sea state
cci

Algorithm Theoretical Basis Document (ATBD)

version 2.0, 06.10.2024

Contents

Author	Approved	Signature	Date
Andey Pleskachevsky, Björn Tings, Guillaume Dodet	Ellis Ash Fabrice Ardhuin		2024-11-06
ESA Acceptance	Digitally signed by Sarah Connors Date: 2024.11.25 17:19:05 Z		

Issue	Date	Comments

List of Acronyms

ACDC	Amplitude Compensation and Dilation Compensation
ADP	Algorithm Development Plan
ALES	Adaptive Leading Edge Subwaveform
ATBD	Algorithm Theoretical Basis Document
BH	Brown-Hayne (model)
cci	Climate Change Initiative
DC	Dilation Compensation
DD	Delay-Doppler
DtC	Distance to Coast
E3UB	End-to-End ECV Uncertainty Budget
ECV	Essential Climate Variable
EMF	Empirical Model Function
FFT	Fast Fourier Transform
GC	Geometric Corrections
GCCM	Grey-level Co-occurrence Matrices
GDR	Geophysical Data Record
GTS	Global Telecommunication System
L4	Level 4
LR-RMC	Low Resolution with Range Migration Correction
LRM	Low Rate Measurement
LUT	Look-Up Table
MAD	Median Absolute Deviation
MLE	Maximum Likelihood Estimator
MQE	Mean Quadratic Error
MWR	Microwave Radiometer
NM	Nelder Mead (algorithm)
NRCS	Normalized Radar Cross-Section
OSTST	Ocean Surface Topography Science Team

PHCP	Percentage of High Correlation
PLRM	Pseudo Low Rate Measurement
PTR	Point Target Response
PVASR	Product Validation and Algorithm Selection Report
RA	Radar Altimeters
RMC	Range Migration Correction
RR	Round Robin
R.m.s.	Root mean square
RMSE	Root mean square error
S3A	Sentinel-3A
S3B	Sentinel-3B
SAR	Synthetic Aperture Radar
SGDR	Sensor Geophysical Data Records
SI	Scatter Index
SSH	Sea Surface Height
SSP	Sea State Processor
STD	Standard Deviation
SWH	Significant Wave Height
WHALES	Wave Height Adaptive Leading Edge Subwaveform (retracker)
w.r.t	with respect to
WV	Wave (mode for SAR)

1. Introduction

This document presents the Algorithm Theoretical Basis Document (ATBD) for **Sea_State_cci**, phase two. The objective of this document is to define the Algorithm Theoretical Basis for all the algorithms developed and used in the framework of the SS_cci phase 2 project. The algorithms include processings for Altimetry and Synthetic Aperture Radar (SAR).

2. Algorithms for Satellite Altimetry Processing

2.1 Nadir altimetry retracking algorithm

2.1.1 Function

The Low Resolution Mode (LRM) waveforms are characterised by a rising leading edge that becomes less steep as the SWH increases, and a slowly decreasing trailing edge. The advent of the Adaptive Leading Edge Subwaveform (ALES, Passaro et al. 2014) retracker has already showed the way to keep the quality of the retrievals in the open ocean while improving the data quality and quantity in the last ~25 km from the coast, where waveforms are particularly corrupted by heterogeneous backscattering from land and sheltered water. Nevertheless, besides still using the full echo to retrieve parameters that are located on the leading edge, the standard retracking method is still affected by a suboptimal distribution of the residuals in the fitting process, which results in high level of noise in the estimations.

WHALES is designed as a unified way to solve these and other problems currently affecting the standard product, and is based on three principles:

1) The application of a weighted fitting solution, whose weights are adapted to the SWH in order to guarantee a more uniform distribution of the residuals during the iterative fitting. This guarantees significantly more precise estimations.

2) A subwaveform strategy to focus the retracking on the portion of the signal of interest, avoiding heterogeneous backscattering in the trailing edge (partially inherited from the ALES retracker). This guarantees efficiency in the coastal zone and a better representation of the oceanic scales of variability.

3) The decorrelation between SWH (Significant Wave Height) and sea level estimation, which corrects for possible covariant errors and increases furthermore the precision.

Moreover, a revisiting of the look-up tables used to correct for the Gaussian approximation of the Point Target Response in the Brown model ensures that the accuracy in the estimation is tailored to the new retracking solution.

2.1.2 Algorithm Definition

Input

- Waveform data: Sensor Geophysical Data Record (SGDR)
- Mission
- Instrumental Correction
- Weights

Output

- SWH
- σ_0

- Quality Flag

2.1.3. Mathematical statement

WHALES is a two-pass retracker. The retracking of each waveform follows the procedure described in the following flow diagram:

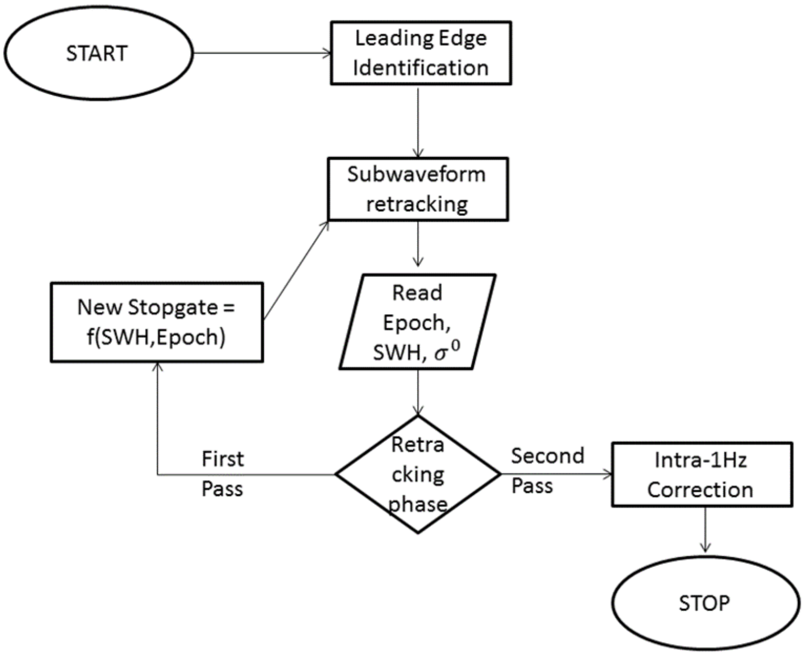

Figure 2.1.1: Flow diagram of WHALES

The functional form used to fit the real waveforms is the Brown-Hayne model as described in the following section. The original waveforms of any altimeter mission are discretized in elements called "gates". In WHALES, the first gate number is identified as 0 and the x-axis of a waveform is sampled in time. For example for Jason-3:

$$x = [0,1 * \tau, 2 * \tau ..., 103 * \tau] \quad \textbf{(2.1)}$$

Where τ is the spacing between two consecutive gates in time (3.125 ns in Jason-3).

The Leading Edge identification includes also the normalisation of the waveform and is performed following these substeps:

1) The waveform is normalised with normalisation factor *N*, where *N* = 1.3 * median(waveform)
2) The leading edge starts when the normalised waveform has a rise of 0.01 units compared to the previous gate (startgate)
3) At this point, the leading edge is considered valid if, for at least four gates after startgate, it does not decrease below 0.1 units (10% of the normalised power).
4) The end of the leading edge (stopgate) is fixed at the first gate in which the derivative changes sign (i.e. the signal starts decreasing and the trailing edge begins), if the change of sign is kept for the following 3 gates. For the specific case of SARAL, the waveform was smoothed before detecting startgate and stopgate.

The scope of the normalisation is indeed to take as reference power a value close to the maximum of the leading edge and, in the case of oceanic waveforms with standard trailing edge noise, the proposed factor *N* is a good approximation.

The first pass of WHALES involves a subwaveform that goes from startgate to stopgate+1. It is therefore a leading-edge-only subwaveform retracking. The vector of weights is filled with 1s. The convergence is therefore found by means of an unweighted Nelder-Mead estimator (see next section). The unknowns and the corresponding initial conditions applied are:

$$\tau = startgate - 1;\ \sigma_c = \frac{stopgate - startgate}{2 * \sqrt{2}};$$
$$P_u = 2 * mean(D[startgate\colon stopgate]) \qquad \textbf{(2.2)}$$

Where *D* is the normalised waveform. In case convergence is not reached, a new attempt is performed extending the subwaveform by two gates, until convergence or until the waveform limit.

After the first pass, the WHALES coefficients are applied to extend the subwaveform. As explained in the next section, the issue is one of defining an appropriate new stopgate for the second pass retracking based upon the SWH estimates from the first pass. For Jason-3, the following coefficients are used:

Stopgate = Ceiling(Tracking point + 3.89 + 3.86*SWH) **(2.3)**

These coefficients were recomputed specifically for the current purposes as explained in the section 2.1.4 “WHALES Coefficients”.

Using the new limits of the subwaveform, a second NM estimation is performed using the same initial conditions of the first pass. This time, the SWH estimation of the first pass is also used to identify the proper set of weights (see next section). WHALES therefore is adaptive in both the subwaveform width and the weights.

The SWH estimated in the second pass is instrumentally corrected by means of a look-up table that takes into account the bias due to the Gaussian approximation of the point target response in the BH model. The estimated P_u is converted in dB and corrected by atmospheric correction and scaling factor, whose fields are contained in the mission data. This constitutes the σ_0 output of the algorithm (backscatter coefficient). The epoch is not provided in the output since its precision and accuracy has not been verified.

The “Fitting Error on the leading edge” (Err) is used as a quality measure for the fitting. It is computed as the RMS difference between the fitted and the real waveform, considering only the gates of the leading edge. When Err>0.3, the quality flag is set to 1, i.e. the quality of the fitting is bad.

Once the full SGDR has been retracked, a further step can be applied to remove the “intra-1Hz correlation” between τ and SWH. WHALES will be applied in the Round Robin with and without this additional step.

2.1.4 Definitions

Brown-Hayne model

WHALES is based on the Brown-Hayne (BH) functional form that models the radar returns from the ocean to the satellite. The BH theoretical ocean model Brown (1977), Hayne (1980) is the standard model for the open ocean retrackers and describes the average return power of a rough scattering surface (i.e. what we simply call waveform). The return power is modelled as follows (equations as reported in Passaro et al., 2014):

$$V_m = a_\xi P_u \frac{[1+erf(u)]}{2} exp(-\nu) + T_n \qquad \textbf{(2.4)}$$

where:

$$a_\xi = exp(\frac{-4sin^2\xi}{\gamma}); \qquad \gamma = sin^2(\theta_0)\frac{1}{2ln(2)};$$

$$u = \frac{t-\tau-c_\xi\sigma_c^2}{\sqrt{2}\sigma_c}; \qquad \nu = c_\xi(t-\tau-0.5c_\xi\sigma_c^2);$$

$$\sigma_c^2 = \sigma_p^2+\sigma_s^2;\ \sigma_s = \frac{SWH}{2c};$$

$$c_\xi = b_\xi a;$$

$$a = 4c\ /\ [\gamma h(1+\frac{h}{R_e})]\ ;b_\xi = cos(2\xi) - \frac{sin^2(2\xi)}{\gamma}$$

where $erf(u)$denotes the error function, c is the speed of light, h the satellite altitude, $R_e$the Earth radius, $\xi$the off-nadir mispointing angle, θ_0 the antenna beam width, τ the Epoch with respect to the nominal tracking reference point, σ_c the rise time of the leading edge (depending on a term σ_s linked to SWH and on the width of the radar point target response σ_p), P_u the amplitude of the signal and T_n the thermal noise level.

Nelder-Mead Algorithm

The Nelder–Mead (NM) algorithm is a simplex optimisation method that does not use the derivatives of its cost function, whilst it searches for the minimum in a many-dimensional space. Specifically, considering m parameters to be estimated, given that a simplex of dimension m is a polytope of the same dimension and with *m* + 1 vertices characterised by *m* + 1 cost function values, NM generates at each step a new point whose cost function is compared with its value at the vertices. If it is smaller, the point becomes a vertex of the new simplex and a new iteration is generated (Nelder and Mead, 1965). Convergence is reached when the diameter of the simplex is smaller than a specified tolerance.
In WHALES, the objective function to be minimised is:

$$C = \sum \quad [W * R^2] \qquad \textbf{(2.5)}$$

where *W* is the vector of weights and the residual *R* is the difference between the real and the fitted waveform. NM is applied using the Python package scipy.optimize.minimize.

