CLIMATE-SPACE - THEME I - B. ADDITIONAL ESSENTIAL CLIMATE VARIABLES (ECVS) - NEW ECV PRODUCTS

SAGE CCI

(Sea Ice Age and Drift)

Algorithm Theoretical Basis Document (ATBD)

Prime & Science Lead: Signe Aaboe
Met Norway, Norway

Technical Officer: Sarah Connors
ESA ECSAT, United Kingdom

Consortium: Norwegian Meteorological Institute (MET Norway)
Science and Technology AS (S&T)
Nansen Environmental and Remote Sensing Center (NERSC)
University of Bremen (UB)
Université catholique de Louvain (UCLouvain)
University of Hamburg (UH)
University of Manitoba (UM)

Document Change Log

Issue	Author	Date	Change	Status
0.5	D. Fantin, S&T	16/09/2025	Document created	
1.0	UB: Hannah Niehaus, Gunnar Spreen. NERSC: Anton Korosov. MET Norway: Signe Aaboe, Emily Down, Thomas Lavergene. S&T: Jacob Hay.	15/04/2026	V 1.0 consolidated	Released to ESA

Document Approval

Prepared by	Hannah Niehaus, Lead Author, IUP Bremen	
Issued by	Daniele Fantin, Project Manager, S[&]T	Digitally signed by Daniele Fantin Date: 2026.05.11 15:41:07 CEST
Checked by	Signe Aaboe Science Leader, MET Norway	2026.05.11
Approved by	Sarah Connors ESA Technical Officer	Digitally signed by Sarah Connors Date: 2026.05.07 14:48:26 +01'00'

Table of Contents

Acronyms and Abbreviations

AARI	Arctic and Antarctic Research Institute (RUS)
ADP	Algorithm Development Plan
AMSR2	Advanced Microwave Scanning Radiometer 2
AMSR-E	Advanced Microwave Scanning Radiometer for EOS
AR	Annual Review
ASCAT	Advanced Scatterometer
ATBD	Algorithm Theoretical Basis Document
AWI	Alfred-Wegener-Institute for Marine and Polar Research
BGEP	Beaufort Gyre Exploration Project
BT	Backtracking-to-known-icetype
C3S	Copernicus Climate Change Service
CAR	Climate Assessment Report
CCI	ESA's Climate Change Initiative
CDR	Climate Data Records
CFOSAT	Chinese-French Oceanography Satellite
CIS	Canadian Ice Service
CMCC	Continuous Maximum Cross-Correlation
CMEMS	Copernicus Marine Service
CMIP	Coupled Model Intercomparison Project
CM SAF	The Climate Monitoring Satellite Application Facility
CMUG	Climate Modelling User Group
CNN	Convolutional Neural Network
CP	Communication Package
CRDP	Climate Research Data Package
CRG	Climate Research Group
DA	Data Assimilation
DAL	Distance Along the Line
DMI	Danish Meteorological Institute (DK)
DNN	Diffusion Neural Network
DOI	Digital Object Identifier
E3UB	End-to-End ECV Uncertainty Budget
ECV	GCOS Essential Climate Variable
ECCC	Environment and Climate Change Canada (CA)
ECMWF	European Center for Medium-Range Weather Forecasts
EO	Earth Observation
ERA5	ECMWF Reanalysis ver. 5
ERS	European Remote-Sensing Satellite
ES	Executive Summary
ESA	European Space Agency
EUMETSAT	European Organization for the Exploration of Meteorological Satellites
FCDR	Fundamental Climate Data Record
FM	Final Meeting
FP	Final Presentation
FR	Final Report
FYI	First-Year Ice
GCOS	WMO/ICO/UNEP Global Climate Observing System
HY-2	Haiyang-2

HR	Horizontal Range
IABP	International Arctic Buoy Programme
IAP	Ice Analysis Polygon
ICESat-2	Ice, Cloud and land Elevation Satellite 2
ICDR	Interim Climate Data Record
IPAB	International Programme for Antarctic Buoys
IPCC	Intergovernmental Panel on Climate Change
IPS	Ice Profiling sonar
ITT	Invitations to Tender
IUP	Institute of Environmental Physics
IWV	Integrated Water Vapour
JAMSTEC	Japan Agency for Marine-Earth Science and Technology (JP)
JAXA	Japan Aerospace Exploration Agency (JP)
KNMI	Royal Netherlands Meteorological Institute
KO	Kick Off
MEMLS	Microwave Emission Model of Layered Snowpacks
METNO or MET Norway	Norwegian Meteorological Institute
LWP	Liquid Water Path
MIZ	Marginal Ice Zone
ML	Machine Learning
MPR	Monthly Progress Report
MS	MileStone
MYI	Multiyear Ice
MYIF	Multiyear Ice Fraction
NASA	National Aeronautics and Space Administration
NERSC	Nansen Environmental and Remote Sensing Center
NIC	National Ice Center
NorESM	Norwegian Earth System Model
NSIDC	National Snow and Ice Data Center (US)
Obs4MIPS	Observations for Model Intercomparison Projects
OEM	Optimal Estimation Method
OSI SAF	The Ocean and Sea Ice Satellite Application Facility
OW	Open Water
PM	Progress Meeting, Project Manager
PMP	Project Management Plan
PMW	Passive Microwave
PSC	Polar Science Center
PDF	probability density function
PSD	Product Specification Document
PSH	Project Scientific Highlights
PSSM	Polynya Signature Simulation Method
PUG	Product User Guide
PVASR	Product Validation and Algorithm Selection Report
PVIR	Product Validation and Intercomparison Report
PVP	Product Validation Plan
QRS	Quarterly Status Reports
QuikSCAT	Quick Scatterometer Mission

RCM	Radarsat Constellation Mission
RID	Review Item Discrepancy
RMSD	Root Mean Square Difference
RMSE	Root Mean Square Error
RRDP	Round Robin Data Package
SAGE	Sea Ice Age and Drift
SAR	Synthetic Aperture Radar
SCAT	Scatterometer
SIC	Sea Ice Concentration
SID	Sea Ice Drift
SIMIP	Sea Ice Model Intercomparison Project
SMMR	Scanning Multichannel Microwave Radiometer
SND	Snow Depth
SRD	System Requirement Document
SoW	Statement of Work
SSD	System Specification Document
SSMI,SSM/I	Special Sensor Microwave - Imager
SSMIS	Special Sensor Microwave - Imager/Sounder
SST	Sea-surface Temperature
SYI	Second-Year Ice
S&T	Science and Technology AS
T2m	2 Meter Air Temperature
TB	Brightness Temperature
UB	University of Bremen
UCLouvain	Université Catholique de Louvain
UH	University of Hamburg
ULS	Upward-Looking Sonar
UM	University of Manitoba
URD	User Requirement Document
UWR	User Workshop Report
WAI	Warm Air Intrusion
WBS	Work Breakdown Structure
WMO	World Meteorological Organization
WP	Work Package
WPD	Work Package Description
WSP	Wind Speed
YI	Young Ice
YI	Year Ice, acronym used in figure 2.3

1 Introduction

1.1 Purpose and Scope

This document contains the Algorithm Theoretical Basis Document (ATBD) for the SAGE project for CLIMATE-SPACE - THEME I - B. ADDITIONAL ESSENTIAL CLIMATE VARIABLES (ECVS) - NEW ECV PRODUCTS, in accordance with the contract [AD1], SoW [AD2], and proposal [AD3-AD10].

This is the first issue of the ATBD document for the SAGE project, which is part of the ESA Climate Change Initiative. The document provides a detailed description of the algorithmic principles underlying the products of sea-ice type and age that were selected for further development following the PVASR [RD-2] and RRDP intercomparison exercise, and the supporting drift component. It forms the basis for subsequent algorithm refinement and product generation as part of the SAGE project.

The purpose of this document is to:

- Explain the underlying physics exploited for satellite-based sea-ice characterisation,
- Introduce primary and auxiliary data sets used for sea ice type and age retrievals,
- Outline how machine learning can be used for sea-ice type classification,
- Explain the key methodology of the algorithms that are planned to be used for the blended SAGE sea-ice age climate data record

1.2 Document Overview

This document is structured as follows:

- Chapter 1 introduces this document.
- Chapter 2 presents the physical background, including microwave interaction with sea ice and the relationship between ice drift and age.
- Chapter 3 describes the satellite and ancillary input data.
- Chapter 4 outlines radiometric and backscatter-based sea-ice type retrieval algorithms, including key implementations such as EC-ICE, ASIMIR, and the C3S Bayesian approach.
- Chapter 5 details the Lagrangian sea-ice age algorithm, including mesh generation and advection methods.
- Chapter 6 introduces the machine learning approach for sea-ice type classification.
- Chapter 7 describes the sea-ice drift algorithm, including retrieval methods, optimisations, and uncertainty estimation.
- Chapter 8 presents initial concepts for the blended SAGE sea-ice age product, outlining first ideas and approaches to be further developed in upcoming R&D activities.

1.3 Applicable Documents

No	Doc. Id	Doc. Title	Date	Issue/ Revision/ Version
AD-1	4000147560/25/I-LR	ESA Contract No.	12/03/2025	NA
AD-2	ESA-EOP-SC-AMT-2024-36	Statement of Work and Annexes and Appendexes	31/07/2024	1.2
AD-3	METNO-ESA-SAGE-CL-001	SAGE Cover Letter	8/11/2024	1.0
AD-4	METNO-ESA-SAGE-TPROP-001	SAGE Technical Proposal	8/11/2024	1.0
AD-5	METNO-ESA-SAGE-IPROP-001	SAGE Implementation Proposal	8/11/2024	1.0
AD-6	METNO-ESA-SAGE-MPROP-001	SAGE Management Proposal	8/11/2024	1.0
AD-7	METNO-ESA-SAGE-FPROP-001	SAGE Financial Proposal	8/11/2024	1.0
AD-8	METNO-ESA-SAGE-CPROP-001	SAGE Contractual Proposal	8/11/2024	1.0
AD-9	METNO-ESA-SAGE-BF-001	SAGE Background and Facilities	8/11/2024	1.0
AD-10	METNO-ESA-SAGE-CV-001	SAGE Curricula Vitae	8/11/2024	1.0
RD-1	METNO-ESA-SAGE-CCI-URD-001	SAGE User Requirements Document	6/11/2025	1.0
RD-2	METNO-ESA-SAGE-CCI-PVASR-001	SAGE Product Validation and Algorithm Selection Report	1/04/2026	1.0
RD-3	METNO-ESA-SAGE-CCI-E3UB-001	End to End ECV Uncertainty Budget	17/04/2026	1.0

Note: If not provided, the reference applies to the latest released Issue/Revision/Version

2 Physical background for satellite-based sea-ice type and age

Sea-ice type directly reflects the physical state of the ice, including its roughness, salinity, and mechanical strength. The classification into types, such as young ice (YI), first-year ice (FYI), second-year ice (SYI), and multiyear ice (MYI), is closely linked to the thermodynamic history. In particular, the number of summer melt seasons an ice parcel survives strongly influences its structure and properties. Consequently, sea-ice type can be interpreted as the "apparent age" of the ice; it describes how old the ice appears based on its observable characteristics.

Sea-ice age, derived from tracking ice parcels over time, represents the number of years the ice can be tracked from formation and through successive summer melt seasons, potentially plus the days since the last summer melt. It therefore provides a discrete measure of age, rather than type categories. Still, it differs from the "true age", or exact time since the ice formed, since the tracked sea-ice age is incremented by one year at every summer season (meaning that a 1-year increment is not necessary 365 days).

Also see discussion on sea-ice age in the User Requirements Document [URD, RD-1].

2.1 Microwave interaction with sea ice

Satellite-based retrieval of sea-ice type relies fundamentally on the interaction between microwave radiation and the physical properties of the sea-ice cover. In SAGE, both passive microwave radiometry (PMW) and active microwave scatterometry (SCAT) are exploited, as they provide complementary information on the thermodynamic and structural state of the ice, which again often is related to the age of the ice.

2.1.1 Dielectric properties

The physics of microwave remote sensing lies in the complex permittivity ($\varepsilon = \varepsilon' - j\varepsilon''$) of ice and snow. The real part (ε') determines the speed of wave propagation and the reflection coefficients, while the imaginary part (ε''), or dielectric loss, represents absorption. Sea ice is a heterogeneous mixture of pure ice, liquid brine and air bubbles. Brine pockets, which are typically vertical, needle-like inclusions in FYI, create dielectric anisotropy, whereby the response varies with polarisation. As the temperature rises, the volume of brine increases non-linearly, which has a drastic effect on both permittivity and loss. In MYI, summer desalination replaces conductive brine with air bubbles, which act as the dominant volume scatterers. Snow cover further complicates this system: its low thermal conductivity can raise the temperature of the underlying ice surface, increasing the ice's dielectric constant and affecting the overall microwave signature.

2.1.2 Emissivity

The emissivity (e) is defined as the ratio of the radiance emitted by a real surface to that of a black body at the same temperature, T. In the microwave regime, where the Rayleigh–Jeans approximation applies, this is expressed as $e = T_B/T$, where T_B is the brightness temperature and T is the actual physical temperature of the surface. For sea ice, emissivity is considered an effective or apparent quantity because the radiation originates from an emitting layer whose depth varies from millimetres to decimetres, depending on the frequency and the ice's microphysical properties, rather than just from the surface. Sea-ice emissivity is

highly dependent on ice type and frequency. Figure 2.1 shows the emissivity of sea-ice types and open ocean measured for different frequencies and for vertical and horizontal polarisation. FYI is characterised by high salinity, which results in a shallow penetration depth and high, relatively stable emissivity (green crosses in Figure 2.1). Similar is observed for the mixed FYI and MYI in late summer (cyan triangles in Figure 2.1). In contrast, MYI undergoes desalination as it ages, replacing brine pockets with air bubbles. These air bubbles act as volume scatterers, causing a significant decrease in emissivity at higher frequencies (red diamonds in Figure 2.1), such as 37 GHz and 89 GHz (Spreen, 2008; Lee 2018). Open ocean inversely shows an increase in emissivity with higher frequencies (blue crosses in Figure 2.1) and additionally shows large difference between the vertical and horizontal polarisation.

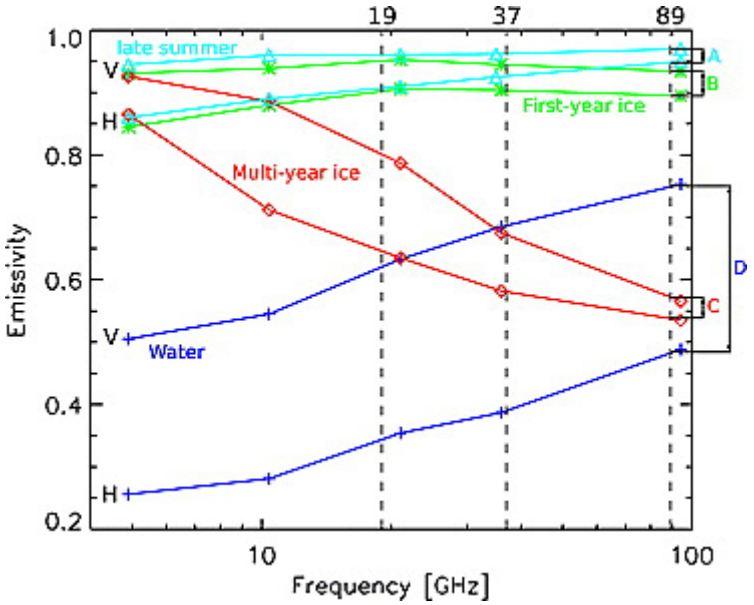

Figure 2.1: Emissivity measurements of sea ice and ocean water at different frequencies with vertical (V) and horizontal (H) polarisation and at an incident angle of 50 °. Figure from Spreen et al. 2008.

This surface-type-specific frequency dependency is commonly used in algorithms for retrieving ice type in the form of the spectral gradient ratio (GR) of the brightness temperature (TB) at two different frequencies (freq1, freq2) using the same polarisation:

$$GR(freq1,\ freq2,\ pol) = \frac{T_B(freq2, pol) - T_B(freq1, pol)}{T_B(freq2, pol) + T_B(freq1, pol)} \quad (2.1)$$

Emissivity also exhibits a strong difference in dependence on polarisation. With vertical polarisation, emissivity is generally higher than with horizontal polarisation, while the degree of difference varies depending on the type of ice. This can be explained by the Brewster effect. The incidence angles typically used by radiometers are close to the Brewster angle, at which point reflection of vertically polarised waves is minimised. This maximises emission according to Kirchhoff's law (e = 1 − R). Therefore, the polarisation ratio (PR), defined as the ratio of the vertically and horizontally brightness temperatures at the same frequency, is very good at distinguishing water versus ice, and is also a suitable parameter for ice age retrieval algorithms:

$$PR(freq) = \frac{T_B(freq,V) - T_B(freq,H)}{T_B(freq,V) + T_B(freq,H)} \quad (2.2)$$

The emissivity contrast between ice types forms the basis for most PMW-based classification algorithms. However, this contrast is strongly reduced during melt conditions, when the presence of liquid water at the surface masks intrinsic emissivity differences, see Figure 2.1. Therefore, most sea-ice type classification algorithms fail to produce a summer product unless they include other data to rely on during summer months, such as drift or model data.

2.1.3 Backscatter

Active microwave sensors measure the backscatter coefficient (σ_0), which represents the proportion of electromagnetic energy reflected directly back to the sensor. σ_0 is typically expressed in decibels (dB) and is an area-normalised average power return. However, in heterogeneous footprints, the measured signal is a weighted average of linear power, and even small areas of high-scattering features, such as pressure ridges, can dominate a pixel's total backscatter. The dielectric properties and microphysical structure of the sea ice and snow determine this interaction, with the resulting signature being an integration of surface scattering at interfaces and volume scattering within the media.

The dominant backscattering mechanism varies significantly by ice type and frequency. In the more saline FYI the dielectric loss is high and penetration depth shallow, causing the signal to be dominated by surface scattering at the air-ice or snow-ice interface, which typically results in lower backscatter. For the more desalinated MYI, the penetration depth increases relative to FYI because the dielectric loss is reduced, and with an increased amount of air bubbles acting as volume scatterers in MYI, the backscatter is enhanced. This enhancement is highly frequency-dependent, as the efficiency of volume scattering increases significantly when the wavelength of the radiation more closely matches the dimensions of these internal inhomogeneities. In multiyear ice, these air bubbles typically range from 1 to 4 mm in diameter. This makes them more effective at scattering shorter wavelengths (higher frequencies). While C-band (~5 GHz) radiation is similarly responsive to surface roughness across ice types, Ku-band (~13 GHz) is much more sensitive to the volume scattering within MYI, and therefore Ku-band scatterometry is facilitating better discrimination between deformed FYI and perennial ice classes. This behaviour is illustrated in Figure 2.2, which shows the seasonal evolution of Ku-band backscatter, with MYI consistently exhibiting higher backscatter than FYI and both ice types displaying a pronounced seasonal dependence.

The backscatter coefficient is highly sensitive to the incidence angle (θ), typically decreasing as this angle increases. This relationship is quantified as a slope, *m*. During winter, FYI typically exhibits a steeper slope than MYI because isotropic volume scattering in MYI decays more slowly with angle than surface scattering in FYI. Like emission, backscatter is influenced by the Brewster effect; vertically polarised waves experience higher transmission into the medium, potentially increasing the relative volume scattering contribution.

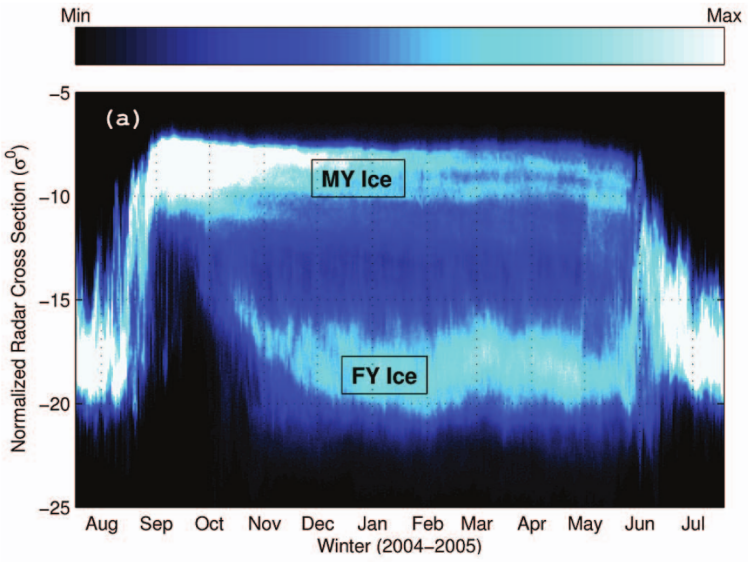

Figure 2.2: Time series of histograms of Ku-band backscatter from QuikSCAT over Arctic sea ice for the winter of 2004–2005. The backscatter differs between ice types and is seasonally dependent. Figure from Swan and Long, 2012.

Whereas MYI typically exhibits stronger backscatter than FYI due to its increased surface roughness and internal volume scattering, this relationship is not always valid. Certain ice conditions, such as deformed FYI or pancake ice (also a FYI type) in the marginal ice zone, can exhibit roughness characteristics similar to MYI (e.g. Remund et al. (2000)) and thus lead to misclassification where extra awareness or correction schemes are required.

