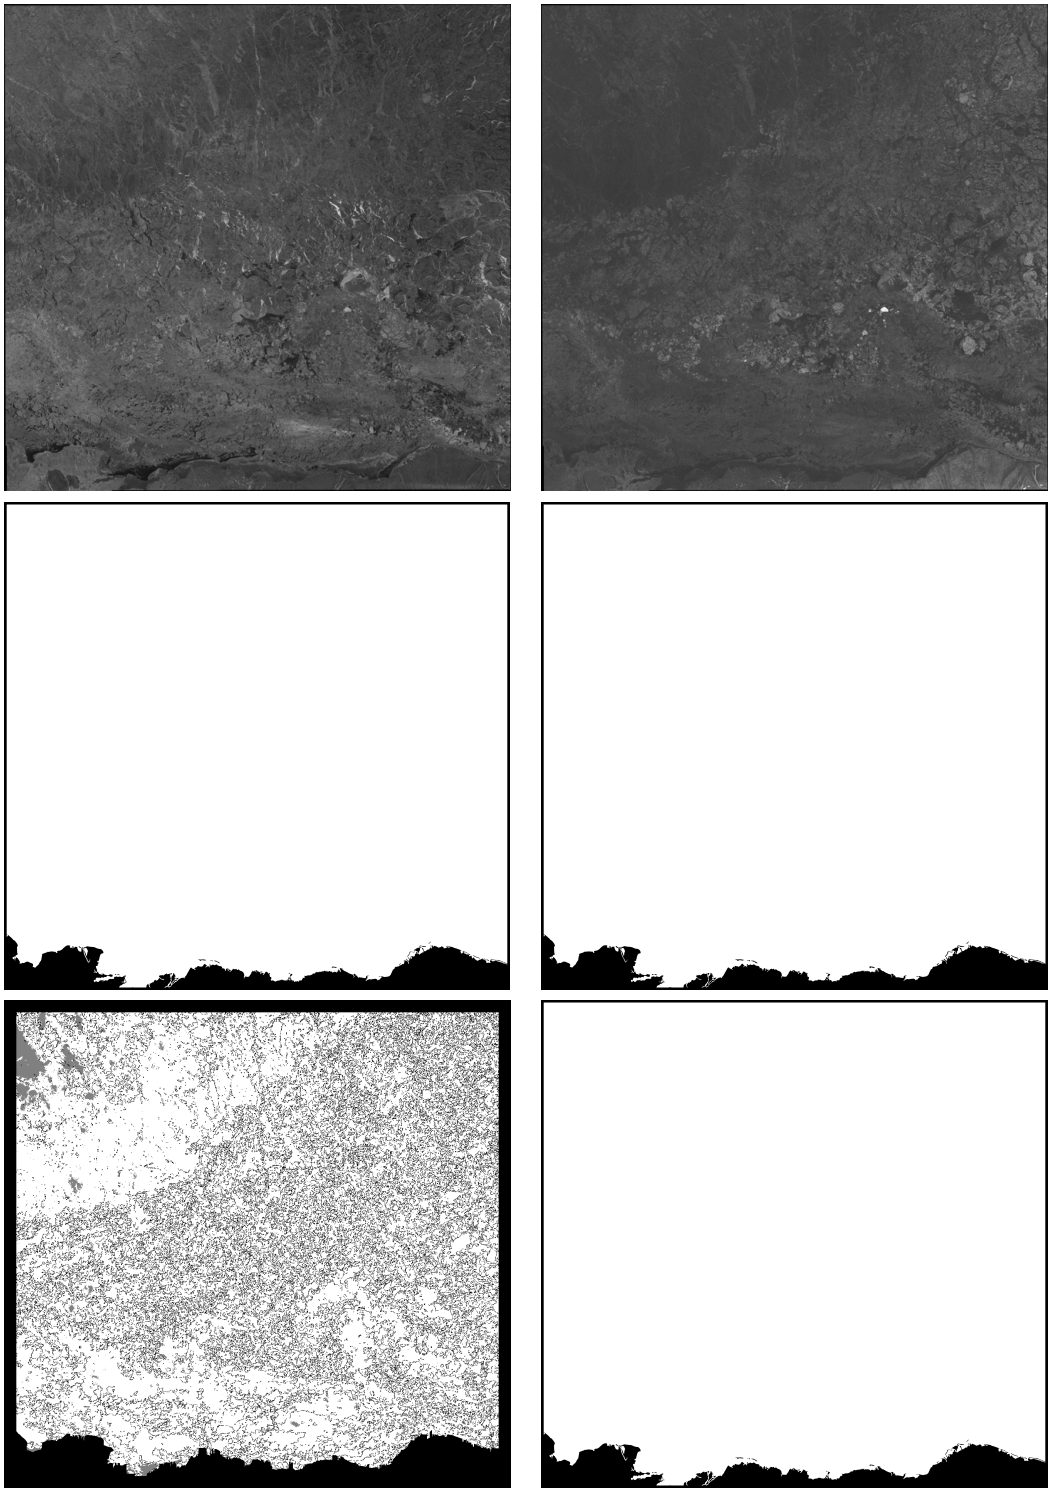

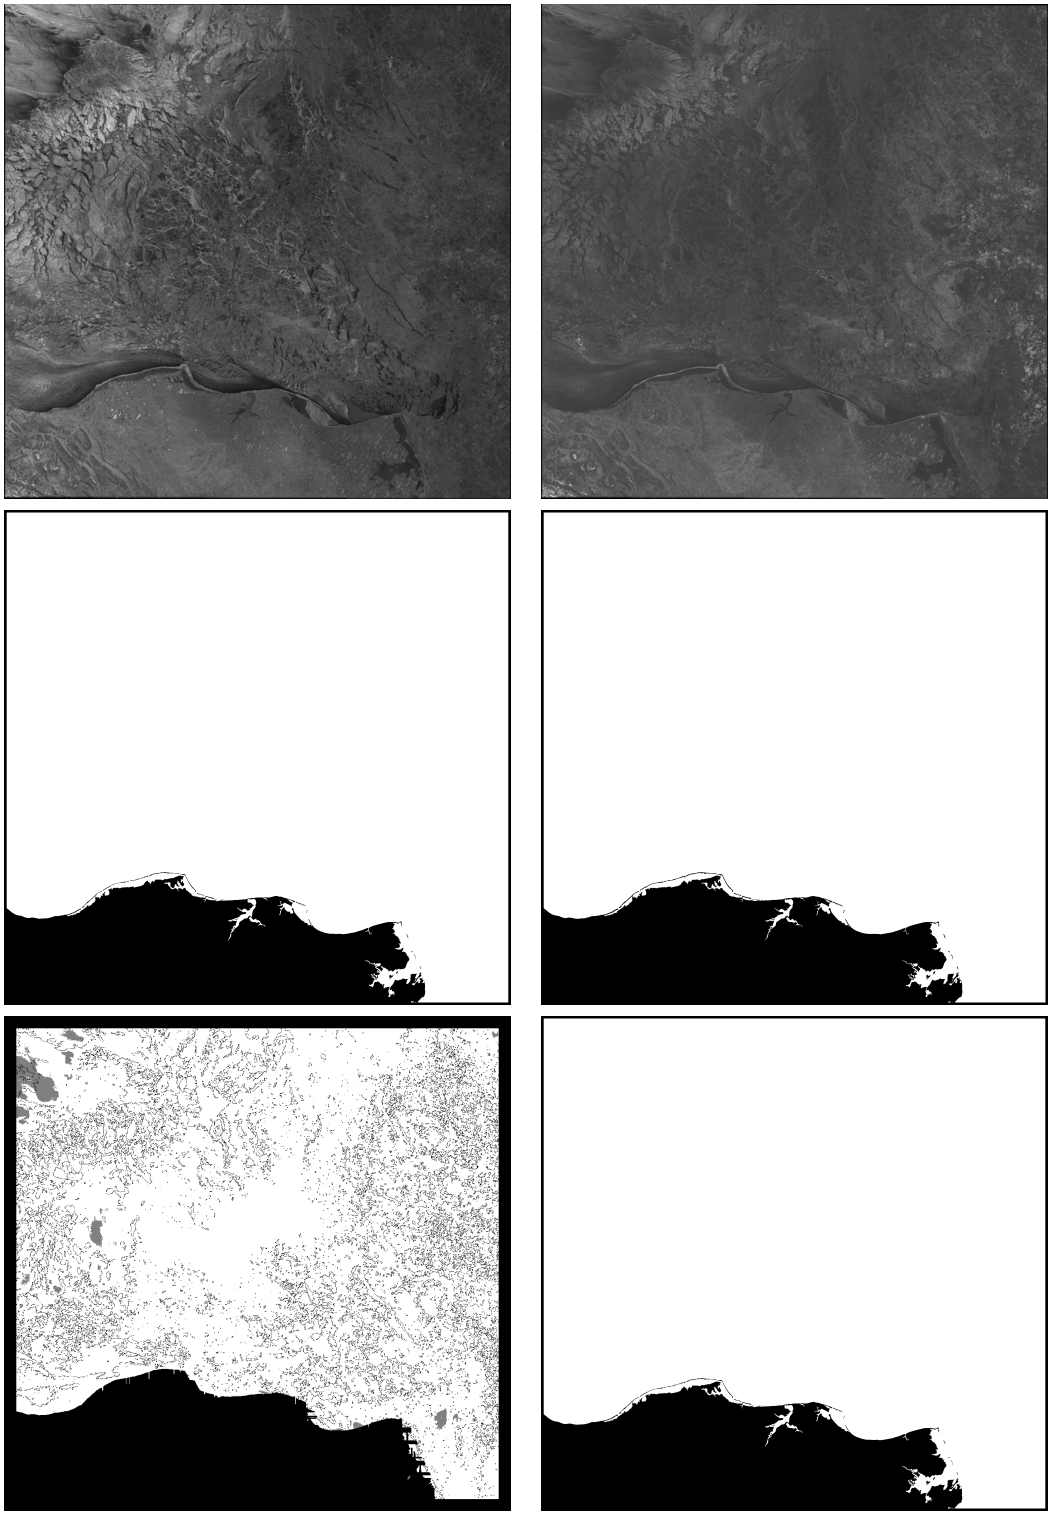

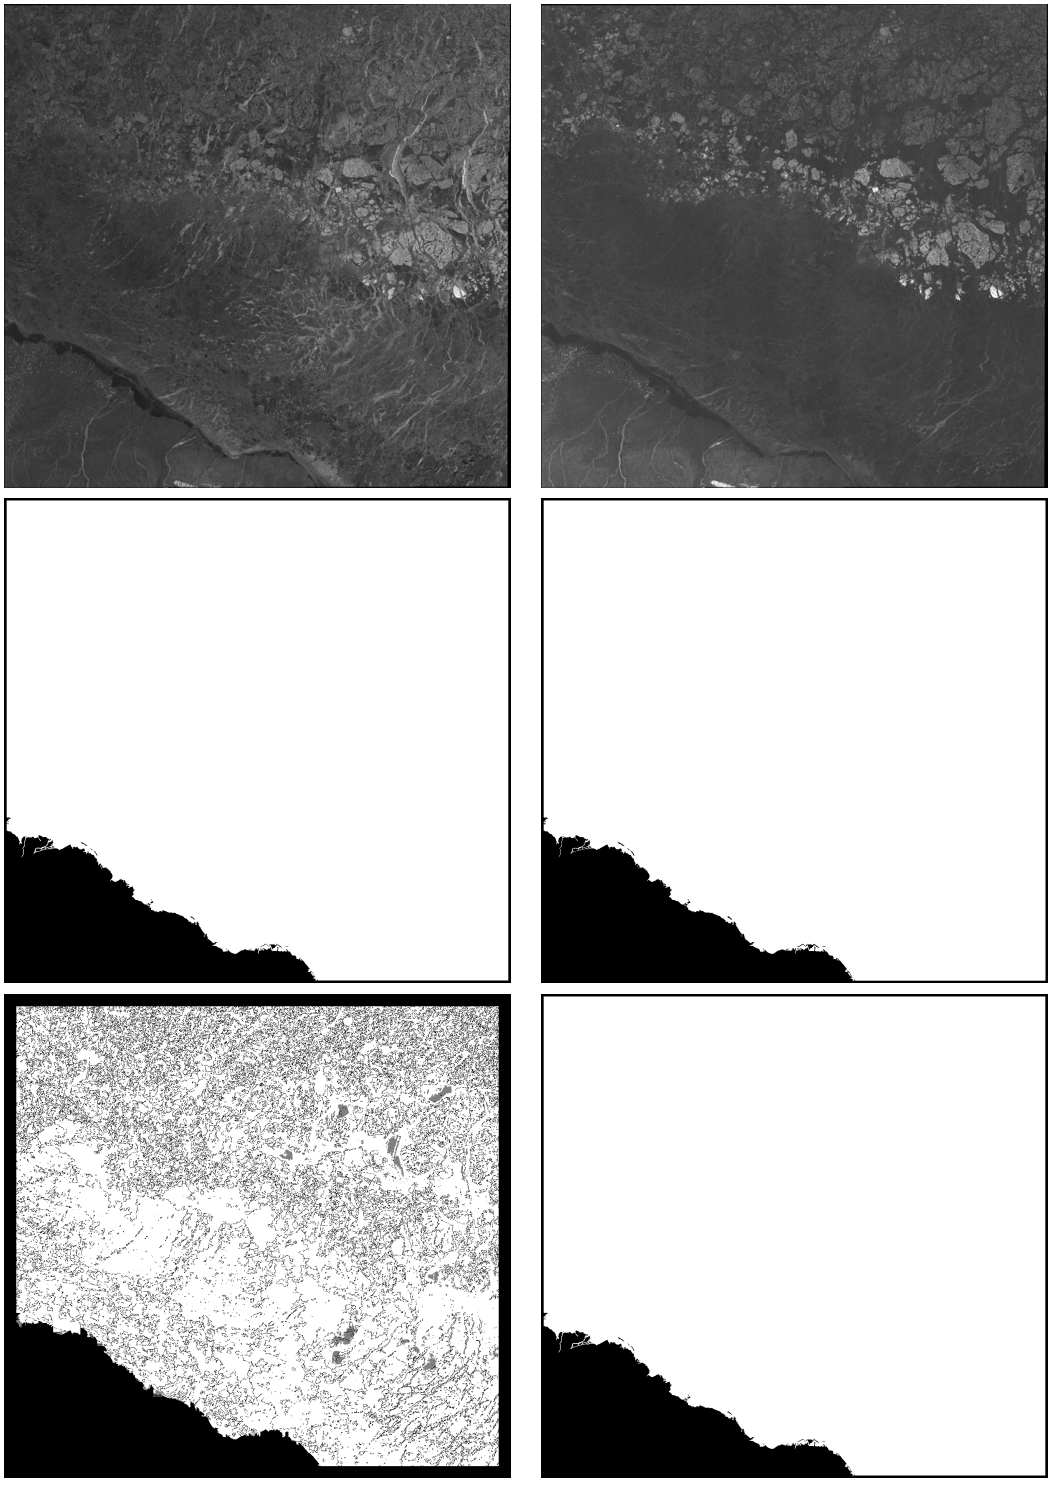

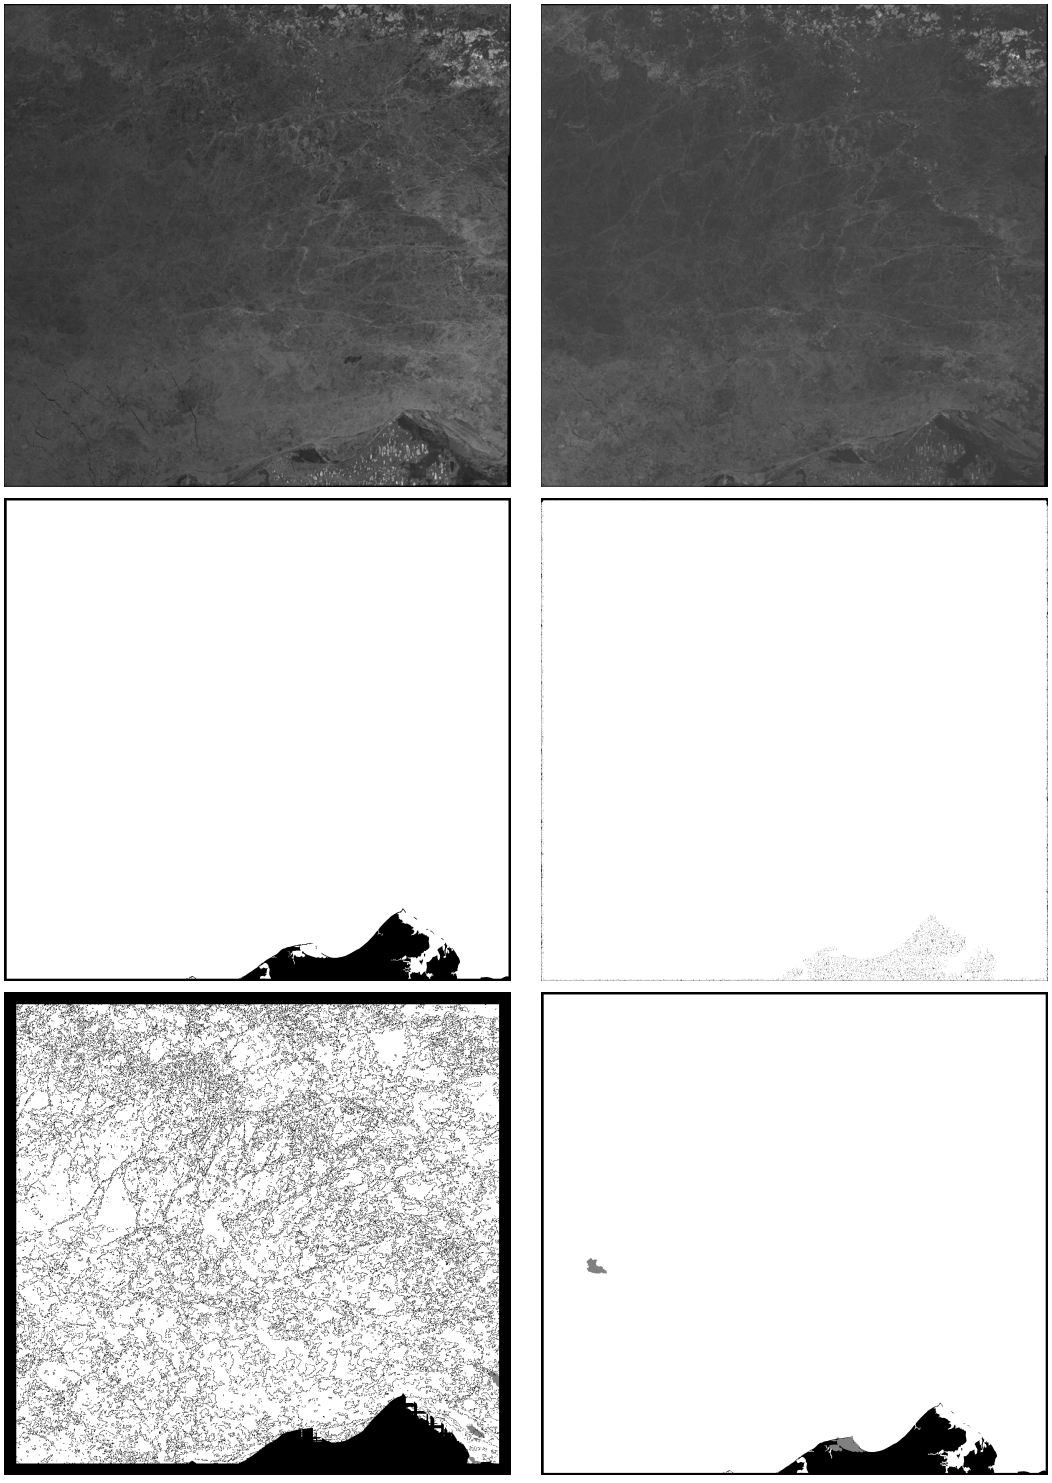

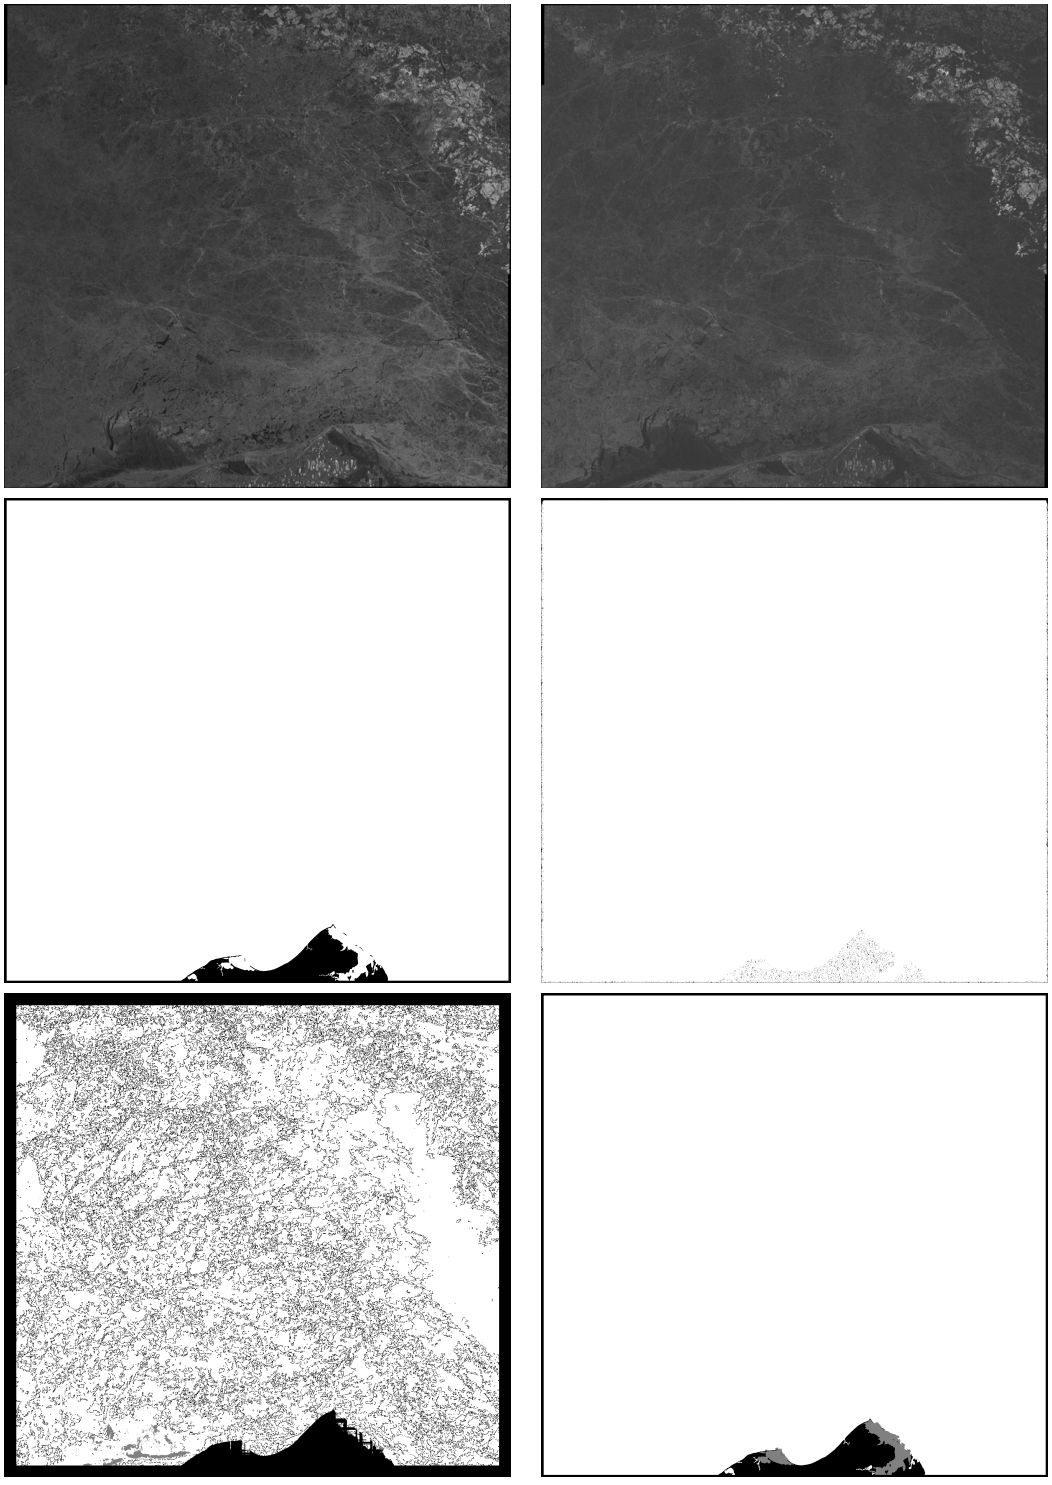

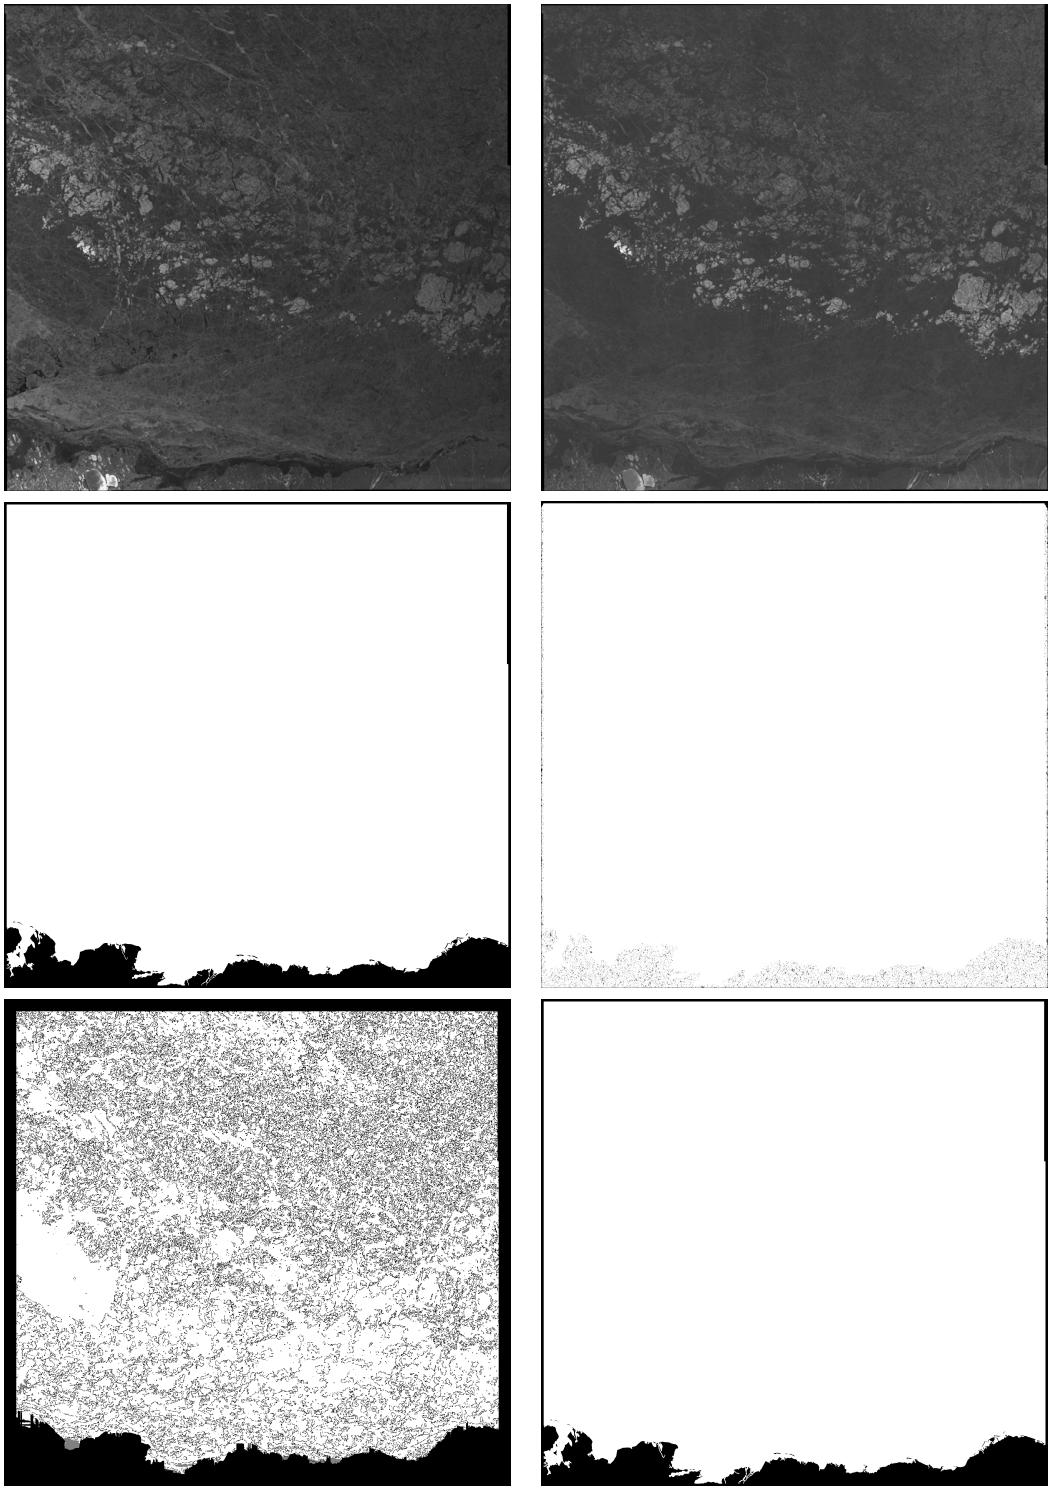

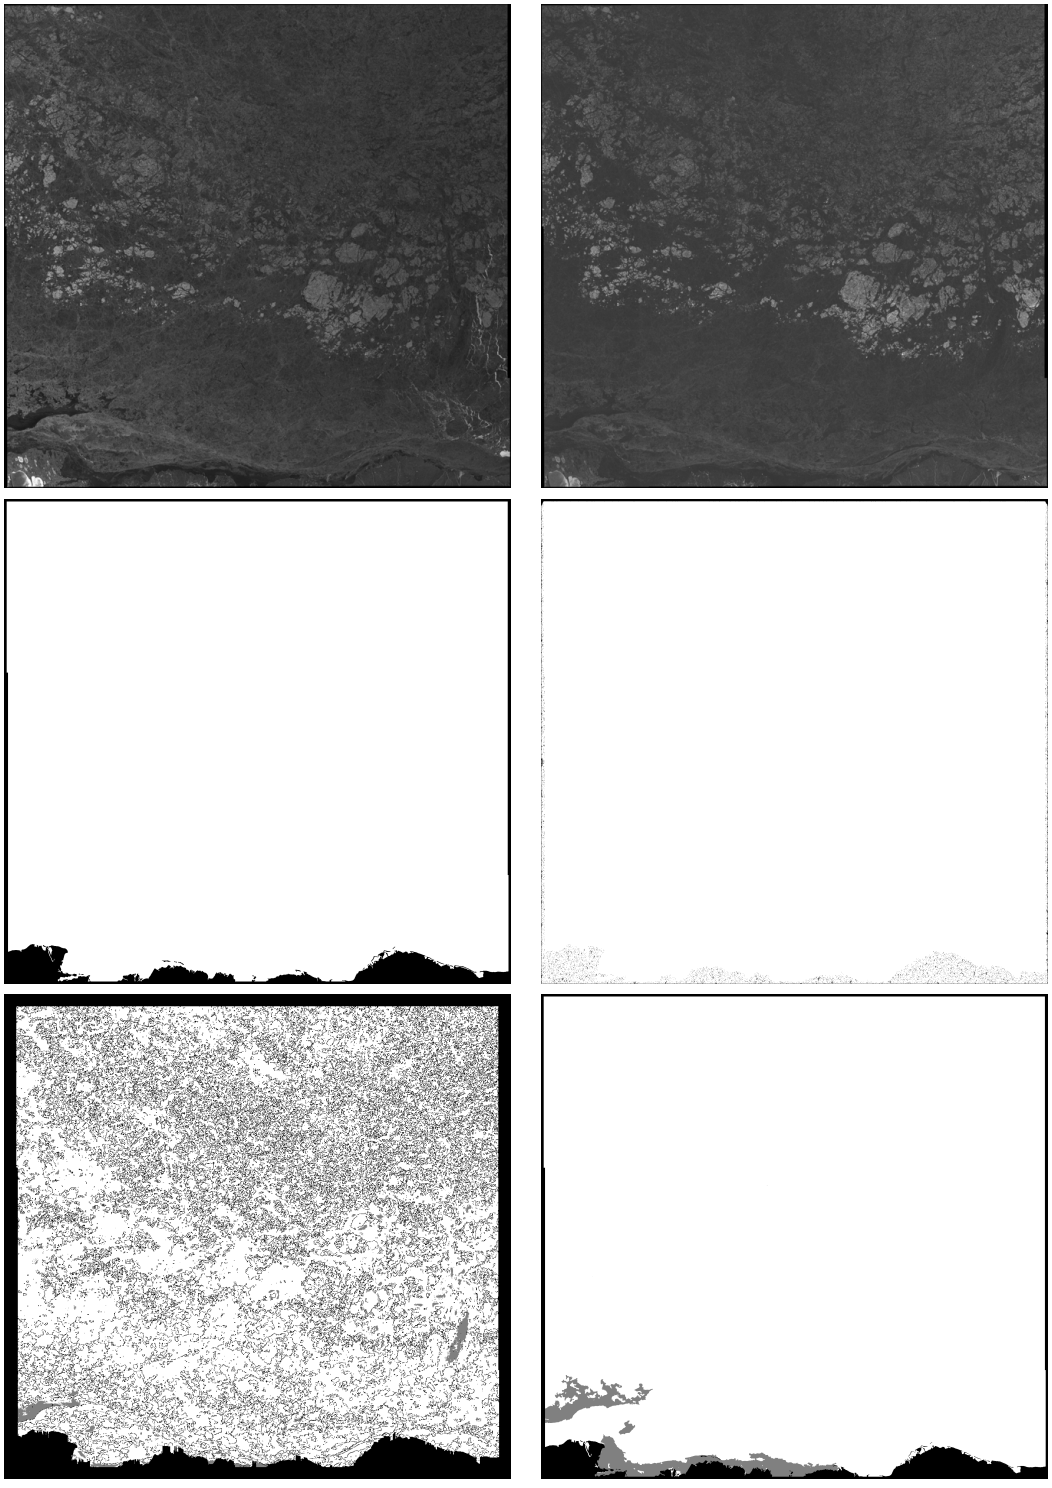