WHALES coefficients

The key concept of the WHALES subwaveform is that a leading-edge only retracker, although also providing results waveforms that do not conform to the BH model, has worse noise performances than a full-waveform retracker and therefore would not guarantee the homogeneity of the result. For best accuracy the subwaveform width for the second pass must be optimised such that it fully includes all gates comprising the leading edge, but with minimal contribution from the trailing edge, where artefacts such as bright target responses may prevent the BH model from accurately describing the shape. Defining startgate and stopgate the first and last gate of the subwaveform of choice, in effect the issue is one of defining an appropriate stopgate for a given SWH. The relationship between SWH and stopgate was derived from Monte Carlo simulations. For each value of SWH ranging from 0.5 to 10 m in steps of 0.5 m, 10000 echoes were simulated with the BH model adding realistic Rayleigh noise, and then averaged to create a simulated high-rate waveform. The resulting waveforms were retracked over the entire waveform, and then over sub-waveform windows with startgate=1 and variable stopgate, and the RMS errors (RMSE) were computed.

The difference of the RMSEs between the "full waveform" estimate and the subwaveform estimates is displayed as a function of the stopgate position in the figure below **Figure 2.1.2** (upper panel). The x axis is, in practice, the width of the sub-waveform, expressed as number of gates from the tracking point to the stopgate. The results for each SWH level are coded in different colours. For all three parameters, the curves converge asymptotically to the full waveform estimates, as expected for this idealised case of "pure-Brown" response of the ocean surface. The relation the second retracking step of WHALES is shown in the panel below and is obtained by setting a tolerance in the RMSE difference of the SWH. In order for WHALES to optimise the need to retrieve signals whose trailing edge is corrupted, the tolerance bar was set to 2 cm at 20 Hz, i.e. 0.45 cm at 1 Hz.

Weights

To derive the adaptive set of weights that is used to find convergence in WHALES, a Monte Carlo simulation is performed. The objective is to base the weighting on the uncertainty of the fitting along the leading edge of the waveform. As an estimation of the uncertainty, the standard deviation (STD) between the simulated echo and a large number of fitted waveforms is used.

Simulated echoes are generated according to the BH model as previously reported, SWH at steps of 0.5 m from 0 to 10 m (10000 waveforms per SWH value). Each echo is retracked with a BH retracker that finds the convergence through an unweighted Nelder-Mead. For each SWH value:

- the value of the residuals between the simulated echo and the fitted waveform is stored
- the position of the start and the end of the leading edge is stored

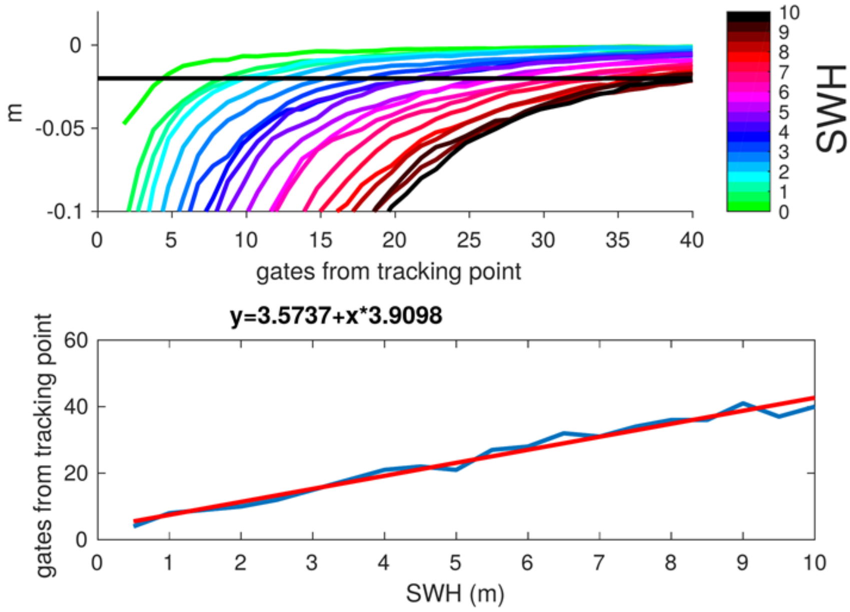

Figure 2.1.2: (Up) Difference of the RMSEs between the "full waveform" estimate and the subwaveform estimates as a function of the stopgate position. (Below): linear relationship obtained by setting a tolerance in the RMSE difference of the SWH in the upper plot. In order for WHALES to optimise the need to retrieve signals whose trailing edge is corrupted, the tolerance bar was set to 2 cm at 20 Hz.

For each SWH value, there will be 10000 value of residuals at each waveform gate, and we can therefore compute their std. The weights to be used in the WHALES retracking will be the inverse of this std, i.e. the so-called “Statistical Weighting”, which in statistics is a recommended choice when the uncertainties of the different points to fit are very different from each other (Wolberg, 2006).

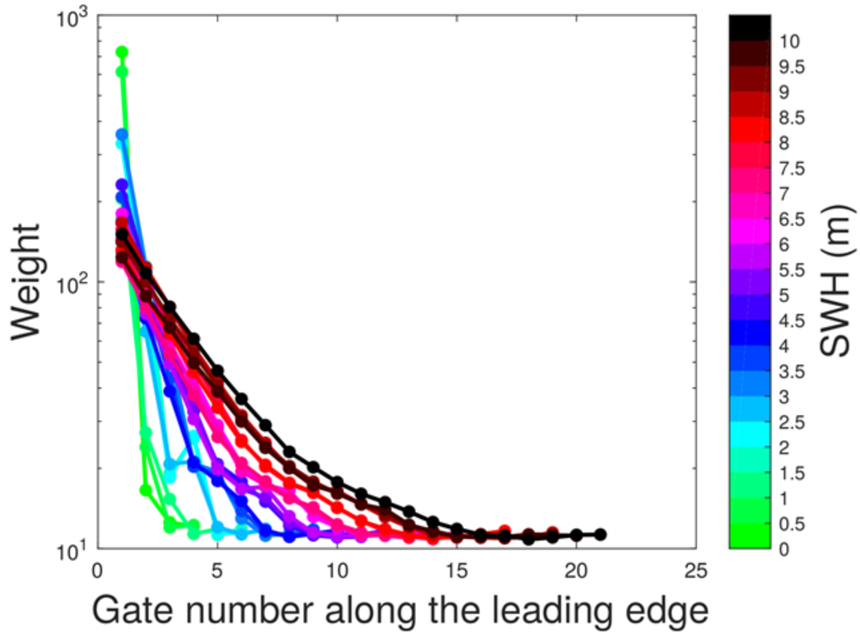

Figure 2.1.3: Variable weights used by WHALES depending on Significant Wave Height and gate number along the leading edge of the waveform.

2.2 Nadir altimetry compression algorithm / resolution choice

The CCI Sea State Dataset v3 provides 1 Hz SWH measurements. These 1Hz measurements are calculated by averaging groups of 20 Hz (40 Hz for Saral):

- for Envisat, CryoSat-2, Saral, Jason-1, Jason-2, Jason-3 missions, the groups of 20 Hz measurements used to calculate the 1Hz values are exactly the same as in the source Agency's GDR & SGDR products. Both CCI and Agency files can be compared one to one, have the same number of measurements, and the same latitude, longitude, time for each 1 Hz measurement.

- for Sentinel3-A, the 20 Hz measurements were grouped by seconds to calculate the 1 Hz values: starting from the first measurement in the file, the next measurements separated by less than one second from the first one are grouped together, and so on. The nominal number of 20 Hz measurements per 1 Hz value is not constant and varies between 19 and 20. It can be less in case of missing data, end of orbit, etc...

The method used to average the 20 Hz measurements into 1 Hz values is the same for all altimeters. For each group of 20 Hz measurements:

1. 20 Hz measurements flagged as bad by the retracker are discarded. 20 Hz measurements are flagged as bad, for altimeters retracked with WHALES, when the fitting error is greater than 0.3.
2. 20 Hz measurements located on land are discarded. The land location is determined from the latitude & longitude provided for each measurement using OpenstreetMap coastline, with a precision of about 50m.
3. 20 Hz SWH measurements not in the -0.5 to 30 m range are discarded
4. the remaining 20 Hz SWH outliers are discarded. The outlier detection scheme is based on the maximum absolute deviation for a 3-sigma criterion, meaning
    - only 20 Hz SWH values within [median(SWH) - 3 * MAD(SWH), median(SWH) 3 * MAD(SWH)] interval are kept
    - with: MAD(SWH) = 1.4286 * median(abs(SWH - median(SWH)))
5. the 1Hz SWH value is computed as the median of the remaining 20Hz SWH values that have not been discarded during the 4 previous steps. Moreover, the swh_num_valid parameter is computed as the number of remaining 20Hz records and the swh_rms parameter is computed as the root mean square of the deviation from the median, considering only valid 20Hz records

The same method is used to calculate 1 Hz sigma0 values from the 20 Hz sigma0 (except no test on the validity range of the sigma0 is performed).

2.3. Nadir altimetry 1Hz SWH quality level and rejection flags

Quality control of individual altimeter measurements is performed with checks over the calculated 1 Hz values and ancillary variables. As a result, the SWH comes with a quality level provided in the swh_quality variable. Its meaning is defined as follows:

value	***meaning***	***description***
0	*undefined*	*the measurement value is not defined or relevant (missing value, etc…), no quality check was applied.*
1	*bad*	*the measurement was qualified as not usable after quality check.*
2	*acceptable*	*the measurement may still be usable for some application or the quality check could not fully assess if it is a bad or good value (suspect).*
3	*good*	*the measurement is usable.*

The tests performed to set SWH measurements to a specific quality level are provided below:

meaning	***reasons***
undefined	• *the 1 Hz SWH or SWH RMS could not be calculated (the number of remaining 20Hz SWH after all check and outlier tests is equal to zero)*
bad	• *the number of valid 20Hz SWH measurements averaged to compute the 1 Hz value is lower than the selected threshold (12 for Saral, 6 for all other altimeters)* • *the 1 Hz SWH value is not in]0., 30.] meter range* • *the 1 Hz SWH RMS is equal to 0. or NaN* • *sea ice concentration is greater than 10%* • *the 1 Hz location is on land (using OpenstreeMap coastline, 50m precision). Since the 1 Hz location is an average of 20 Hz measurements, this may include mixed land/sea areas.* • *the 1 Hz SWH value did not pass the SWH RMS test (described below)* • *the 1 Hz SWH value did not pass the outlier test*
acceptable	• *sea ice concentration is non zero but lower than 10%*
good	• *all remaining measurements*

The quality flag is provided in the swh_rejection_flag variable.

Algorithm description for computing SWH RMS Look Up Table

As documented in Sepulveda et al. (2015), SWH measurements derived from radar altimeter measurements can be contaminated by the presence of land in the footprint, strong rain events, and so-called “sigma0 blooms” due to weak winds or surface slicks. According to these authors, the standard deviation of 20Hz SWH values over one second (hereinafter called SWH

RMS) is one of the most relevant parameters to detect erroneous values of SWH. Since the SWH RMS level strongly depends on SWH, constant threshold values are not adequate to efficiently remove SWH RMS outliers. Therefore, Sepulveda et al. (2015) and Queffeulou (2016) proposed a methodology to set a statistical threshold on SWH RMS that depends on SWH, and which can be used a posteriori to filter out erroneous SWH measurements. Since the SWH RMS for a given narrow SWH presents a log-normal distribution, these authors propose to estimate the upper threshold for the logarithm of SWH RMS as the sum of the mean value and thrice the standard deviation. Moreover, in order to reduce discontinuity in the SWH RMS threshold function for large SWH values, where the number of records is too low to derive robust statistics, a second-order polynomial function is fitted.