2.2 Drift and age relationship

2.2.1 Conceptual link between ice motion and age classification

Sea-ice drift and sea-ice age are intrinsically linked because age is determined by the cumulative transport history of the ice cover. Sea-ice drift describes the horizontal motion of ice driven by winds, ocean currents, and internal ice stresses, while sea-ice age represents the time elapsed since the ice was formed or since it last fully melted. In a Lagrangian perspective, individual ice parcels are followed as they move through the Arctic or Antarctic domain, and their age increases continuously as long as they survive successive seasons. The spatial distribution of ice age therefore emerges from the pathways, residence times, and survival probabilities associated with the drift field rather than from local conditions alone.

This relationship means that realistic sea-ice age estimation requires accurate representation of both advection and ice loss processes. Drift transports newly formed ice away from source regions, recirculates older ice within long-lived reservoirs, and exports ice through gateways where it may melt or leave the basin. Convergence and divergence modify concentration and thickness distribution, indirectly influencing the likelihood of survival through summer melt. Consequently, regions containing old ice generally correspond to areas of persistent retention and repeated seasonal survival, whereas younger ice is associated with recent

formation, rapid export, or frequent melt. Sea-ice age can therefore be interpreted as an integrated diagnostic variable that combines information on ice motion, thermodynamic forcing, and long-term memory of the sea-ice system.

2.3 Sea ice age nomenclatures

The classification and naming of sea-ice age and type are not standardised across scientific and operational communities, leading to inconsistencies in terminology and interpretation. Within SAGE, resolving these differences has been a key effort, as consistent class definitions are essential for the PVASR intercomparison and the construction of the RRDP, see [RD-1] and [RD-2]. This issue is illustrated in Figure 2.3, which compares two established nomenclatures (green and blue) and the SAGE-derived nomenclature (red).

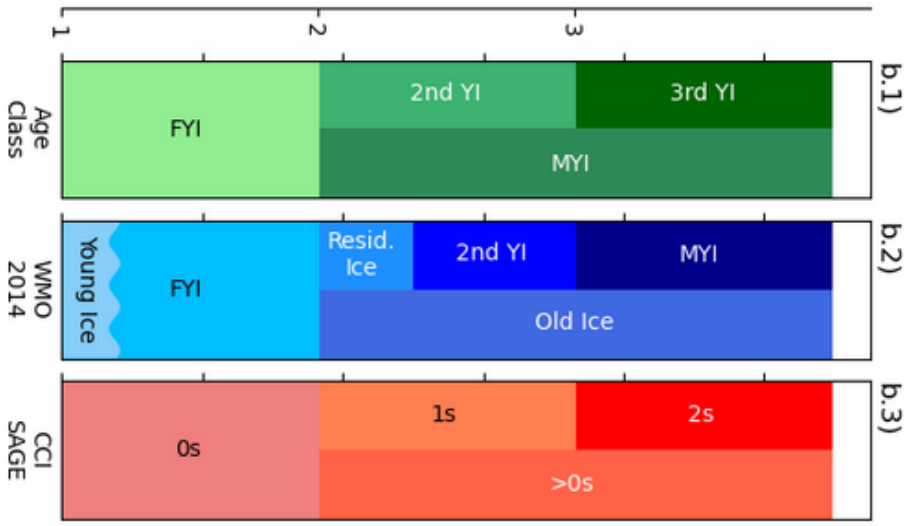

Figure 2.3: Illustration of different sea-ice nomenclatures: b.1) Age Classes, b.2) WMO Sea Ice Nomenclature (2014), b.3) the nomenclature developed in CCI SAGE.

Figure 2.3 panels b.1 - b.3) illustrate the current issue with sea-ice **nomenclature**. The three panels illustrate how sea-ice age would be named in three different conventions. On b.1), in greens, the "Age Class" nomenclature, including FYI, second-year ice (2YI), 3YI, and higher-order age classes, and with MYI defined as all ice of 2YI and older. This is in wide use in the climate / remote sensing communities. On b.2), in blues, the nomenclature of mariners (WMO Sea Ice Nomenclature, 2014) that introduces several additional labels, e.g. Young Ice and Residual (First-Year) Ice, where the latter is characterised as older ice that has survived at least one summer melt. Note that according to the WMO nomenclature, MYI is ice that has survived at least two summers' melt. Because these first two nomenclatures share the same labels (e.g. "2nd YI" and "MYI") but with different definitions, the community has a nomenclature issue.

The third nomenclature on b.3), in reds, is the new nomenclature we defined in the CCI SAGE PVASR exercise. This last nomenclature has no ambiguities (since it uses new labels). This nomenclature would solve some of

the ambiguities with the other two nomenclatures, but it would take time and efforts to have it widely accepted. At this stage, it remains to be seen if CCI SAGE should push for introducing a new nomenclature, or rather acknowledge the issue with existing nomenclature and explicitly adopt one (and only one) of them. In this report, we have so far used the “Age Class” nomenclature when referring to classes of FYI and MYI.

3 Input data

3.1 Satellite data

3.1.1 ASCAT backscatter

The Advanced Scatterometer (ASCAT) is an active microwave sensor carried on the European Meteorological Operational (MetOp) series of polar-orbiting satellites, including MetOp-A, MetOp-B, and MetOp-C. Operational since 2007, the sensor measures the normalised radar backscattering cross section (σ0) at C-band (5.3 GHz) utilising vertical (VV) polarisation. The instrument's measurement geometry uses two 550 km-wide swaths on both sides of the satellite track, sampling backscatter from three different azimuth directions to characterise the surface.

3.1.2 AMSR brightness temperatures

The Advanced Microwave Scanning Radiometers from the Japan Aerospace Exploration Agency (JAXA) measure dual-polarised - vertical (V) and horizontal (H) - brightness temperatures at various frequencies. Table 3.1 presents an overview of the different AMSR sensors and their specifications. They are all operated on satellites in near-polar, sun-synchronous orbits at an altitude of about 660 to 700 km. The data is divided into ascending and descending swaths with a combined 29 to 30 files per day.

Table 3.1: Properties of the AMSR instrument series.

Instrument	AMSR-E	AMSR2	AMSR3
Satellite	Aqua	GCOM-W1	GOSAT-GW
Central Frequencies (GHz)	6.9, 10.7, 18.7, 23.8, 36.5, 89.0	6.9, 7.3, 10.65, 18.7, 23.8, 36.5, 89.0	6.9, 7.3, 10.25, 10.65, 18.7, 23.8, 36.5, 89.0, 165.5, 183.3
Polarisation	H & V	H & V	H & V up to 89 GHz V only for 165.5, 183.3 GHz
Footprint size	43x75 km (6GHz) 29x51 km (10.65 GHz) 16x27 km (18.7 GHz) 14x21 km (23.8 GHz) 9x14 km (36.5 GHz) 4x6 km (89 GHz)	35x62 km (6.9 and 7.3 GHz) 24x42 km (10.65 and 18.7 GHz) 11x19 km (23.8 GHz) 7x12 km (36.5 GHz) 3x5 km (89 GHz)	34x58 km (6.9 and 7.3 GHz) 22x39 km (10.25 and 10.65 GHz) 12x21 km (18.7 GHz) 14x24 km (23.8 GHz) 7x11 km (36.5 GHz) 3x5 km (89 GHz and higher)
Lifetime	May 2002 - Oct 2011	May 2012 - present	June 2025 - present

3.1.3 SMMR, SSM/I and SSMIS data

SMMR was a prototype of multi-channel conical-scanning MW imagers flown on the last satellite in NASA's Nimbus satellite program. SMMR-NI07 has frequency channels from 6.6 to 37 GHz in both V and H polarisations, but no near-90 GHz imagery.

The SSM/I and SSMIS instruments are conically-scanning multi-frequency microwave radiometers on board the DMSP platforms of the US Department of Defence (DoD), with imaging frequencies ranging from 19.3 GHz to near-90 GHz (85.5 GHz for SSM/I and 91.1 GHz for SSMIS). The two near-90 GHz channels (at V and H polarisations) have the highest spatial resolution (13 x 14 km -3dB diameter of the instantaneous field of view). The SSM/I sensors are installed on the F08, F10, F11, F13, F14 and F15 missions, and SSMIS sensors on the F16, F17 and F18 missions. The first SSM/I mission on F08 had an early failure of the near-90 GHz channels (Fennig et al., (2020)). The V polarisation channel noise strongly increased towards the end of 1987 until the instrument was switched off. The H polarisation channel noise steadily increased from deployment until early 1990, when it failed for the first time, though it recovered for a short period at the end of 1990 before failing again.

Table 3.2: Properties of the SMMR-NI07, SSM/I and SSMIS instruments 1978 - present.

Instrument	SMMR-NI07	SSMI-F08, SSMI-F10, SSMI-F11, SSMI-F13, SSMI-F14, SSMI-F15	SSMIS-F16, SSMIS-F17, SSMIS-F18
Satellite	Nimbus-7	DMSP-F08, DMSP-F10, DMSP-F11, DMSP-F13, DMSP-F14, DMSP-F15	DMSP-F16, DMSP-F17, DMSP-F18
Central frequencies (GHz)	6.6, 10.7, 18.0, 21.0, 37.0	19.35, 22.235, 37.0, and 85.0(compromised on SSMI-F08)	19.35, 22.235, 37.0, various from 52.8 - 63.3, 91.655, 150, 183.31
Polarisations	V & H	V & H (except 22.235 GHz V only)	V & H (22.235 GHz V only, near-60 GHz either H or V or RC, >100 GHz H only)
Thermal resolution	0.9 - 1.5 K	0.37 – 0.73 K	0.21 – 1.23 K
Footprint size	95x160 km (6.6 GHz) 60x100 km (10.7 GHz) 35x60 km (18.0 GHz) 30x50 km (21.0 GHz) 17x29 km (37.0 GHz)	45x68 km (19.35 GHz) 40x60 km (22.235 GHz) 24x36 km (37.0 GHz) 11x16 km (85.0 GHz)	42.4x70.1 km (19.35 and 22.235 GHz) 27.5x44.2 km (37.0 GHz) 17.5x25.8 km (50 - 64 GHz) 13.1x14.4 km (>90 GHz)

Pixel size	40x26 km (6.6 GHz) 20x26 km (10.7 GHz) 20x26 km (18.0 GHz) 20x26 km (21.0 GHz) 10x26 km (37.0 GHz)	25 x 12.5 km (up to 37 GHz) 12.5 x 12.5 km (85.0 GHz)	25 x 12.5 km (up to 37 GHz) 37.5 x12.5 km (50 - 60 GHz) 75.0 x12.5 km (60 - 64 GHz) 12.5 x 12.5 km (>90 GHz)
Global coverage	Every 2 days	Every day	Every day
Time coverage	1978-10-25 – 1987-08-20	1987-07-09 – 2008-12-31	2005-11-20 – present (unsupported)

3.2 Sea ice drift products

Sea-ice drift data are used across several components of the SAGE system. In the Lagrangian sea-ice age algorithm, sea-ice drift constitutes the primary input, as it drives the advection of ice parcels and thereby determines the evolution of sea-ice age. In other components, drift data are used as ancillary information, for example, in training, evaluation, and post-processing.

Two main long-term sea-ice drift climate data records (CDRs) are currently available, provided by the EUMETSAT Ocean and Sea Ice Satellite Application Facility (OSI SAF) and the National Snow and Ice Data Center (NSIDC). Within the selected type and age algorithms in SAGE, both CDRs are used. Additionally, the OSI SAF near-real-time (NRT) drift product is used to extend to the present time. Note that an ICDR (OSI-435) has recently been released by OSI SAF to supplement the existing CDR (OSI-455). This will replace the NRT product in future sea-ice type and age retrievals.

These products are based on observations from PMW sensors such as SSM/I (>= F10) and SSMIS (CMSAF FCDR), AMSR-E (NSIDC) and AMSR2 (JAXA). The NRT product also utilises C-band scatterometer data from the ASCAT missions. Both products employ the same sea ice motion-tracking methodology, continuous maximum cross-correlation (CMCC), which involves a fractional-pixel pattern matching of the brightness temperatures. These brightness temperatures are pre-processed with Laplacian filters, and the vectors are then post-processed with correction schemes as described in Section 7 below. The climate products also use the ERA5 wind reanalysis together with a free drift model to gap-fill in the summer season where the early satellite data is not sufficient to overcome the challenges of the surface melt and high atmospheric water content.

Table 3.3: Overview of sea-ice drift products and their availability.

Product ID	Record	Description	Source
OSI-455	CDR	Global low-resolution sea-ice drift CDR based on PMW data from SSM/I, SSMIS, AMSR-E and AMSR2, plus ERA5 wind reanalysis for gap-filling of the summer season. Product covers: 1991-2020.	EUMETSAT OSI SAF doi: 10.15770/EUM_SAF_OSI_0012
OSI-435	ICDR	Global low-resolution sea-ice drift CDR based on PMW data from SSMIS and AMSR2, plus ERA5 wind reanalysis for gap-filling of the summer season. Product covers: 2021-present. (Soon to be released)	EUMETSAT OSI SAF doi: TBC
OSI-405 (OSI-405-b, OSI-405-c, OSI-405-d)	NRT	Daily provided sea-ice drift vectors based on passive and active microwave observations. Available since 2009. Different versions, depending on time, are indicated by appending a “b”, “c” or “d”.	EUMETSAT OSI SAF doi: 10.15770/EUM_SAF_OSI_NRT_2007
NSIDC-0116	CDR	Polar Pathfinder Daily 25 km EASE-Grid Sea Ice Motion Vectors v4. Data from Northern and Southern Hemisphere, 1978-2022	NASA NSIDC doi:10.5067/INAWUWO7QH7B (Tschudi et al. (2019))

3.3 Sea ice concentration

As for sea-ice drift, sea-ice concentration (SIC) data is used across several components in the SAGE system: as primary input for the Lagrangian sea-ice age algorithm, and as ancillary data for the sea-ice type algorithms.

The primary source of sea-ice concentration is CDRs provided by OSI SAF. These are all derived from PMW data but are produced as separate sensor-based CDRs: The coarse long-term CDR/ICDR are based on SMMR, SSM/I and SSMIS, while the medium-resolution CDR/ICDR are based on AMSR-E and AMSR2.

All OSI SAF concentration products are on a 25 km EASE2 grid and they are retrieved using the SICCI3LF algorithm based on the ~19 GHz and ~37 GHz imagery. The products also contain information about the land/water mask in the status flag which was obtained from the status flag of the OSI SAF SIC product which is used in the Lagrangian sea ice age product.

Table 3.4: Overview of sea ice concentration products and their availability.

Product ID	Record	Description	Source
OSI-450-a1	CDR	Global Sea Ice Concentration CDR v3.1, based on SMMR/SSMI/SSMIS. 1978-2020.	EUMETSAT OSI SAF doi: http://dx.doi.org/10.15770/EUM_SAF_OSI_0023
OSI-430-a	ICDR	Global Sea Ice Concentration ICDR v3, based on SSMIS. 2021-2025.	EUMETSAT OSI SAF doi: http://dx.doi.org/10.15770/EUM_SAF_OSI_0014
OSI-438	ICDR	Global Sea Ice Concentration ICDR v3, based on AMSR2. 2021-present. (Soon to be released)	EUMETSAT OSI SAF doi:10.15770/EUM_SAF_OSI_0024

3.4 Ancillary data

3.4.1 ERA5 Reanalysis data

ERA5 is the fifth-generation reanalysis from the European Centre for Medium-Range Weather Forecast (ECMWF), providing a globally complete and consistent dataset of climate and weather variables from 1940 to the present (Hersbach, 2020). This product utilises data assimilation to combine physical model data with worldwide observations, resulting in hourly estimates for a vast array of atmospheric, ocean-wave, and land-surface variables. The data is provided on a global 0.25-degree regular latitude-longitude grid and includes an underlying 10-member ensemble at three-hourly intervals to provide sampled uncertainty estimates.

Table 3.5: Auxiliary variables used from ERA5 reanalysis.

Auxiliary product type	Record	Description	Source
Wind speed, T2m, TCLW, TCWV, SST	CDR	ERA5: Global 1-hourly atmosphere reanalysis fields of 10m wind speed, 2m air temperature, total column liquid water (same as liquid water path LWP), total column water vapour, and sea-surface temperature	C3S CDS (doi:10.24381/cds.bd0915c6) (produced by ECMWF, see Hersbach et al. (2020))
Wind speed, T2m, TCLW, TCWV, SST	ICDR	ERA5T: the same data stream as ERA5 but ERAT5 is made available earlier (generally 5 days latency) however without the quality control going into ERA5.	C3S CDS doi:10.24381/cds.bd0915c6 (produced by ECMWF, see Hersbach et al. (2020))

3.4.2 Merged Warren-AMSR2 snow depth climatology

The merged Warren-AMSR2 climatology is a monthly snow depth product that integrates the long-term Warren climatology derived from historical Soviet drifting station measurements (Hendricks and Paul (2022), Warren et al. (1999)) with modern daily snow depth estimates over first-year sea ice from AMSR2 satellite observations (Rostosky et al. (2018)).

3.5 Sea ice climatology

A climatological maximum sea-ice extent is used within the processing to remove grossly erroneous sea-ice detections outside regions where sea ice is expected to occur for a given time of year. In the C3S sea-ice type this climatology is inherited from the OSI SAF sea-ice concentration CDR, OSI-450-a (see Section 3.3). Its basis is the monthly varying maximum sea-ice extent climatology from the NSIDC sea-ice concentration CDR v3, covering the period 1987–2007. The OSI SAF version includes further modifications, mainly manual editing of selected regions using information from operational ice charts, and an expansion with a buffer zone of 150 km in the Northern Hemisphere and 250 km in the Southern Hemisphere. The larger expansion in SH is to cope with the more dynamic behaviour of the sea ice and the slightly positive trends in the SH sea-ice extent (Parkinson, 2019).

4 Radiometric and backscatter-based sea ice type algorithms

4.1 Overview

The selection of candidate radiometric and backscatter-based algorithms is based on the round-robin intercomparison study [RD-2], where methods were evaluated consistently against the RRDP. The results show that established passive microwave and combined approaches outperform both simple channel combinations and purely scatterometer-based methods. Among all evaluated algorithms, C3S, ECICE, and ASIMIR form the top-performing group in terms of FYI/MYI discrimination accuracy.

In the following, each core algorithm is described in a dedicated section, while common processing elements, such as pre-processing and post-processing, are described separately due to their related implementation across the algorithms.

4.2 Pre-processing

Pre-processing steps are commonly applied in sea-ice type retrieval algorithms prior to the core classification in order to minimise atmospheric contributions and land-spillover effects, and to ensure consistency across sensors. These steps are typically applied to Level-1 and Level-2 brightness temperature data.

Land spill-over correction is applied in coastal regions to account for mixed land-sea footprints. In the C3S pre-processing this is performed by a linear combination of the land and sea surface contributions. The observed brightness temperature T_B can be expressed as:

$$T_B = (1 - \alpha)T_{sea} + \alpha T_{land} \quad (4.1)$$

Where α represents the fractional land contribution within the sensor footprint. The corrected sea-only brightness temperature, T_{sea}, is used in subsequent processing.

For the EC-ICE and ASIMIR algorithms, a land mask is used to exclude all pixels that are partially covered by land from the processed data.

Atmospheric correction is performed using a radiative transfer model, RTM, driven by auxiliary data from numerical weather prediction or reanalysis, including total column water vapour, cloud liquid water, near-surface air temperature, and wind speed. The RTM used in C3S follows the model function from Wentz (1983) for SMMR and Wentz (1997) for SSM/I and SSMIS, and Wentz and Meissner (2000) for AMSR-E and AMSR2.

The latter is also used in the ASIMIR algorithm, but there it is implemented in the core model because atmospheric variables are also retrieved.

For EC-ICE, no atmospheric correction is applied.

All channels are then collocated onto a common grid, ensuring spatial consistency between frequencies and polarisations.

4.3 EC-ICE

The EC-ICE algorithm is a mathematical optimisation algorithm designed to retrieve total and partial sea ice concentrations from a combination of active and passive microwave remote sensing data. In principle, several channels of radiometer (brightness temperature) and scatterometer (normalised radar backscatter cross-section) measurements can be used to distinguish as many surface types as there are input channels.

4.3.1 Input data

The implemented version of the retrieval (in line with the implementation setup at the Institute of Environmental Physics (IUP), University of Bremen, https://data.seaice.uni-bremen.de/MultiYearIce/MYIuserguide.pdf) uses AMSR2 brightness temperatures and ASCAT scatterometer data introduced in chapter 3. The AMSR2 swath data is combined and interpolated into daily maps on a polar stereographic grid with nominal grid spacing of 12.5 km. The ASCAT data are resampled to the same grid.

4.3.2 Algorithm scheme

The EC-ICE algorithm distinguishes between four surface types: Open Water (OW), Young Ice (YI), First-Year Ice (FYI) and Multiyear Ice (MYI). These ice types can be separated due to their internal structure and surface composition. The radiometric distinction between MYI and FYI is driven by the lower salinity and a more porous structure of MYI due to brine drainage, leading to increased radar backscatter coefficients and lower microwave emissivity compared to FYI. Microwave signatures are primarily influenced by the ice surface composition rather than bulk ice properties. This is particularly significant given that snow cover on sea ice undergoes continuous and rapid metamorphosis due to steep temperature gradients and changing atmospheric conditions causing microwave signatures to overlap and thus leading to radiometrically diverse surfaces even within a single ice category. To account for this natural variability, EC-ICE utilises empirical probability density distributions of the radiometric values derived from manually sampling homogeneous areas of known ice types instead of using single values as tie points.