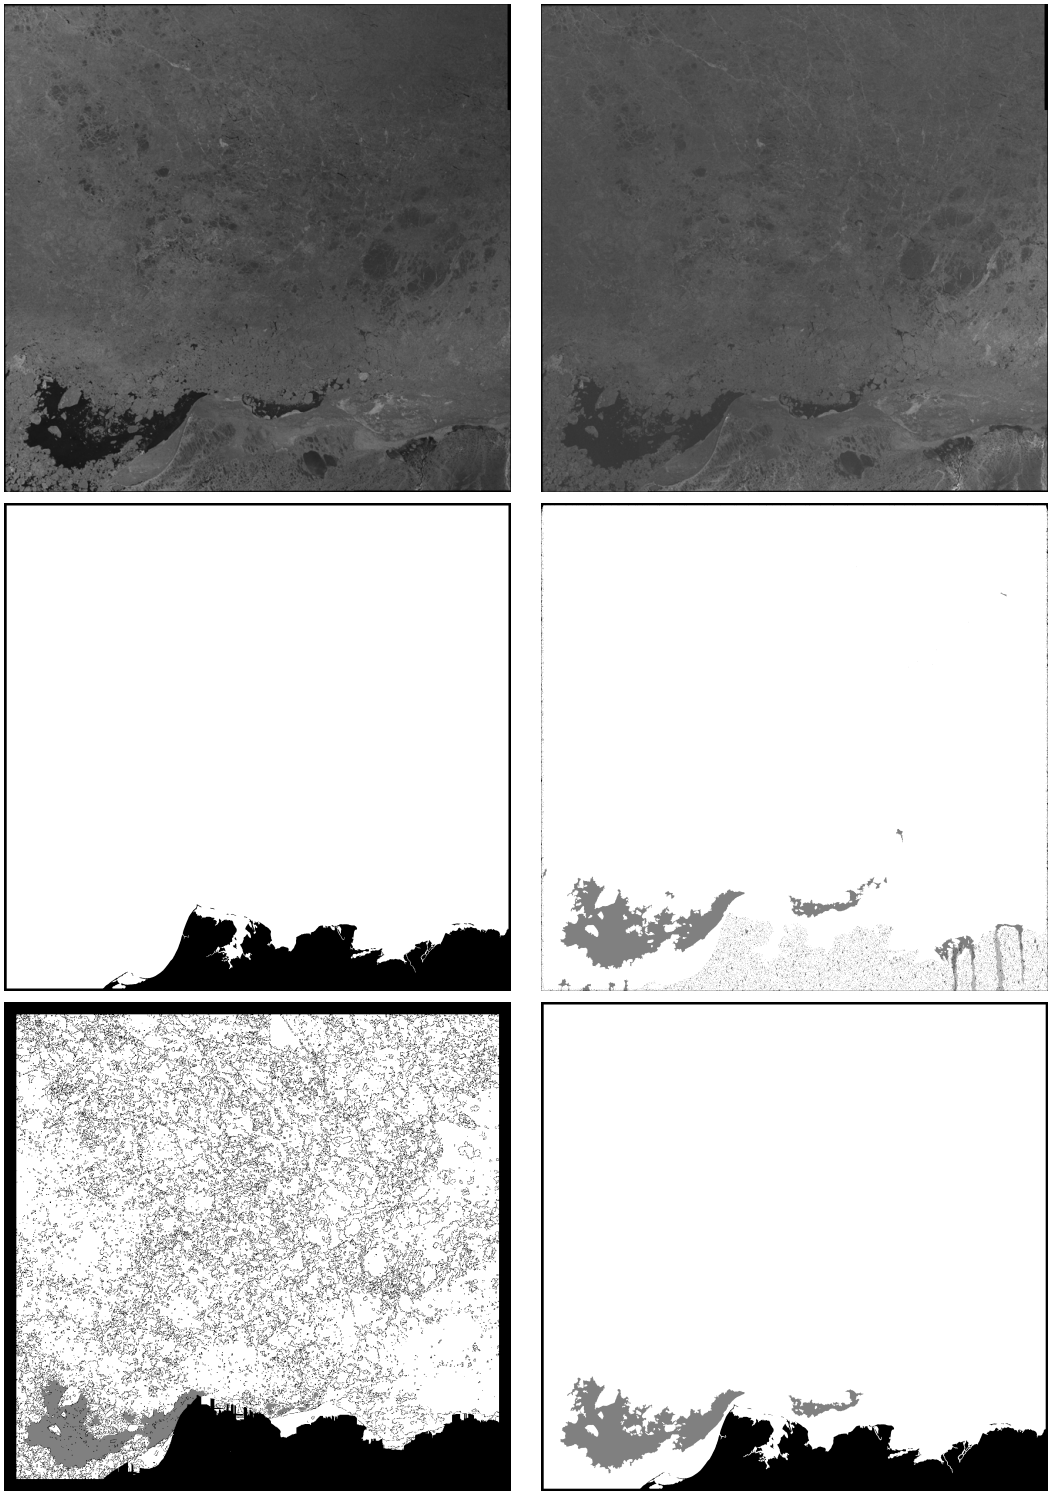

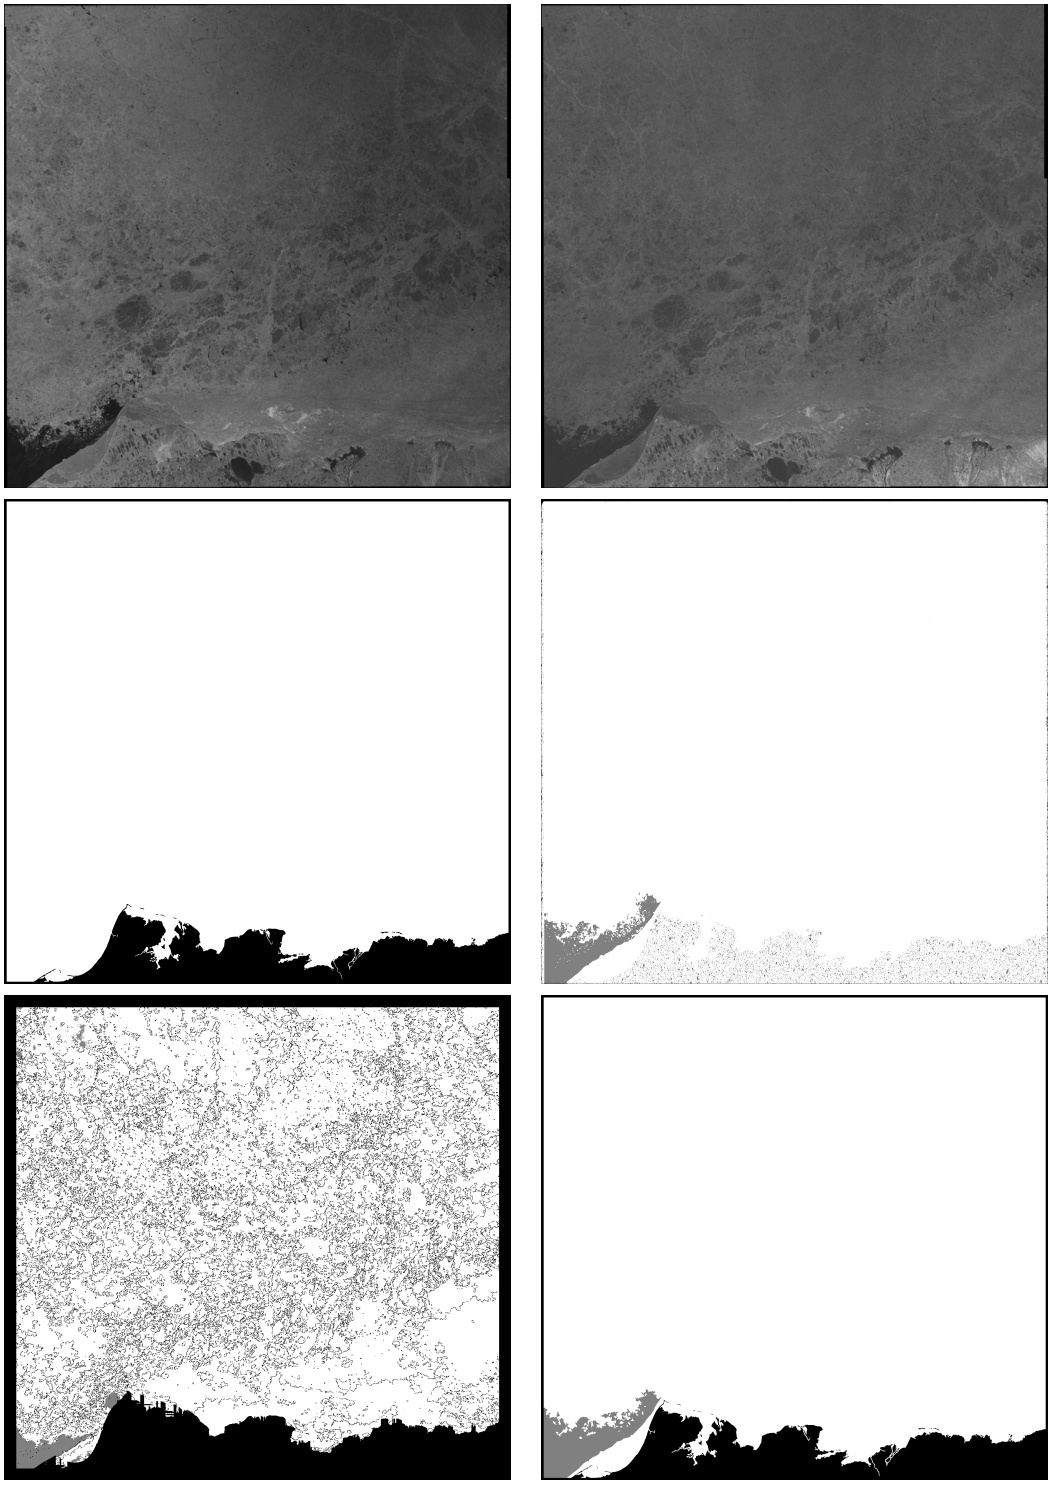

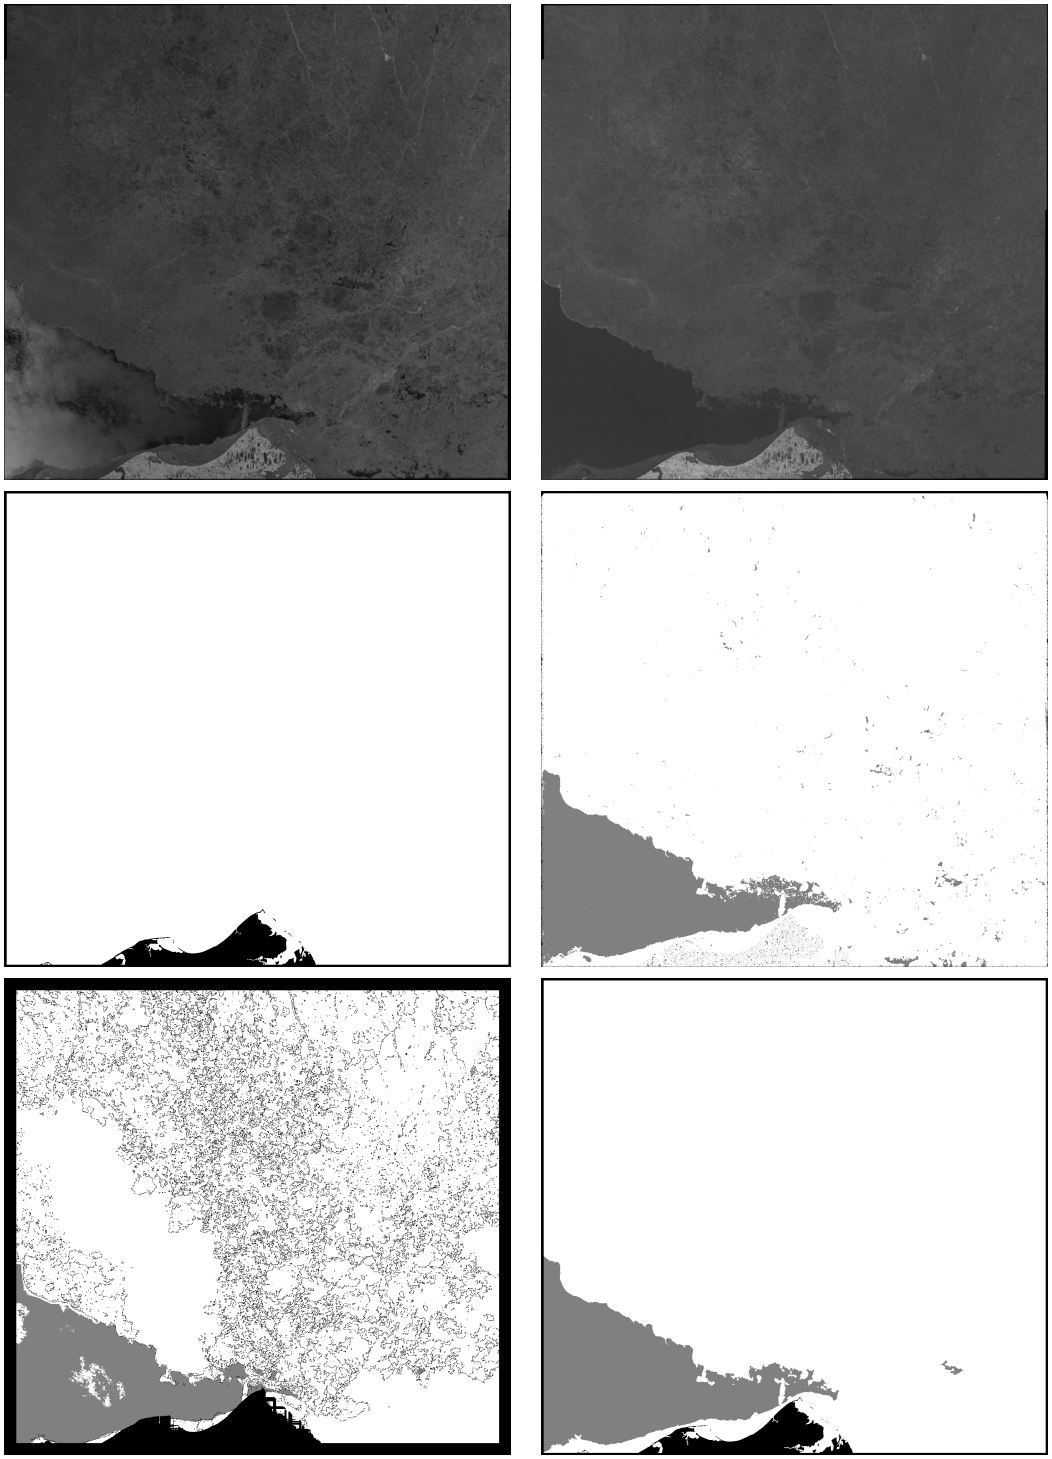

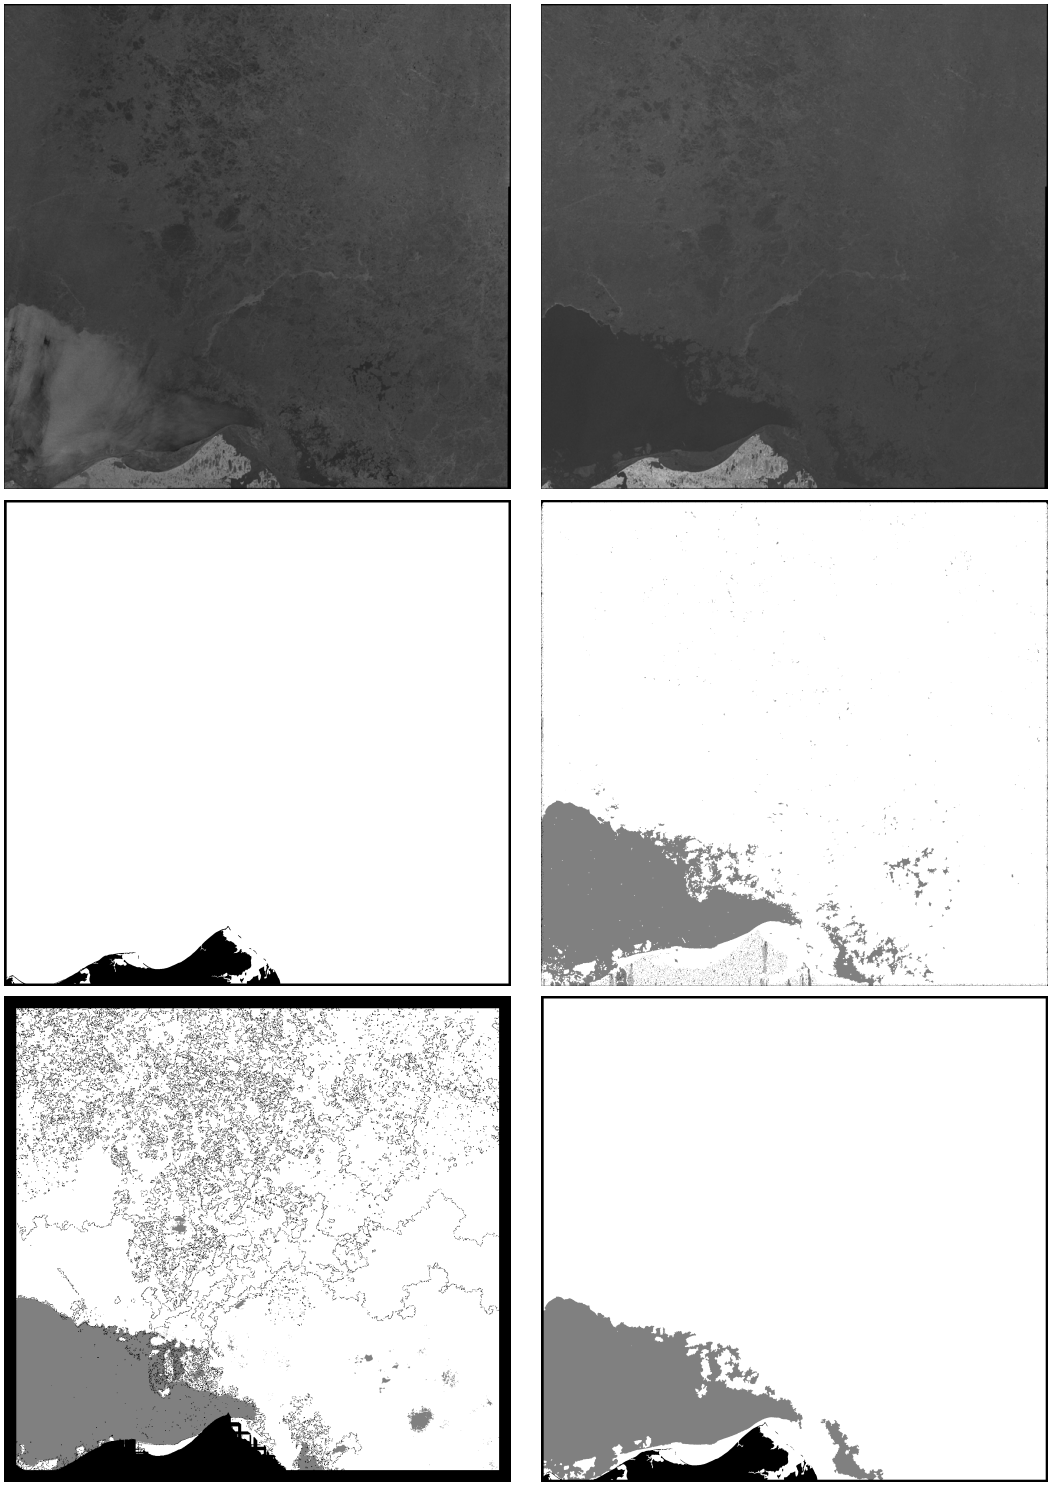

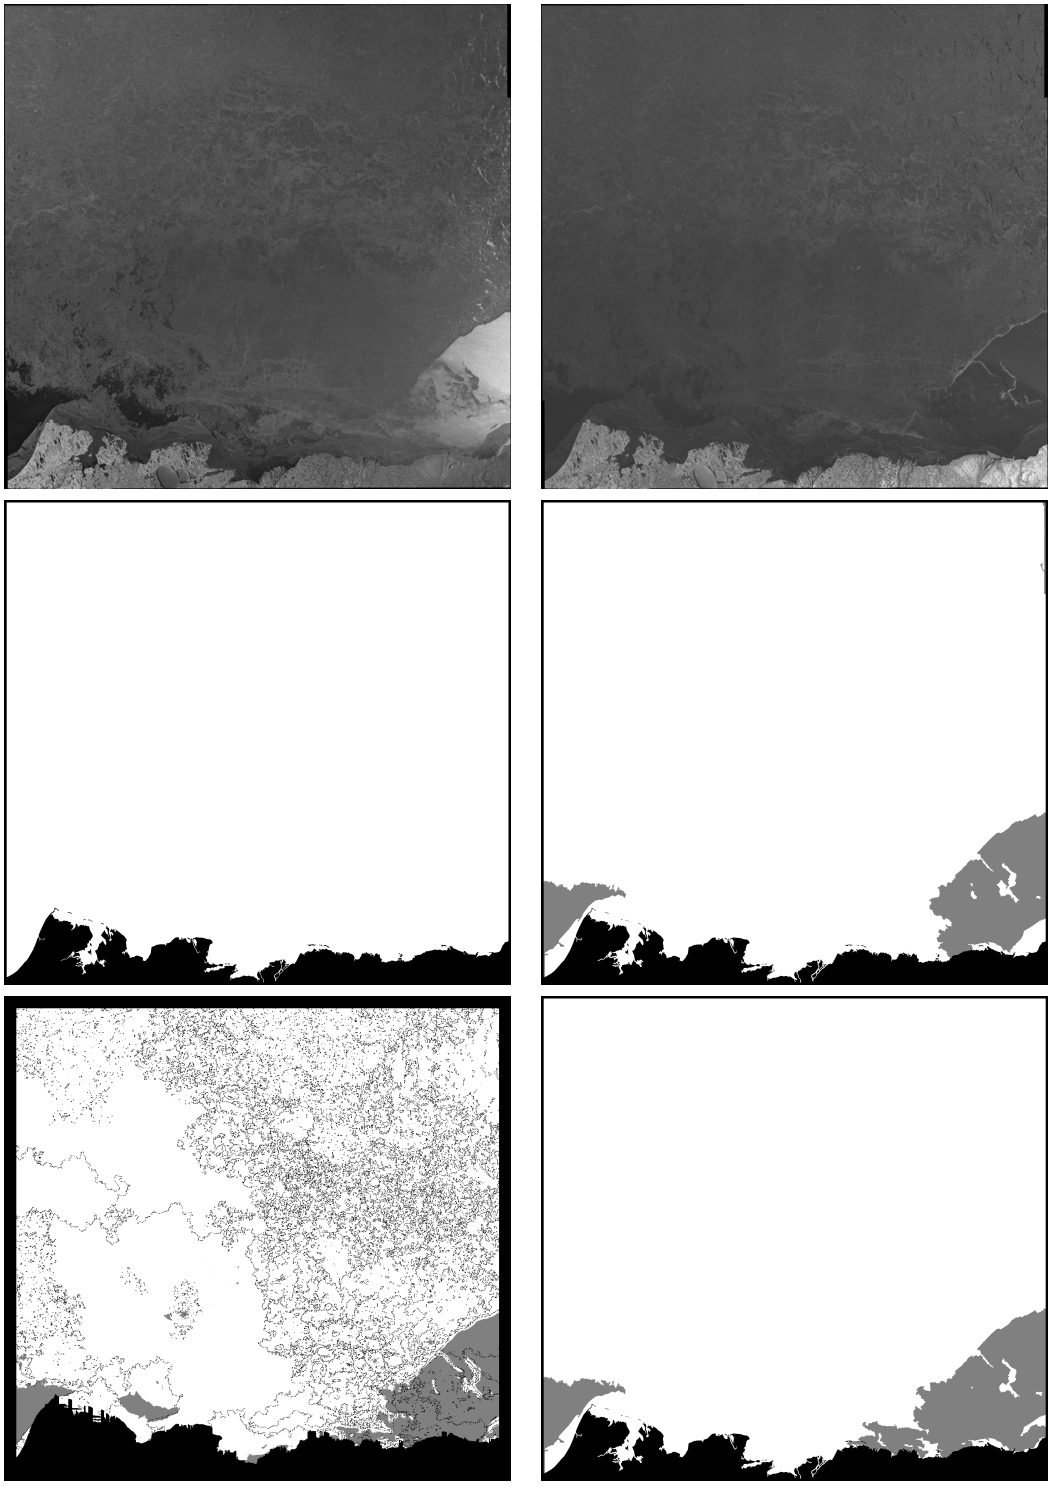

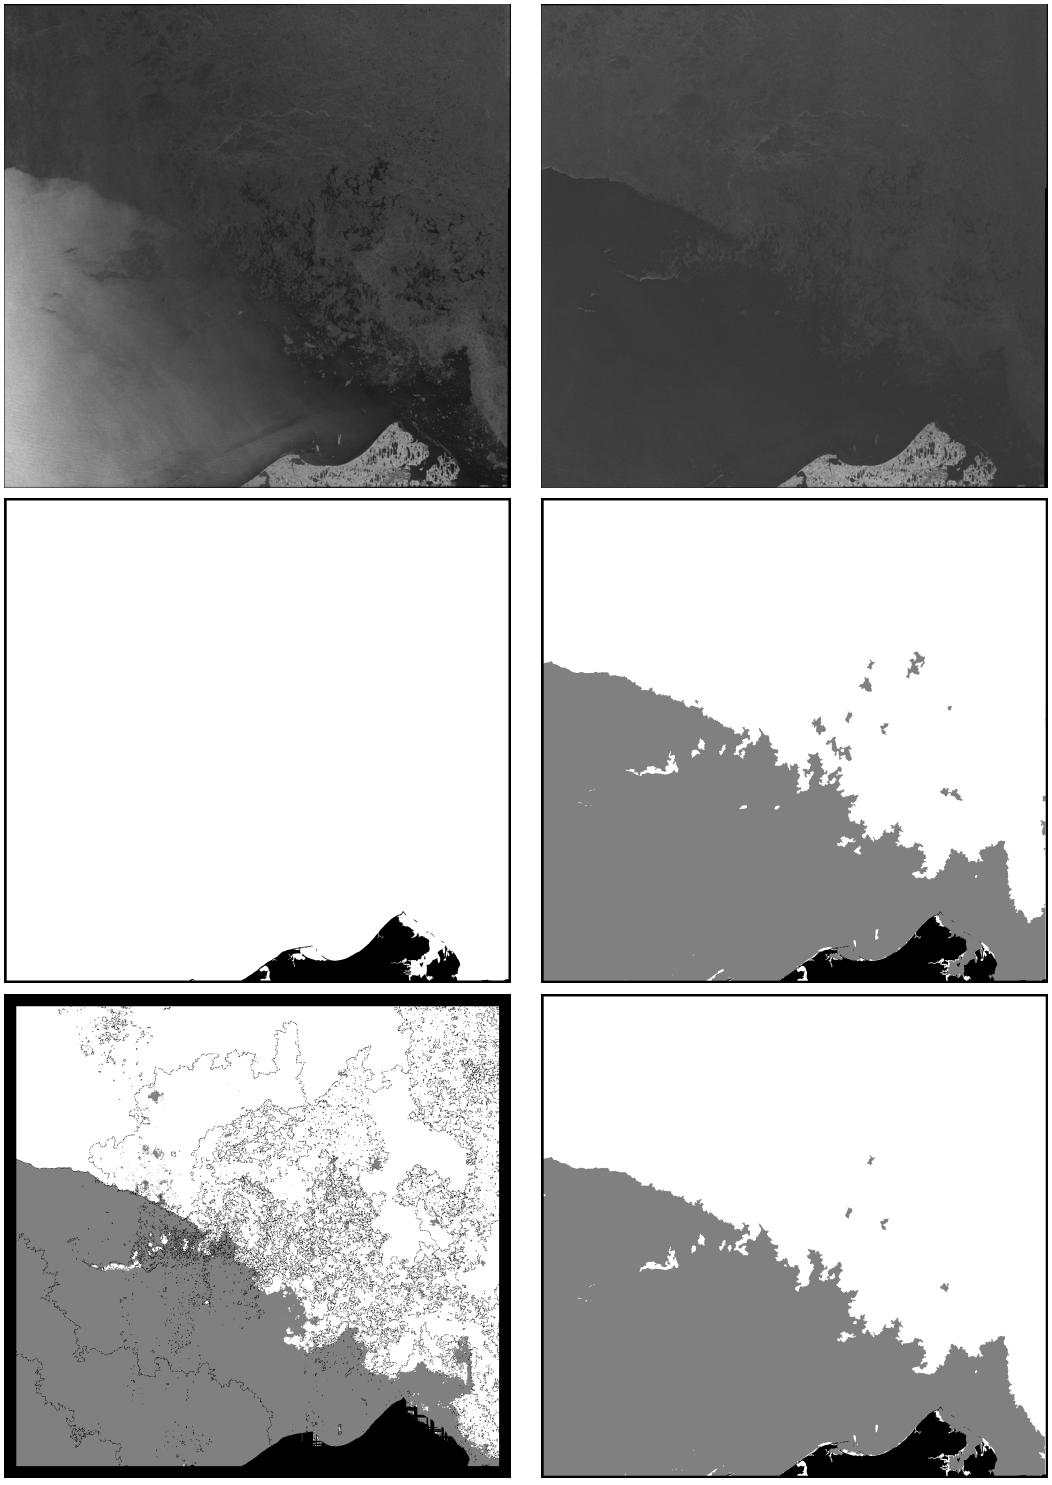

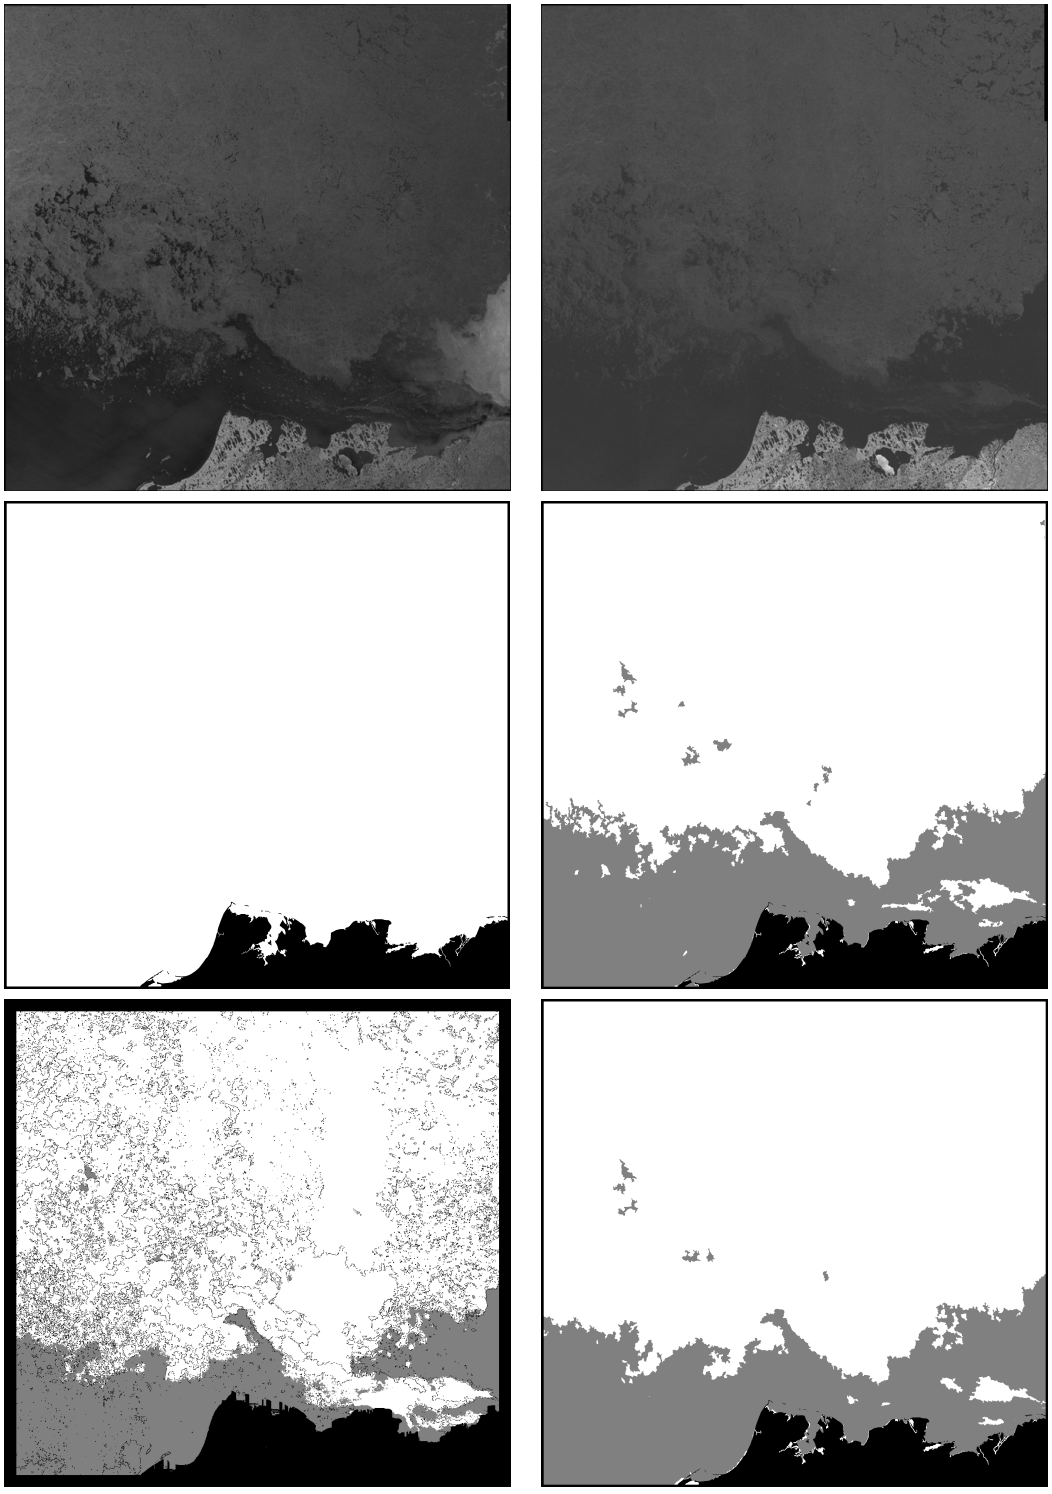

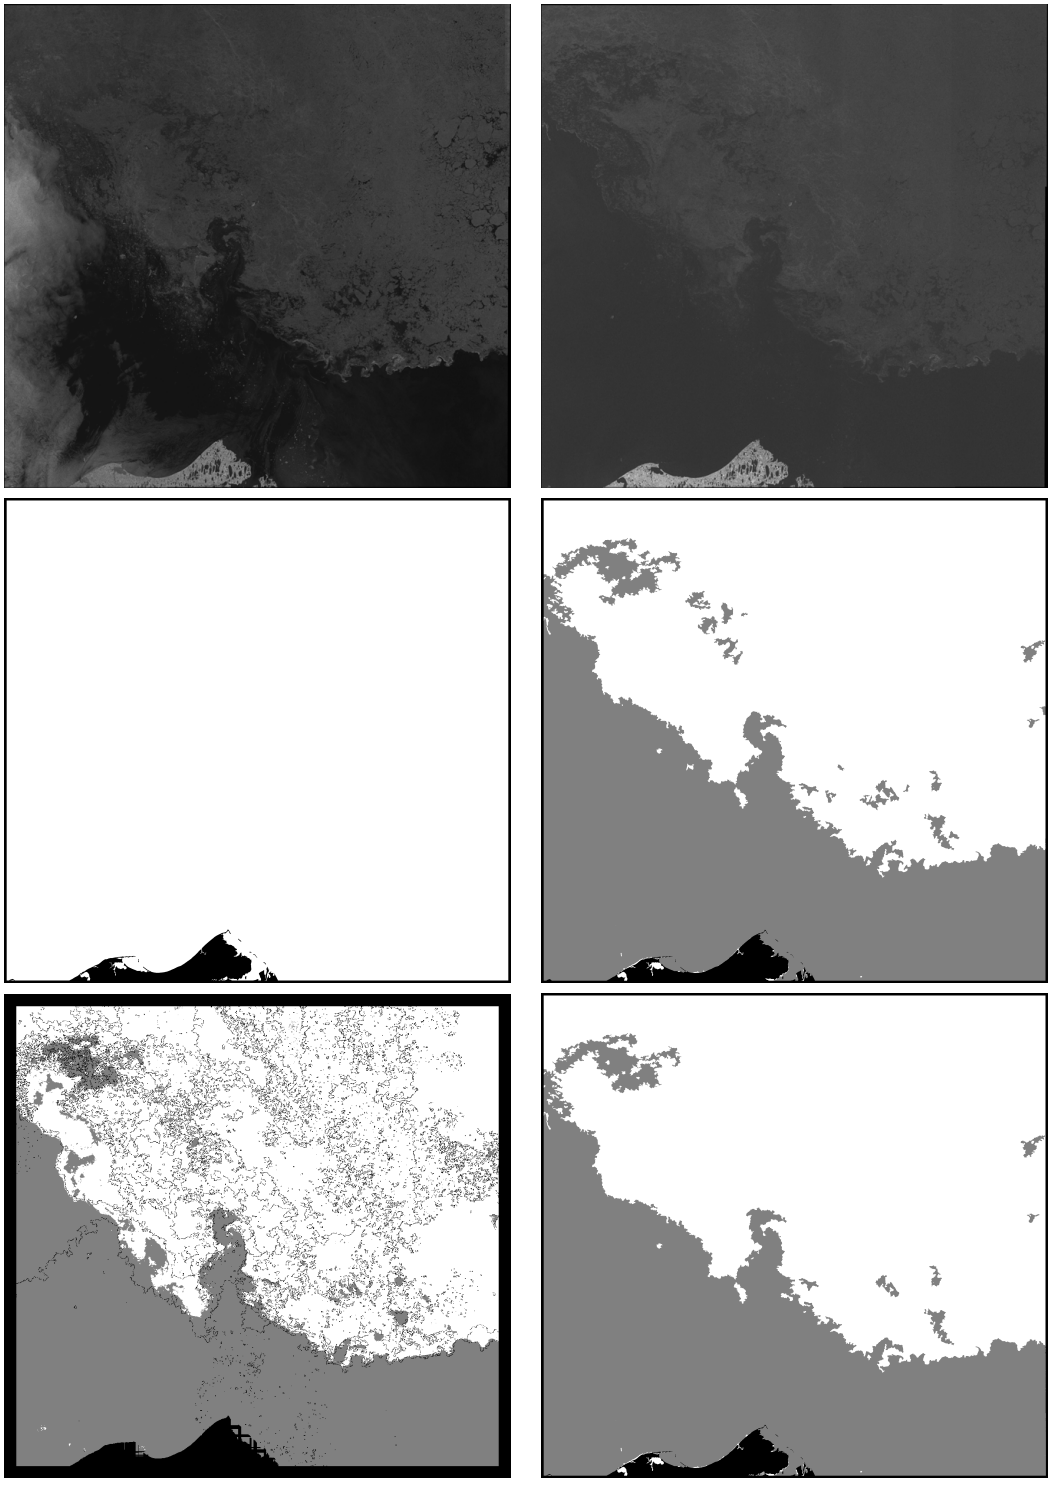