The methodology implemented to determine a SWH RMS LUTfor each mission of the Sea State CCI dataset can be described as follows:

1. One cycle of SWH and SWH RMS measurements are sampled and invalid measurements are rejected based on the following quality flags: land mask, SWH range;
2. the mean and standard deviation of log(SWH RMS) are computed for SWH bins of 0.5 m width, ranging from 0 to 15 m, with a 0.05 m increment. Only bins with more than 100 values are considered;
3. upper threshold on SWH RMS are computed for each SWH bin as:

exp(mean(log(SWH RMS))+3*std(log(SWH_RMS)));

4. a second-order polynomial function is fitted to the SWH RMS threshold function for SWH values comprised between 3 and 10m and is extrapolated up to SWH = 12m;
5. finally, a look-up table (LUT) of SWH RMS thresholds is derived by combining the SWH RMS threshold values for SWH lower than 3m and the values of the polynomial function for SWH higher than 3m; for SWH above 12 m, a constant value equal to the SWH RMS threshold at 12m is taken;

Note that the selection of the lower and upper bounds (3-10m) used to estimate the second-order polynomial function is mission specific, and will be determined for each method from visual inspection of the SWH_RMS = f(SWH) function.

Figure 2.3.1 shows the distribution of SWH RMS (left panel) and log(SWH RMS) (right panel) values as a function of SWH using S3A PLRM measurements during cycle 27.

Three distinct patterns can be identified:

1) the main axis of the scatter where SWH RMS presents a quasi-linear increase with SWH;
2) a set of large scattered SWH RMS values, above the red curve, for SWH range about 1–7 m, extending beyond the upper limit of the figure (most data of this second pattern correspond to along-track, almost isolated, erroneous spikes);
3) a nonlinear behavior of SWH RMS values at low sea state (SWH below 2 m).

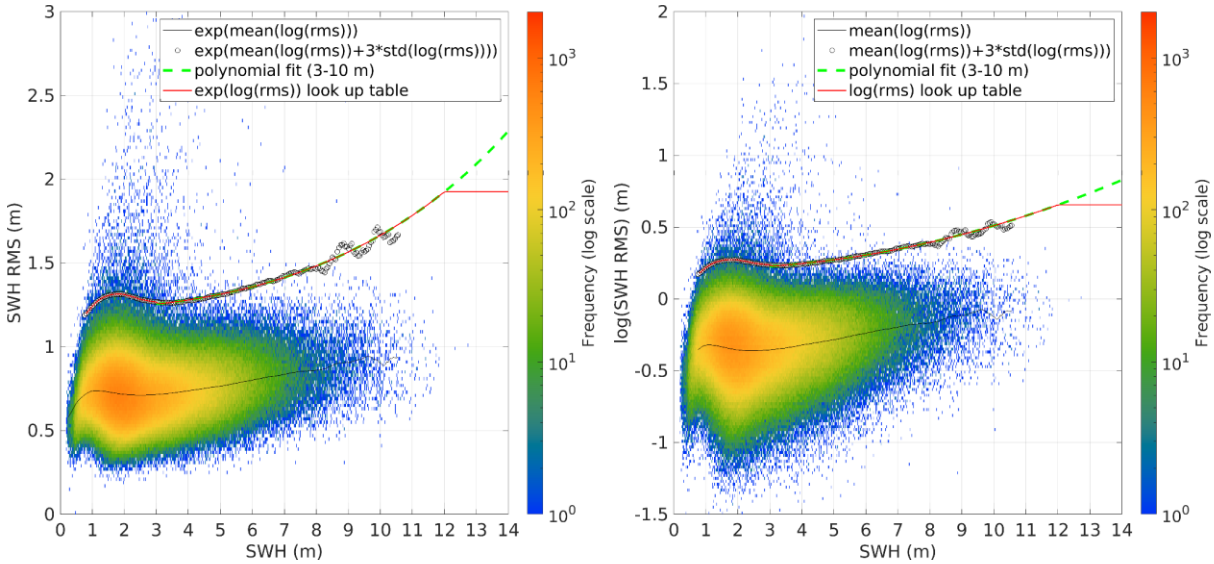

Figure 2.3.1: 2D histograms of SWH RMS against SWH (left panel), and log(SWH RMS) against SWH

In order to investigate the sensitivity of the SWH LUT values to the selected cycles, a set of 8 LUT were computed considering eight orbit cycles. Results indicate limited variability between cycles, particularly for SWH below 8 m. We see that individual SWH LUT start diverging for SWH larger than 8 m, which indicates that the mean LUT is therefore more representative of overall conditions than individual LUT computed on a single cycle. The final SWH LUT is therefore computed as the mean of an ensemble of 8 SWH LUT estimated from height independent cycles. This represents an improvement in comparison to the methodology implemented during the phase 1 of the Sea State CCI project, where 90 consecutive days of data were used to estimate the SWH LUT.

2.4. Nadir altimetry SWH intercalibration

Following the methodology developed during phase 1 (see Dodet et al., 2020), reference altimeter missions calibrated against in situ data are used to intercalibrate the remaining altimeter missions. For the version 3 dataset (phase 1), the Jason-2 mission was selected as reference because its time extent overlapped all the other missions included in the dataset. For version 4, the four missions of the Jason series (J1, J2, J3, J-CS) are selected as reference in order to cover a >20-year time extent and overlap with the majority of the missions, as well as to expand the range of sea state conditions used to derive the absolute calibration function. Indeed, in phase 1 the limited number of matchups between Jason-2 and in situ measurements prevented the derivation of correction below 1 m. Moreover, these four missions present tandem phases between subsequent missions which allows a robust inter-calibration between them. Once inter-calibrated, this reference “metamission” is calibrated against in situ data, considered as the ground truth. In a second step, the remaining missions (namely, ERS-1, ERS-2, Envisat, Saral, CryoSat-2, S3A and S3B) are calibrated against this reference “metamission”. Finally, an independent verification of this calibration methodology is performed by comparing the statistical metrics between altimeter and in situ data and between altimeter missions before and after the calibration corrections are applied.

The adjusted significant wave height SWH is provided in swh_adjusted variable.

2.4.1. Selection of in situ platforms

For the absolute calibration of the altimeter reference mission, we have selected a set of 72 quality controlled in situ platforms from the reprocessed product (WAVE-REP) of the CMEMS INSTAC database (http://www.marineinsitu.eu/). This selection corresponds to offshore platforms (>100km from the coast) with more than 1000 matchups with altimeter data. This selection contains stations mostly from the US National Data Buoy Center (60%), the Canadian Marine Environmental Data Service (15%) and the UK Met Office (8%). The remaining stations (17%) come from other US and European organisms.

The CMEMS In Situ Thematic Assembly Center (CMEMS INSTAC) is a component of the Copernicus Marine Service and its role is to ensure consistent and reliable access to a range of in situ data for service production and validation. For this purpose, CMEMS INSTAC collects multi-source/multiplatform data, and performs consistent quality control before distributing the data in a common format to the CMEMS Marine Forecasting Centres (MFC). The data can be found at http://www.marineinsitu.eu/. For the validation of altimeter wave data, we have considered all in situ platforms located at a minimum distance of 50 km from the shore at depth larger than 100 m. In total, 99 wave buoys were selected, with distance to the coast comprised between 50 and 2000 km. In order to ensure that only good quality measurements are selected, a number of tests have been applied:

- stationarity: rejecting SWH buoy measurements having a constant SWH value over more than 48 h;
- location: rejecting SWH buoy measurements over unrealistic position change within a short period of time;
- CMEMS quality check: using the CMEMS native quality flags on time, location and SWH to reject bad or suspect SWH buoy measurements;
- precision: detecting and rejecting SWH buoy measurements having a precision equal or larger than 0.5 meters;
- range: rejecting SWH buoy measurements out of the 0-30 meter range.

2.4.2. Data pairing method

Matchups between altimeter and in situ measurements

Matchups between altimeter and in situ platform measurements were computed within a 100-km distance and 30-minute time window. The along-track altimeter and in situ SWH observations were smoothed with 50-km and 1-h running average, respectively, in order to reduce high-frequency variability and match the 0.5 degree model resolution. Matchups associated with in situ SWH measurements below 0.05 were discarded, assuming altimeter measurements unreliable in such low sea state conditions.

Cross-overs between two altimeter missions

A cross-over between two altimeter missions is identified when the two altimeter ground tracks cross each other within a one-hour time window. The altimeter cross-over data are estimated as the 50-km average of the along-track SWH in order to filter out high-frequency variability.

Data pairing during tandem phases between altimeter missions

During tandem phases, the two altimeter missions are sharing the same orbit with a slight time delay (less than 1 min) and the so-called “cross-over” data simply correspond to the SWH data recorded at the same location, after applying a 50-km along-track moving average.

2.4.3. Calibration method

In order to account for the non-linear error characteristics of radar altimeter measurements, in particular for low sea state conditions, the calibration methodology implemented in v4 uses the data binning technique developed during phase 1 (see **Figure 2.4.1**), so that the calibration correction can accommodate for swh- and mission-dependent error structure. First, the median of the residuals between uncorrected and reference data (the reference data being in situ match-ups for Jason-2 calibration and the calibrated Jason-2 data at cross-over for the remaining missions) is computed for 0.20m bin width, with a 0.05 m increment over the full SWH range. Then, a robust linear regression (green line on **Figure 2.4.1** top middle panel) is fitted through the median of the residuals (magenta dots on **Figure 2.4.1** middle panel) over a mission-specific SWH range to extrapolate the error function at high SWH. The threshold and range are estimated in order to ensure a smooth transition between the non-linear corrections for low-to-medium sea state conditions and the linear corrections for medium-to-high sea state conditions. Finally a 5-point moving average is applied to the corrections in order to reduce small scale oscillations.

As opposed to phase 1, the calibration technique implemented in phase 2 is using the results of a spectral wave model hindcast as an intermediate dataset between the two altimeter missions to intercalibrate. A correction look up table is first computed between the reference mission and model outputs interpolated in space and time at the altimeter ground track location (blue circles on **Figure 2.4.2**). Then a second correction look up table is computed between the mission to be intercalibrated with the reference mission and the model outputs (red circles on **Figure 2.4.2**). Finally, the first LUT is subtracted to the second one in order to derive the final correction between the two missions (black dots on **Figure 2.4.2**). In comparison to the original method (black circles on **Figure 2.4.2**), the model-based approach is able to align the two altimeter missions over a much wider SWH range (see differences between black circles and black dots for SWH < 1m) since a much larger number of “crossovers” is affordable.

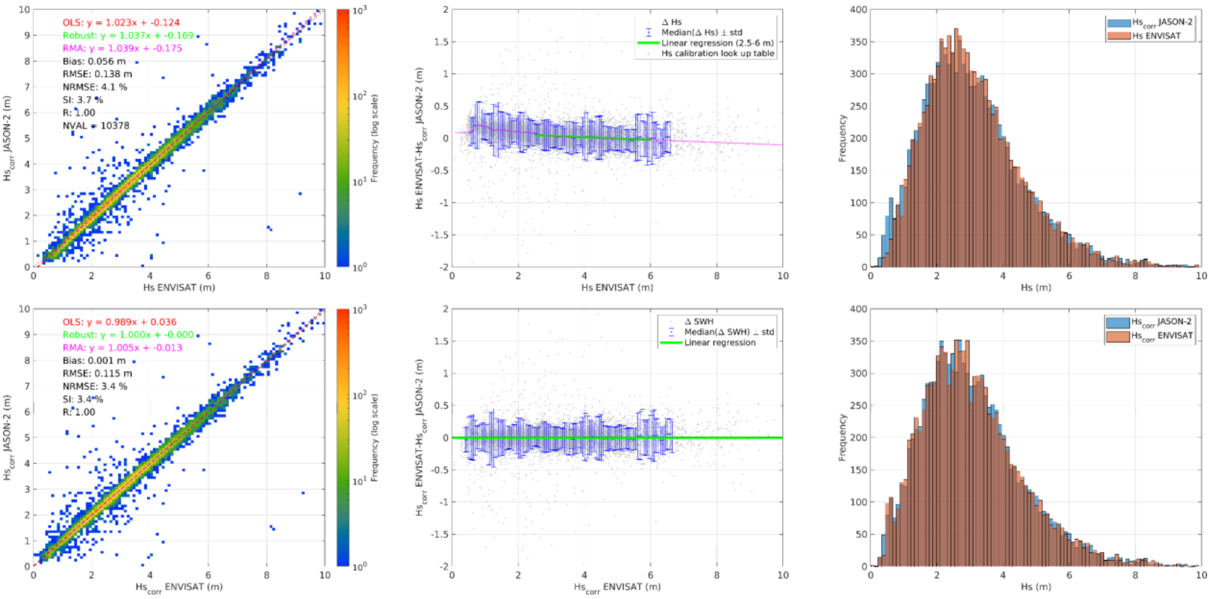

Figure 2.4.1: Scatter diagram between Jason-2 and Envisat (left); residual error between Jason-2 and Envisat measurements at crossover (middle); histograms of Jason-2 and Envisat SWH measurements before (top) and after (bottom) intercalibrationis applied.