To resolve the four surface types, the algorithm utilises an observation vector I_{obs} composed of four independent variables:

- σ^0: backscatter cross-section at 5.3 GHz VV, incidence angle-adjusted at 40 °
- **GR(37V,19V):** vertical polarisation Gradient Ratio of 36.5 GHz and 18.7 GHz (from AMSR-E or AMSR2), defined following Equation 2.1
- **TB(37V)** and **TB(37H)**: brightness temperatures at 36.5 GHz, in vertical and horizontal polarisation

The core of EC-ICE is a linear mixing model which assumes that a single satellite observation, integrated over the sensor's footprint, can be decomposed into a linear mixture of components produced by the various ice types within that footprint. Following this linear mixture approach, the expected observation vector I_{obs} of a mixed footprint is

$$I_{obs} = c_{ow} I_{ow} + c_{yi} I_{yi} + c_{fyi} I_{fyi} + c_{myi} I_{myi} \quad (4.2)$$

where c_i is the concentration of surface type i and I_i is the representative signature vector of the surface type i containing tie points for σ^0, GR37V19V, TB37V and TB37H. To ensure physical consistency, the concentrations c_i are restricted to values between 0 and 1, and the sum of all concentrations must equal 1. Given these constraints, the algorithm determines the concentration vector $C = [c_{ow}, c_{yi}, c_{fyi}, c_{myi}]$ by minimising a normalised quadratic cost function that is the sum of squared differences between observed and estimated radiometric values across n (=4) variables. As EC-ICE uses typical distributions (Section 4.3.3) of the radiometric quantities which can be interpreted as probability density functions instead of fixed tie points to account for the natural variability, a Monte Carlo approach is used. With this approach, a set of m=1000 possible combinations of typical brightness temperatures, gradient ratio and backscatter is randomly drawn from the input distributions. The algorithm executes m independent trials producing a distribution of possible concentration vectors for a single observation. The final concentration for each surface type is determined as the median of the resulting distribution, which is then normalised to ensure the total concentration is exactly 100%.

4.3.3 Selection of PDFs

The only difference in applying the EC-ICE algorithm to the Arctic and Antarctic is in the selection of representative brightness temperature and backscatter signal distributions for the retrieved surface types. Despite using the same classes, their signatures strongly differ because Antarctic sea ice is different from Arctic sea ice (Melsheimer et al. (2023)).

Figure 4.1 shows the distributions used for the Arctic. The FYI and YI distributions were primarily compiled using independent "truth" data from the Canadian Ice Service (CIS) Radarsat-1 image analysis, which provides delineated Ice Analysis Polygons (IAPs) characterized by a uniform distribution of ice types and concentrations. To ensure integrity, only satellite footprints that were located entirely within these homogeneous polygons were sampled. The MYI distributions were derived from stable winter conditions in the central Arctic Basin in order to capture the distinct low-emissivity and high-backscatter signatures that are typical of desalinated, porous, perennial ice. OW distributions were sampled from many homogeneous footprints across the Arctic associated with near-zero surface wind velocity.

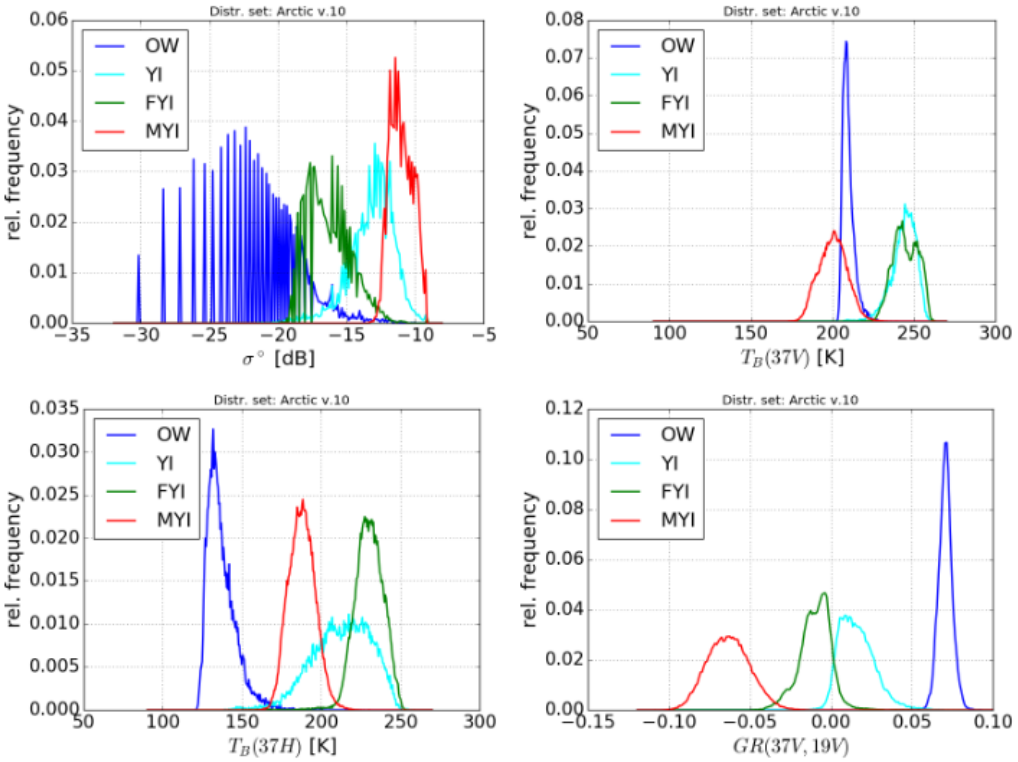

Figure 4.1: Typical distributions in the Arctic for the four separated surface type classes: open water (OW, dark blue), young ice (YI, cyan), first-year ice (FYI, green) and multiyear ice (MYI, red). Distributions are shown for the four independent retrieval parameters: backscatter (top left), brightness temperature at 37 GHz at vertical (top right) and horizontal (bottom left) polarisation and gradient ratio (GR) between 37 GHz and 19 GHz at vertical polarisation. Figure from Melsheimer et al. (2022).

A different approach was used to derive distributions for the Antarctic because pure, homogeneous areas of a single ice type are more difficult to locate through visual analysis alone. The temporal evolution of sea ice and specialised satellite datasets were used to identify representative samples. The MYI distribution was established by sampling daily ASI sea ice concentration maps at the beginning of the cold season, specifically in February and March. At the onset of freeze-up, any ice surviving the summer, primarily in the Weddell Sea, was defined as MYI. It is important to note that these samples were taken only at the beginning of the season, so they may not fully represent the radiometric evolution of MYI throughout the winter. To derive the signature for YI, the algorithm sampled regions later in the freezing season that were located far from the known MYI domains of the early season. This ensured that the sampled ice was seasonal, having grown entirely within the current year. The YI distributions were obtained using an Antarctic polynya satellite dataset from 2018, derived from the Polynya Signature Simulation Method (PSSM) (Markus and Burns (1995), Kern et al. (2007)). Specifically, 'thin ice' up to 20 cm thick was sampled from active polynyas in the western Ross Sea. OW distributions were taken from sample areas in the Southern Ocean in August 2018 and from the ice-free part of the Ross Sea in March 2018. These samples were collected over a period of several

days and covered a large area, ensuring a wide range of atmospheric and oceanic conditions were represented. The resulting distributions are shown in Figure 4.2.

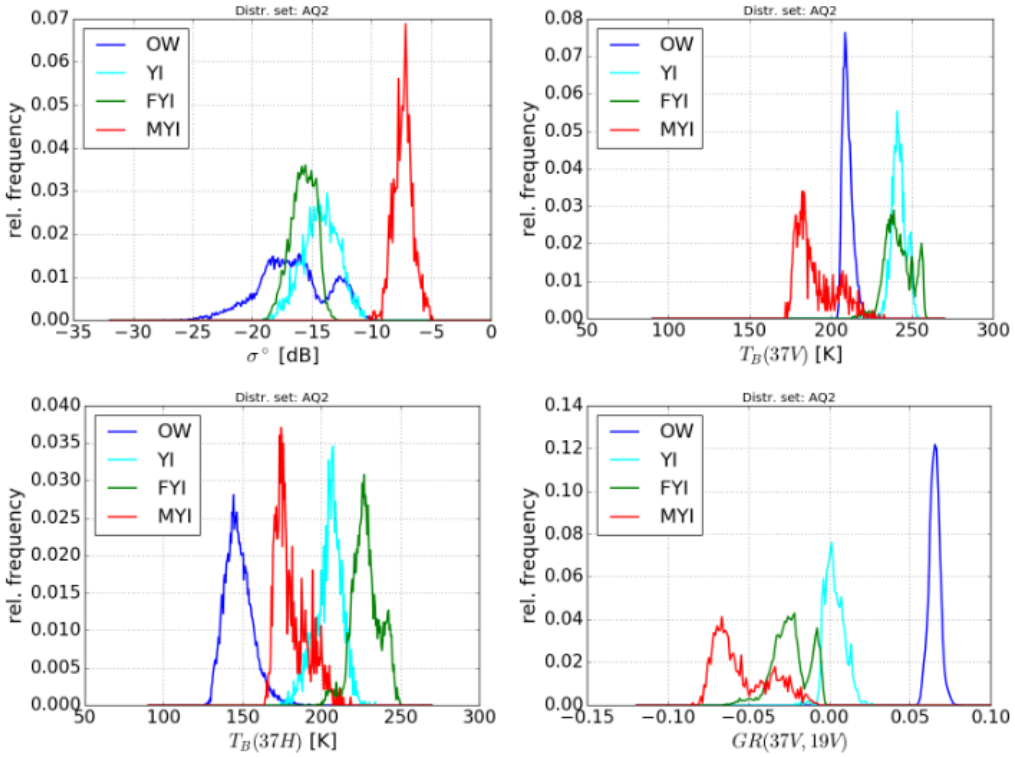

Figure 4.2: Typical distributions in the Antarctic for the four separated surface type classes: open water (OW, dark blue), young ice (YI, cyan), first-year ice (FYI, green) and multiyear ice (MYI, red). Distributions are shown for the four independent retrieval parameters: backscatter (top left), brightness temperature at 37 GHz at vertical (top right) and horizontal (bottom left) polarisation and gradient ratio (GR) between 37 GHz and 19 GHz at vertical polarisation. Figure from Melsheimer et al. (2022).

4.3.4 Algorithm uncertainty

The algorithm quantifies internal retrieval stability by calculating a confidence level CL for each pixel, which is derived from the statistical spread of the 1000 Monte Carlo realisations. This metric is mathematically defined as:

$$CL = 1 - \frac{MAD}{AD_{max}} \tag{4.3}$$

where MAD is the mean absolute deviation of the 1000 trial solutions from their median and AD_{max} is the maximum absolute deviation observed within those trials. The resulting confidence values range from 0.0 to 1.0, with lower values signifying a wider distribution of potential solutions. This approach serves to propagate the uncertainty inherent in the characteristic radiometric parameter value distributions through the optimisation process. More details on the algorithm and its implementation can be found in https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4683347.

To mitigate physical ambiguities in microwave signatures caused by weather and surface metamorphosis, the IUP processing chain incorporates two specific post-processing correction schemes, one based on temperature and the other on drift tracking. Both of them are described in chapter 4.6 on general post-processing steps.

4.4 ASIMIR

The ASIMIR (Multi-parameter Atmosphere, Surface, Ice Microwave Inversion Retrieval) algorithm, developed by Rückert et al. (2023), is designed to simultaneously retrieve atmospheric and surface parameters from satellite microwave radiometer data (AMSR-E and AMSR2) under Arctic freezing conditions. The algorithm is based on the inversion of a physical forward model F which simulates the upwelling brightness temperatures T_B observed by a satellite at the top of the atmosphere:

$$T_B = F(X) \tag{4.4}$$

where X is the state vector describing the surface and atmospheric state. By combining a multi-layer microwave emission model for snow and sea ice with established parameterisations for the Arctic atmosphere and open ocean, the model is able to account for the high variability of Arctic surface emissions while simultaneously considering atmospheric effects and enables the self-consistent retrieval of a set of nine geophysical parameters:

Ice surface parameters
- Sea ice concentration (SIC)
- Multiyear ice fraction (MYIF)
- Snow depth (SND)

Atmospheric parameters
- Integrated water vapour (IWV)
- Liquid water path (LWP)

Temperature parameters
- Snow-ice interface temperature (T_{si})
- Snow-air interface temperature (T_{sa})
- Sea-surface temperature (SST)

Oceanic parameters
- Wind speed (WSP)

4.4.1 Input data

The algorithm runs on AMSR brightness temperatures, as described in Section 3.1.2, using frequencies ranging from 6.9 to 89 GHz. Additionally, a priori data for the retrieved geophysical parameters in the state vector are required. These are mostly extracted from ERA5 reanalysis data , described in Section 3.4.1. For the snow depth, the merged Warenn-AMSR2 climatology described in Section 3.4.2 is used. For the MYI fraction, the output from EC-ICE is used to generate a monthly average.

4.4.2 Physical forward model

The brightness temperature T_B observed by the satellite at a given frequency f and viewing angle (55° for AMSR2 and AMSR-E) can be expressed as a combination of three components: the radiation emitted directly by the surface $T_{B,surf}$ the upwelling atmospheric emission $T_{B,U}$, and the downwelling radiation scattered upwards by the surface $T_{B,\Omega}$, which includes both the atmospheric downwelling term and the surface reflectivity. It can be written as

$$T_B(\theta, f) = T_{B,U}(\theta, f) + \tau(\theta, f)(T_{B,\Omega}(\theta, f) + T_{B,surf}(\theta, f)) \quad (4.5)$$

with the total atmospheric transmittance τ that attenuates the surface-emitted and reflected radiation. The surface contribution from $T_{B,surf}$ and parts of $T_{B,\Omega}$ is treated as a linear combination of sea ice $T_{B,ice}$ and open water $T_{B,ow}$ brightness temperatures, weighted by the sea ice concentration (SIC).

The atmospheric and oceanic parameterisations within $T_{B,U}$ and $T_{B,\Omega}$ are calculated using established parameterisations that simulate upwelling/downwelling emission and surface reflectivity based on the retrieved water vapor, liquid water, wind speed, and sea surface temperature (Wentz and Meissner (2000), Meissner and Wentz (2004,2012)). The contribution of sea ice to the surface signal $T_{B,ice}$ is modeled using the Microwave Emission Model of Layered Snowpacks extended to sea ice (MEMLS_ice), which simulates the emission and scattering of a physical, multi-layer snow and ice pack. This physical model treats the surface as a stack of planar homogeneous layers (usually four for snow and five for ice) and uses six-flux theory to calculate radiative transfer (Wiesmann and Mätzler (1999), Mätzler and Wiesmann (1999), Tonboe (2006,2010)). This approach explicitly accounts for internal reflections (Fresnel equations), absorption and scattering. Essential physical parameters are incorporated for each layer, including thickness, density, salinity and exponential correlation length. The model also treats the snow-air (T_{sa}) and snow-ice (T_{si}) interface temperatures as independent variables, accurately representing non-equilibrium thermal gradients. The values of the fixed model parameters that are implemented are summarised in Rückert et al. (2023, Table 1).

4.4.3 Inversion Scheme

The inversion of the forward model $F(X)$ is performed using a Bayesian Optimal Estimation Method (OEM) which finds the most likely state vector X to describe the geophysical state given the measurement vector Y (consisting of the observed brightness temperatures). The scheme minimises a cost function χ^2 that weights the residual between observed and simulated brightness temperatures against the deviation of the state vector from a priori knowledge, using their respective covariance matrices. This minimisation is executed through the Levenberg-Marquardt algorithm, an iterative procedure that adjusts the state vector based on the Jacobian matrix of the forward model. The a priori state is defined using monthly climatological data from ERA5 reanalysis and the NASA Team algorithm (Cavalieri et al. (1984,1997)) that are linearly interpolated to the specific time and location of the satellite observation. A diagonal covariance matrix is used, assuming no correlations.

4.4.4 Algorithm uncertainty

The ASIMIR algorithm provides a mathematically rigorous estimate of the retrieval uncertainty for each geophysical parameter through its Bayesian optimal estimation framework. These uncertainties are represented by the retrieved covariance matrix $\hat{S}$, where the square root of its diagonal entries defines the 1-sigma uncertainty for each variable. The calculation of $\hat{S}$ integrates the forward model's sensitivity (represented by the Jacobian matrix) with the a priori covariance matrix S_a and an effective measurement uncertainty matrix S_e. S_e accounts for both, forward model uncertainty S_b and sensor-specific radiometric noise S_y, the latter of which is assumed to be 1 K for all channels. The model uncertainty was quantified through a sensitivity analysis involving 1000 MEMLS_ice simulations per frequency tuple, capturing the variability lost by fixing certain physical parameters like layer density or salinity. The resulting uncertainties are highly frequency-dependent with values from approximately 1.2 K to 3.9 K at 6.9 GHz and up to 7.0 K at 89 GHz.

4.4.5 Additional uncertainty sources

The atmospheric Model: The atmospheric sub-model was not specifically developed for the polar atmosphere, which introduces inherent but unquantified model uncertainty.
Melting surface conditions: The retrieval is limited to freezing conditions only and therefore only applied in the Arctic winter months (October to April). Occasional melting periods can significantly increase the retrieval uncertainty, as the forward model assumes zero wetness and cannot radiometrically distinguish between surface meltwater and the open ocean. Consequently, these unsimulated conditions often lead to retrieval failures due to non-convergence or produce biased geophysical parameters.
Model constraints: The number of layers for the snow and ice pack simulation is fixed and therefore the possibility to resolve gradients of temperature or salinity is limited.

4.5 C3S Bayesian sea-ice type classification algorithm

The C3S sea-ice type product is based on a Bayesian classification framework applied to passive microwave radiometer observations. The algorithm estimates, at each grid cell, the probability of belonging to predefined surface classes and assigns the class with the highest probability. The retrieval is a multi-stage processing chain, as illustrated in the flow-chart in Figure 4.3. Starting from Level-1 brightness temperature swath data, a series of pre-processing steps (see Section 4.2) are applied, followed by the derivation of Level-2 geophysical classification parameters.

At Level-3, the data are gridded and combined across sensors, and the core computation of Bayesian class probabilities (Section 4.5.1) is performed. This step relies on probability density functions (PDFs) derived from training data over known surface classes (Section 4.5.2). Finally, at Level-4, the classification is refined through post-processing steps, including gap filling, correction schemes (Section 4.6), and the application of masks and filters. Uncertainty estimates are derived directly from the Bayesian probabilistic output.

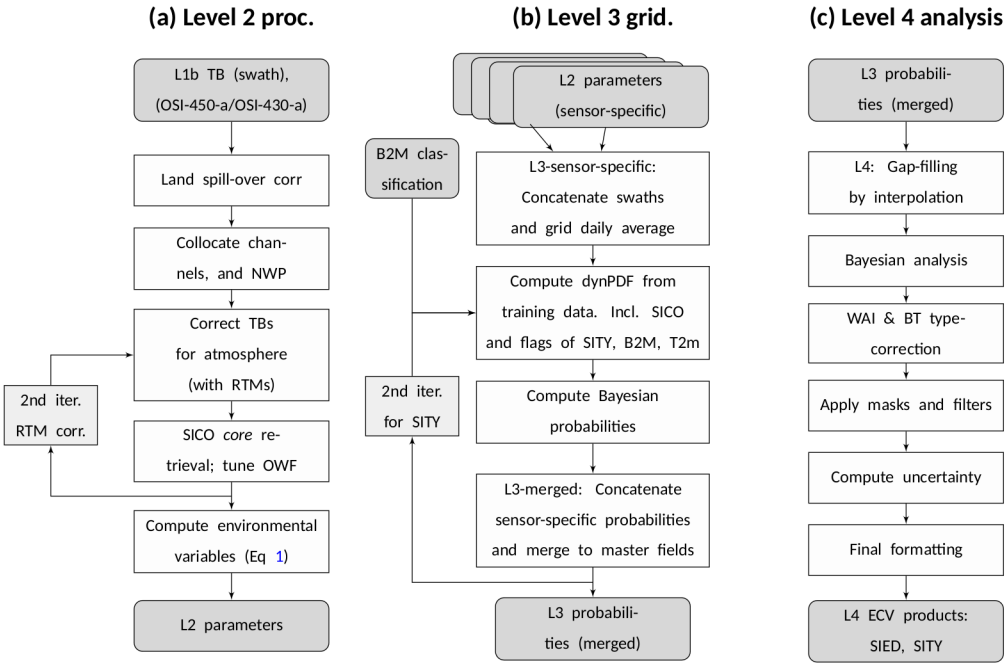

Figure 4.3: The main processing elements in the daily calculations of sea-ice edge and type are divided into 3 sub-routines: (a) Level 2, (b) Level 3 (sensor-specific and sensor-merged), and (c) Level 4 analysis. The darker grey boxes (rounded corners) represent data input and output, the white boxes represent the processing steps. Level 2 are the inputs to Level 3, and the Level 3 gridded fields are the inputs to the Level 4 analysis. The light grey boxes indicate iterative processing steps. Figure from C3S ATBD for sea-ice edge and sea-ice type CDRs (Aaboe et al. (2025)).

The primary classification parameter derived from the observations is the gradient ratio GR1937V, complemented by polarisation ratios such as PRnear90 and PR19, as defined in Eqs. 2.1 and 2.2, for defining the edge and open-ice zone between open water and closed ice. The selection of frequencies near 19 GHz and 37 GHz ensures consistency across the full time series back to 1978, while higher-frequency channels near 90 GHz are used, where available since 1991, to enhance spatial resolution, particularly in the sea-ice edge regions.