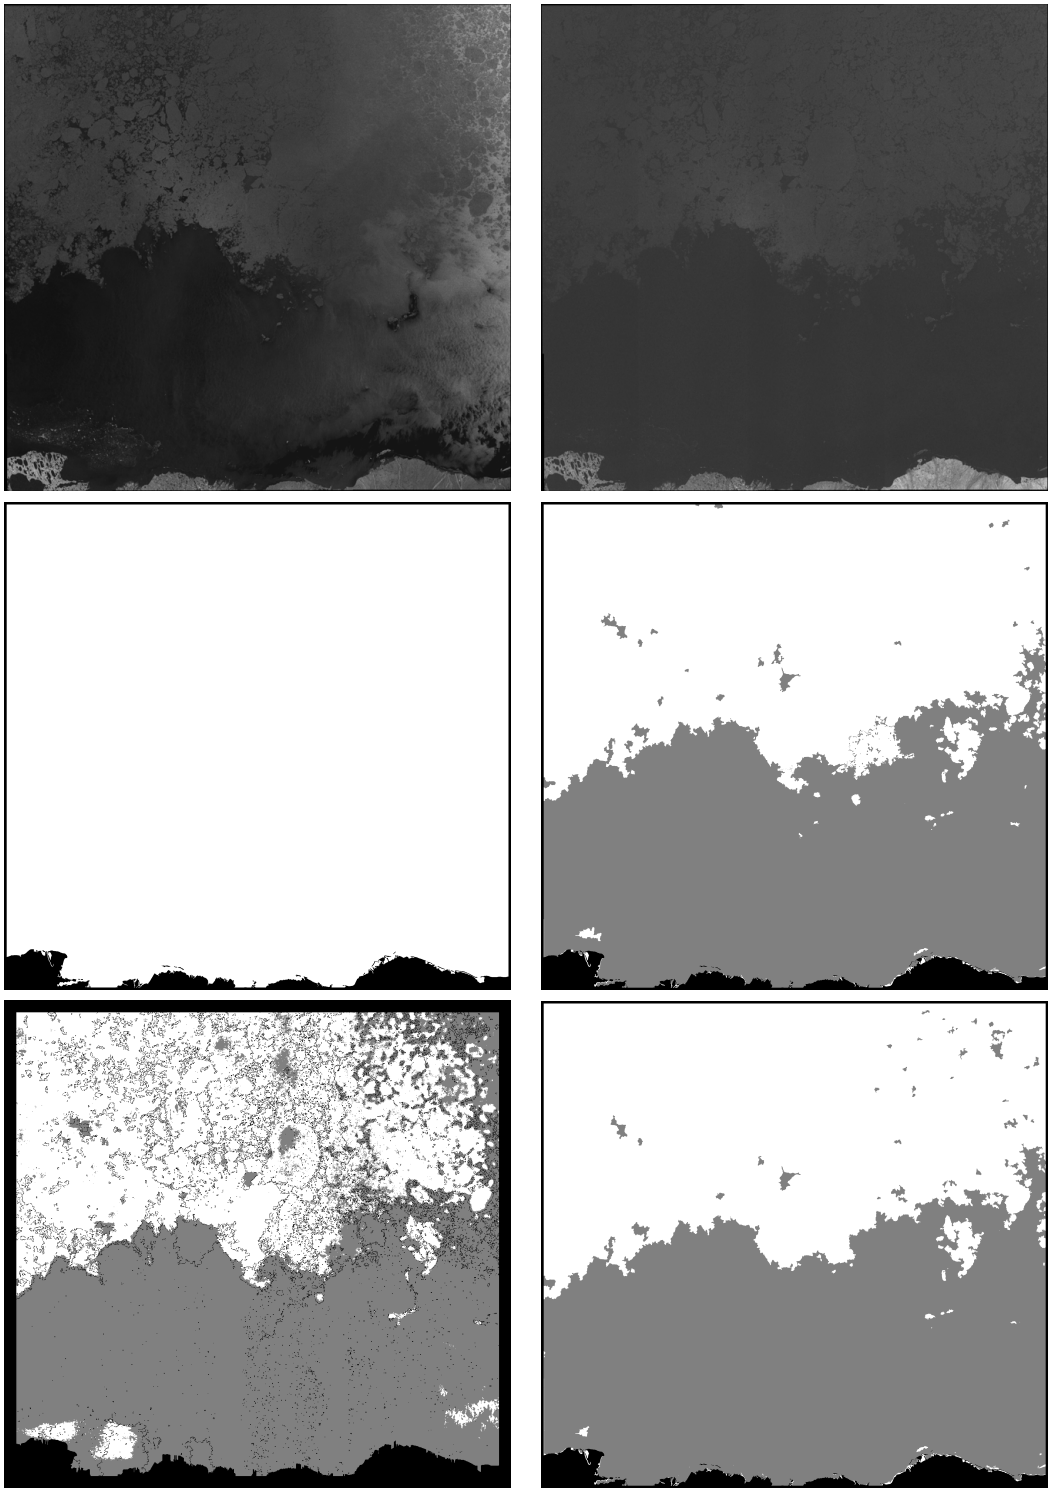

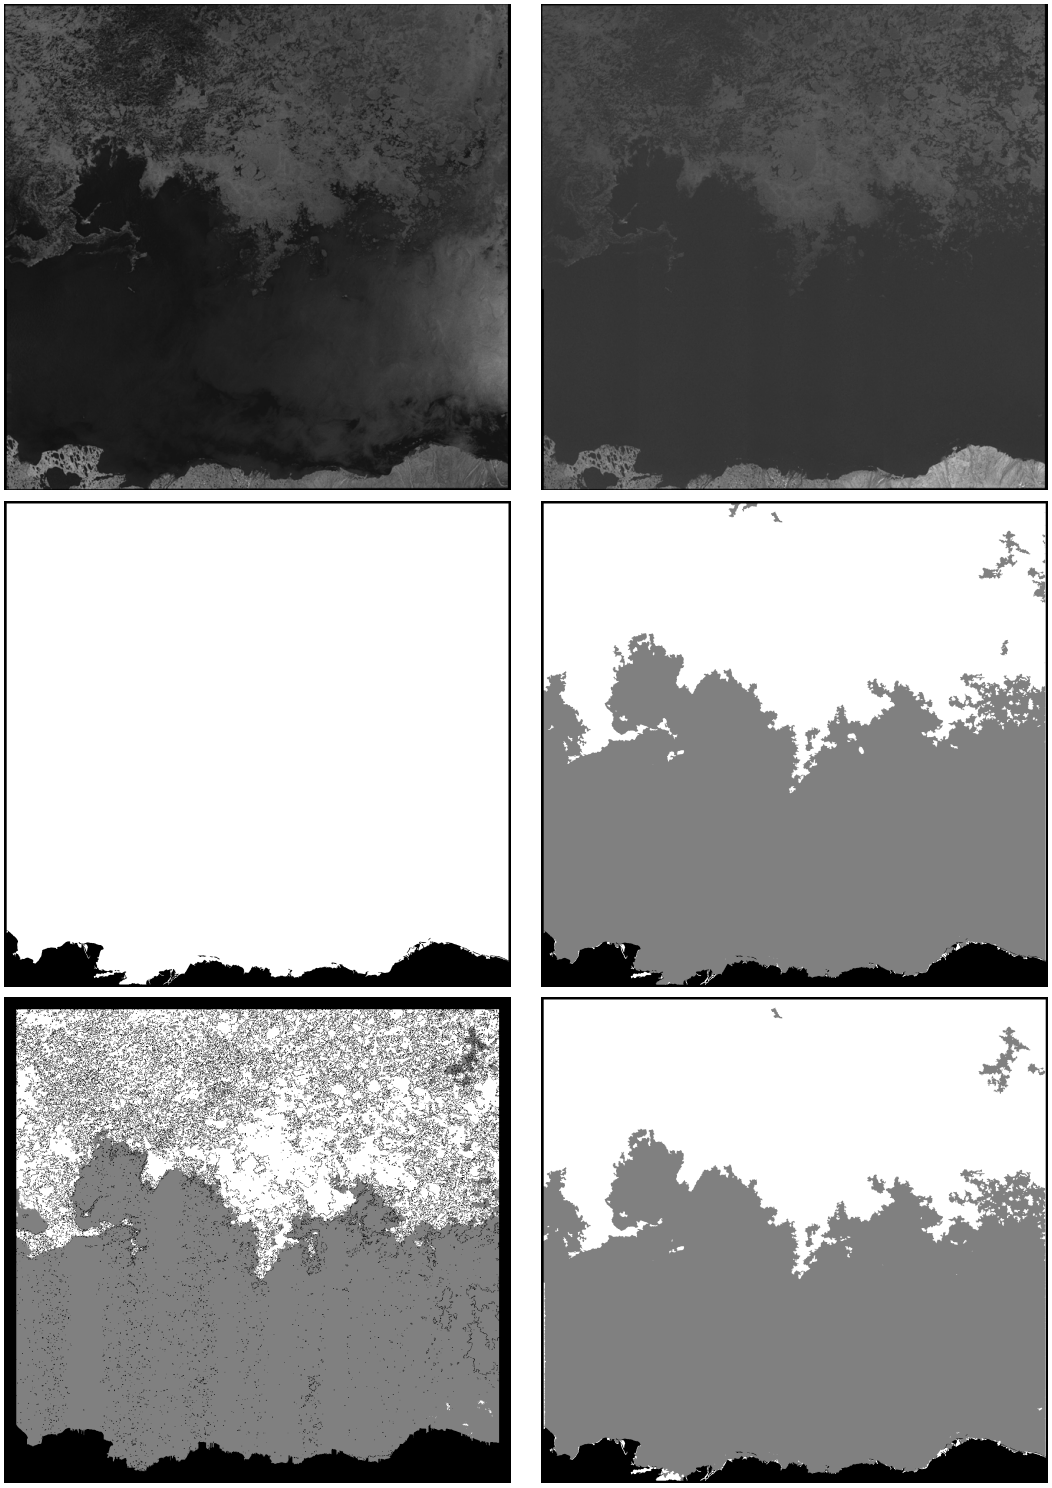

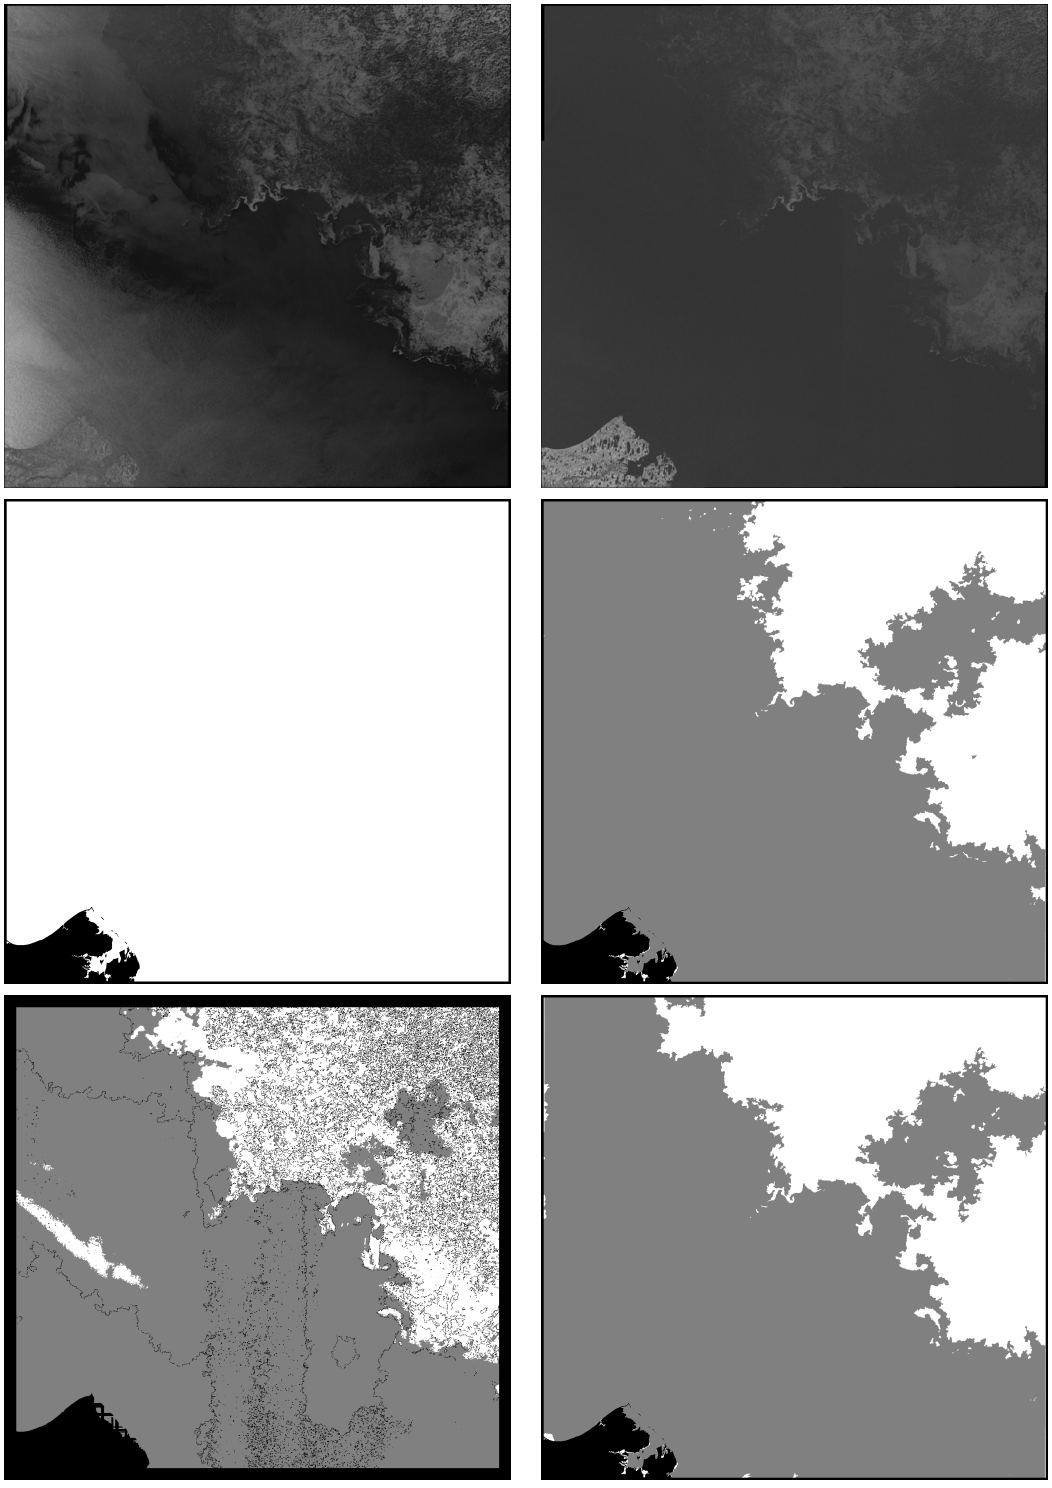

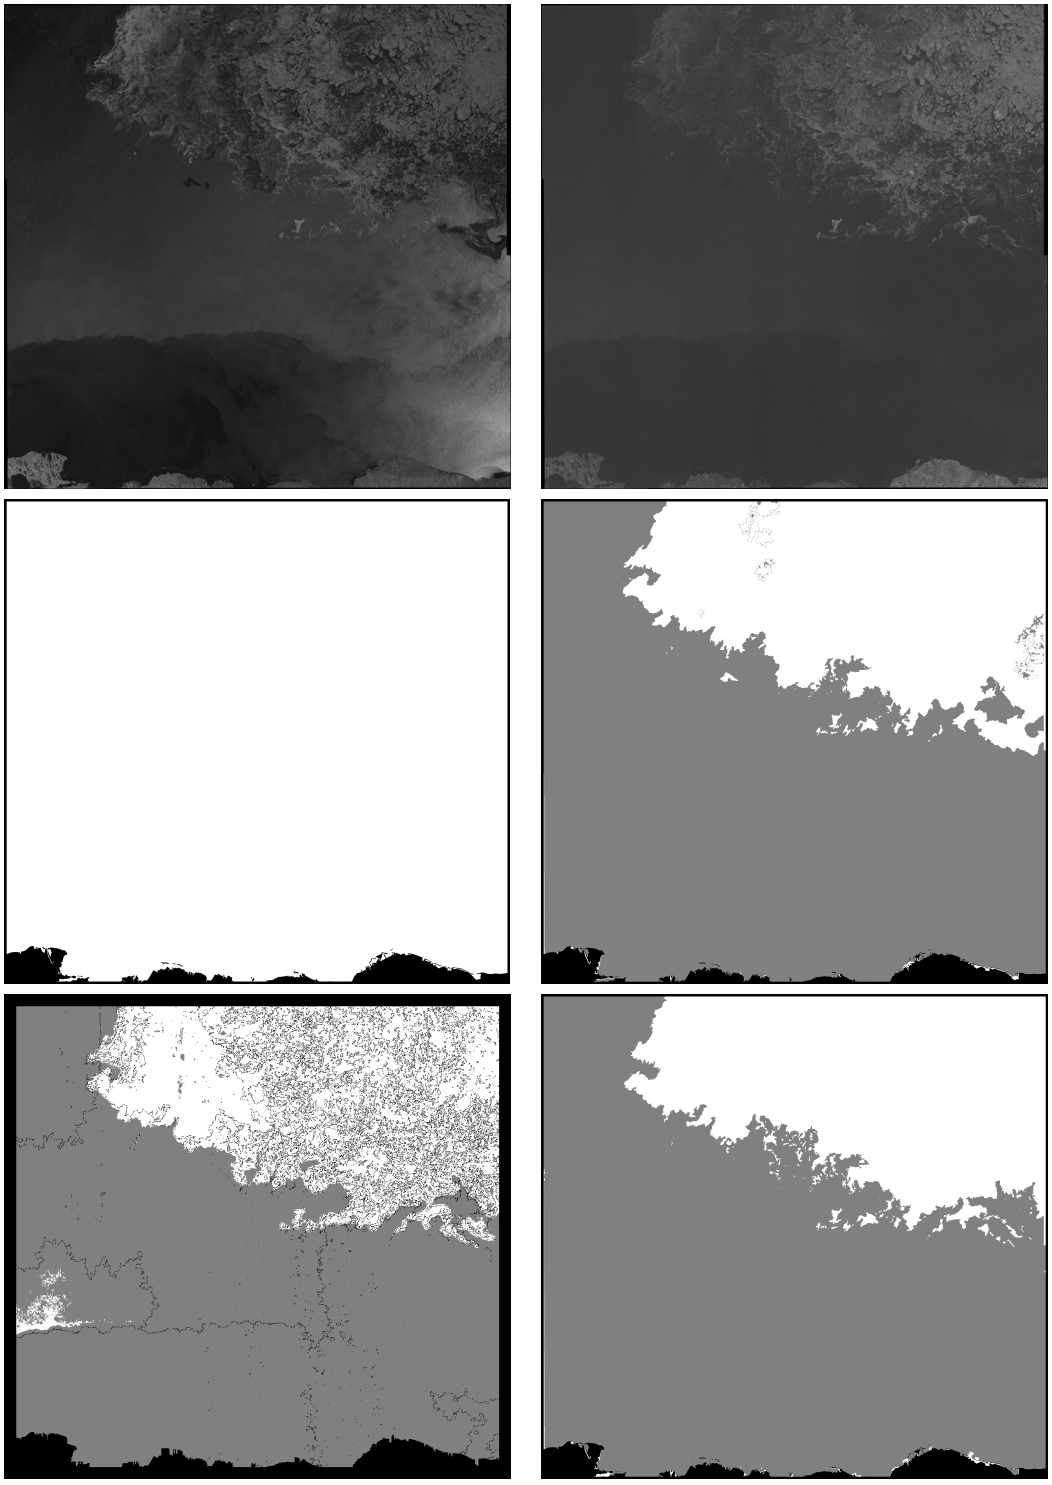

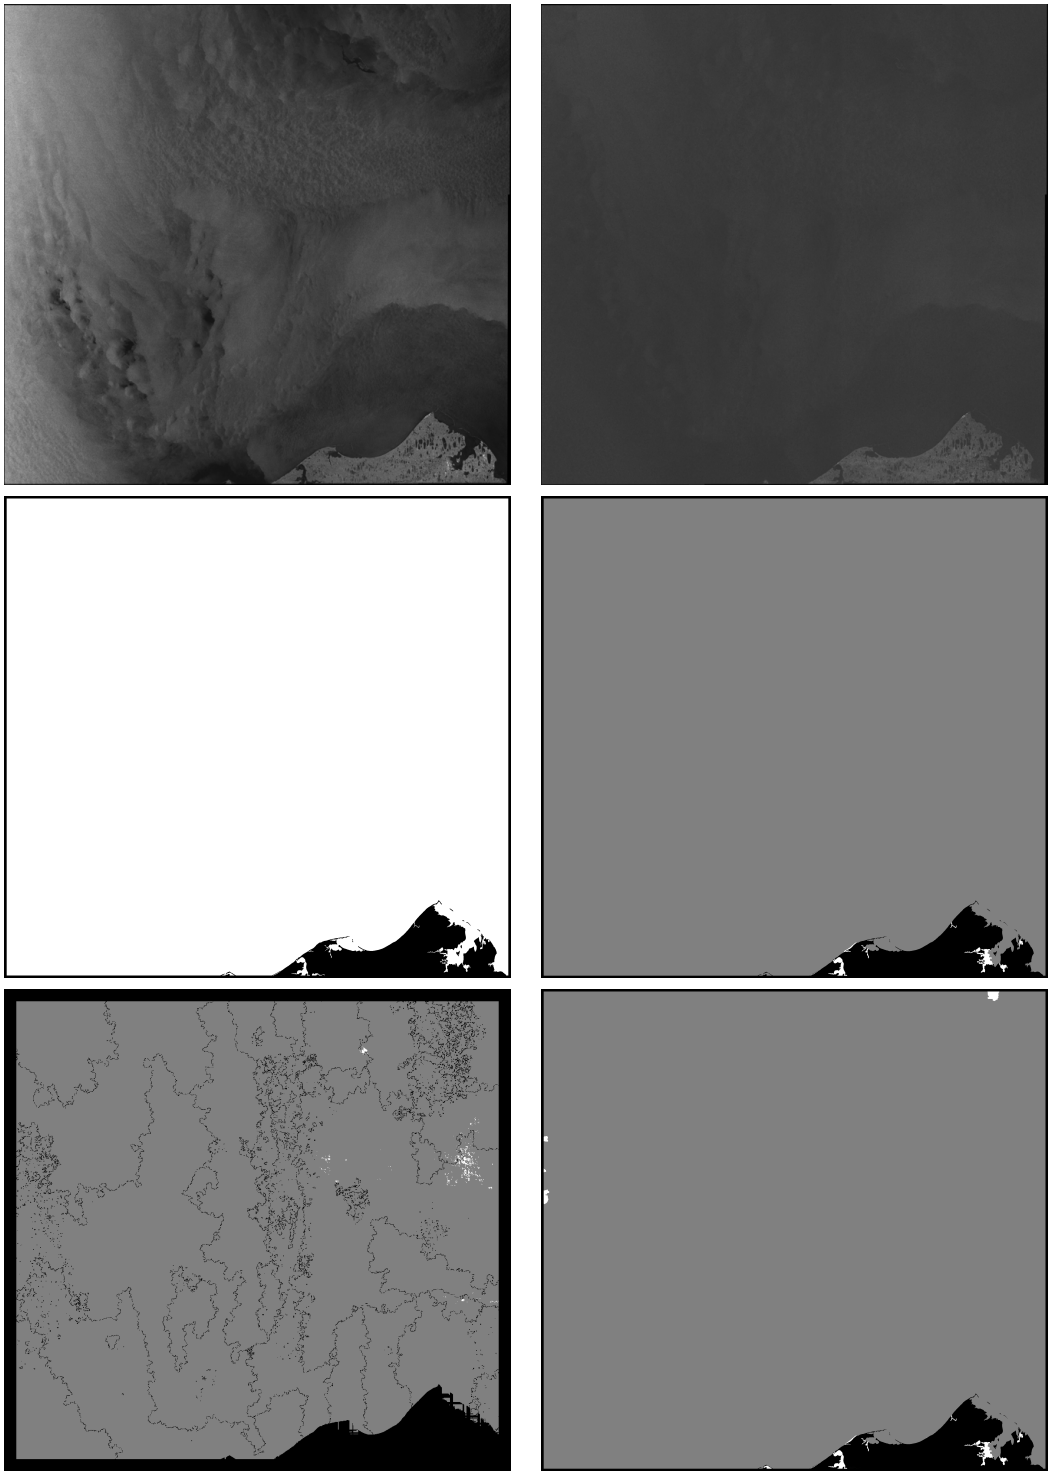

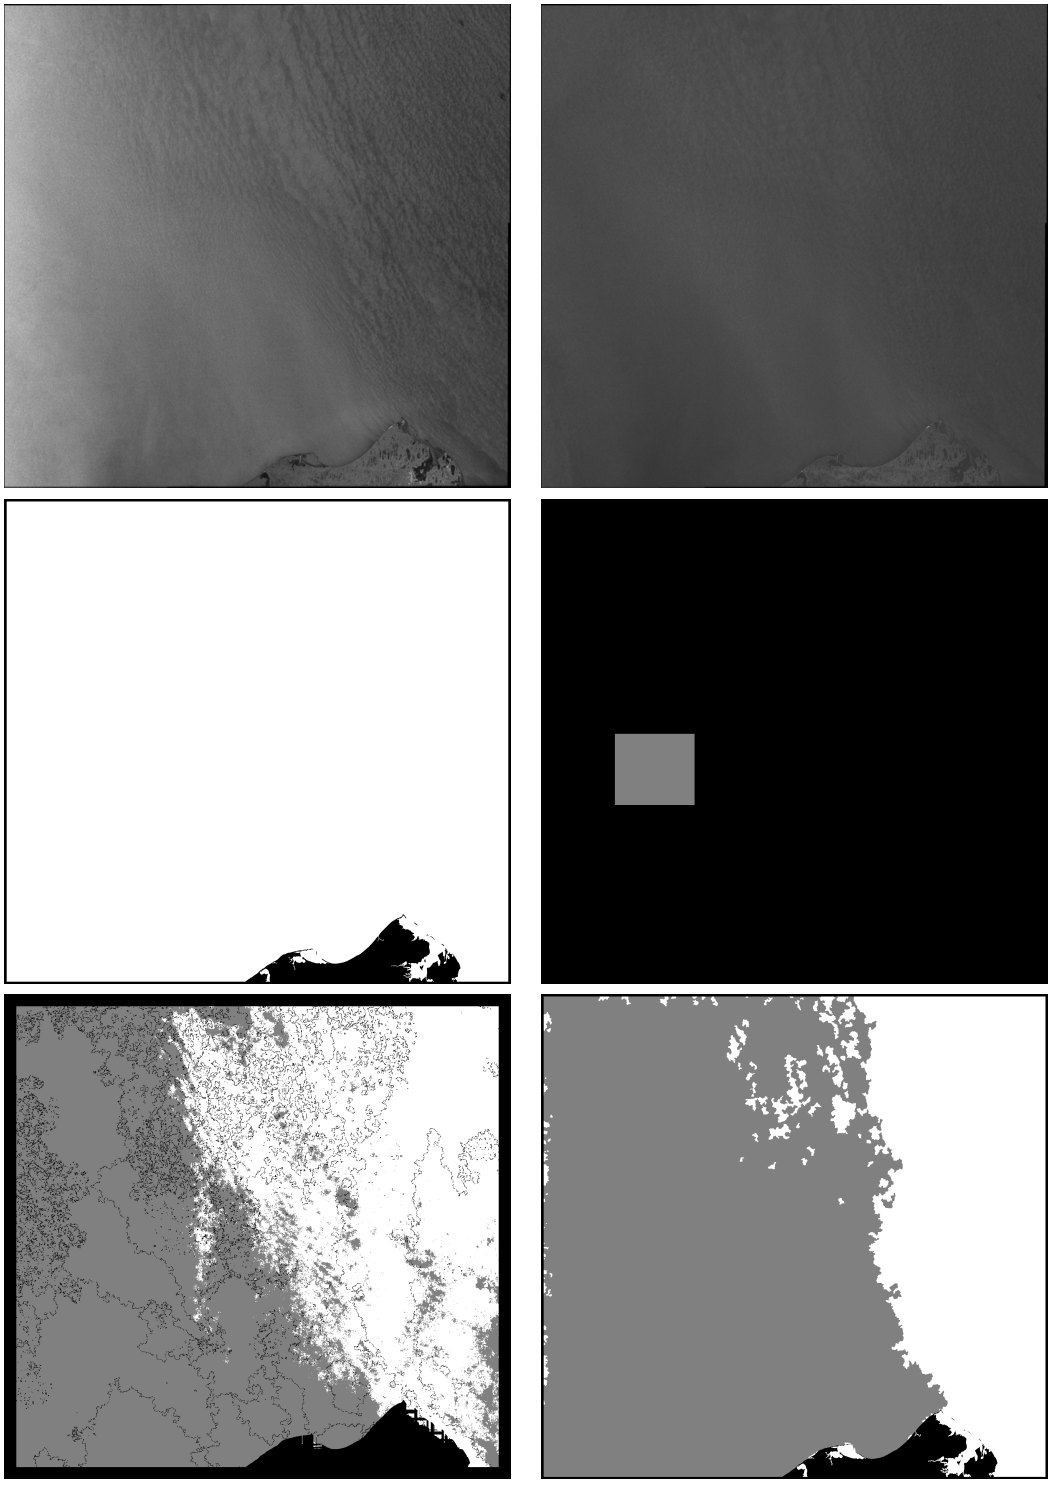

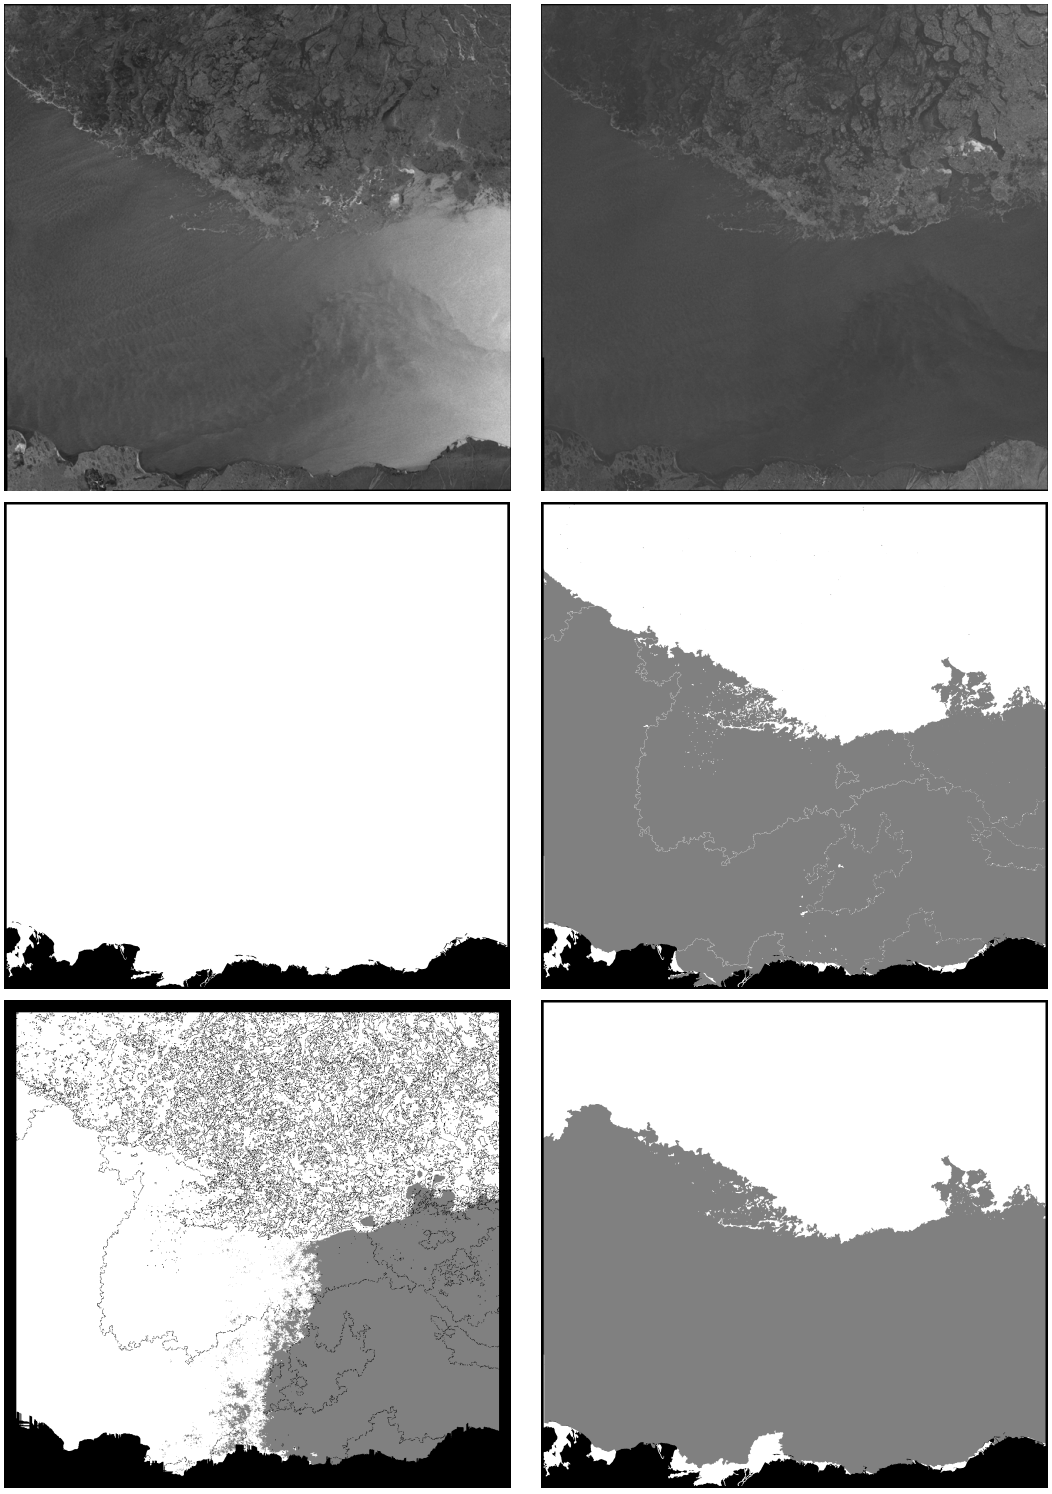