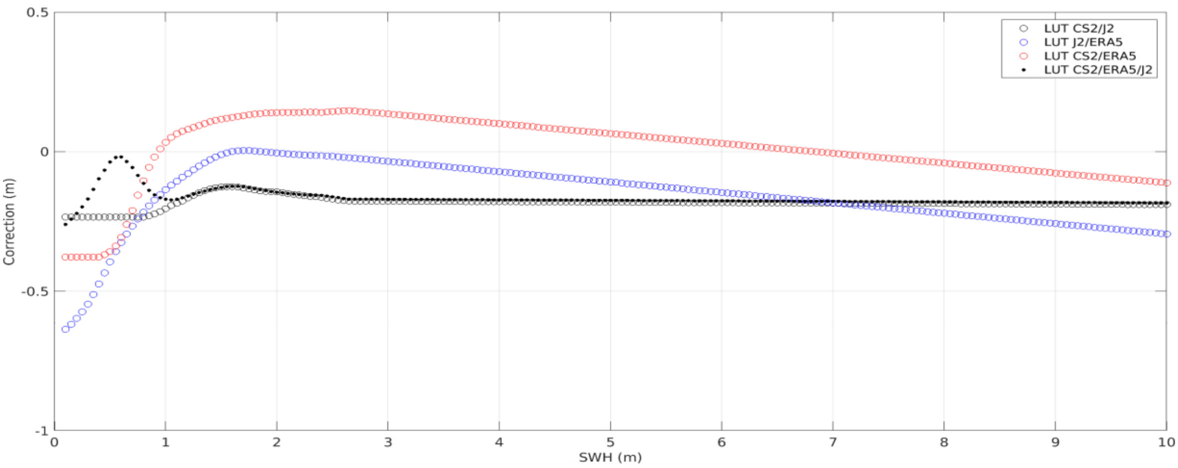

Figure 2.4.2: Calibration Look Up Tables derived from direct comparison between Cryosat-2 and Jason-2 (black circles) or from indirect comparisons (black dots) based on Cryosat-2 versus ERA-5 (red circles) and Jason-2 versus ERA-5 (blue circles) comparisons.

2.5 Nadir altimetry Sigma0 Dual-Frequency Calibration

There are two algorithms to be applied to the sigma0 data, whether these come from a proprietary CCI algorithm such as WHALES or from the standard distributed GDRs (Geophysical Data Records). The first is a non-linear correction to adjust for errors in the settings of the AGC (Automatic Gain Control); at present this correction only seems pertinent for Jason-2. The second correction is to bring all altimetry records into alignment, so that universal algorithms may be applied such as for wind speed or slick detection, with those methods and results being independent of the particular altimeter data in question. That correction is one major adjustment for the specific mission plus minor adjustments as a function of time to allow for gradual changes in instrument performance e.g. degradation of the antenna or amplification circuitry.

2.5.1. Non-linear AGC correction

For Jason-2 it has been noted that although the AGC settings are given to 0.01 dB, the actual change in attenuation in switching from one value to another does not precisely match the difference in their nominal values. Having determined and tabulated the correct increments, these need to be applied cumulatively to give the correct AGC setting (apart from an undetermined solitary bias). As the 1Hz data simply uses a combination of these discrete AGC settings, the correction is an exact linear interpolation between that for each AGC setting, and consequently it can be simply applied as piecewise linear interpolation (see **Figure 2.5.1.**)

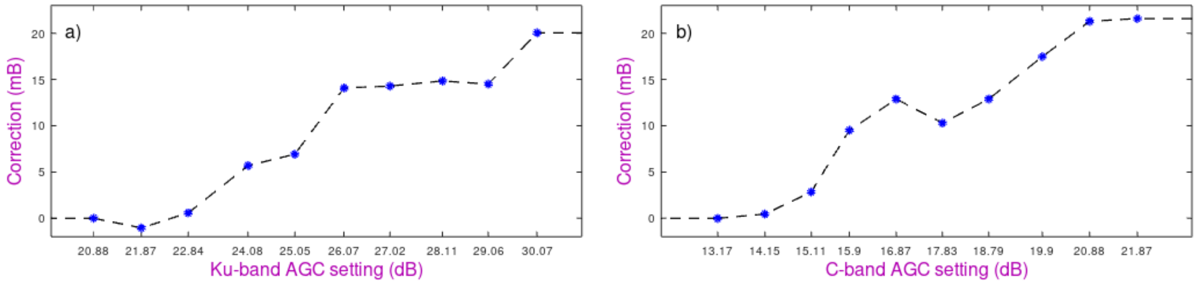

Figure 2.5.1: Adjustment to sigma0 values for incorrect AGC values on Jason-2. a) Ku-band, b) C-band. The correction is to be determined from piecewise linear interpolation using the stated 1 Hz AGC values and then subtracted from the sigma0 value. Note there are a wide range of AGC settings, as the instrument aims to cover surfaces with very different reflectivity (ocean, ice and land); the correction has only been determined for the 10 AGC settings most frequently encountered over the open ocean, with constant value extrapolated beyond those spans.

2.5.2. Instrument-specific time-series

A time-varying set of corrections is to be applied uniformly to all sigma0 values, firstly correcting for the definition of sigma0 for that altimeter, and the uncertainty in the attenuation provided by the waveguides, and secondly for gradual changes in instrument performance that have not been adequately resolved by the use of the internal CAL-1 and CAL-2 loops. The constancy of σ^0-σ^0 curves enables this correction to be determined to fine temporal resolution (e.g. 10 days) for dual-frequency altimeters, but corrections will be much less well resolved in time for instruments operating at Ku-band only.

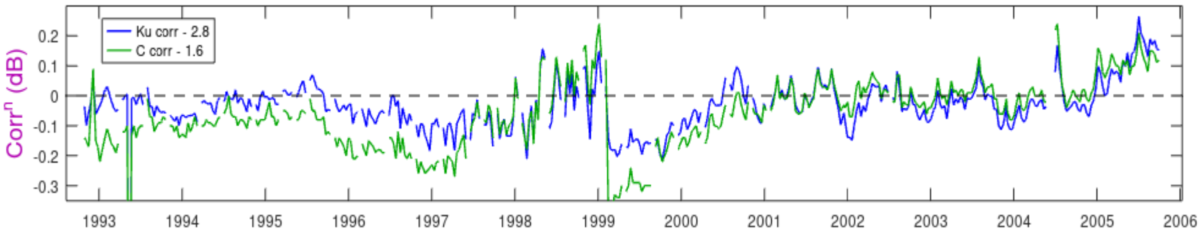

Figure 2.5.2: Changes in Ku- and C-band corrections for the TOPEX instrument. Note that degradation became noticeable around cycle 150 (late 1996), with an eventual switch to the redundant Side B after cycle 235 (Feb. 1999). (Some of the observed 59-day periodicity may later be removed by an allowance for SST.)

3. Algorithms for Synthetic Aperture Radar Processing

3.1. Algorithm SAR_SeaStaR for estimation of integrated sea state parameters

3.1.1. General SAR_SeaStaR description

The empirical algorithm SAR_SeaStaR (SAR Sea State Retrieval) was developed at German Aerospace Center DLR, Maritime Safety and Security Lab Bremen, and adopted for different satellites Sentinel-1 (S1) and TerraSAR-X (TS-X) and modes:

- S1 Wave Mode (WM) Level-1 (L1) products
- S1 Interferometric Wide Swath Mode (IW)
- S1 Extra Wide (EW)
- TerraSAR-X (TS-X) StripMap (SM)

SAR_SeaStaR bases on a combination of the linear regression (LR) functions CWAVE_EX (Pleskachevsky et al., 2022) and a machine learning (ML) approach using the support vector machine (SVM) technique (Pleskachevsky et al., 2024). A combination of the methods means the SAR_SeaStaR algorithm includes both: CWAVE_EX functions and SVM models. This means, by processing SAR scenes:

- the parameters are derived by CWAVE_EX in the first stage of the processing
- a series of parameters (e.g. SWH, TM2, see **Tab.3.1.1**) have an improved SVM solution based on first guess SWH from CWAVE_EX. The use of the SVM is optional, it can be skipped if timeliness is preferred over result quality.

Tab.3.1.1: Current state of the state parameters estimated by SAR_SeaStaR

№	Parameter	unit	symbol	LR solution CWAVE_EX	ML solution			
					S1-VW	S1-IW	S1-EW	TSX-SM
1	Total significant wave height	m	SWH	+	+	+	+	+
2	Mean wave period (Tm0-1)	s	TM0	+		+		
3	First moment wave period	s	TM1	+		+		
4	Second moment wave period	s	TM2	+		+	+	+
5	Wave height swell dominant	m	SW1	+		+		
6	Wave height swell secondary	m	SW2	+				
7	Wave height windsea	m	SWW	+		+		
8	Mean period windsea	s	TMW	+				

Comment: *the azimuthal cutoff (one of used internal SAR features and not a sea state parameter) is not is stored in DLR ocean products. However, it was stored in intermediate results (used SAR features, total ca.50 features) for each ID and can be additionally extracted.*

The scenes are processed in raster and result in continuous sea state fields, with the exception of S1 WV, where averaged values for sea state parameters for along-orbit imagettes of 20 km × 20 km are presented. Outputs are total significant wave height SWH, wave heights of dominant and secondary swell SW1, SW2, wave height of windsea SWW, mean wave period

TM0, first and second moment wave periods TM1, TM2 and periods of windsea TMW (see **Tab.3.1.1**). **Fig.3.1.1** shows an example of sea state processing for all eight sea state parameters from S1 IW scene covering ca. 1600 km × 200 km, acquired during a strong storm in the North Atlantic on 2020-02-14 at 18:45 UTC with SWH reaching ca. 13 m. Processing in a 5 km raster results in ca. 1500 subscenes for each individual IW image. The isolines depict the results of WFWAM model at 18:00 UTC (excepting TM1 period which is not provided by this model).