4.5.1 Bayesian approach

The mathematical expression for the Bayesian probability (p) of the ice surface class (I) and with the measurement of GR1937V is given as:

$$p(I|GR1937V) = \frac{p(GR1937V|I)}{p(GR1937V|FYI) + p(GR1937V|MYI)} \quad (4.6)$$

In Equation 4.5, the expression is a simplified form of the general expression, where multiple independent parameters can be combined; here the implementation is limited to GR1937V. Another simplification is that the classification of ice is divided into two categories, either FYI or MYI. In addition, the prior probabilities of each class are set equal, thus avoiding bias towards a priori ice distributions and thereby cancelling each

other out in Equation 4.5. The $p(GR|I)$ terms are the conditional probabilities that describe the behaviour of the observations (here GR based on TB) over a certain known surface class.

The final ice type at each grid cell is defined as the class with the highest posterior probability. If the maximum probability does not exceed a predefined confidence threshold (typically around 75%), the classification is considered uncertain and the pixel is assigned to an ambiguous ice class.

4.5.2 Dynamic PDFs

The properties of sea ice, and consequently the measurements of its emissivity, change throughout the seasons. Therefore, dynamic conditional probabilities are calculated from a daily-updated training dataset covering a 15-day period centred on the specific day.

In order to select training data in a dynamical framework, several selection criteria are applied:

Geographical selection is the first step in sea-ice type classification, where training data are collected from regions with a high likelihood of each ice class. In the Arctic, the oldest ice is typically found north of Greenland and Canada, while newer ice dominates the Siberian shelf. In the Antarctic, most multi-year ice is located in the Weddell Sea. These regions are defined broadly to ensure sufficient training data throughout the seasonal cycle and across the climate record.

A priori knowledge of ice types' distribution is required to select reliable training data to avoid contamination between FYI and MYI. This is addressed using (1) the most recent sea-ice type product (e.g. from the previous day) and (2) an independent classification based on sea-ice drift, concentration, and backtracking to the minimum ice extent (Section 4.5.3).

A temperature threshold filter avoids using training data over ice that is under exposure to some warm event. HereNWP data of T2m is used .

A concentration threshold filter of 70% concentration is used to select only training data over closed ice regions to reduce influence from open water.

Assumption of Gaussian distribution

For the selected training data, the probability density functions (PDFs) of each classification parameter (e.g. *GR1937V*) conditioned on a surface class (*I*) are assumed to follow Gaussian distributions. Under this assumption, each class-specific PDF is fully described by its mean and variance, which are estimated from the training data. With this assumption, the conditional probabilities can be expressed as:

$$p(GR|I) = \frac{1}{\sigma} exp(-\frac{(GR-\mu)^2}{2\sigma}) \quad (4.7)$$

where μ and σ are the mean and standard deviation estimates of the measured *GR* for each class. μ and σ are computed from the daily-updated training dataset per sensor, and we refer to the conditional probability as the dynamic PDF (dynPDF).

Two sets of dynPDFs are generated daily, one for each hemisphere. Figure 4.4 shows examples for the Northern Hemisphere, including the corresponding Gaussian fits. A Gaussian assumption is used for simplicity and generally provides a good representation of the training data. While asymmetric distributions may offer improved fits in some cases, no suitable alternative has yet been established.

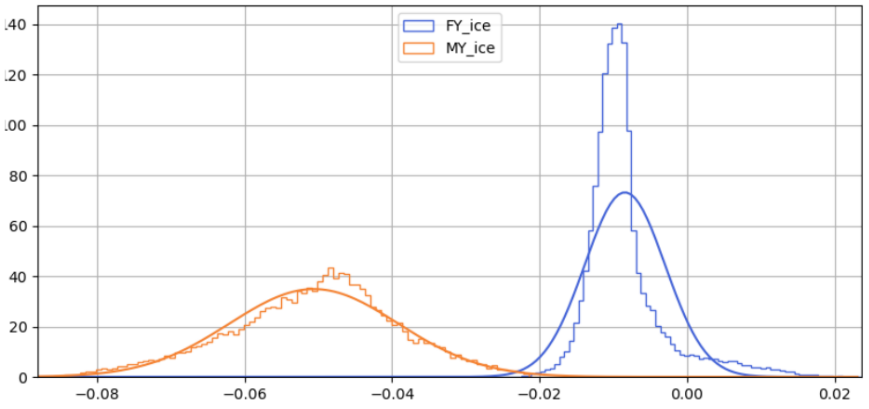

Figure 4.4: Density plot of the GR1937V for the sea-ice type classes of FYI (blue) and MYI (orange). Data are collocated from 15 days of SSMIS data centred on 2013-03-01. The stair plots show the distribution of the data, and the thin lines show the corresponding Gaussian fit to the data. Figure from C3S ATBD for sea-ice edge and sea-ice type CDRs (Aaboe et al. (2025)).

4.5.3 Backtracking ice-type classification

An independent estimate of sea-ice type is derived from sea-ice drift and concentration fields through backtracking of ice parcels. For this, sea-ice concentration OSI-450-a and OSI-438 (Section 3.3) and sea-ice drift NSIDC-0116 and OSI-405 (Section 3.2) are used. The backtracking classification is performed in two ways: for selecting training data (Section 4.5.2) and for post-processing the classification (Section 4.6.1).

Backtracking to the minimum extent is based on the assumption that all ice present at the time of minimum sea-ice extent becomes multiyear ice (MYI), while ice formed afterwards is first-year ice (FYI). Ice parcels are tracked backwards in time using sea-ice drift fields until they reach the date of minimum sea-ice extent. If the tracked trajectory leads to open water en route to the minimum date, the pixel is classified as FYI; otherwise, it is classified as MYI. The method is applied only to pixels with sufficiently high ice concentration and provides an independent estimate of the large-scale MYI distribution. This independent ice-type classification is used to sort out target data (Section 4.5.2).

Backtracking to known ice type involves tracking ice parcels backwards over shorter time periods to a known ice-type field. This approach reduces dependence on long, continuous time series of concentration and drift data and enables estimation of the expected evolution of ice type under drift. Backtracking periods of different lengths (e.g. several days to months) are used to derive a probabilistic estimate of the likely ice type. This information is primarily used in post-processing to identify and correct inconsistent or spurious classifications (Section 4.6.1).

4.5.4 Level-4 gap filling and masks

The Level-4 processing involves post-processing correction schemes to remove physically implausible classifications. These include primarily the temperature-based and drift-based corrections, described in Section 4.6, and climatological maximum sea-ice extent masks (Section 3.5) to suppress spurious ice detection outside expected regions.

Data gaps arising from missing swath coverage or the polar observation hole are filled using a two-step interpolation approach. First, temporal interpolation is applied using neighbouring days, followed by spatial interpolation based on distance-weighted averaging.

4.5.5 Uncertainty estimates

The uncertainty in the C3S sea-ice type classification is a purely statistical uncertainty and is derived directly from the Bayesian classification output. For each grid cell, the posterior probabilities of the predefined surface classes provide a quantitative measure of classification uncertainty.

Low uncertainty is associated with cases where one class clearly dominates the posterior probability distribution, whereas higher uncertainty arises when probabilities are more evenly distributed between classes, indicating ambiguity in the observed microwave signal. This typically occurs under conditions such as melt, in marginal ice zones, in transition zones between ice classes, or during atmospheric disturbances.

In addition to this probabilistic measure, pixels affected by post-processing corrections (e.g. temperature- or drift-based adjustments) are explicitly flagged and assigned increased uncertainty to reflect reduced reliability.

The resulting uncertainty is provided on a per-pixel basis and reflects the intrinsic ambiguity of the Bayesian classification and the impact of known limitations in the retrieval; however, it captures only part of the total uncertainty and does not include all sources propagating through the processing chain.

4.6 Post-processing

The ASIMIR algorithm does not include sea-ice-type-specific post-processing steps, as it is designed as a multi-parameter retrieval, not specifically designed for ice-type detection only. In contrast, the EC-ICE and C3S algorithms rely to a large extent on post-processing schemes to reduce noise and improve classification, particularly during challenging conditions. These schemes make use of independent auxiliary information, such as near-surface air temperature and sea-ice drift history. For EC-ICE and C3S, the two major post-processing schemes are explained here.

4.6.1 Drift correction

Drift correction is a spatiotemporal post-processing approach designed to enforce physical consistency in satellite-derived sea-ice type products, particularly for MYI. It addresses a fundamental limitation of microwave-based retrievals: the fact that instantaneous radiometric observations alone are often insufficient to reliably distinguish MYI from FYI under certain environmental conditions.

In the EC-ICE post-processing (following Ye et al. (2016b)), it is based on the physical principle that the area of MYI cannot increase during the freezing season after the initial freeze. Its total area can only decrease through export, deformation or melting. Retrieval errors usually occur when physical processes such as snow grain metamorphosis, surface deformation or late-winter melt-refreeze cycles cause MYI to resemble the high-backscatter and low-emissivity signature of FYI.

The algorithm uses daily ice motion vectors to constrain the spatial evolution of the MYI domain via a sequential predictive process. First, the MYI domain for each day is established by identifying all pixels where the retrieved concentration exceeds a specific threshold (e.g. 20%). The scheme then applies daily drift vectors to these pixels to calculate their expected displacement, thereby defining a physically plausible MYI domain for the following day. To account for uncertainties in satellite-derived ice drift products, the predicted domain is then expanded by a one-pixel buffer. Any MYI concentration retrieved outside this expanded boundary is considered physically impossible.
Furthermore, the scheme includes a consistency check for pixels within the plausible domain to identify misclassifications caused by rapid snow metamorphosis. It monitors for anomalous changes in MYI concentration that occur concurrently with sharp decreases in the 37 GHz horizontal brightness temperature ($\Delta T37 < -20$ K) or the horizontal range (HR):

$$TB_{37H} <- 20\,K \qquad HR = TB_{19H} - TB_{37H} < -10\,K \qquad (4.8)$$

When either of these criteria is met, the anomalous concentration is replaced with the previous day's value. This ensures temporal continuity and a more realistic, monotonic decrease in MYI area throughout the freezing season.

Following a principle similar to EC-ICE above and to Ye et al. (2016b), C3S post-processing also uses tracking of the ice-type field to correct false or physically unrealistic classifications between FYI and MYI. In this approach, the backtracking-to-known-ice-type (BT) classification (Section 4.5.2) is used to identify inconsistencies between the retrieved ice type and the drift-based history.

The BT information is applied in two ways: (1) to fill remaining gaps in MYI classification not resolved by the temperature correction (Section 4.6.2) and (2) to mask spurious MYI occurrences that are not supported by the drift history. This is particularly relevant in regions where sudden or persistent MYI misclassification is observed.

The effectiveness of the drift correction depends on the backtracking period and the persistence of the anomalous signal. All pixels affected by the BT-based correction or masking are explicitly flagged in the output product.

4.6.2 Temperature correction

Microwave-based sea-ice type classification is sensitive to short-term atmospheric variability, particularly warm-air intrusions (WAI) associated with cyclonic activity (e.g., Tonboe et al. (2003)). This physical ambiguity

arises because warm air advection introduces liquid water into the snowpack on MYI. Under these conditions of wet snow, the microwave emissivity of the snowpack increases while radar backscatter decreases. This causes the radiometric signature of MYI to approach values that are typical of FYI. This results in a sudden, non-physical drop in MYI concentrations, which typically lasts from a few days to weeks, until cold temperatures resume and the signature is restored.

To address this limitation, temperature-based correction schemes have been developed. These methods use near-surface air temperature to identify anomalous warming events and correct the resulting non-physical changes in ice type. The general approach follows the methodology of Ye et al. (2016), which forms the basis for the implementations used in both EC-ICE and C3S.

The temporal post-processing module (Ye et al. (2016a)) is a correction scheme designed to mitigate the misclassification of MYI as FYI during atmospheric warm spells in the transition seasons. The correction procedure uses 2 m air temperature (T2m) data from meteorological reanalysis (described in Chapter 3.2.3) to identify warm episodes lasting up to 10 days. An episode is triggered when the air temperature rises above a threshold (T1), and ends when temperatures fall below another threshold (T2).

In EC-ICE implementation, these threshold values are -1 °C and 2 °C, respectively. Additionally, if the retrieved MYI concentration during this identified period drops by more than a specified threshold (set at 10%) and then recovers afterwards, the intermediate values are flagged as anomalous. To restore the physically plausible MYI concentrations, the scheme replaces these erroneous values using linear interpolation between the MYI concentrations observed immediately before and after the warm episode. This restorative approach ensures the temporal continuity of the MYI area during transient autumn and spring melting events, effectively correcting the dips in the time series that would otherwise lead to an underestimation of the perennial ice cover.

A similar correction approach is used in the C3S algorithm, which is based on the same methodology. However, instead of applying interpolation, the C3S implementation distinguishes between reliable and ambiguous cases based on the temporal behaviour of the ice-type signal and the strength of the temperature anomaly. In reliable cases, MYI is restored during the identified gap period, whereas in more ambiguous situations, the affected pixels are assigned to an intermediate or ambiguous class.

All corrected pixels in the C3S product are explicitly flagged, and their associated uncertainties are increased to reflect reduced confidence in the classification during such events.

These temperature-based corrections reduce non-physical variability in MYI classification and improve temporal consistency, particularly during short-lived warming events that affect the microwave signal.

5 Lagrangian sea ice age algorithm

5.1 Overview of the sea ice age algorithm

The sea ice age is computed from sea ice concentration and drift in the following way. A pan-Arctic field of concentration is taken from satellite observations at the end of the melt season, when the ice extent is minimal, and all ice is assumed to be multi-year ice (MYI, C_{MYI}) (Figure 5.1, Step 1). The C_{MYI} field is repetitively advected (morphed) using daily satellite-observed sea ice drift fields (Figure 5.1, Step 2). Changes of concentration due to ice deformation or melting are accounted for during the advection process (see Section 5.3 below).

Some time after the initialisation of the MYI field advection (e.g. on 1 Jan 1992, as shown in Figure 5.1), the advected field represents a concentration of MYI, which is lower than the total observed concentration and is denoted as C_{A0} (i.e., advected for less than one year). The difference between the total and advected fields yields the concentration of the first-year ice (Figure 5.1, Step 3):

$$C_{1Y} = C_{TOT} - C_{A0} \quad (5.1)$$

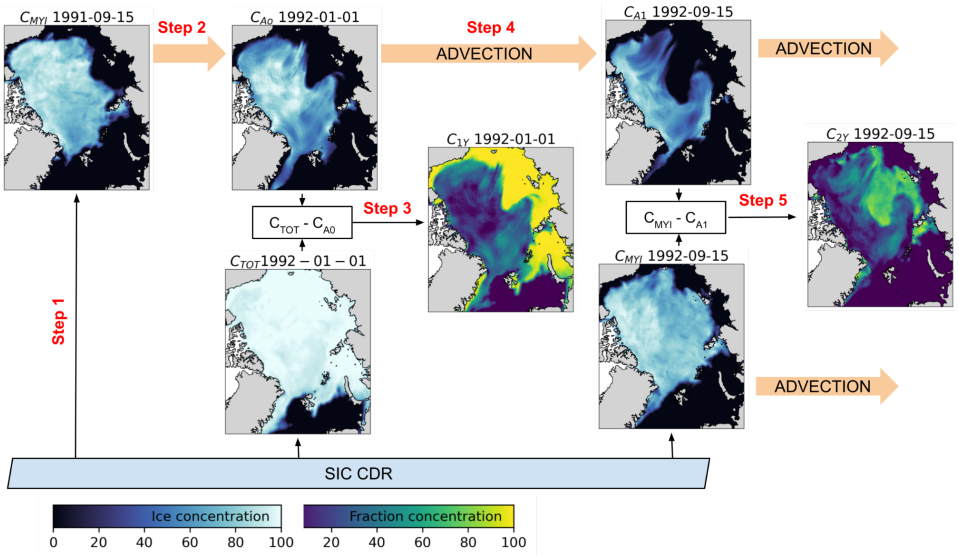

Figure 5.1: A scheme of sea ice age fraction computation within one year. The fields of MYI and TOT concentration originate from SIC CDR (Ice concentration colour map). These fields are advected using sea ice drift from the SID CDR (advection is shown by thick orange arrows). The advected fields of MYI are labelled as "C_{Ai} YYYY-MM-DD", where i denotes the number of years of advection and YYYY-MM-DD denotes the date of the advected field. MYI and advected fields are used to compute sea ice age fractions (Fraction concentration colourmap), labeled as "C_{jY} YYYY-MM-DD", where j denotes the age of the ice fraction, i.e. C_{1Y} stands for the first-year ice and C_{2Y} stands for the second-year ice.

Hereafter, we define the first-year ice (C_{1Y}) as ice formed during the ongoing freezing season and that has not yet experienced melting.

One year after the initialisation, the advected field (Figure 5.1, Step 4) represents the concentration of sea ice which is at least two years old and is denoted C_{A1} (advected for one year). At that time, the total observed concentration again reaches a minimum, representing the concentration of multi-year ice. C_{MY} has higher values than the advected C_{A1} as it also contains a fraction of the second-year ice (Figure 5.1, Step 5):

$$C_{2Y} = C_{MY} - C_{A1}. \quad (5.2)$$

Hereafter, we define the second-year ice (C_{2Y}) as ice formed during the previous freezing season and that has survived one melting season. It should be noted that according to the nomenclature of the World Meteorological Organisation (Sea Ice Nomenclature, WMO-259), the first-year ice (FYI) that survives the summer minimum is called “residual ice”, and it turns into second-year ice only on 1 January of the coming winter. Nevertheless, in this algorithm, we define that all survived FYI turns into the second year after 15 September.

As shown in Figure 5.2, both C_{A1} and the current C_{MY} are advected further using the sea ice drift and after one more year they become C_{A2} and C_{A1}, i.e. ice fractions advected for two years and for one year. Since C_{A1} contains C_{A2} and, similar to the previous year, the new C_{MY} contains C_{A1} we can compute fractions of the second- and third-year ice:

$$C_{2Y} = C_{MY} - C_{A1}$$
$$C_{3Y} = C_{A1} - C_{A2} \quad (5.3)$$

This workflow is repeated, and the ice age fraction can be computed using a generic formula:

$$C_{NY} = C_{A(N-2)} - C_{A(N-1)}, \quad (5.4)$$

Where N is an integer number indicating the age of sea ice fraction.

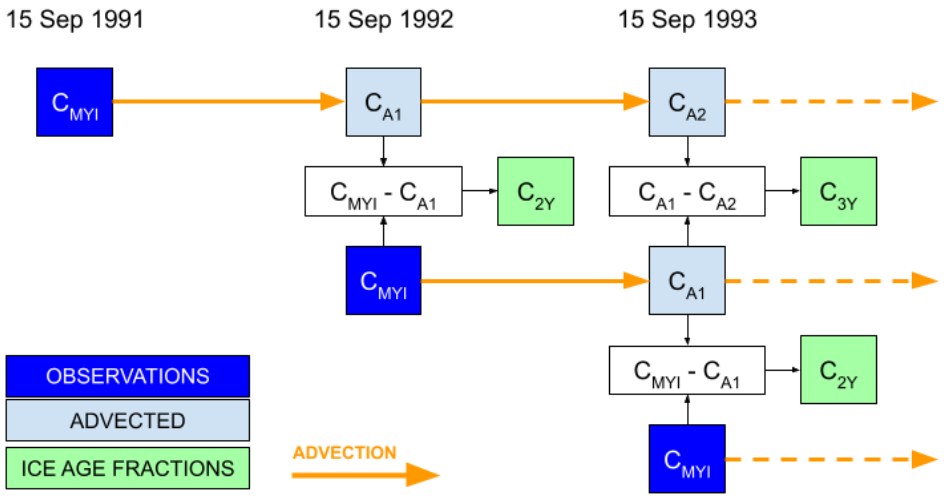

Figure 5.2: A scheme of sea ice age fraction computation for multiple years. The field of multi-year ice from observations is shown as dark blue blocks, advected fields are shown as light blue blocks, and the computed ice age fractions are shown as green blocks. Orange arrows show advection.

In this algorithm, we use a Lagrangian scheme, where nodes of a triangular mesh are advected using the ice drift vectors. Triangular elements in this mesh are advected without diffusion unless a remeshing occurs. The sections below provide a detailed description of the advection scheme.

5.2 Generation of initial mesh

The initial triangular mesh is created for the region of interest, including the Arctic Ocean above 60 °N and excluding the Baffin Bay and Canadian Archipelago, by triangulating points on a regular grid in EASE2 projection (Brodzik and Knowles, 2011) with ≈ 25 km spacing. Points located on land more than 150 km from the coastline are excluded from triangulation. The mesh is optimised with the Laplace method using the GMSH library (Geuzaine and Remacle, 2009). The nodes of the mesh located on land are marked as fixed: they cannot be moved by the advection or remeshing procedures.

5.3 Advection of mesh nodes and remeshing

The daily ice drift vectors from the OSI SAF product are linearly interpolated on the nodes of the triangular mesh and position of the nodes is updated: $X_{n+1} = X_n + U_n$, where X_n is the initial position of nodes, U_n is ice drift velocity in km/d and X_{n+1} is the new node position.

After the advection, some elements (triangles) of the mesh are critically distorted and require remeshing. The following criteria are set to detect distorted elements:

- Edge of the element is shorter than 13 or longer than 38 km;
- Element has an angle below 15°;
- Element area is smaller than 20 km²;

- Element is flipped.

The end results are not very sensitive to the mesh size. These parameter values are chosen to keep the area of the mesh elements below that of the destination grid elements with a spatial resolution of 25 km, while ensuring the elements are large enough for efficient advection and, especially, to avoid time-consuming remeshing.