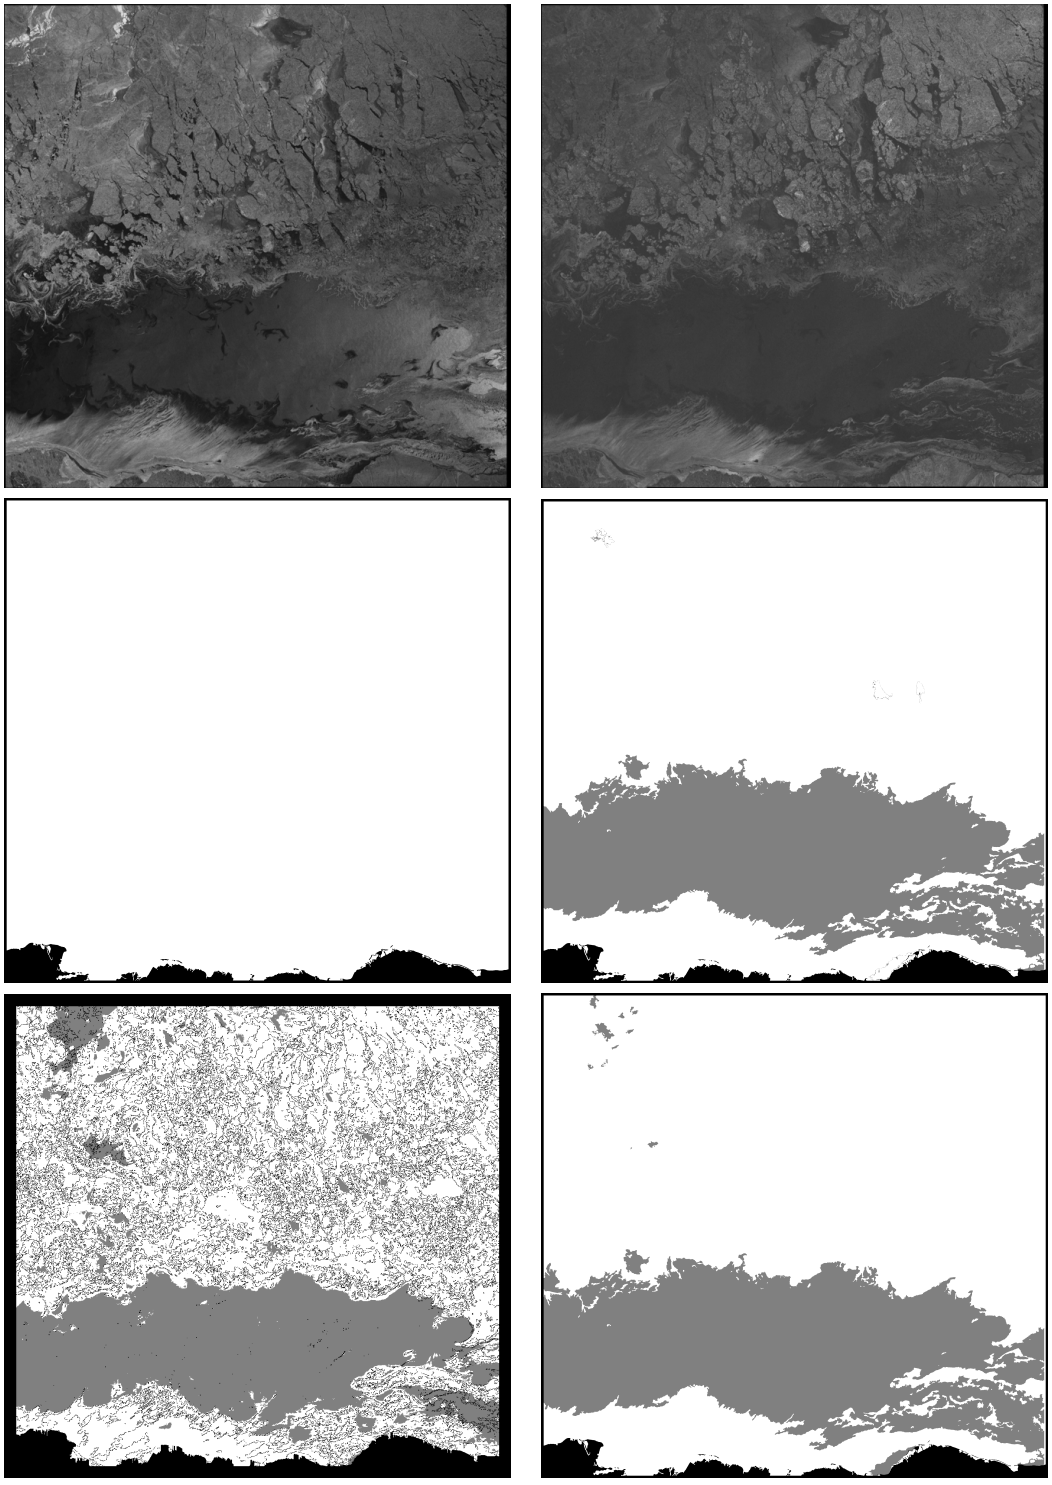

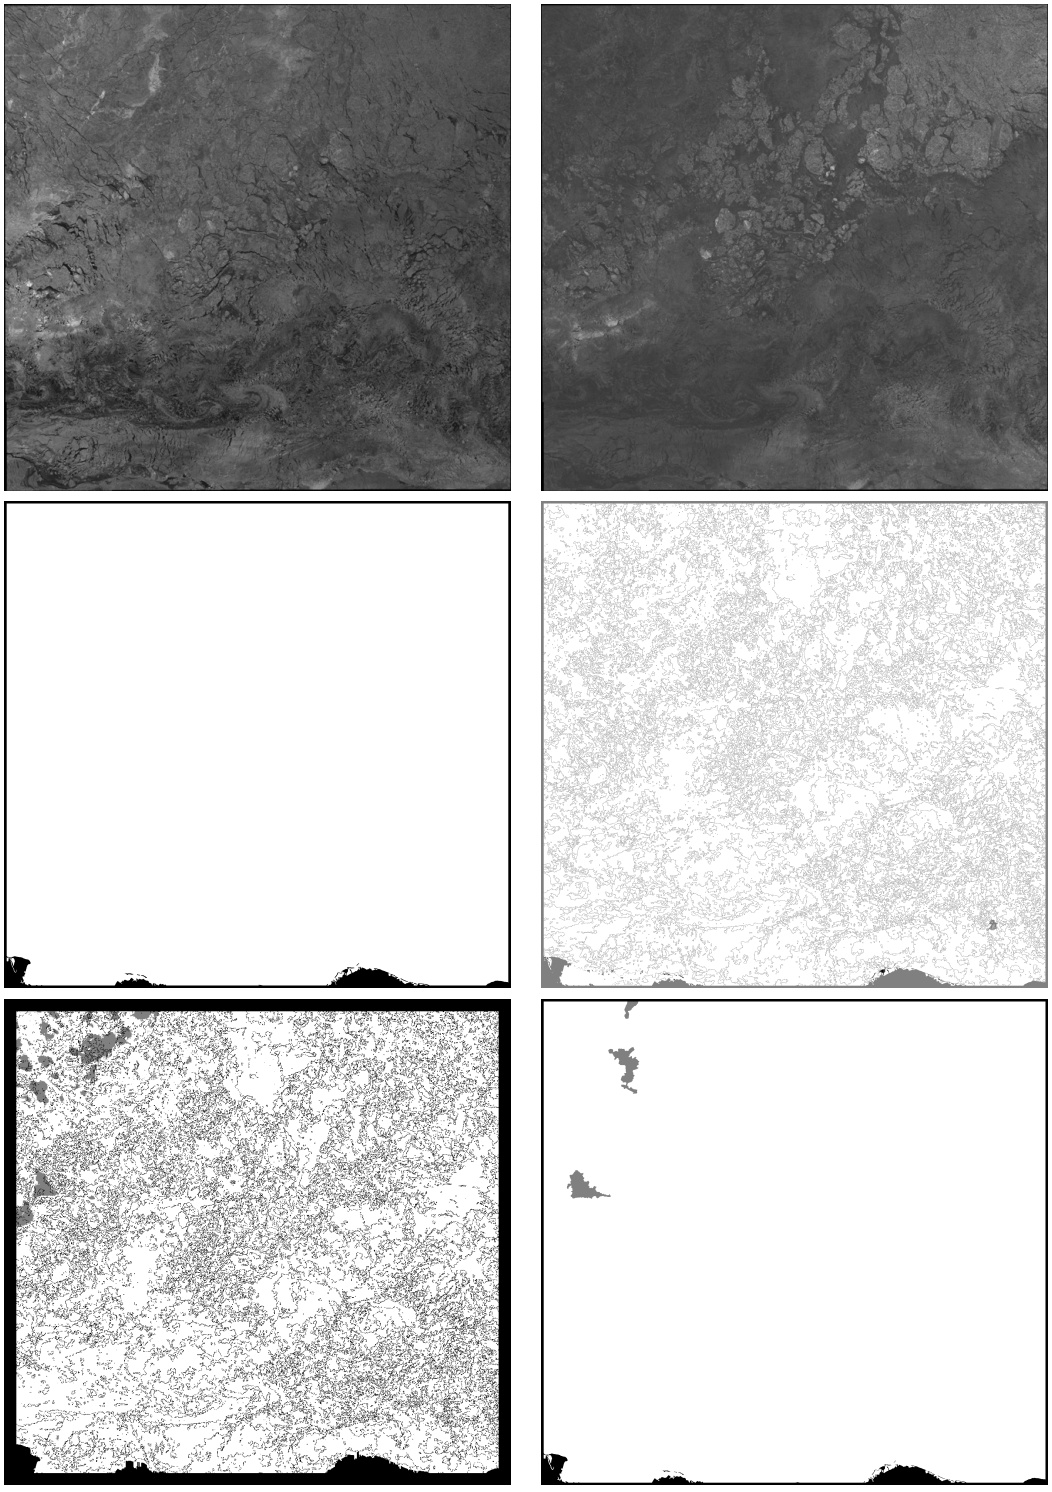

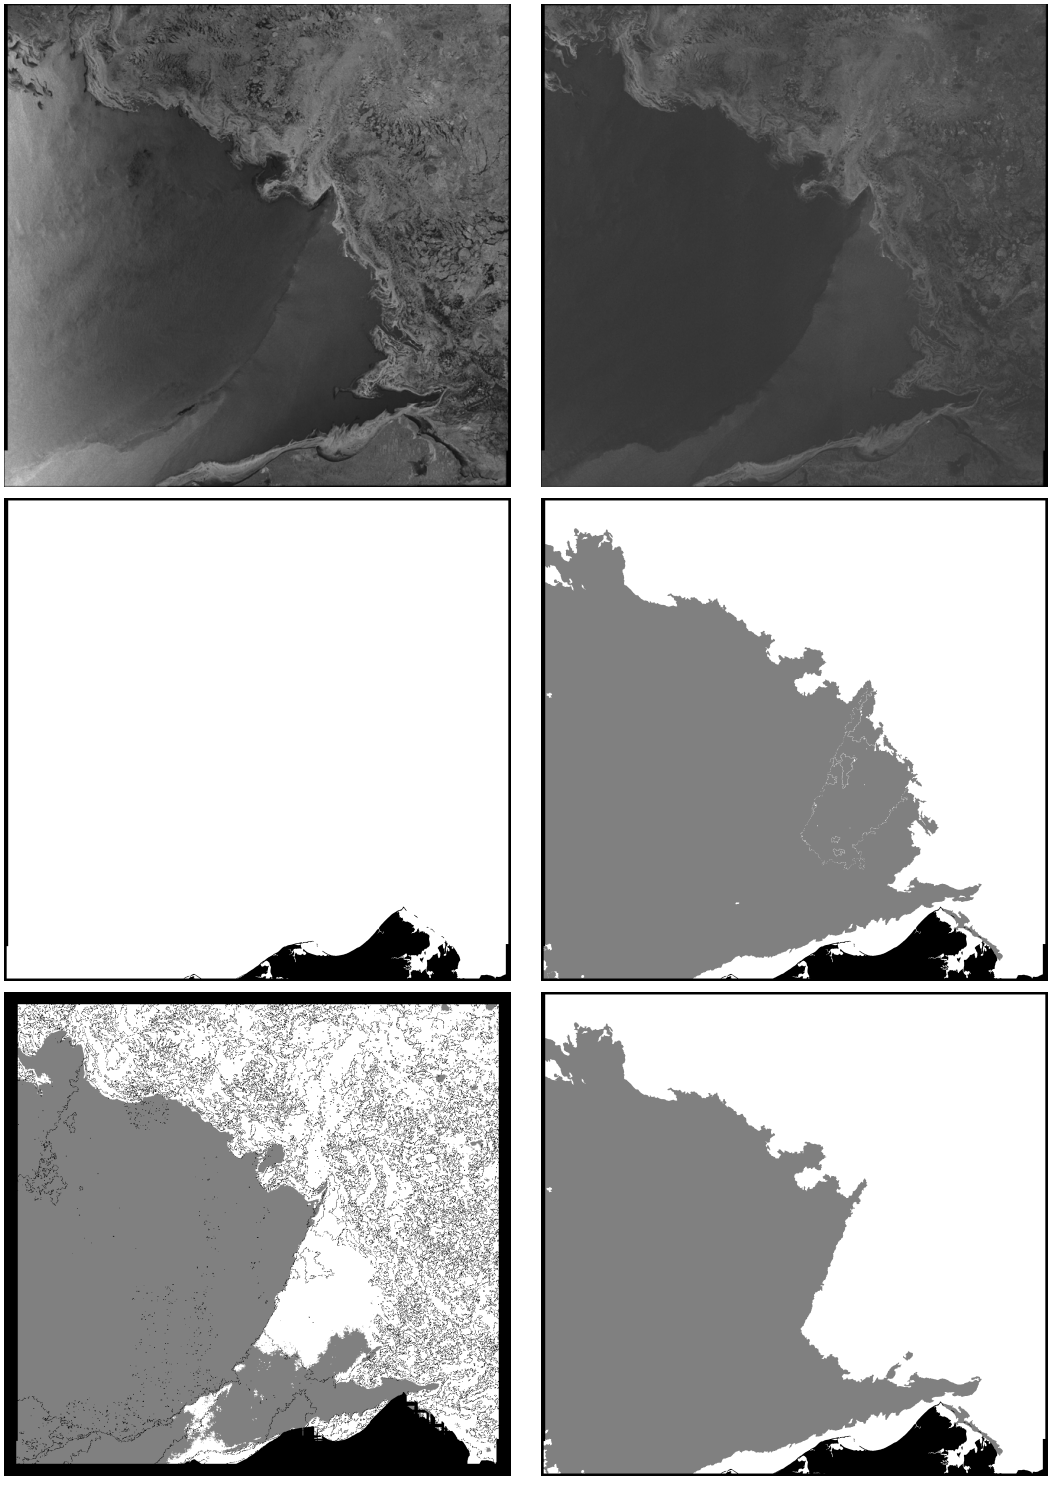

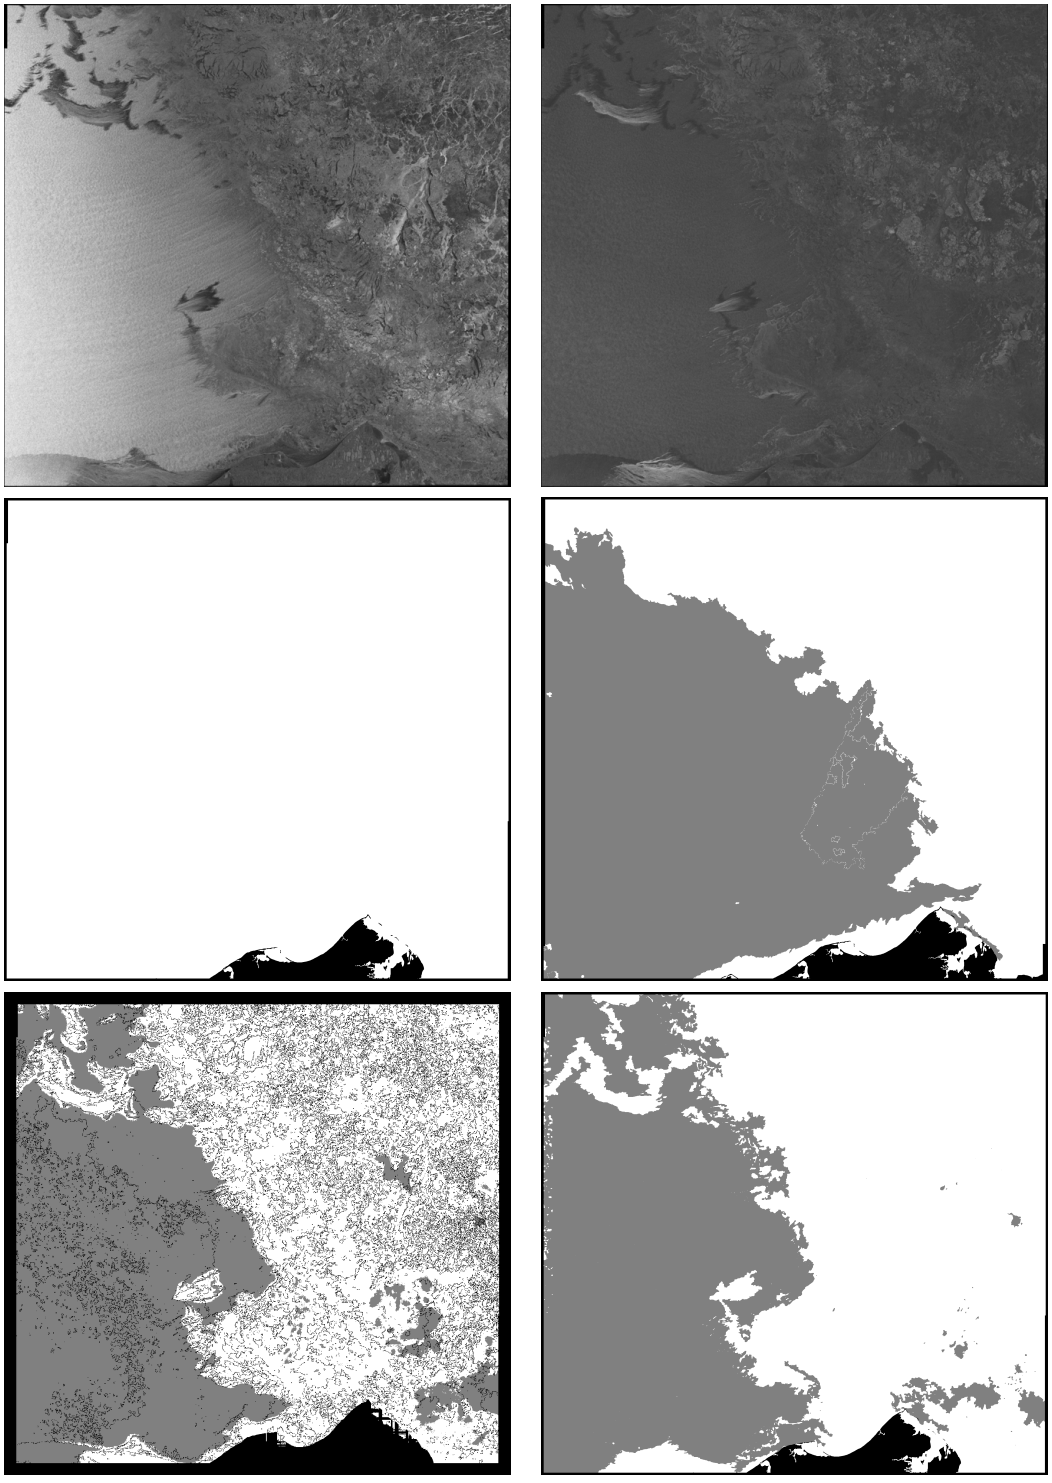

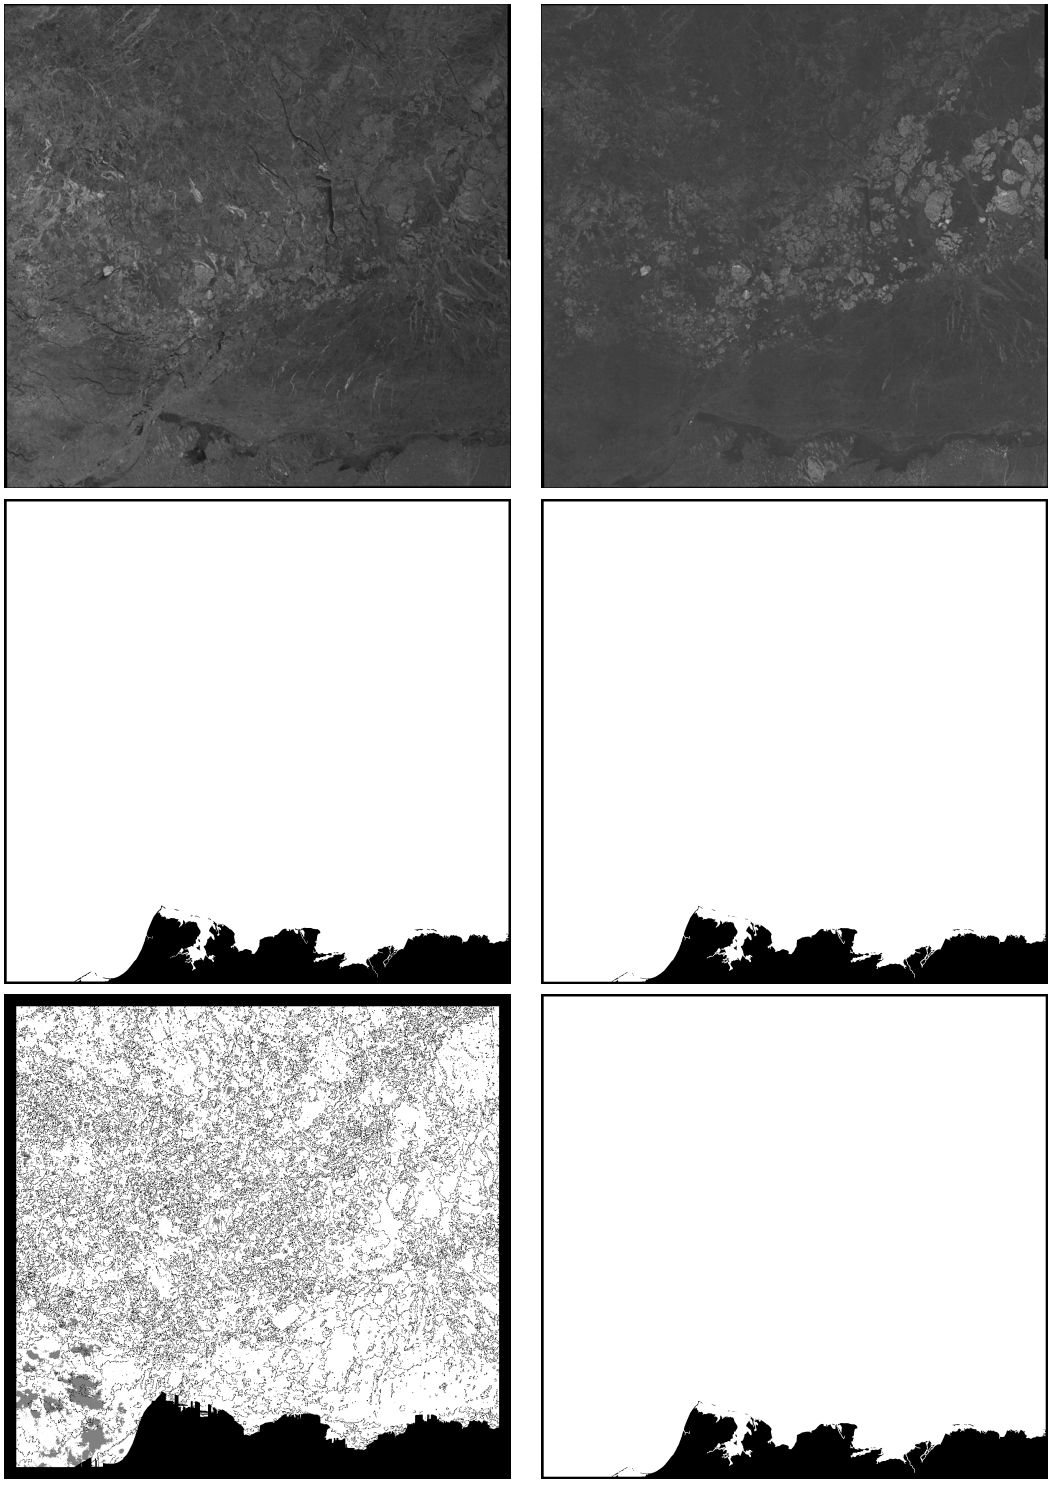

Appendix B

How to Generate Results in MAGIC

How to generate an SVM model

1. organize data set in LAN format
   (a) ground truths for each image are also needed
   (b) use GDAL for image format conversions if necessary `http://www.gdal.org/ogr/`
2. configure path to data set in SVM section of MAGIC codebase
3. configure SVM feature search code to operate on correct number of scenes
4. compile MAGIC in release mode (debug mode is too slow for this)
5. run the exe and select SVM feature search
6. LOO SVM models for each scene are output to the exe directory

How to generate a final result

1. copy svm_model.svm to image directory (may have to rename svm file)
2. run MAGIC exe
3. load dualpol LAN image
4. select HV pol
5. Segment/Classification->Glocal->Create Autopolygons->Watershed based
   (a) 12 by 12
6. wait to finish (should take ~10 sec)
7. Segment/Classification->Glocal->Set global labels

(a) add 4 labels

8. Segment/Classification->Segment Polygons and Save to Glue Later
   (a) use only HV pol
9. wait to finish (should take ~2 min)
10. click save
11. Segment/Classification->Glue Previous Segmentation
    (a) glue using IRGS
    (b) 6 classes
    (c) use only HV pol
12. wait to finish (should take ~10 sec)
13. click save
14. Labeling->Ice/Water Labeling->IRGS regionbased Labeling
15. wait to finish (should take ~25 min)
16. click save