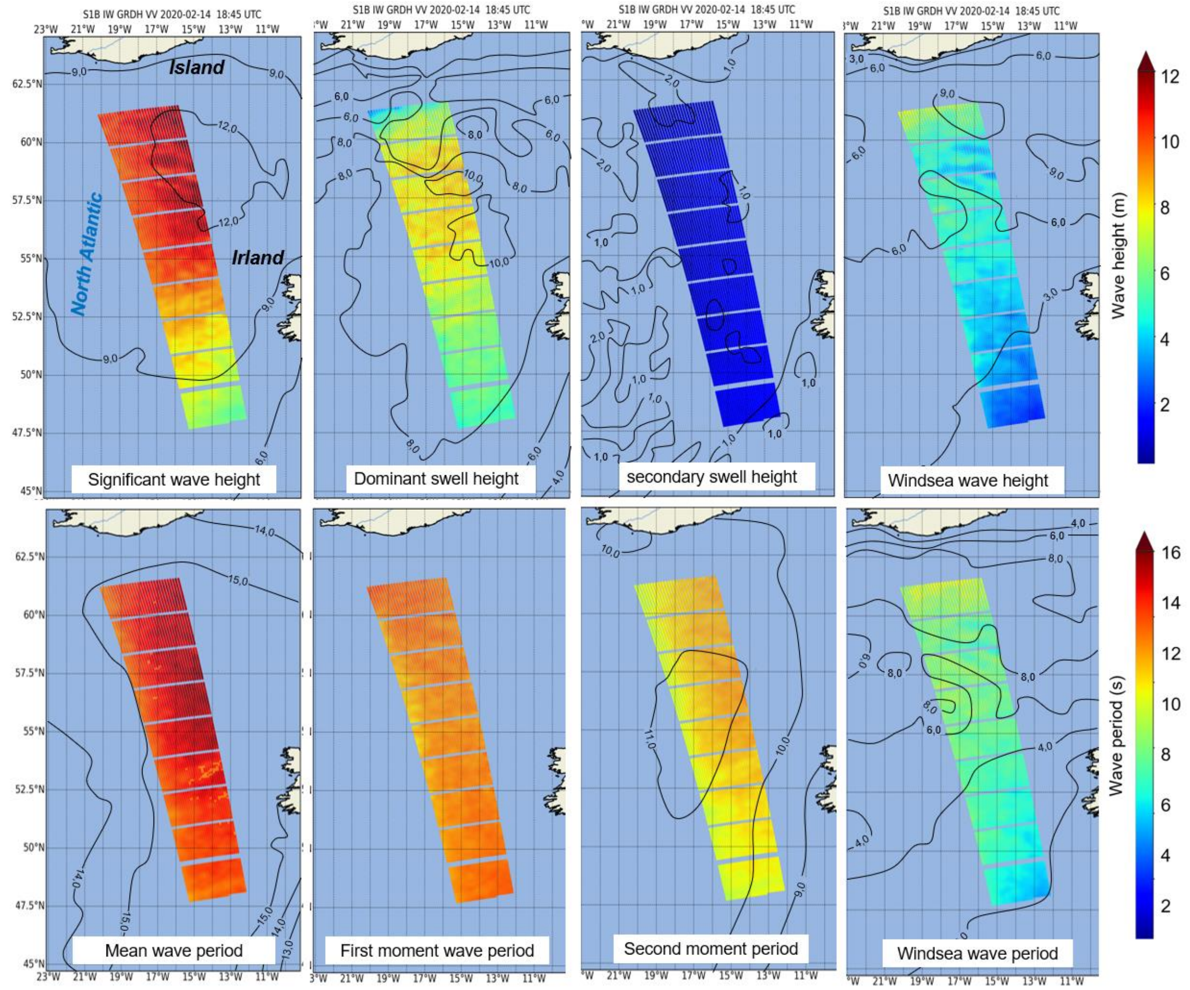

Figure 3.1.1: Example of eight sea state fields processed from a S1 IW scene with ca. 1600 km × 200 km coverage acquired during a strong storm in the North Atlantic on 2020-02-14 at 18:45 UTC with SWH reaching ca. 13 m. Processing in a 5 km raster results in ca. 1500 subscenes (approximately ~30×50) for each individual IW image. Isolines shows the results of forecast WFWAM at 18:00 UTC (excluding T_{m1} not provided by Copernicus Marine Environment CMEMS).

The priority of the algorithm design is an automatic, fast and robust raster processing of SAR acquisitions independent from wave patterns, which means it is also applicable when only clutter is visible in the SAR images. The SSP runs daily at the ground station for Sentinel-1 IW scenes in North and Baltic Sea. The model functions are based on the spectral analysis of subscenes in the wavenumber space and allow direct significant wave height (*SWH)* estimation from image spectra without transferring into wave spectra. The model functions are based on integrated image spectra parameters as well as local wind information estimated by the CMOD geophysical model functions. Additionally, a texture analysis based on Grey Level Co-occurrence Matrices (GLCM) is performed and the extracted textural features are integrated in the model function. The processing includes three main steps:

1. Data preparation: reading, calibration and artefact/outlier pre-filtering
2. Feature extraction & model function application
3. Control of results

The detailed descriptions can be seen in Pleskachevsky et al., 2022 and in Pleskachevsky et a., 2024.

3.1.2. SAR_SeaStaR processing details

The SAR_SeaStaR includes the complete processing chain with a series of steps each needed to reach high accuracy:

1. filtering image artefacts (e.g. ships, wakes, offshore windfarms constructions, etc.),
2. resampling and denoising (e.g. for S1 IW resampling from 10 m to 2.5 m pixel spacing)
3. SAR features estimation and control
4. model functions (linear regression and machine learning models) for estimation of sea state parameters
5. control of results using filtering procedures.

Steps 1-3 are described in (Pleskachevsky et al., 2016) and (Pleskachevsky et al., 2019) in detail. A series of CWAVE_EX LR model functions were developed for each satellite/mode and each sea state parameter. Detailed information on development of latest model functions and the control-of-results procedures can be found in (Pleskachevsky et al., 2024). MLs were trained for two H_s and T_{m2} for TS-X SM, S1 IW and S1 EW. For S1 IW also the model functions for Tm0-1 were developed.

In a classic way, the estimation of sea state parameters is based on an analysis of normalized radar cross-section (NRCS) of subscenes. One of the basic variables represents the SAR image spectrum obtained using fast Fourier transformation (FFT) applied to the ground range detected, radiometrically calibrated, filtered, denoised, land-masked and normalised subscenes with a size of 1,024 × 1,024 pixels in wave number k domain as introduced in (Pleskachevsky et al., 2019). The primary SAR features estimated from a subscene are of five different types:

- NRCS and NRCS statistics (variance, skewness, kurtosis, etc.), in total 9 features.
- Geophysical parameters (wind speed using CMOD-5 algorithms for C-band, 1 feature
- Grey Level Co-occurrence Matrix (GLCM) parameters (homogeneity, dissimilarity, etc.), in total 8 features
- Spectral parameters based on image spectrum integration of different wavelength domains (0-30 m, 30-100 m, 100-400 m, etc.) and spectral width parameters (Longuet-Higgins, Goda), in total 17 features

- Spectral parameters using products of normalized image spectrum with orthonormal functions (CWAVE approach) and cutoff wavelength estimated using autocorrelation function (ACF), in total 21 features

The detailed description of all features can be found in Pleskachevsky et al. (2019) and Pleskachevsky et al. (2022). In total, 54 primary and the 77 most significant secondary features are applied in CWAVE_EX. Secondary features are combinations of primary features in quadratic and inverse forms. The input for new SVM functions are primary features complemented with:

- first-guess H_s from linear regression solution from CWAVE_EX
- incidence angle with an accuracy of third decimal place
- flag identifying the satellite (S1-A or S1-B for Sentinel-1 and TS-X or TD-X for TerraSAR-X).
- pixel spacing (different for e.g. TS-X SM radiometric enhanced (RE) and spatially enhanced (SE) products)
- flag identifying polarisation (HH or VV)

In a series of several analyses used for the control of results (Pleskachevsky et al., 2022). One of the most acting features is the so-called Rosenthal parameter, named after Wolfgang Rosenthal (RIP 2016), who suggested its introduction. The Rosenthal parameter *E*r indicates the image spectrum energy integrated with an additional scaling of 1/*k*. In contrast to a direct integration, the *Er* integration amplified the longer wavelength signals in the image spectrum. In this way, objects such as ships or ship wakes in a subscene significantly increase the value of *Er* and can thus be filtered out. For each satellite/mode the threshold value of *Er* was experimentally established using statistics of processed scenes indicating the validity of the subscene for sea state estimation.

3.1.3. SAR products for processing using SAR_SeaStaR

For S1, the method was trained and applied for three SAR modes WV, IW and EW (state-of-the-art SAR_SeaStaR for 2024). For TerraSAR-X, the SAR_SeaStaR was trained for StripMap mode (SM).

For S1 WV, the Single Look Complex (SLC) products have been used. Per a day around 60 S1 WV products (ascending or descending tracks) each with around 120 individual imagettes for both S1-A and S1-B are acquired. Each individual S1 WV product of is 2–10 GB (ca. 5 TB/month).

S1 WV acquires two parallel tracks with incidence angles of around 23° (wv1) and around 36° (wv2) with **imagettes** (small images with an approximate footprint of **20 km × 20 km** acquired every 200 km along each wv1 and wv2 track with a 100 km offset and distance of 100 km between wv1 and wv2 tracks. Using both wv1 and wv2 this means along-track imagettes **each 100 km**. The length of a track (relative orbit with ID) varies from around 1,000 km (10 imagettes) to 12,000 km (120 imagettes). The nominal spatial pixel resolution of S1 WV is around 3 m depending on the local incidence angle. The scenes can be acquired in vertical (VV) or horizontal (HH) co-polarization. However, more than 95% of the data were acquired in VV polarization. **Pixel spacing** differs between **3-5 m**. **WV products are acquired exclusively in open ocean**

For S1 IW and EW, the GRD (Ground Range Detected) L1 products have been used, available in single (HH or VV) or dual polarization (HH+HV or VV+VH). For sea state estimation, the VV or HH polarization data are used, with priority given to VV products. Both IW and EW images were often acquired in strips of up to 2,000 km along the orbit covering the ocean surface divided later into individual images with separated ID.

- **S1 IW** mode combines a large swath width with a moderate geometric resolution. The individual IW images cover approximately 200 km in azimuth and **250 km in the range** direction with a **pixel spacing of 10 m**. The L1 products are available in single (HH or VV) or dual polarization (HH+HV or VV+VH). For sea state estimation, the VV or HH polarization data were used, with priority given to VV products. **In ocean, IW products are mostly acquired in coastal-near areas**.

- **S1 EW** mode is similar to the IW mode, but the EW mode acquires data over a wider area than for IW mode using five sub-swaths. The EW mode acquires data over ca. **400 km swath** width with a coarser **pixel spacing of 40 m** (GRDM products). In the open ocean, EW products are mostly acquired in polar regions. However, a number of EW images were still acquired in ice-free regions and around the equator and near Madagascar Island

The modes are acquired worldwide, however, two different mods can't be acquired at the same time: S1 flips from one mod to an another. For example, in the North Atlantic there is a large area in which not a single WV were acquired, since only IW are acquired there**. Fig.3.1.2.** shows an example of the daily acquisitions by S1-A and S1-B for 2020-09-01. As can be seen, during one day the following images were acquired:

- 54 S1 WV tracks/orbits with 6480 individual imagettes,
- 894 S1 IW individual images,
- 229 S1 EW individual images,
- 10 S1 SM individual images

Fig.3.1.3 (Pleskachevsky et al., 2022) shows an example of Sentinel-1 WV archive processing for one day (right half) and one month (left half). A data gap in North Atlantic is quite visible.
For TerraSAR-S, StrimMap (SM) SM L1 Multi-Look Ground Range Detected (MGD) standard products has been used.

- **TS-X SM** mode with ca. **30 km × 50 km** and 3 m resolution, **pixel spacing 1.2 – 4.5 m** differs for spatially enhanced (SE) and radiometric enhanced (RE) products. **In ocean**, TS-X SM are mostly acquired **in costal areas**, only individual images can be found in open ocean. More than **70%** of all ocean-included images are acquired over harbours and include **less than 5 km water from landline**.