The following operations are applied recursively to the mesh, changing one edge at a time (see Fig. 5.3 for illustration):

- If an edge is shorter than a threshold or an angle in the element is below a threshold, the shortest edge of the element is collapsed: two nodes that belong to the shortest edge are replaced with the one node between them, and the element that had the shortest node is removed (Figure 5.3, A).
- If an edge is longer than a threshold, the edge is split in two: a node is added in the middle of the edge, and the initial large element is replaced with two smaller elements sharing one new edge (Figure 5.3, B).
- If an element is flipped, the edge over which it is flipped is removed, and the created quadrangle is bisected by the edge connecting the two other nodes (Figure 5.3, C).

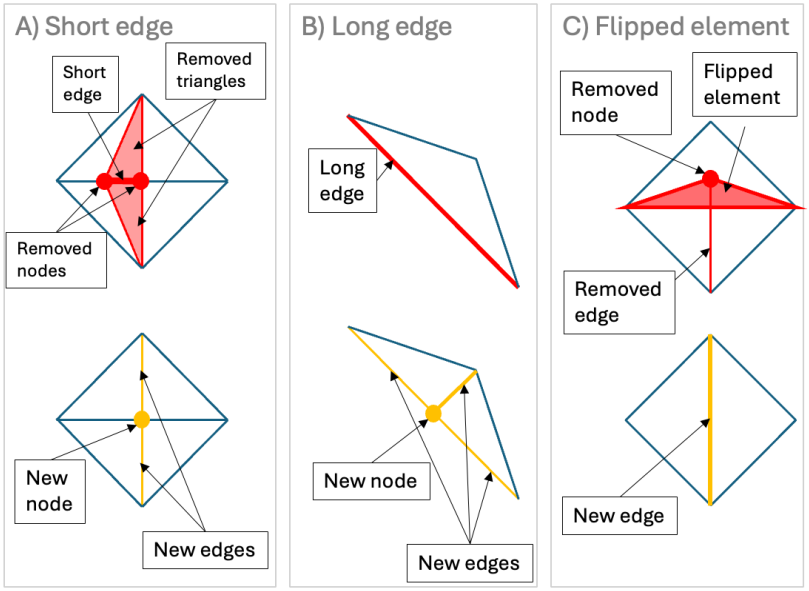

Figure 5.3: Scheme of three types of remeshing: collapsing of a short edge (A), splitting of a long edge (B), removing a flipped element (C). Edges and elements before remeshing are shown in red in the upper row, and the new mesh is shown in yellow in the lower row.

5.4 Mapping between advected meshes

Most of the elements of the mesh at step *n+1* are not distorted enough to require remeshing and have a corresponding element on the mesh at step *n*. For these elements, a conservative mapping is used – the entire content from the element on the previous step is transferred to the corresponding element on the next step:

$$A_{i,n+1} = A_{i,n}, \quad (5.5)$$

where *A* is an areal (i.e., not voluminal) content, e.g., area of sea ice in an element.

Since the area of the non-remeshed element may change, the concentration (*C*) of an areal content (e.g., sea ice concentration) also changes:

$$C_{i,n+1} = k_{i,n+1}\, C_{i,n} \quad (5.6)$$

where $k_{i,n+1}$ is a factor equal to the ratio of the changed element area a_i:

$$k_{i,n+1} = a_{i,n} / a_{i,n+1} \quad (5.7)$$

For the elements that were remeshed, we find all elements on the advected mesh that they intersect with (see Figure 5.4). The total areal content of the remeshed element is equal to the weighted average of the intersecting elements, where the weight is proportional to the area of intersection. The concentration of an areal content is, therefore, equal to the weighted average of concentrations adjusted by the factor accounting for change in the element area:

$$C_{i,n+1} = \sum_i (w_j\, k_{i,j,n+1}\, C_{i,j,n}) / \sum_j w_j \quad (5.8)$$

where w_j is the weight of the j-th intersecting element. Note that this equation is generic and can also be used for conservative mapping with $j = 1$ and $w_1 = 1$.

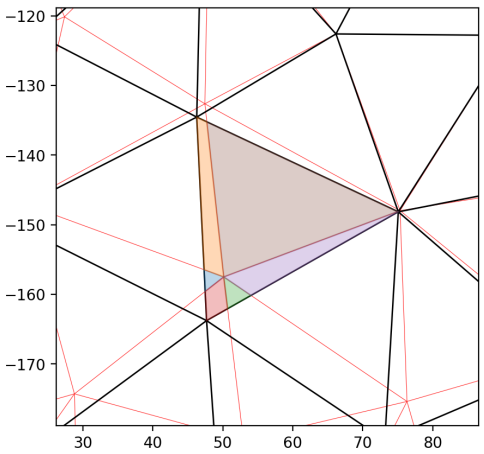

Figure 5.4: Illustration of computing weights for the intersecting mesh elements. The previous (advected) mesh is shown by red colour. The new (remeshed and optimised) mesh is shown in black. Intersections of the central element from the new mesh with elements from the previous mesh are shown by polygons of various colours.

After each advection, the indices of the elements on the source mesh, along with their weights, are saved together with the new mesh to facilitate quick mapping between meshes. The advection and remeshing process described above is applied to the whole time series of sea ice drift data once to produce daily meshes and mappings between them. These meshes and mappings are then used to advect the sea ice concentration field.

5.5 Advection of the sea ice concentration field

Prior to advection, the initial sea ice concentration field ($C_{OBS,n}$) for a selected date is linearly interpolated from the OSI SAF gridded product to the centres of elements of the corresponding mesh. For the next day, the corresponding mesh and the mapping are loaded, and the advected concentration (C_{n+1}) is computed using Equation 5.4. The OSI SAF concentration for the next day ($C_{OBS,n+1}$) is also interpolated on the same mesh. The observed concentration conditions the advected concentration:

$$\hat{C}_{n+1} = min(C_{n+1}, C_{OBS,n+1}). \quad (5.9)$$

Conditioning is necessary in cases where the element area decreases due to convergence, resulting in a concentration that exceeds the observed value. Conditioning can be interpreted as ice ridging – the excess of sea ice area in a compressed element increases the sea ice thickness, which we cannot account for, as we don't have accurate enough observations of sea ice thickness during the entire period of SIC and SID observations.

The observed concentration can be lower than the advected one, also due to errors (e.g., an unaccounted atmospheric impact on brightness temperatures of passive microwave data). In such a case, the advected

concentrations are also conditioned, and the uncertainty of C_{OBS} is used for computing the uncertainty of the advected field as described below.

5.6 Initialisation of the MYI concentration field

The concentration of 15 September can be assumed to be minimal and taken as the concentration of multi-year ice. However, in the central part of the Arctic, freezing starts earlier, and by 15 September, the first-year ice is already present in some regions. To avoid including FYI in the MYI map, the annual initialisation procedure is implemented as follows. The observed total concentration fields on all days between 5 September and 14 September are taken and advected until 15 September independently of each other without capping. Then, the minimum concentration from the advected fields is taken into the MYI concentration field:

$$C_{MYI} = min(C_{A,i}), \quad (5.10)$$

where $C_{A,i}$ is the advected concentration starting from i-th date (i ranging between 5 and 14 September).

6 Machine learning for sea ice type

For classification and segmentation of image-type data within the geospatial domain we have, over the past few years, seen an increase in both usage and performance of machine learning models. Specifically the usage of deep learning methods like Mask-RCNN, UNets, and similar convolutional neural network architectures. These model architectures have remained relevant and acted as benchmarks for state of the art (SoTA) over the last decade. We support our models with the use of pre-training tasks (Vats et al. (2024)), data augmentation (Xie et al. (2020)), data mixing and pseudolabeling (Kage et al. (2026)), and conditional random fields (Jiao et al.(2020)).

Putting it simply: we will be using a segmentation model which outputs results that are spatially co-located with the inputs. The model will act as a simple lens through which the multi-variable inputs over a geographical region is reduced to a single map over ice types in that same geographical space (or equivalently a likelihood/probability distribution over the ice types in the area of interest).

6.1 Pipeline steps

1. Data preprocessing
2. Sampling / mixing strategies
3. Pre-training + Training
4. Post-processing
5. Model selection
6. Deployment

Steps (1,2) primarily focus on loading, normalising, and sampling of the data. We also account for whether to include a normalisation layer for the inputs, which may affect the strategies for pre-conditioning. Different sampling strategies may be applicable to the utilised snapshot data, where we consider the coverage and/or weighting of sample snapshots. We grid the input images and use samples from subsections of the grid as input to our training pipeline. To ensure better generalisation from the trained network, we consider perturbation and augmentation strategies that may apply. Since there is a canonical orientation of the input data, it seems likely that the network will encode some information from the position on the map. This encoding may contain information on hidden states that are statistically significant, and we may find that rotation or flipping of the input results in a poorer overall performance. We will likely therefore focus only on augmentation methods that do not perturb the orientation of the data, and focus on methods that are coherent with the variations of the data, e.g. added noise, delabeling and mixing.

Step (2) also contains an option for mixing strategy, whereby one mixes two inputs in a given batch to introduce synthetic delineation. Here we will experiment with including a CowMix-strategy (French et al. (2020)), which uses a variation of the CutMix algorithm that is sampled from gaussian noise, and then smoothed to form masks that resemble the patterns on a cow (Figure 6.1). These have previously shown to be more useful than a straight CutMix when dealing with organic or naturally occurring boundaries. We include these to pairwise mix between samples in an input batch in order to ensure a stronger learning-signal

from the boundaries between ice types (ITY), which also opens the opportunity to use delabeling of uncertain boundaries in step (1).

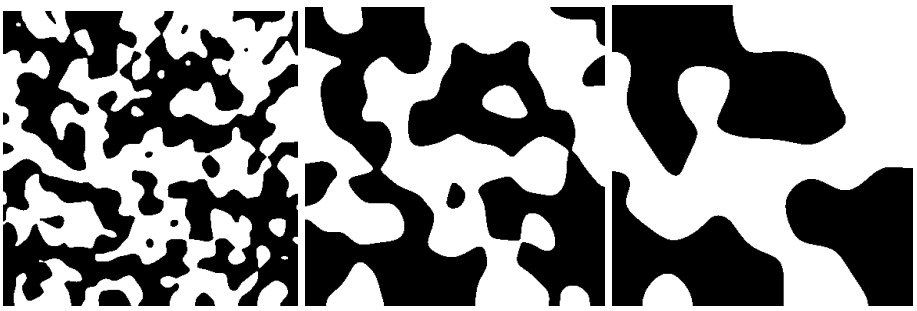

Figure 6.1: CowMask figures from French et al. (2020).

Step (3) contains the overall training steps, with an optional pre-training. Pre-training the network has been shown to moderately improve the performance (Vats et al. (2024)) when training on predicting known, derivable values from the input. An easy target for pre-training can be the season, or other known meta-data, noting that the pre-training is primarily used as a tool to train the network on the structure of the data and provide a level of adjacency to the target task.

The training itself consists of providing the network with batches of training samples, and comparing the output to the target output. Here we will use a variation of the cross-entropy loss known as Focal Loss, which has an additional hyperparameter, gamma, that allows for a higher weighting on pixels that the network consistently gets wrong and a down-weighting of pixels that are consistently correct. We use a standard AdamW optimiser for the updating of the weights and backpropagation of the gradients, with a sensible selection of parameters based on the current understanding of best practices (Orvieto et al. (2025)).

The I/O dimensions of the model is [B, N, H, W] / [B, D, H, W], where B is the batch-size, and N is the amount of input channels. The output channels D correspond to the different ice-types that our model will be classifying, and ultimately reduce to a single channel during post-processing. H and W refer to the height and width of the samples as determined in Step (1).

Step (4), post processing, contains the aforementioned reduction of predicted classes to a single label. This step can be done using a naive ArgMax, i.e. picking the class with the highest likelihood, or by setting threshold values for each class. We will also include an option of conditional random field (CRF) - post processing, which has been shown to sharpen the edges between classes in some cases.

Step (5), model selection, is performed based on a randomised holdout of the training data. For this we retain a portion of the training data that is used to compare the performance of models with different tunable parameters and inclusions of e.g. CRF. Once the parameters have been selected we train a final version of the network that uses all the available training data before handing over for comparison with the RRDP.

Step (6), deployment, here refers to the delivered model for RRDP comparison and potential future implementation of the model as part of the overall family of algorithms that are used for the prediction of ice types. The model is saved and bundled so that it may be rerun with arbitrary input data, i.e. it can take any data of the shape [C, H, W], where C is the channels which it is trained on and is therefore locked, and (H, W) are the input width and height, only limited at a minimum that is determined by the bottleneck size of the UNet.

6.2 Data

6.2.1 Quality assurances

For all AI / ML applications the quality of data matters. It is generally accepted that neural networks act as function approximations, which means that the quality of input data dictates the function that will be approximated in the end. We assume the data we train our model on will remain representative of any future data, when accounting for small variations in noise. Similarly we assume the orientation of the input data will remain consistent with what is trained on, which allows the network to encode some bias of the underlying topology in the domain of interest. If the model will be used for both Northern and Southern hemispheres, then the arguments of Section 6.1 (1) will need to be revisited, and there would be a stronger case for introducing rotation or flipping of the data as augmentations to decouple the model from spatial or geographical information and bias.

6.2.2 Input data

We use subsamples of the available RRDP-snapshot data [RD-2] as our inputs, i.e.:

- All available AMSR channels
- ERA5 data
- ASCAT

All / any data augmentations are applied equally across channels and inputs. Noting that any added noise will still be channel-dependent, and should only simulate what is likely to be seen in the future, rather than universally applied. The label data will not be subjected to noise, but may be unlabeled in a buffer zone around the edges when using mixing, in order to artificially introduce sharper boundaries in the synthetic recomposition of the image-pairs.

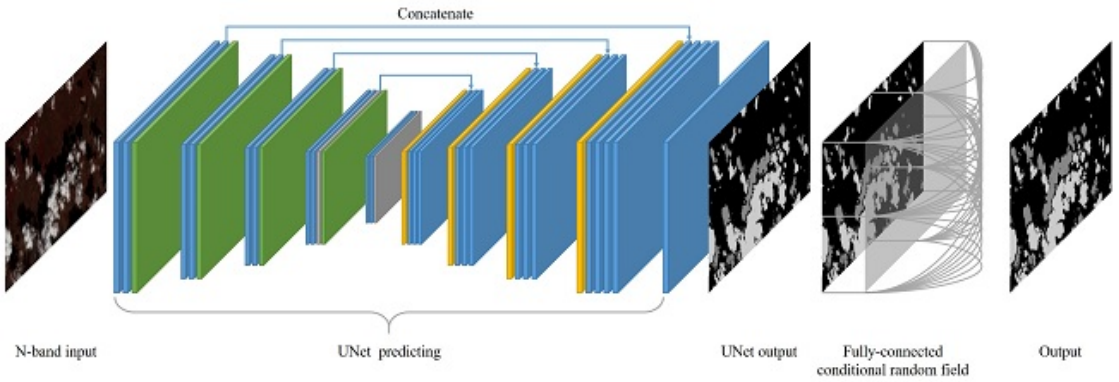

Figure 6.2: Example U-net architecture with conditional random field (CRF). Figure from Jiao L, et. al. (2020).

6.3 Network / Model

We will be using a UNet model, which has a conv-net encoder, and an attention-based decoder. The encoder portion will include drop-out, which has been shown to yield models with probabilistic outputs coherent with bayesian models (Gal et al. (2016)), allowing us to estimate the predictive uncertainty of our models. Further, we will be using a Focal Loss as replacement for a plain entropy-based loss-function, as this ensures the model trains better on the outliers of its current performance (Lin et al. (2018)). UNet based models have been developing since 2015, and remain a standard within all forms of semantic segmentation, particularly within the medical and remote sensing fields.

The models remain competitive with Vision Transformers and other methods used today in remote sensing, and are still used to benchmark new approaches or architectures. Due to their residual CNN structure they provide a gradual receptive field that is reminiscent of how our own vision works, where more weight is placed in the immediate proximity of a potential feature.

6.4 Results

Resulting models will be reported in an aggregate format, where a statistical performance is measured based on multiple runs to provide some insight into the stability of the model. For this we will instantiate each model used multiple times, and perform a full training run to estimate both the mean and variance of the performance. Final results will be reported both as confusion matrices and F1 scores (macro-averaged), which should provide a good indicator of performance over segmentation / classification tasks.

7 Sea ice drift algorithm

The SAGE project requires a sea-ice drift product with as long a time series as possible, and a dedicated R&D work is therefore required to extend the EUMETSAT OSI SAF global low-resolution sea-ice drift Climate Data Record (CDR) v1 product (OSI-455) back as far as possible in time from 1991.

7.1 Input satellite and ancillary data

7.1.1 Satellite data

The original Climate Data Record (CDR) v1, stretching from 1991 to 2020 inclusive, used data from SSMI-F10 (beginning January 1991), SSMI-F11, SSMI-F13, SSMI-F14, SSMI-F15, SSMIS-F16, SSMIS-F17, SSMIS-f18, AMSR-E and AMSR2-GW1. In order to extend back to 1978, data from SMMR-NI07 (October 1978 - August 1987) and SSMI-F08 (September 1987 - December 1991) are considered. See Section 3.1.3 for more information about the input data. Due to the limitations of the input data to be used for the CDR back-extension (SMMR-NI07 lacking the near-90 GHz channel, and SSMI-F08 having an early failure in these channels), these data were omitted from the original CDR, and must be processed carefully, focussing on the lower-frequency 37 GHz channels.

7.1.2 Ancillary data

Sea-ice concentration is used for ice/ocean/landmasks (see Section 3.3), and the Free Drift model uses the ERA5 wind reanalysis (10m u- and v- components of wind - see Section 3.4.1 for details about the ERA5 reanalysis).

The Free Drift Model gap-filling retrievals are dependent on a training dataset of drift vectors in the summer period derived from AMSR-E and AMSR2-GW1 data with the Continuous Maximum Cross-Correlation method. See input data Section 3.1.2 for information about AMSR-E and AMSR2. The training data is retrieved from the 36.5 GHz imagery (V and H polarisations) as it offers the best compromise between retrieval accuracy and spatial resolution (7 x 12 km). The 89 GHz channels have higher resolution, but are also more affected by atmospheric liquid water path and surface melting, and are consequently not used.

There has not as yet emerged a solid methodology to retrieve uncertainties from the Continuous Maximum Cross-Correlation algorithm for retrieval of sea-ice drift, and instead the common method used in the community is to validate the product against in-situ data from buoys and derive uncertainties from such validation statistics, scaled with the duration of the drift vector. Several different programs have on-ice drifting buoys which have data in the early period of the CDR, including:

- The International Arctic Buoy Programme (IABP) was initiated by the Polar Science Center (PSC), Applied Physics Laboratory at the University of Washington and began its operations in early 1979. It is now the reference programme for exchange of Arctic Ocean buoy trajectories (as well as air temperature and pressure data). The IABP itself does not deploy buoys, but receives quality-controls, formats and re-distributes buoy trajectories from many other institutions. The archive of 3-hourly trajectories covering 1979 to 2016 was accessed.

- The Atlas of Antarctic Sea Ice Motion that was developed at the Karlsruhe Institute of Technology (KIT) in 2004 and covers the period 1979 to 2000. Some of these buoys originate from the International Programme for Antarctic Buoys (IPAB). The Atlas is now hosted by AWI (https://data.meereisportal.de/eisatlas/).

7.2 Retrieval algorithms

7.2.1 Processing Chain

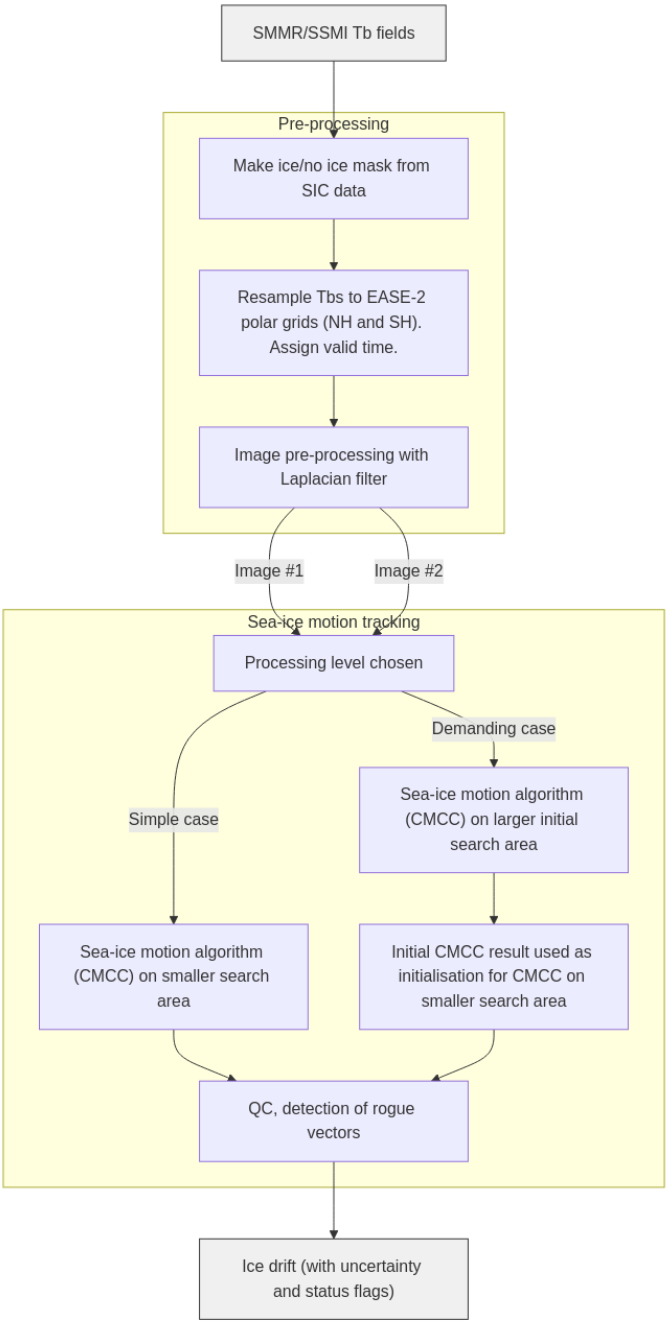

Figure 7.1: Flowchart of the ice drift retrieval.

7.2.2 Laplacian pre-processing

The remapped brightness temperatures are not used directly as input to the sea-ice motion tracking step. Instead, a Laplacian filter is first applied to the remapped imagery to enhance and stabilise intensity patterns. There are excellent ready-made functions providing Laplacian filters, for example `ndimage.laplace` from the `scipy` toolbox (Virtanen et al. (2020)). It is sufficient to use one of these standard functions, and therefore the mathematical description of this step is not described here.

7.2.3 Continuous Maximum Cross-Correlation

The algorithm to retrieve sea-ice drift operates on a pair of consecutive images. In the CDR, these images are consecutive daily brightness temperature maps, so the time separation is 1 day. In the case of the early SMMR-NI07 data, these data are only every second day, and the time separation becomes 2 days.

The first step of the algorithm is to retrieve the locations where the CMCC algorithm will be applied. For each pixel, there is a sub-image or block used for the algorithm, defined by the radius within which the algorithm operates. The pixels within these sub-images in both the start and end images must be fully over sea ice, and must not contain missing data, for that pixel to be processed by the algorithm.

$L_0(x,y)[i]$ denotes the $[i]$th pixel of the start sub-image centred on point (x,y), where the coordinates (x,y) are defined according to the image projection. N is the total number of pixels in the sub-image, and depends on the input diameter of the sub-image. While L_0 refers to the start image, L_1 refers to the end image, and $L_1(u,v)$ to the end sub-image centred at (u,v), where u and v can be fractional numbers, and can be also expressed with the shifts, i.e.

$$u = x + \delta x$$
$$v = y + \delta y \qquad (7.1)$$

The pixels of each candidate sub-image $L_1(u,v)$ used for comparison with the start sub-image are computed from bi-linear interpolations of the pixels of L_1 as follows:

$$\begin{aligned} L_1(u,v)[i] = {} & (1-\epsilon_u) \times (1-\epsilon_v) \times L_1(\bar{u},\bar{v})[i] \\ & + (1-\epsilon_u) \times \epsilon_v \times L_1(\bar{u},\bar{v}+s_v)[i] \\ & + \epsilon_u \times (1-\epsilon_v) \times L_1(\bar{u}+s_u,\bar{v})[i] \\ & + \epsilon_u \times \epsilon_v \times L_1(\bar{u}+s_u,\bar{v}+s_v)[i] \end{aligned} \qquad (7.2)$$

where

$$\begin{aligned} \overline{u} &= Trunc(u) \\ \epsilon_u &= |u - \overline{u}| \\ s_u &= \frac{u}{|u|} \end{aligned} \tag{7.3}$$

and the equivalent for v.
The match between a start and end image can then be evaluated by the correlation metric

$$\rho(x, y, \delta x, \delta y) = \frac{1}{\sigma(L_0(x,y))\sigma(L_1(x+\delta x, y+\delta y))} \sum_{i=1}^{N} (L_0(x,y)[i] - \langle L_0(x,y)\rangle) \cdot (L_1(x+\delta x, y+\delta y)[i] - \langle L_1(x+\delta x, y+\delta y)\rangle) \tag{7.4}$$

where the mean of the sub-image is

$$\langle L_0(x,y)\rangle = \frac{1}{N}\sum_{i=1}^{N} L_0(x,y)[i] \tag{7.5}$$

and the standard deviation of the sub-image is

$$\sigma(L_0(x,y)) = \sqrt{(\langle L_0^2(x,y)\rangle - \langle L_0(x,y)\rangle^2)} \tag{7.6}$$

The correlation metric $\rho(x, y, \delta x, \delta y)$ takes values between -1 and +1, where high values indicate a good match between the sub-images, indicating that the drift vector $(\delta x, \delta y)$ from this retrieval is optimal.

This can be expressed mathematically as

$$max\{(x,y) \in D\}\ \rho(x, y, \delta x, \delta y) \tag{7.7}$$

where D is a validity domain for $(\delta x, \delta y)$. This maximisation problem is solved at all grid positions where the motion vector is searched for, each optimisation being conducted independently from the others.

So far, this methodology has described the cross-correlation between a single pair of images. However, multiple image pairs may be available, taking into account different frequencies (37 GHz and 85 GHz) and different polarisations (V and H). Following Lavergne et al. (2010), an implicit merging of the imaging channels selected is implemented by maximising a sum of the cross-correlation functions:

$$max\{(x,y) \in D\}\ \frac{1}{M}\sum_{c=1}^{M} \rho^c(x, y, \delta x, \delta y) \tag{7.8}$$

where M is the number of channels.

This maximisation problem is solved by the Nelder-Mead algorithm (Nelder & Mead (1968)). This algorithm is simple to implement, and has good convergence and computational properties in problems with low dimensionality (Lagarius et al. (1998)). First, starting points for the optimisation are sampled on a regular

length-angle grid centred around $(0, 0)$. The length increment here is 10 km and the angle increment is 45 °. The circle has radius R, the maximum drift distance, which defines the validity domain D.

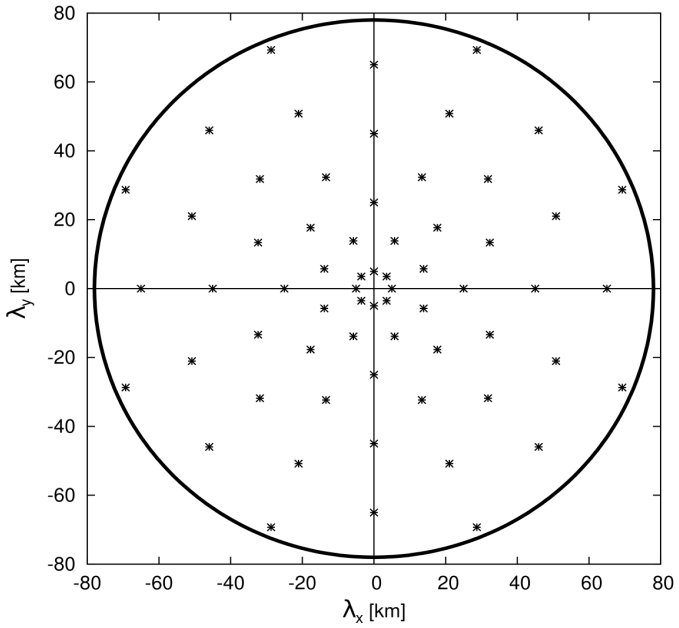

Figure 7.2: Initialisation points for the Nelder-Mead algorithm, chosen on a regular length-angle grid around point (0, 0).

The correlation metric $\rho(x, y, \delta x, \delta y)$ is calculated at each of these initialisation points, and the best three vertices are kept for optimising the algorithm. Termination and convergence is tested via a relative difference of function values at the current best and worst vertices, f_b and f_w. The condition for convergence is

$$\left|f_b - f_w\right| < (f_b + f_w) \times \tau + \epsilon \qquad (7.9)$$

where τ and ϵ are small, positive values, acting as thresholds to determine when the solution has stabilised. Once the difference between the best and worst of the three vertices is small, the solution is determined to have converged. The maximum number of iterations is set to 1000.

The search region is restricted by a-priori knowledge to a disc D of radius R centred on (xc, yc), to prevent the algorithm for exploring too large an area. The condition for the drift vector $(\delta x, \delta y)$ to belong to domain D is given by

$$(\delta x, \delta y) \in D_{xc,yc} \Leftrightarrow d(xc, yc, \delta x, \delta y) < R \qquad (7.10)$$

where $d(xc, yc, \delta x, \delta y)$ is the distance between the centre of the disc and the tip of the drift vector. Although (xc, yc) is the centre of the domain D, this does not mean that drift vectors will be preferentially retrieved close to $(0, 0)$, instead the whole domain D is explored. R is computed as the maximum expected

speed of the ice (v_{max}) multiplied by the time separation of the images, $T_1 - T_0$, and so R is the maximum expected distance that could be covered. Here a v_{max} of 0.45 m s^{-1} is used.

The domain constraint cannot be used as is, since this results in abrupt and nonlinear behaviour, and so instead it is implemented as a soft constraint based on a mono-dimensional sigmoid function as

$$W(d) = \frac{1}{1+e^{k(d-R)}} \quad (7.11)$$

where k is a parameter controlling the steepness of the sigmoid around the cut-off radius R. By construction, $W(R) = 0.5$. By using a large enough value for k, the W can be made arbitrarily close to the Heaviside step function, yet remaining smooth and continuous.

This penalty is applied to the correlation function as

$$\rho^{c,D}(x, y, \delta x, \delta y) = (\rho(x, y, \delta x, \delta y) + 1) \times W(d(x, y, \delta x, \delta y)) - 1 \quad (7.12)$$

where $\rho^{c,D}$ is the penalised correlation function. Solving for the maximum of $\rho^{c,D}$ is taken as a proxy for solving the original, constrained, optimisation problem, and so now

$$max\{(x, y) \in D\} \frac{1}{M} \sum_{c=1}^{M} \rho^{c,D}(x, y, \delta x, \delta y) \quad (7.13)$$

is used in the Nelder-Mead algorithm.

Figure 7.3 shows a mono-dimensional example of applying a sigmoid penalty function to a synthetic correlation function. Evaluations for x lower than L are dominated by the correlation value $\rho(x)$ while those occurring outside the domain (x larger than L) return very bad scores of close to -1.

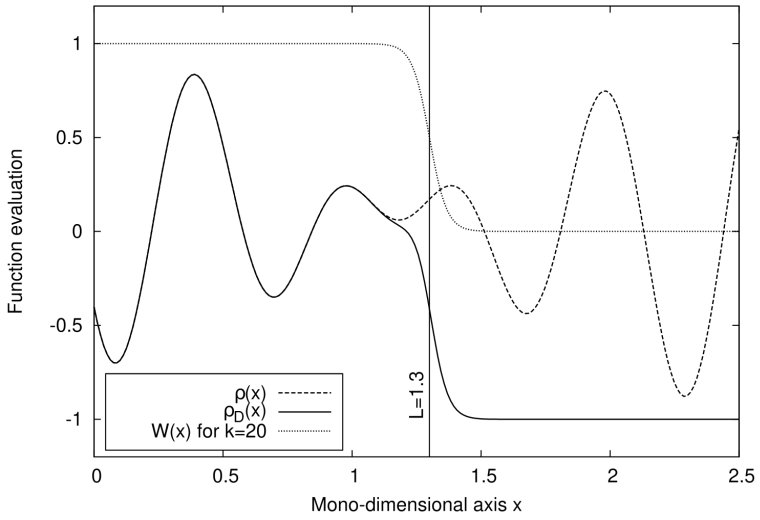

Figure 7.3: Example soft constraint implemented with a sigmoid penalisation function W and the application to a synthetic, mono-dimensional correlation signal ρ. Here, L = 1.3 and k = 20.

7.2.4 Detection and correction of rogue vectors

Once the CMCC algorithm has been applied, it is essential to apply a filtering step to detect and correct or remove erroneous ("rogue") vectors. Such erroneous vectors can arise due to convergence of the Nelder-Mead algorithm in a local maximum, noise in the sub-images, or edge effects in the sub-images.

The filtering step is based on the distance from individual displacement vectors to the average of its neighbouring vectors. If this distance is less than a fixed threshold, the displacement vector being tested is validated and another vector is tested. Otherwise, a new CMCC motion tracking optimisation is triggered, in which the Nelder-Mead algorithm is initialised and run like in the previous section, except that the validity domain D is altered to constrain the result for this vector.

Δ^{avg} is the distance between the tip of the current drift vector $(\delta x, \delta y)$ and the zonal average drift vector $(\delta x^{avg}, \delta y^{avg})$ calculated from the 8 closest neighbouring drift vectors. The local domain D then used for the new optimisation is then the disc with centre $(\delta x^{avg}, \delta y^{avg})$ and radius Δ_{max}^{avg}, where Δ_{max}^{avg} is set to 10 km. Neighbouring vectors with a maximum correlation value of less than 0.5 are not used, to avoid degrading the average drift field with possibly wrong estimates.

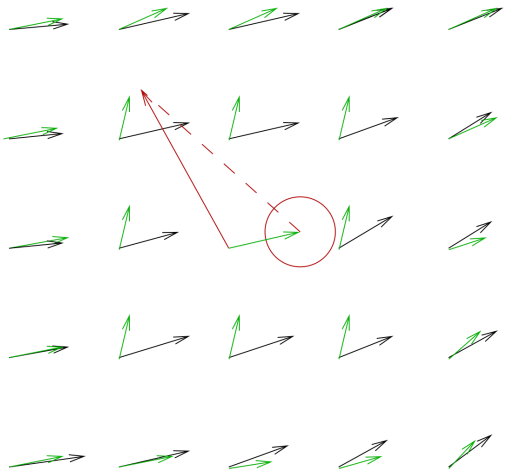

Figure 7.4: Example case where there is an erroneous drift vector (red) in the initial field of vectors (black) from the CMCC algorithm. The locally-averaged vector field is plotted in green. Δ^{avg} is the length of the dashed red line. The red disc shows the validity domain D that is used to re-optimise the erroneous vector, with centre $(\delta x^{avg}, \delta y^{avg})$ and radius Δ_{max}^{avg} (10 km).

If a satisfactory maximum correlation is found during this second CMCC optimisation limited to the new domain, it is kept and the surrounding average vectors are immediately updated, as well as their lengths. If the constrained optimisation does not converge or if the new vector does not have a good enough maximum correlation value, both the old and new vectors are discarded, and the average vectors at the neighbouring locations are updated.

This method works well when isolated vectors are erroneous, but can sometimes fail when several erroneous vectors are in close proximity. This happens especially when noise dominates the signal in a large region of one of the images. In these cases, the order in which the vectors are corrected has an influence on the final efficiency for the filtering.

To minimise this influence, the vectors are first sorted from the largest to the shortest Δ^{avg} and the filtering is first applied to the vector exhibiting the longest Δ^{avg}. Since changing a vector has an influence on its direct neighbours, the sorting is repeated after each correction. A mechanism is used to avoid the computation entering an infinite loop. This strategy ensures that the good vectors around an erroneous initial vector are not modified before the rogue vector is processed through the filter.

In the case where the new optimisation does not lead to an acceptable CMCC value (if the value is below the threshold or if the algorithm does not converge with the new constraint), the vector position is recorded as non-feasible in the status flags, and the fields of drift vectors get a fill value. This processing is done until all vectors are either validated, corrected, or flagged as non-feasible.

7.2.5 Use of a 2-step CMCC

Since the earlier SSMI data and the SMMR-NI07 data is coarser resolution and lower quality than the data from 1991 onwards used in the currently existing CDR, a refinement is considered for improving the CMCC algorithm. When defining the size of the search radius for the CMCC process, there is a trade-off in that the bigger the search domain, the more smoothing there is in the resulting initial vector field, as each vector is based on more of the same data that the neighbouring pixels are based on. Therefore, using a smaller search domain allows retrieval of smaller scale features in the resulting vector field, but it also results in more rogue vectors that need to be corrected. Therefore, picking the search radius is a compromise between field smoothing and number of erroneous vectors.

A 2-step CMCC configuration would use a broader matching window to retrieve the drift more accurately but without the smaller-scale details. There would then be a second round of CMCC with a smaller search radius centred on the tip of the vector retrieved from the first step, which may improve the resolution of the retrieved field without adding excess rogue vectors. This 2-step CMCC is currently in R&D and will be further described in future versions of this ATBD, if selected.

7.2.6 Free-drift Model based on wind

In the summer period, the sea ice drift from satellite motion tracking is less reliable due to surface melting and increased liquid water and vapour content in the atmosphere. The free-drift method is an alternative measure of sea-ice drift from a theoretical model based on wind speed when the stresses within the sea ice are neglected. This method is less valid in the winter season, as the internal sea-ice stresses can be too high. It can also be less reliable in case of low wind speeds.

Table 7.1: Definition of symbols.

Symbol	*Definition*
h	sea ice thickness
ρ	sea ice density
$\vec{u}$	sea ice velocity
t	time
f	Coriolis parameter
$\vec{k}$	vertical unit vector
$\nabla \cdot \vec{\sigma}$	internal sea ice stress
$\vec{\tau}_a$	wind stress
$\vec{\tau}_w$	water stress

g	gravitational acceleration
η	sea surface elevation
p_a	atmospheric pressure
U_{wg}	geostrophic water velocity
U_a	surface wind velocity
ρ_a	air density
ρ_w	water density
C_a	air drag coefficient
C_w	water drag coefficient
θ_a	air turning angle
θ_w	water turning angle

The free drift model of ice drift is based on the 2-dimensional equation of motion of sea ice when internal stresses are neglected (Mozo (2017), Leppäranta (2011)). The symbols used are defined in Table 7.1.
The 3-dimensional equation of motion integrated through the thickness of the ice to obtain a 2-dimensional motion equation is

$$h\rho((\partial\vec{u}/\partial t) + (\vec{u}\cdot\nabla\vec{u}) + (f\vec{k}\times\vec{u})) = \nabla\cdot\vec{\sigma} + \vec{\tau}_a + \vec{\tau}_w - h\rho g\nabla\eta - h\nabla p_a \quad (7.14)$$

The physical quantities that the parts of this equation represent are summarised in Table 7.2.

Table 7.2: Breakdown of the equation of motion of sea ice.

Equation part	*Representing*
$h\rho(\partial\vec{u}/\partial t)$	Local acceleration
$h\rho(\vec{u}\cdot\nabla\vec{u})$	Advective acceleration
$h\rho(f\vec{k}\times\vec{u})$	Coriolis acceleration
$\nabla\cdot\vec{\sigma}$	Internal sea ice stress
$\vec{\tau}_a$	Wind stress
$\vec{\tau}_w$	Water stress
$h\rho g\nabla\eta$	Sea surface gradient
$h\nabla p_a$	Air pressure

The local and advective accelerations and the air pressure are negligible. In the free drift model, which is a good approximation for individual ice floes and non-compact ice fields, the internal sea-ice stress is set to 0. With the further approximation that the Coriolis and sea surface gradient term can be neglected, which is valid in the case of moderate to high surface winds, this reduces simply to

$$\vec{\tau}_a + \vec{\tau}_w \approx 0 \qquad (7.15)$$

The Coriolis effect causes the stress to not act in the direction of the wind or the water current. This effect can be parameterised by turning angles, so the wind stress can be written as

$$\vec{\tau}_a = \rho_a C_a |U_a| e^{i\theta a} \vec{U}_a \qquad (7.16)$$

Due to the stronger coupling of the water and ice, the water stress is calculated using the difference between the geostrophic current and the sea-ice velocity. The water stress also includes a turning angle and can be expressed as

$$\vec{\tau}_w = \rho_w C_w |U_{wg} - u| e^{i\theta w} (\vec{U}_{wg} - \vec{u})$$

With the addition of an error term $\vec{e}$, the sea-ice drift can be expressed by

$$\vec{u} = A\vec{U}_a + \vec{U}_{wg} + \vec{e} \quad \text{where} \quad A = \sqrt{(\rho_a C_a / \rho_w C_w)}\, e^{-i(\theta w - \theta a)} \qquad (7.17)$$

For ease of use, this can be expressed as a matrix

$$\vec{u} = G m_e \quad + \vec{e} \quad \text{where} \quad G = [\vec{U}_a \quad 1] \quad \text{and} \quad m_e = \begin{bmatrix} A \\ \vec{U}_{wg} \end{bmatrix} \qquad (7.18)$$

Both sea ice drift and wind observations are used to solve for the model parameters m_e in an inverse problem (parameter tuning). A least squares estimation to this problem where $\vec{e}$ is zero (Menke and Menke (2012)) is

$$m_e = [G^T G]^{-1} G^T \vec{u} \quad (7.19)$$

where G^T is the transpose of G.

The error term e which provides information on the validity of the model is solved as

$$e = (I - G[G^T G]^{-1} G^T) \vec{u} \quad (7.20)$$

where I is the identity matrix.

In order to generate monthly maps of wind-model parameters, sea-ice drift data computed with the CMCC algorithm from the 37 GHz imagery of the AMSR-E and AMSR2 sensors for spring-summer-autumn periods, as well as ERA5 wind data are used. The tuning of the free-drift model thus covers only the period 2002-2020, since the near-90 GHz imagery of the SSM/I and SSMIS missions do not have high enough quality retrievals in the summer periods to be included (due to surface melt and atmospheric wetness), and neither do their 37 or 19 GHz imagery (too coarse resolution).

For implementation reasons, parameter tuning is solved using complex numbers, with the x-component of the vectors being the real part of the complex number and the y-component of the parameter being the imaginary part. A series of masking steps then take place. First, the parameter fields are masked with a landmask and a monthly climatological maximum extent mask. Second, a maximum sea-ice drift mask is applied. It defines grid cells and regions within the climatological sea-ice cover where drift vectors are never calculated from satellites, typically close to coasts. This mask is created from accumulating all the drift vectors from the single-sensor satellite products over the climate data record 1991-2020 and defines the maximum extent of the satellite-based drift vectors.