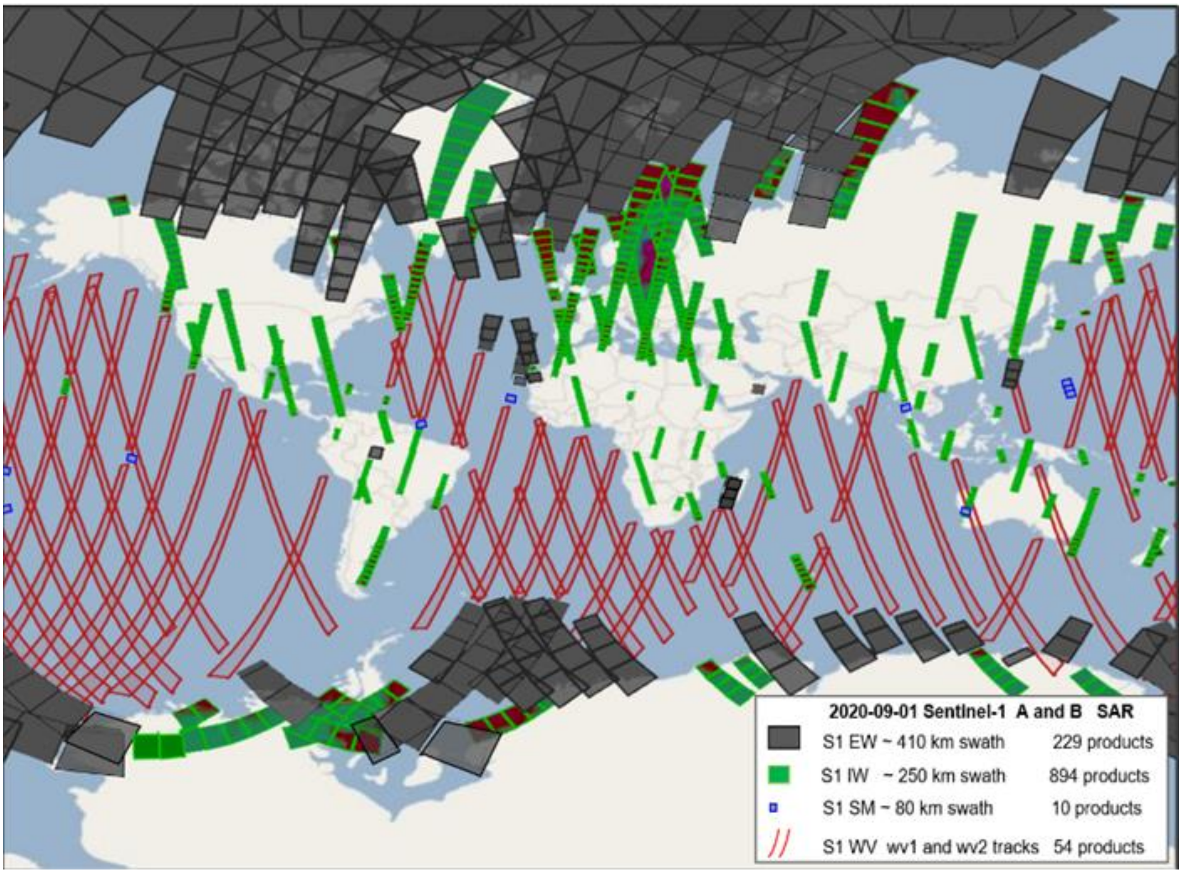

Figure 3.1.2: An example of S1-A and S1-B worldwide acquisitions on 2020-09-01. There are 54 S1 WV tracks (red), 894 S1 IW images (green), 229 S1 EW images (grey) and 10 StripMap images (blue).

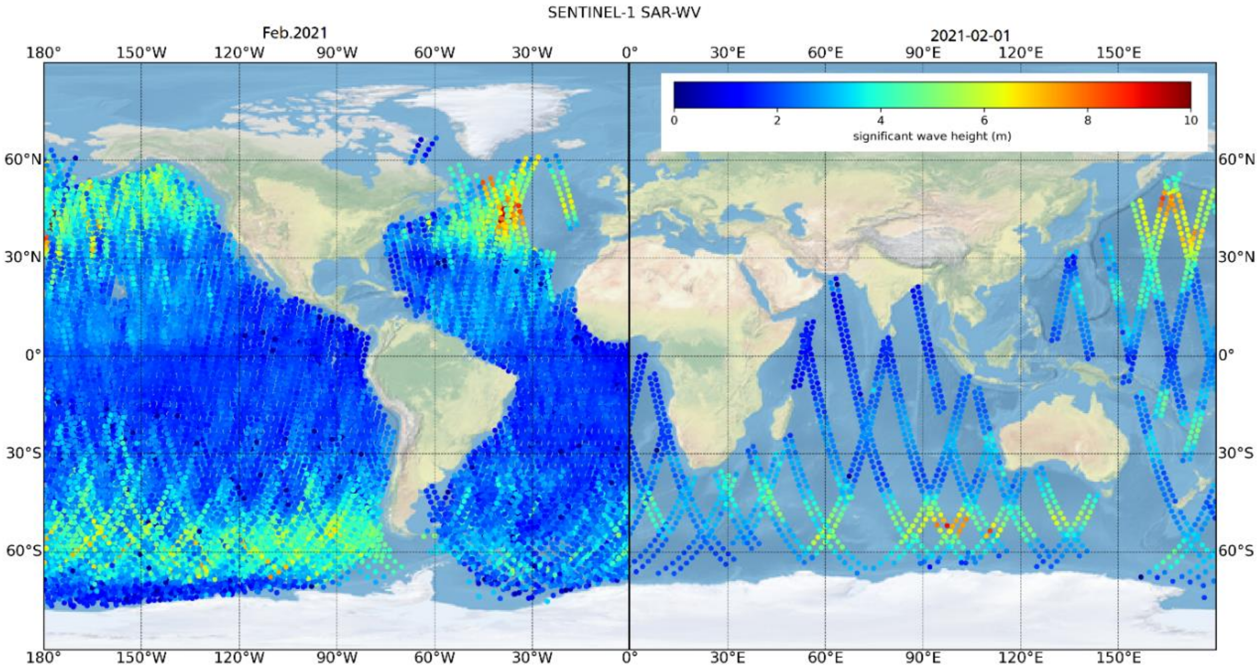

Figure 3.1.3: Example of Sentinel-1 Wave Mode WV archive processing. On the right half of the globe only one-day of acquisitions is displayed, on the left half all data acquired during February 2021.

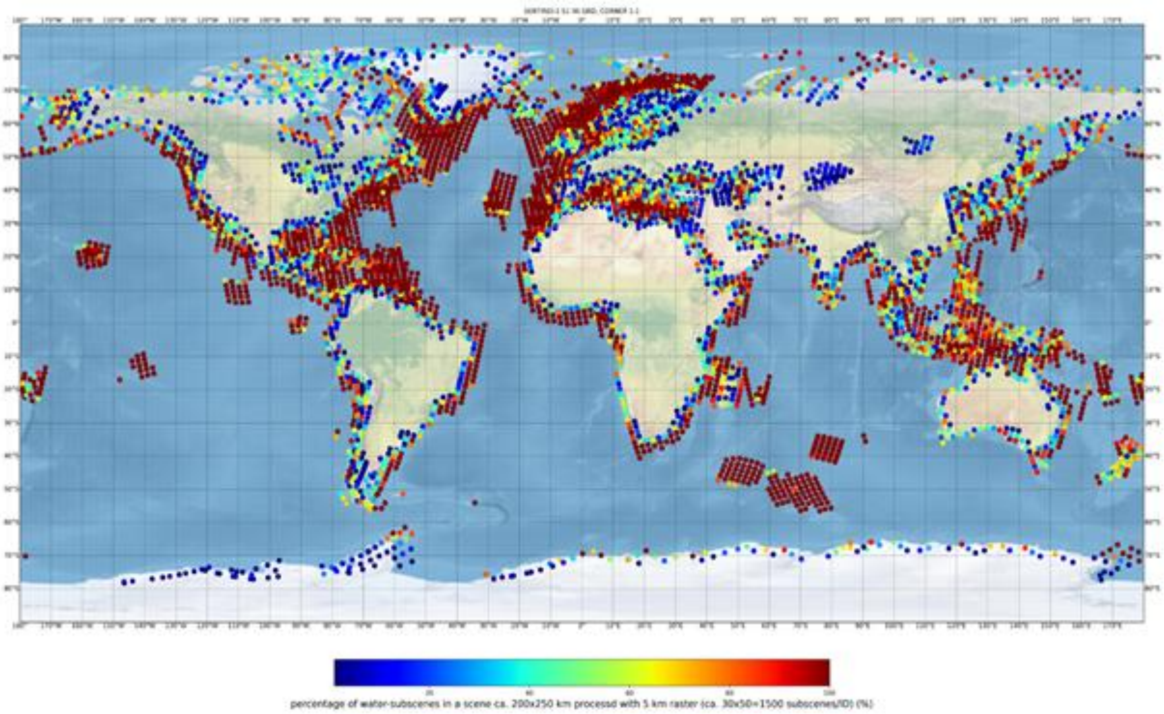

Figure 3.1.4: Example coverage of marine scenes S1 IW in 2020. Shown are upper-left corner coordinates of each ID product, colour shows percentage of processed points (to filtered out, land) for each ID.

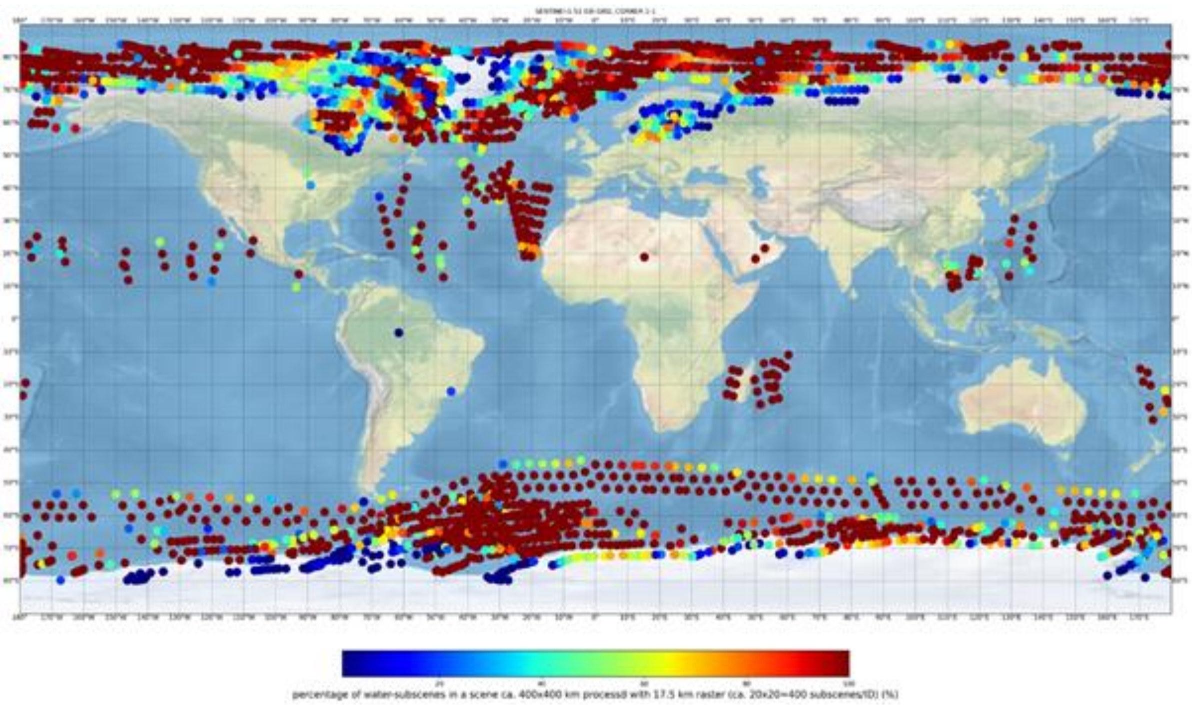

Figure 3.1.5: Example coverage of marine scenes S1 EW in 2020. Shown are upper-left corner coordinates of each product ID, color means percentage of water points (to filtered out, land) for each ID.