Then, a gap-filling by extrapolation step is applied. This extrapolation is necessary because the wind-driven model is tuned on the later part of the timeseries (AMSR-E and AMSR2) but the model is applied to the earlier portion of the CDR as well, when sea ice covered more area in the Arctic. The grid cells that neither have inverted parameters, land, the ”no-ice” region or the ”no-drift” region are gap-filled by extrapolation.
Finally, the parameters are smoothed using a Gaussian weighting resampling method with a smaller radius to remove small speckle. These maps of parameters and residuals are saved in monthly parameter files.

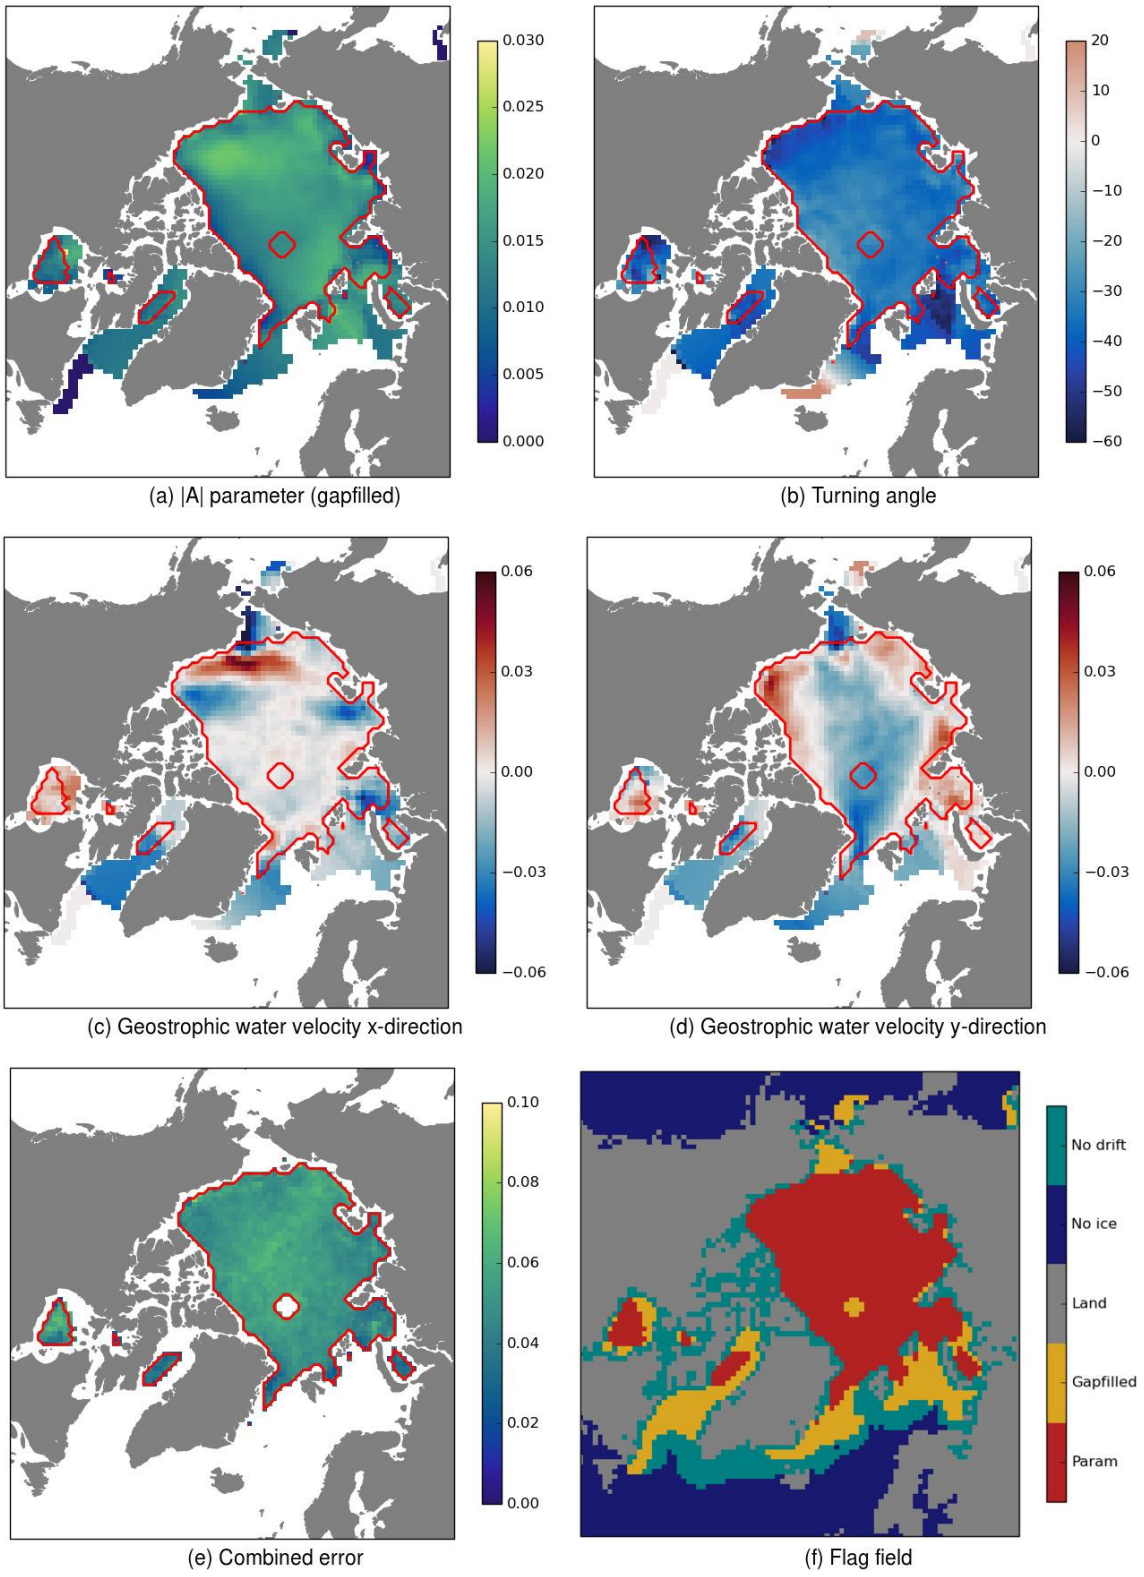

Figure 7.5: Example of parameter fields, July in the Northern hemisphere. (a) |A| (b) θ (c) U_{wgx} (d) U_{wgx} (e) Mean of e_x and e_y errors (f) flag field. The red contours enclose the non-gapfilled region.

Using the calculated model parameter fields, the sea-ice drift is then calculated directly from wind velocity data. For each date, two ice drift maps are calculated from the ERA5 wind field, using the parameter files from the two bracketing months. The two ice drift maps are combined in a weighted mean, according to the day of the month - a linear interpolation is used to combine the two ice drift maps in proportion to the time period between the current date and the two bracketing months. This linear interpolation is to avoid sudden changes in sea-ice drift vectors at the transition between months.

The final step in the wind-drift calculation is to apply the sea-ice mask for the current date, to ensure that wind-driven vectors are only reported over sea ice. The ice mask is created from the daily single-sensor satellite drift products from the same day.

7.3 Optimisations

7.3.1 Sensor-specific optimisations

This will be described in future versions of the ATBD.

7.3.2 Consistency between products from different sensors

This will be described in future versions of the ATBD.

7.3.3 Data merging

This will be described in future versions of the ATBD.

7.4 Estimation of uncertainties

See the E3UB document in which the methodology is described. Briefly, this follows the same method as used for other sea-ice drift datasets at the moment in comparing the drift vectors to the motion of on-ice drifting buoys to retrieve an uncertainty. The retrieved drift vectors and buoy vectors are collocated in time and space and then the vectors are directly compared to obtain the standard deviation and bias.

7.5 Limitations of the product

The parts of the record including all of the summer seasons which are gap-filled by the ERA5 wind drift vectors should not be used for trend analysis. This is because this model was trained on AMSR-E and AMSR2-GW1 data from 2002-2020, and therefore was trained only on a small part of the CDR and not on the full extent. Moreover, the result of this training was monthly wind files averaged from this entire period, and therefore no time-dependency of the drift is retained from the training data.

7.6 Early version of the back-extension

Figure 7.6 shows early tests of the sea-ice drift fields obtained from SSMI-F08 data from both 37 GHz data and 85 GHz data in 1989, thus two years earlier than the start of the OSI SAF CDR v1. The figure documents that processing the 37 GHz imagery results in much smoother vector fields than when processing 85 GHz, as expected. The 85 GHz data can capture smaller-scale features, but is rarely able to retrieve a full vector field.

It should be noted that these days are picked as good examples, and that there are many days when a lot fewer, or no vectors at all, are retrieved.

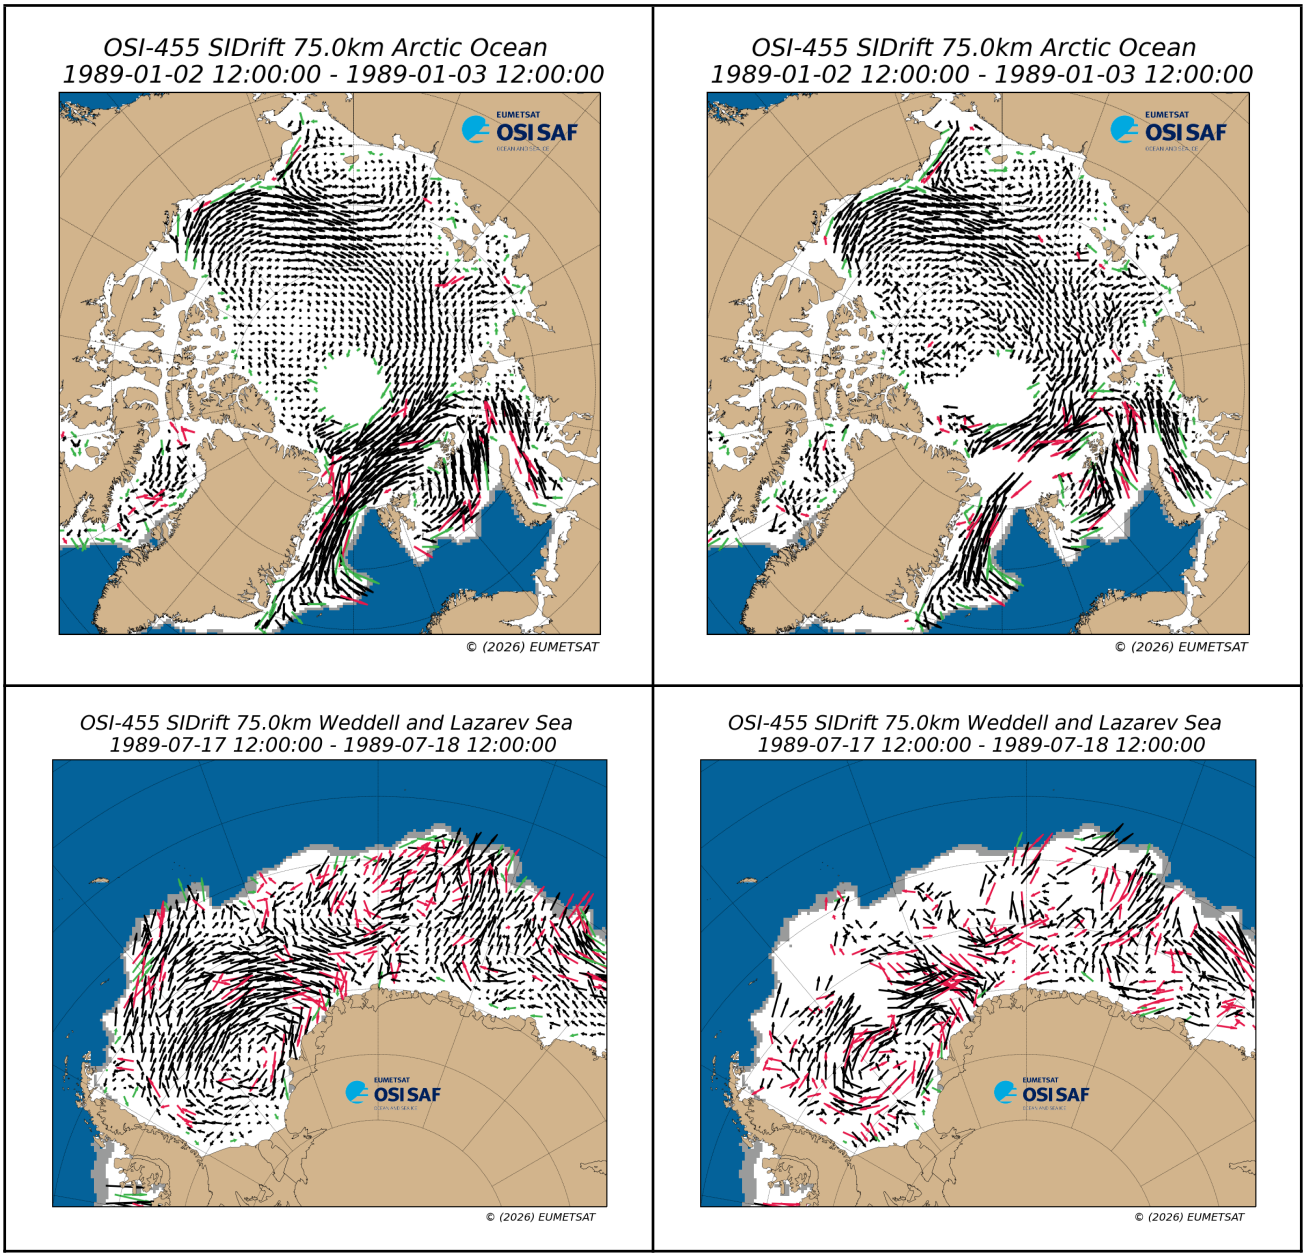

Figure 7.6: Example fields of ice drift from SSMI-F08, showing 37GHz (left) and 85 GHz (right) in combined V- and H-polarisations in 1989. Shows both Northern Hemisphere (top) and Southern Hemisphere Weddell Sea region (bottom).

Figure 7.7 shows early tests of the sea-ice drift fields obtained from SMMR-NI07 data from 37 GHz data. Here the much larger “pole hole” gap around the North Pole can be seen. Again, these days are among the best examples, there are many more days with few to zero vectors retrieved.

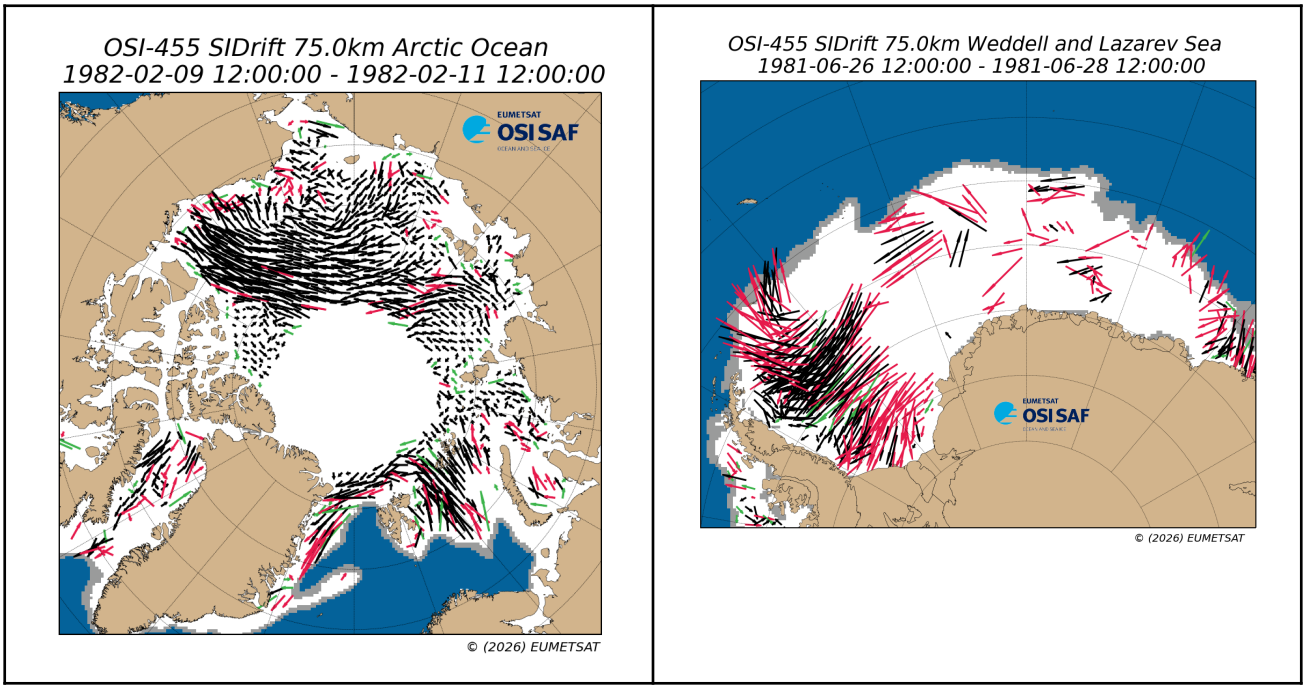

Figure 7.7: Example fields of ice drift from SMMR-NI07, showing 37GHz (left) in combined V- and H-polarisations in 1982 Northern Hemisphere (left) and 1981 Southern Hemisphere Weddell Sea region (right).

A time series of the number of vectors per instrument, frequency and polarisation is shown in Figure 7.8. The seasonal effect of the higher atmospheric water and increased surface water during the summer can be seen, with the number of vectors for most instrument setups dropping to near zero. There are consistently fewer vectors for SMMR-NI07 than the SSMI instruments. For SSMI-F08, the comparison of the purple time series (85 GHz) compared to the blue time series (37 GHz) visualises the failure of the 85 GHz channels, especially the V polarisation.

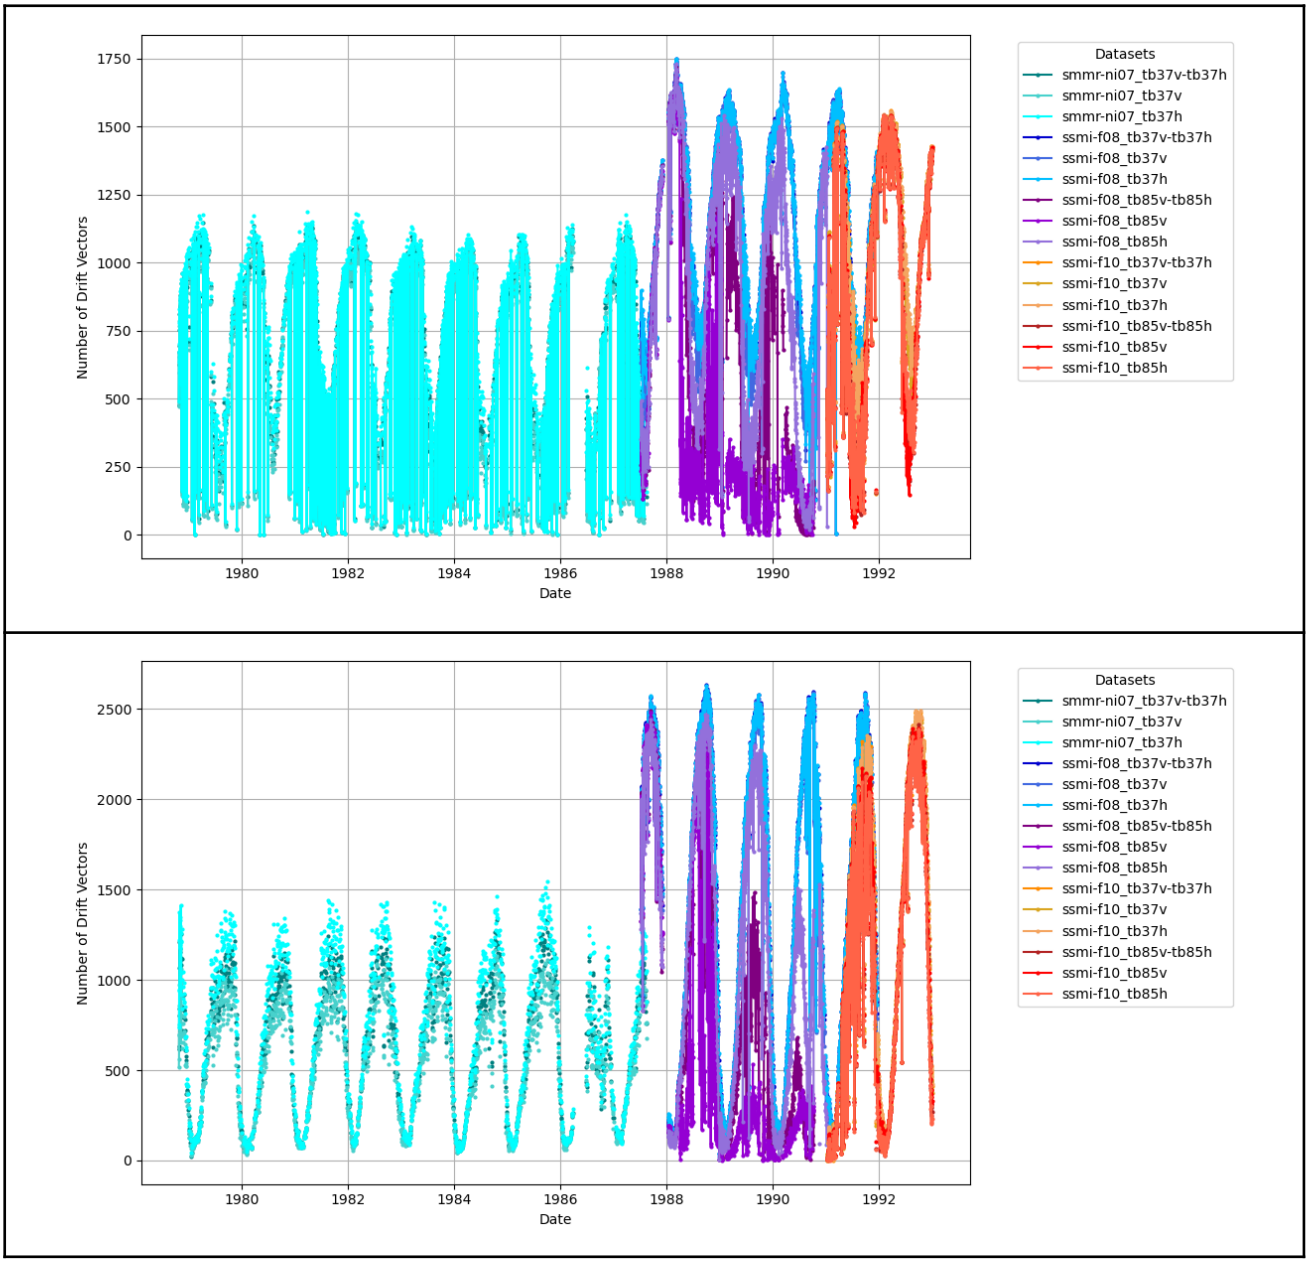

Figure 7.8: Time series of the total number of drift vectors available for the various frequencies and polarisations of instruments SMMR-NI07, SSMI-F08 and SSMI-F10. The top panel shows the Northern Hemisphere and the lower panel the Southern Hemisphere.

Early results from the first validation are shown in the E3UB document [RD-3].

8 Blended SAGE sea ice age product

The development of a blended sea-ice age product is the central objective of the SAGE project. Building on the experience gained from the PVASR intercomparison and the RRDP framework, the SAGE team will now initiate a dedicated research and development phase to design, test, and evaluate approaches to combine complementary information from microwave sea-ice type retrievals, Lagrangian tracking methods, and machine learning techniques.

At this stage, no single approach has been selected. Instead, multiple candidate methodologies are under consideration, reflecting different perspectives and the strengths of the individual algorithm components. These approaches will be further developed, tested, and compared in an extended RRDP framework.

The following sections outline an initial candidate approach for blending, intended as a first proof of concept. It should be understood as an exploratory direction rather than a final algorithm design.

8.1 Motivation

Microwave-interpreted sea-ice type algorithms and Lagrangian sea-ice age approaches provide complementary perspectives on the state and evolution of sea ice. As discussed in Section 2.3, the definition and nomenclature of sea-ice classes are themselves non-trivial, reflecting both observable properties and thermodynamic history, and highlighting the complexity of consistently describing sea ice across different methodologies.

Microwave-based retrievals offer high spatial resolution, comparable to that of passive and active microwave satellite data, and are directly linked to observable physical properties of the ice surface, enabling detailed mapping of ice-type distributions, typically into the main categories of FYI and MYI. However, these methods are limited in their ability to resolve multiple age classes and can be affected by atmospheric conditions, surface melt, and noise in the satellite signal. During summer months or under certain weather conditions (e.g., too moist atmosphere, short local warming events, etc.) this may lead to false classification if no independent data sources are included (see Section 4.6 on post-processing correction schemes).

In contrast, the Lagrangian sea-ice age approach derives ice age by tracking ice parcels over time, providing a physically consistent description of the ice's thermodynamic history and enabling the distinction of multiple age classes and their individual fractional concentration. This approach is largely independent of short-term atmospheric variability but depends on the quality and resolution of the input drift and concentration fields, which can lead to reduced spatial detail and smoothing of fine-scale features. Due to the low resolution of the sea ice drift data used as input to the Lagrangian algorithm and inherent diffusion during remeshing, the resulting spatial distribution of ice age has also low resolution.

The goal of blending the microwave and Lagrangian algorithms is to exploit their respective strengths: to enhance the spatial resolution and observational sensitivity of microwave-based methods while avoiding noise from input data and preserving tracking of many ice age fractions.

One candidate approach currently being considered for blending the two types of algorithms is assimilation of brightness temperatures into the data-driven model of sea-ice age. As described below, the assimilation process can easily control the amount of information which is ingested from satellite data into the ice age tracking system and balance between adding noise or enhancing resolution.

8.2 Algorithm overview

The Lagrangian algorithm moves nodes of a triangular mesh resulting in motion and deformation of the mesh elements and changes of concentration of ice fractions in these elements due to divergence. As such, it can be seen as a numerical model, which is forced by the observed sea ice drift and concentration.

There exist many methods of data assimilation (DA), which blend observations and outputs of numerical models to improve the model state for future simulations. In case of atmospheric or ocean or sea ice models, satellite observations are often used in DA to improve operational forecasts or hindcasts (also called re-analyses). We developed a DA method to utilise data from passive and active microwave satellite observations to improve the resolution of the sea ice age product.

Conceptually, the algorithm works as follows:

1. Concentration of multi-year ice is advected for one day as described in Section 5 (forecasted concentration, C_F).
2. The advected MYI concentration is collocated with all available observations of brightness temperature and backscatter for this data on the same grid (observed brightness temperatures, T_O).
3. An inverse (i.e., from satellite observations to geophysical field) machine learning probabilistic model is trained on a fraction of the collocated data and is evaluated on the independent fraction of the collocated data (ML model, M = Training(C_F, T_O)).
4. The ML model is applied to satellite data to predict a field of purely satellite observation based MYI. The ML model also computes the uncertainty of the predicted field (predicted concentration and uncertainty, C_P, U_P = M(T_O)).
5. The forecasted and predicted ML concentrations are blended using uncertainty as a weighting factor, meaning that the more uncertain the ML prediction, the less its contribution (analysed concentration, $C_A = C_F\ U_P + C_P\ (1 - U_P)$)
6. The analysed MYI concentration is used further in the sea ice age computations and for advection into the next day.

8.3 Some implementation details

First test studies of this blending approach have been carried out. As this methodology is still under active development, at present we do not go into implementation details of each step but only present the key solutions we found.

8.3.1 The ML model and training

We use Random Forest Regressor (RF) as the backbone of our ML model. Configuration of RF allows to control the level of generalisation and accuracy. The following configuration was found to have a good balance of generalisation and accuracy:

- The number of trees in the forest: 20
- The maximum depth of the tree: 20
- The minimum number of samples required to split an internal node: 5

Each tree in the RF is trained with only 10% of randomly sampled input data.

8.3.2 Computation of uncertainty

40% of independent input data is used for evaluation of each tree independently using the r^2 score.
During inference, the outputs of the trees are averaged (C_{AVG}) using the r^2 score as weight and a standard deviation of the prediction (C_{STD}) is computed for each predicted concentration. Variance of the prediction (C_{VAR}) is computed as a ratio of standard deviation to the mean and the uncertainty is computed by scaling the variance: $U_P = \alpha + \beta$ CVAR, where α and β are coefficients found in sensitivity experiments (see below).

8.3.3 Detection and reduction of noise in input data

Strictly speaking, the microwave signal detected by a satellite sensor consists of signals from sea ice, ocean, atmosphere, galaxy and noise from the electronics onboard. However, for our application only the signal from sea ice is considered useful, therefore the other components are grouped into noise. The higher the noise (i.e., contribution from other, non useful signal sources) in the input data, the larger are differences in the predictions by the individual trees in the random forest. In addition, the prediction from each tree has lower accuracy. Therefore, the standard deviation (or rather the variance, i.e. STD normalised by the averaged signal) and the accuracy (estimated using r^2 score) can serve as a measure of the noise either in individual pixels or overall.

As described before, to account and reduce noise in individual pixels we compute the standard deviation, variance and uncertainty and use when blending with the forecasted MYI concentration. In order to reduce the impact of trees that were trained on noisy data, we use the r^2 score to weight the prediction when computing the average from the random forest. In addition, we compute an average of the r^2 scores from all trees and use that to scale the variance and reduce the weight of observations when the quality of retrieval is low: $C_{VAR} = (C_{STD} / S_{AVG}) / <r^2>^2$, where $<r^2>$ is the averaged r^2 score.

8.3.4 Sensitivity experiments

Several parameters of the DA system were tested during the sensitivity experiments including configuration of the RF regressor, α and β coefficients, formulas for computing the variance, etc. In the sensitivity experiments we used the passive and active microwave data from RRDP for assimilation and the manual ice charts for evaluation. The experiments run from Sep 2019 to Dec 2024, when a sufficient number of observations are available.

9 References

Aaboe, S., Down, E.J., Sørensen, A., Lavergne, T., Eastwood, S. (2025). Sea Ice Edge and Sea Ice Type Climate Data Records, Version 4.0: Algorithm Theoretical Basis Document. Copernicus Climate Change Service. Document ref. WP2-FDDP-SIET-v4.0_C3S2-312b_ATBD-of-v4.0-SeaIceEdgeType-products_v4.1. Accessed from C3S CDS DOI: 10.24381/cds.29c46d83

Brodzik, M. J. and Knowles, K.: EASE-Grid 2.0 Land-Ocean-Coastline-Ice Masks Derived from Boston University MODIS/Terra Land Cover Data, Version 1, https://doi.org/10.5067/VY2JQZL9J8AQ, accessed: 2025-06-02, 2011.

Cavalieri, D. J., Gloersen, P., & Campbell, W. J. (1984). Determination of sea ice parameters with the NIMBUS 7 SMMR. *Journal of Geophysical Research*, 89(D4), 5355–5369. https://doi.org/10.1029/JD089iD04p05355

Cavalieri, D. J., Parkinson, C. L., Gloersen, P., & Zwally, H. J. (1997). Arctic and Antarctic sea ice concentrations from multichannel passive-microwave satellite data sets: October 1978-September 1995 User's Guide (Tech. Rep. No. NASA-TM-104647). NASA Goddard Space Flight Center. Retrieved from https://ntrs.nasa.gov/citations/19980076134

Copernicus Climate Change Service, Climate Data Store, (2023): ERA5 hourly data on single levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). doi: 10.24381/cds.adbb2d47

EUMETSAT OSI SAF, Global sea ice concentration climate data record 1978-2020 (v3.1, 2025), OSI-450-a1, doi:10.15770/EUM_SAF_OSI_0023

EUMETSAT OSI SAF, Global Low Resolution Sea Ice Drift data record 1991-2020 (v1, 2022), OSI-455, doi: 10.15770/EUM_SAF_OSI_0012

Fennig, K., Schröder, M., Andersson, A., and Hollmann, R.: A Fundamental Climate Data Record of SMMR, SSM/I, and SSMIS brightness temperatures, Earth Syst. Sci. Data, 12, 647–681, https://doi.org/10.5194/essd-12-647-2020, 2020.

French, G., Oliver, A., & Salimans, T. (2020). Milking cowmask for semi-supervised image classification. arXiv preprint arXiv:2003.12022.

Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning (arXiv:1506.02142). arXiv. https://doi.org/10.48550/arXiv.1506.02142

Geuzaine, C. and Remacle, J.-F.: Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities, International Journal for Numerical Methods in Engineering, 79, 1309–1331, https://doi.org/10.1002/nme.2579, 2009.

Hendricks, S., & Paul, S. (2022). Product user guide and algorithm specification: AWI CryoSat-2 sea ice thickness (version 2.5). (Computer software manual). https://doi.org/10.5281/zenodo.7416200

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., et al. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 1999–2049. https://doi.org/10.1002/qj.3803

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. Neural Networks, 2(5), 359–366. https://doi.org/10.1016/0893-6080(89)90020-8

Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. Machine Learning, 110(3), 457–506. https://doi.org/10.1007/s10994-021-05946-3

Hwang, P. and T. Lavergne. Validation and comparison of OSI SAF low and medium resolution and IFREMER/Cersat sea ice drift products. 2010. EUMETSAT Associate & Visiting Scientist scheme report.

Jiang D, Li X, Zhang K, Marinsek S, Hong W, Wu Y. Automatic Supraglacial Lake Extraction in Greenland Using Sentinel-1 SAR Images and Attention-Based U-Net. Remote Sensing. 2022; 14(19):4998. https://doi.org/10.3390/rs14194998

Jiao L, Huo L, Hu C, Tang P. Refined UNet: UNet-Based Refinement Network for Cloud and Shadow Precise Segmentation. Remote Sensing. 2020; 12(12):2001. https://doi.org/10.3390/rs12122001

Kage, P., Rothenberger, J. C., Andreadis, P., & Diochnos, D. I. (2026). A Review of Pseudo-Labeling for Computer Vision. Journal of Artificial Intelligence Research, 85. https://doi.org/10.1613/jair.1.19656

Kern, S., Spreen, G., Kaleschke, L., De La Rosa, S., and Heygster, G. (2007). Polynya Signature Simulation Method polynya area in comparison to AMSR-E 89 GHz sea-ice concentrations in the Ross Sea and off the Adélie Coast, Antarctica, for 2002: first results, Ann. Glaciol., 46, 409–418, https://doi.org/10.3189/172756407782871585

Lagarias, J. C. , J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9:112–147, 1998.

Lavergne, T., S. Eastwood, Z. Teffah, H. Schyberg, and L.-A. Breivik. Sea ice motion from low resolution satellite sensors: an alternative method and its validation in the Arctic. *Journal of Geophysical Research*, 2010. doi:10.1029/2009JC005958.

Lee, S.-M., B.-J. Sohn, & H. Shi. Impact of ice surface and volume scatterings on the microwave sea ice apparent emissivity. Journal of Geophysical Research: Atmospheres, 123, 9220–9237. https://doi.org/10.1029/2018JD028688, 2018.

Leppäranta, M. The drift of sea ice. Springer-Praxis Books in Geophysical Sciences. Springer-Verlag, 2nd edition, 2011.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2018). Focal Loss for Dense Object Detection (arXiv:1708.02002). arXiv. https://doi.org/10.48550/arXiv.1708.02002

Markus, T. and Burns, B. A.(1995). A method to estimate subpixel-scale coastal polynyas with satellite passive microwave data, J. Geophys. Res.-Oceans, 100, 4473–4487, https://doi.org/10.1029/94JC02278

Mätzler, C., & Wiesmann, A. (1999). Extension of the microwave emission model of layered snowpacks to coarse-grained snow. *Remote Sensing of Environment*, 70(3), 317–325. https://doi.org/10.1016/s0034-4257(99)00047-4

Meissner, T., & Wentz, F. (2004). The complex dielectric constant of pure and sea water from microwave satellite observations. *IEEE Transactions on Geoscience and Remote Sensing*, 42(9), 1836–1849. https://doi.org/10.1109/TGRS.2004.831888

Meissner, T., & Wentz, F. J. (2012). The emissivity of the ocean surface between 6 and 90 GHz over a large range of wind speeds and Earth incidence angles. *IEEE Transactions on Geoscience and Remote Sensing*, 50(8), 3004–3026. https://doi.org/10.1109/tgrs.2011.2179662

Melsheimer, C., Spreen, G., Ye, Y., and Shokr, M. (2023). First results of Antarctic sea ice type retrieval from active and passive microwave remote sensing data. The Cryosphere, 17(1), 105–126. https://doi.org/10.5194/tc-17-105-2023

Melsheimer, C., Spreen, G. (2022). IUP Multiyear Ice Concentration and other sea ice types, Version 1.1 (Arctic)/Version AQ2 (Antarctic) User Guide. https://data.seaice.uni-bremen.de/MultiYearIce/MYIuserguide.pdf

Menke, W. and J. Menke. Environmental data analysis with MatLab. Elsevier, 2012.

Mozo, J.. Investigating the utility of a linear sea ice drift model for the Arctic. PhD thesis, University of Oslo. 2017.

Nelder, J. A. and R. Mead. A simplex method for function minimization. *Computational Journal*, 7:308–313, 1968.

Orvieto, A., & Gower, R. (2025). In Search of Adam's Secret Sauce (arXiv:2505.21829). arXiv. https://doi.org/10.48550/arXiv.2505.21829

Parkinson, C. L. (2019). A 40-y record reveals gradual Antarctic sea ice increases followed by decreases

at rates far exceeding the rates seen in the Arctic. Proceedings of the National Academy of Sciences, 116(29):14414–14423. doi:10.1073/pnas.1906556116.

Remund, Q., D. Long, and M. Drinkwater, "An iterative approach to multisensor sea ice classification," IEEE Trans. Geosci. Remote Sens., vol. 38, no. 4, pp. 1843–1856, Jul. 2000.

Rostosky, P., Spreen, G., Farrell, S. L., Frost, T., Heygster, G., & Melsheimer, C. (2018). Snow depth retrieval on Arctic sea ice from passive microwave radiometers — Improvements and extensions to multiyear ice using lower frequencies. *Journal of Geophysical Research: Oceans*, 123(10), 7120–7138. https://doi.org/10.1029/2018JC014028

Rückert, J. E., Huntemann, M., Tonboe, R. T., & Spreen, G. (2023). Modeling snow and ice microwave emissions in the Arctic for a multi-parameter retrieval of surface and atmospheric variables from microwave radiometer satellite data. Earth and Space Science, 10, e2023EA003177. https://doi.org/10.1029/2023EA003177

Spreen, G., L. Kaleschke, and G. Heygster. Sea ice remote sensing using AMSR-E 89-GHz channels, J. Geophys. Res., 113, C02S03, doi:10.1029/2005JC003384, 2008.

Swan, A. M., & Long, D. G. (2012). Multiyear Arctic sea ice classification using QuikSCAT. IEEE Transactions on Geoscience and Remote Sensing, 50(9), 3317–3326. https://doi.org/10.1109/tgrs.2012.2184123

Tonboe, R., Andersen, S., and Toudal, L. (2003). Anomalous winter sea ice backscatter and brightness temperatures. DMI Scientific Report 03-13, Danish Meteorological Institute. https://www.dmi.dk/fileadmin/Rapporter/SR/sr03-13.pdf

Tonboe, R. T., Heygster, G., Petersen, L. T., & Andersen, S. (2006). Thermal microwave radiation: Applications for remote sensing. Institution of Engineering and Technology.

Tonboe, R. T. (2010). The simulated sea ice thermal microwave emission at window and sounding frequencies. *Tellus A: Dynamic Meteorology and Oceanography*, 62(3), 333–344. https://doi.org/10.1111/j.1600-0870.2010.00434.x

Tschudi, M., Meier, W. N., Stewart, J. S., Fowler, C., and Maslanik, J. (2019). Polar Pathfinder Daily 25 km EASE-Grid Sea Ice Motion Vectors, Version 4. Boulder, Colorado USA. NASA National Snow and Ice Data Center. Distributed Active Archive Center. doi:10.5067/INAWUWO7QH7B .

Vats, A., Völgyes, D., Vermeer, M., Pedersen, M., Raja, K., Fantin, D. S. M., & Hay, J. A. (2024). Terrain-Informed Self-Supervised Learning: Enhancing Building Footprint Extraction From LiDAR Data With Limited Annotations. IEEE Transactions on Geoscience and Remote Sensing, 62, 1–10. https://doi.org/10.1109/TGRS.2024.3391391

Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020. doi:https://doi.org/10.1038/s41592-019-0686-2.

Warren, S. G., Rigor, I. G., Untersteiner, N., Radionov, V. F., Bryazgin, N. N., Aleksandrov, Y. I., & Colony, R. (1999). Snow depth on Arctic sea ice. *Journal of Climate*, 12(6), 1814–1829. https://doi.org/10.1175/1520-0442(1999)012⟨1814:SDOASI⟩2.0.CO;2

Wentz, F. J. A model function for ocean microwave brightness temperatures. Journal of Geophysical Research 88(C3), 1892-1908, 1983.

Wentz, F. J. A well-calibrated ocean algorithm for SSM/I. Journal of Geophysical Research 102(C4), 8703-8718, 1997.

Wentz, F. J., & Meissner, T. (2000). Algorithm theoretical basis document: AMSR ocean algorithm, version 2 (Tech. Rep.). Remote Sensing Systems. Retrieved from https://eospso.gsfc.nasa.gov/sites/default/files/atbd/atbd-amsr-ocean.pdf

WMO Sea Ice Nomenclature (WMO No. 259, volume 1 – Terminology and Codes, Volume II – Illustrated Glossary and III – International System of Sea-Ice Symbols)
by March 2014 (5th Session of JCOMM Expert Team on Sea Ice). Accessed from https://cryo.met.no/sites/cryo/files/IceService_docs/WMO_259-2015_multilingual.pdf

Wiesmann, A., & Mätzler, C. (1999). Microwave emission model of layered snowpacks. *Remote Sensing of Environment*, 70(3), 307–316. https://doi.org/10.1016/S0034-4257(99)00046-2

Xie, Q., Dai, Z., Hovy, E., Luong, M.-T., & Le, Q. V. (2020). Unsupervised Data Augmentation for Consistency Training (arXiv:1904.12848). arXiv. https://doi.org/10.48550/arXiv.1904.12848

Y. Ye, G. Heygster and M. Shokr, "Improving Multiyear Ice Concentration Estimates With Reanalysis Air Temperatures," in *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 5, pp. 2602-2614, May 2016a, doi: 10.1109/TGRS.2015.2503884.

Ye, Y.; Shokr, M.; Heygster, G.; Spreen, G. Improving Multiyear Sea Ice Concentration Estimates with Sea Ice Drift. *Remote Sens.* **2016b**, *8*, 397. https://doi.org/10.3390/rs8050397