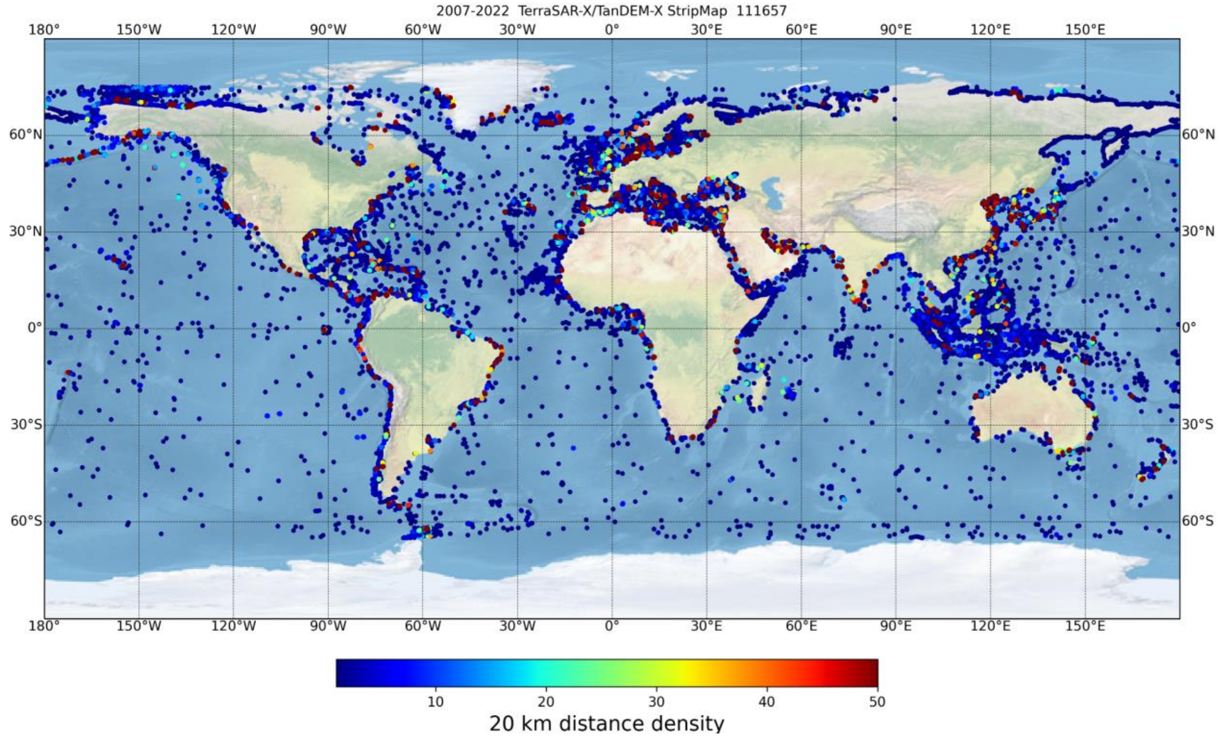

Figure 3.1.6: Example coverage of marine scenes TS-X SM (ca. 30 km×50 km, ca.3 m pixel spacing) in 2020. Shown are upper-left corner coordinates of each product ID, colour shows density for 20 km distance.

3.1.4. SAR_SeaStaR processing workflow

The technical realisation of the **SAR_SeaStaR** algorithm in the Sea State Processor (SSP) software-packet using modular architecture allows to switch SVM functions to a more modern one when expanding the number of training events, or disable it. In this case, the processor gives the results only using linear regression CWAWE_EX.

The pre-filtering of artefacts and outliers is important to obtain uncontaminated image regions, because contaminated regions interfere with the extraction of image features of sea state. A direct application of the Empirical Model Function (EMF) to the features extracted from a subscene often leads to inaccuracies in SWH estimation with outliers in the range of meters. The sources of these errors are in the first place a number of natural and man-made artefacts like current boundaries, wind streaks, ships, wind farm constructions or buoys. They can be divided into two classes: radar echo is much stronger than background backscatter (e.g. ships) or radar echo is much weaker than background backscatter (e.g. oil slicks). A pre-filtering procedure to recognize and possibly remove the signals not produced by sea state before the analysis is applied by replacing the outlier pixels in the current sub-subscene by the mean value of the subscene.

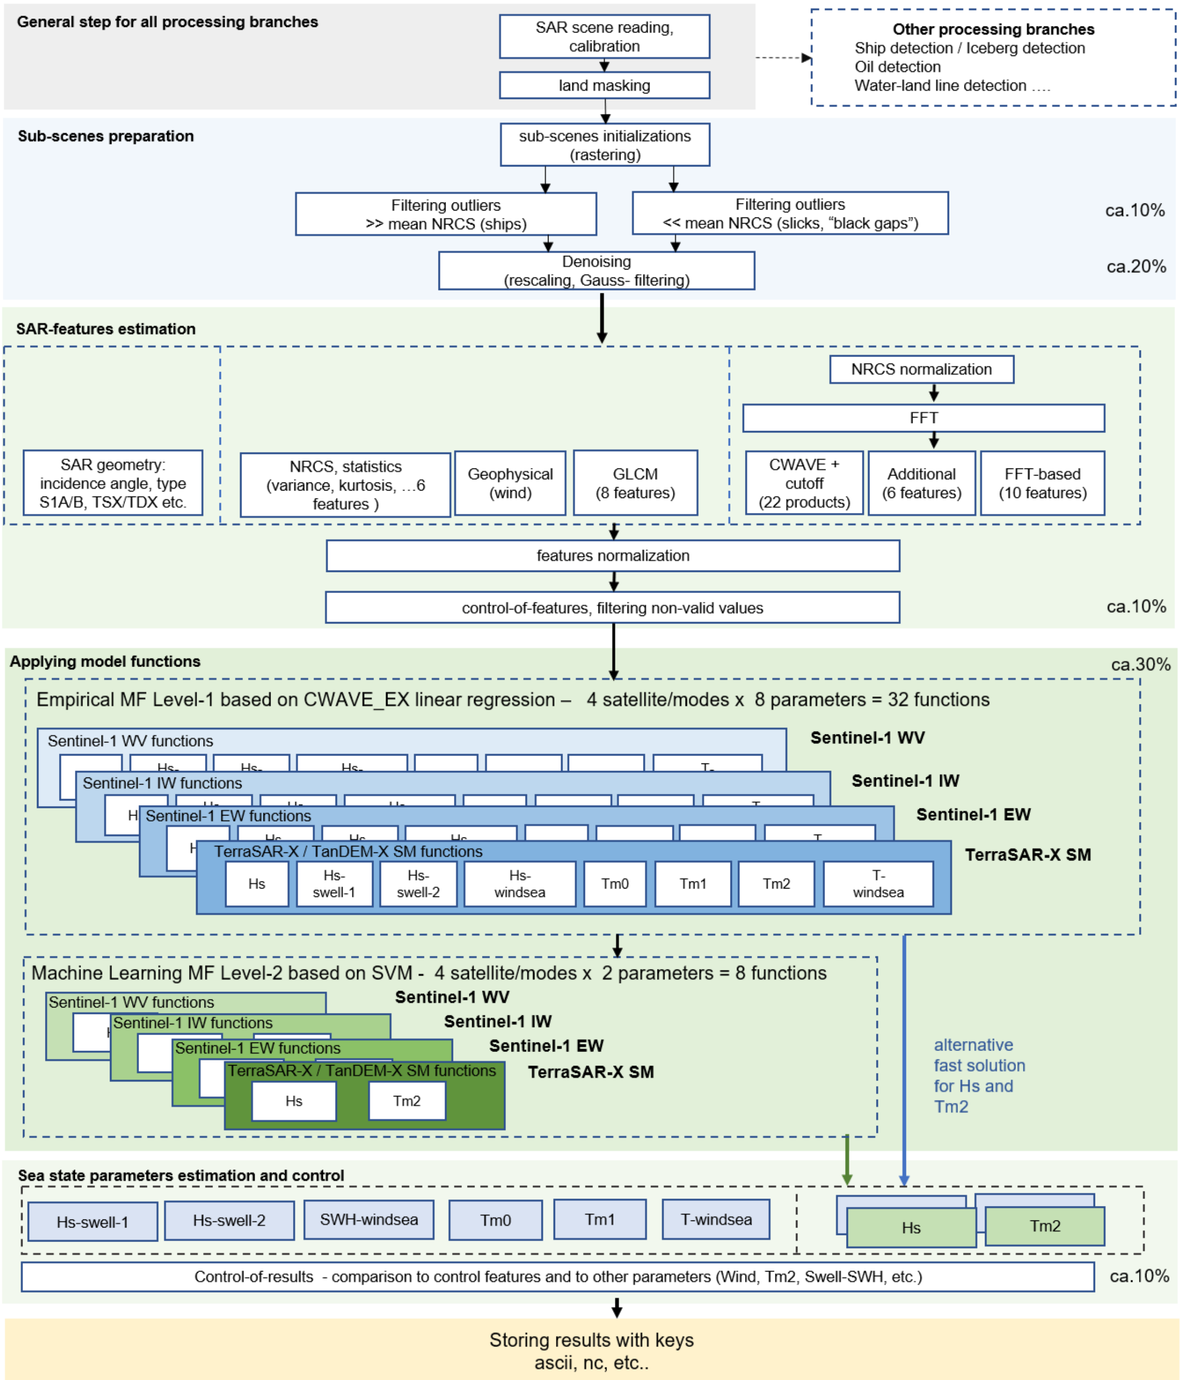

Figure 3.1.7: SAR-SeaStaR algorithm workflow realised in Sea State Processor. For each processing operation, the approximate effects in percent on resulting RMSE in terms of wave height are shown on the right; for the model functions, this is compared to CWAVE method.

3.1.5 Algorithm Definition

Input Data

- S1 WV SLC products
- S1 IW GRD products
- S 1 EW GRD products
- TerraSAR-X SM RE products

Output Data

- Primary:
    - SWH total significant wave height
    - Quality Flag
- Secondary:
    - Tm0-1 mean period (TM0-1)
    - Tm1 first moment period (TM1)
    - Tm2 second moment period (TM2)
    - Tmw windsea period (TMW)
    - SWH dominant swell wave height (SW1)
    - SHW secondary swell wave height (SW2)
    - SWH windsea wave height (SWW)

3.2. Deep learning based algorithm for the inversion of wave spectra from SAR cross spectra

3.2.1 General description

For Phase 2 of the Sea State CCI+ project, ODL is developing a deep learning based approach [Goodfellow et al 2016., Reichstein et al. 2019, Camps-Valls et al. 2021, Quach et al. 2020] to produce and incorporate wave partition parameter estimations from SAR data into the CCI+ Sea State data product.

The estimation of wave partition parameters from SAR data is a complex problem, particularly due to the non-linear nature of the SAR imaging mechanism. Currently, this problem is tackled by a direct inversion of the two-dimensional wave spectra from SAR-derived cross-spectra, followed by a subsequent partitioning of the estimated directional wave spectra and a partition-based estimation of wave partition parameters.

In particular, the current inversion scheme is based on the geophysical and geometrical modeling of the imaging process, and relies on multiple hypothesis and approximations throughout a series of processing stages, which can hinder inversion performance. In particular, it is based on the ocean-to-SAR cross-spectral transform, which can be split into a sum of a quasi-linear and a non-linear term. The inversion model, however, is based solely on the quasi-linear term. As such, the inversion of directional wave spectra requires the removal of the non-linear contribution of wind sea from the SAR cross spectra, by means of the closed-form equation, followed by a quasi-linear inversion of the obtained linear approximation of the cross-spectra to produce estimates of the directional wave spectra. The obtained spectra are then partitioned, via a watershed algorithm, and parameters are estimated for each identified partition.

Based on recent advances in deep learning, ODL is, at this stage, exploring deep learning approaches to improve on the current scheme for the inversion of directional wave spectra from SAR data, before converging towards a final architecture that can be applied to the different SAR missions to be processed for inclusion into the CCI Sea State data product.

The choice of deep learning based methods relies on multiple identified methodological advantages that make neural networks an appealing alternative to tackle the issue of wave directional spectra inversion from SAR data. Indeed, regarding the limitations associated with the SAR imaging mechanism, deep neural networks can accommodate additional data as input with relatively minimal modifications, which is particularly relevant to exploit ancillary data (such as wind). Additionally, problem-specific custom loss functions could also help tackle these issues by enhancing interpretability (Murdoch et al., 2019; Molnar et al., 2020; Gennatas et al., 2020; Li et al., 2022) and allowing for the introduction of physically-informed constraints (Karniadakis et al., 2021; Raissi et al., 2017; Thuerey et al., 2021; Hao et al., 2022). Furthermore, choosing a network architecture exploiting an encoder-decoder structure could help prevent issues regarding the high-dimensionality of SAR images, as well as provide some linearization capabilities to deal with the non-linear nature of the SAR imaging process. Finally, extending and combining neural network architectures with other frameworks could help up further tackle issues related to the problem of wave partition parameter retrieval from SAR imaging data, such as uncertainty characterization (e.g. ensemble methods (Dietterich, 2000), bayesian networks (Goan et al., 2020) and state-time data augmentation (Shorten et al., 2019)), identification of relevant information within SAR scenes (e.g. attention-based mechanisms (Vaswani et al., 2017; Niu et al., 2021)), and retrieval performance enhancement and improved generalization (e.g. GANs (Goodfellow et al., 2020) for data augmentation), among others.

3.2.2 Function

In this context, ODL is developing a deep learning based algorithm focusing directly on the inversion of directional wave spectra from SAR cross-spectra. The computation of wave partition parameters from the obtained directional wave spectra is to be achieved by direct partitioning of the produced directional wave spectra, using a watershed partitioning algorithm, followed by the computation of partition parameters for each detected partition.

As such, the proposed algorithm aims at leveraging machine learning and deep learning techniques to improve on the current model-based inversion, and perform a fully data-driven inversion of the wave directional spectra from SAR cross-spectra. The final objective is thus to provide **estimations** of **wave partition parameters**, including **partition significant wave height** ($H_s^{partition\text{-}n}$), **partition peak wavenumber** ($K_{peak}^{partition\text{-}n}$) and **azimuth**($\theta_{peak}^{partition\text{-}n}$), and **spectral** ($K_{sp}^{partition\text{-}n}$) and **directional spread** ($\theta_{sp}^{partition\text{-}n}$) from SAR imagery, in particular for the wave modes of Sentinel 1, ERS and Envisat.

3.2.3 Algorithm Definition

Input Data

- S1 WV SLC products (for the initial model, dataset release V4)
- Envisat WM ESA WVI products (adaptation of initial model, dataset release V5)
- ERS-1 WM products (ERS format) (adaptation of initial model, dataset release V5)
- ERS-2 WM products (ERS format) (adaptation of initial model, dataset release V5)

Output Data

- Primary:
    - N-th Partition SWH H_s^n
    - N-th Partition peak wavenumber K_{peak}^n
    - N-th Partition peak azimuth θ_{peak}^n
    - N-th Partition wavenumber spread K_{sp}^n
    - N-th Partition directional spread θ_{sp}^n

3.2.4. Mathematical Statement

3.2.4.1. Inversion of wave spectra from SAR cross-spectra

a) Problem formulation

We formulate the inversion of the directional wave spectra from SAR-derived cross-spectra as an image processing problem, where we aim at learning the transfer function f between SAR cross-spectra and the corresponding directional wave spectra.

Given the real and imaginary parts of the SAR-derived cross spectra $[P_{re}(k,\theta), P_{im}(k,\theta)]$, where k and θ represent, respectively, wavenumber and direction, we aim at approximating the non-linear function f that verifies:

$$S(k,\theta) = f(P_{re}(k,\theta), P_{im}(k,\theta)) \qquad \textbf{(3.1)}$$

where $S(k,\theta)$ is the corresponding wave directional spectra.

From this formulation, we propose a deep-learning based approach, relying on deep a neural network, that approximates function f . To this end, the network is trained by taking the pair of real and imaginary parts of the SAR cross-spectra $[P_{re}(k,\theta), P_{im}(k,\theta)]$ as input and producing, as output, the corresponding colocalized wave spectra $S(k,\theta)$, as depicted in **Figure 3.2.1.**

Given that the inversion is formulated as an image processing problem, the proposed deep learning approach will rely on a convolutional neural network, which are particularly well adapted to the processing of two dimensional signals.

The basic building block of convolutional neural networks is the convolutional layer. A convolutional layer takes as input a two dimensional signal, a convolution with a trainable kernel is then applied to the input, a bias term is added to the result of the convolution, and a non-linear activation function is then applied. Mathematically, this can be formulated as:

$$O(x_o, y_o) = f_a(I(x_i, y_i) * K(x_i, y_i) + b(x_o, y_o)) \qquad \textbf{(3.2)}$$

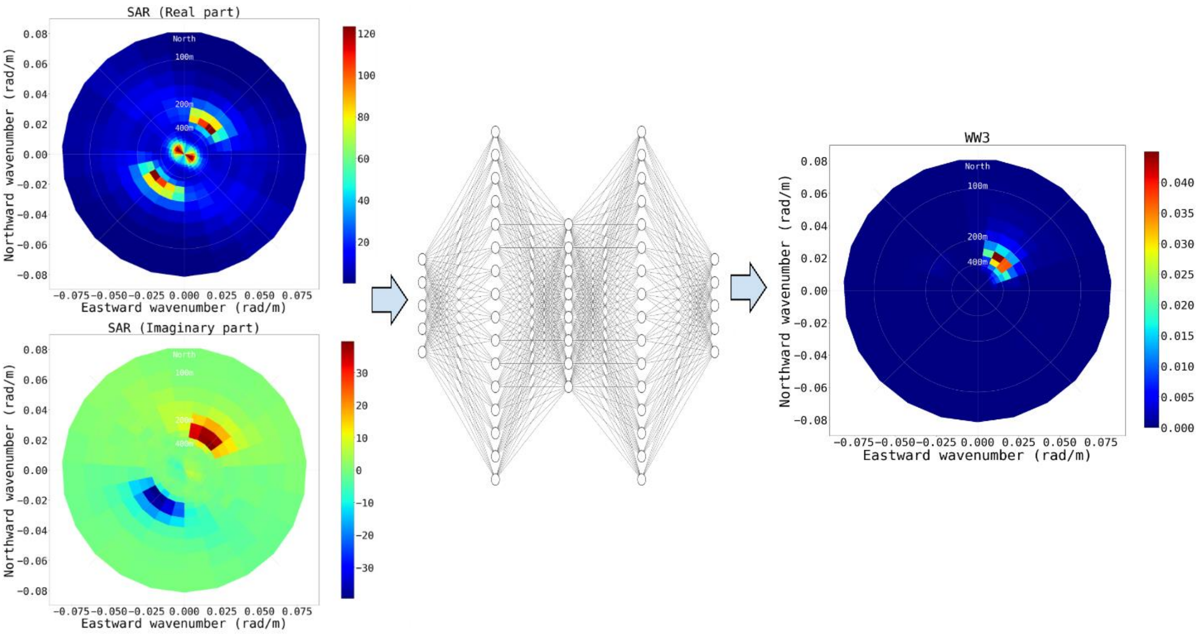

Figure 3.2.1: Schematic representation of the proposed deep learning approach

b) Deep-learning based schemes

where $I(x_i, y_i)$ and $O(x_o, y_o)$ represent, respectively, the layer's input and output, $K(x_k, y_k)$ and $b(x_o, y_o)$ represent, respectively, the trainable convolution kernel and bias term, and f_a is the non-linear activation function.

Importantly, the dimensions of the input, output and kernel may be different, allowing the network to dilate or expand the number of dimensions. Special layers designed to dilate and expand intermediate results, namely MaxPooling (an operator that locally takes the maximum value of the input) and UpSampling (which produces a higher dimensional output by locally repeating the input value along several dimensions), are used to further manipulate intermediate results.

A deep convolutional network consists of multiple convolutional layers stacked and connected sequentially together to produce a complex approximator of non-linear functions. Each layer may also contain multiple channels, consisting of parallel convolutional layers, to increase modeling complexity.

At each layer, the trainable kernels and biases are learned from data by minimizing a loss function. This is done iteratively, by means of a gradient descent based algorithm. The backpropagation approach, in which gradient updates at the output layer are computed and propagated backwards throughout the network, allows for the derivation, at each layer, of the necessary updates to the layer's parameters.

The learning process performs multiple iterations of this minimization. At each iteration, a randomly sampled subset of the learning data (known as a batch) is used. The term epoch is used to signify the processing, by multiple batches, of all of the training data through the network.

As such, multiple architectures of deep convolutional networks may be constructed by adequately choosing the number of layers and channels, the kernel sizes and strides, and the choice of the layer activation functions and the loss function to be minimized globally, among others. The learning process must also be tuned by adequately selecting a number of hyperparameters, including the learning rate, which controls the magnitude of parameter updates with respect to the gradient descent, the number of epochs or stopping criteria used in training, and the gradient-descent based algorithm to perform the iterative minimization of the loss function and associated parameter updates, among others.

c) Proposed network architecture

As previously stated, ODL is currently testing and comparing different network architectures to tackle the inversion of directional wave spectra from SAR cross-spectra. In particular, preliminary results demonstrate the relevance of a modified U-net based architecture (Ronneberger et al., 2015), introducing a softmax layer inside skip connections to select relevant parts of the spectra.

The proposed U-net based architecture relies on stacking convolutional layers to build an encore-decoder architecture. As such, the first part of the network relies on convolutional layers and max-pooling operators to build a contractive network, which effectively reduces dimensionality at each successive layer. In this manner, an information bottleneck is achieved, forcing the network to learn a lower dimensional representation, thus leveraging only the most relevant information within the training dataset. The lower dimensional representation is then fed into the second part of the U-net. consisting of an expansive network exploiting convolutional layers and upsampling operators to expand dimensionality at each successive layer.

Importantly, encoder-decoder architectures are useful to tackle the curse of dimensionality, as previously stated, but are prone to the loss of context information at each layer, which translates into two-dimensional spectral structural information in the case of wave spectra. To tackle this problem, U-nets rely on ‘skip connections’ that feed the output of each layer in the contractive path directly into the equivalent layer in the expansive path. In the proposed implementation, skip connections are further modified to include a soft-max activated layer to select, at each considered scale, relevant parts of the spectra. A schematic representation of the proposed architecture is depicted in **Figure 3.2.2**.

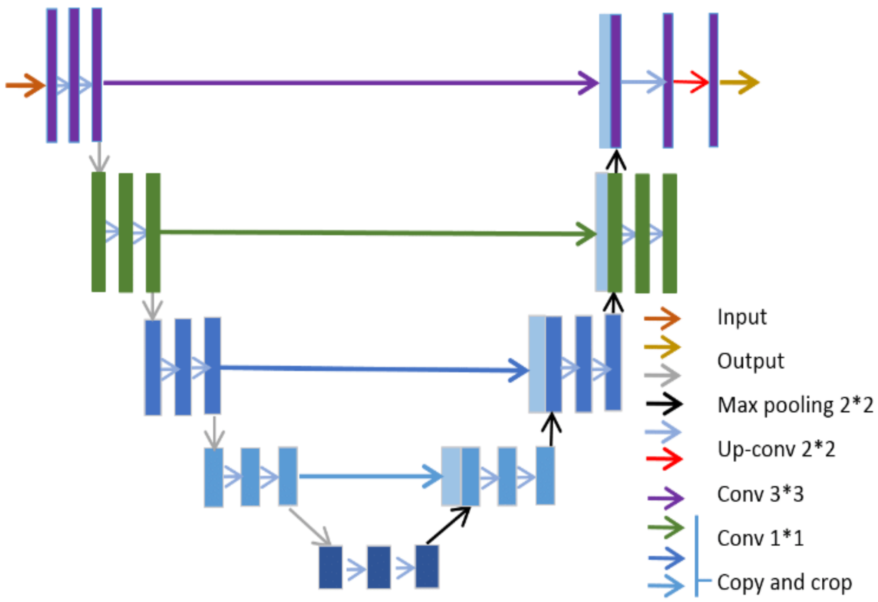

Figure 3.2.2. Schematic representation of the proposed U-net architecture.

The proposed architecture is trained by minimizing a custom loss function, measuring the mismatch between the predicted two dimensional wave spectra $\hat{S}(k,\theta)$ and the ground-truth two dimensional wave spectra $S(k,\theta)$, combining 3 distinct terms:

$$L\left(S(k,\theta),\hat{S}(k,\theta)\right) =$$

$$\lambda_{MAE} \cdot \sum_{k,\theta} MAE\left(S(k,\theta),\hat{S}(k,\theta)\right) + \lambda_{SSIM} \cdot \sum_{k,\theta} \left(1 - SSIM\left(S(k,\theta),\hat{S}(k,\theta)\right)\right) +$$
$$\lambda_F \cdot \sum_k \sum_m \left(RMSE\left(a_m(S(k,\theta)), a_m(\hat{S}(k,\theta))\right) + RMSE\left(b_m(S(k,\theta)), b_m(\hat{S}(k,\theta))\right)\right) \quad \textbf{(3.3)}$$

where λ_{MAE}, λ_{SSIM} and λ_F are user-set hyperparameters used to control the relative importance of each term in the loss function (and were set to λ_{MAE} =0.5, $\lambda_{SSIM} = 0.3$ and $\lambda_F = 0.2$ in our preliminary tests).

The three error terms considered are, respectively:

- A classic error metric, namely the mean absolute error (MAE): $MAE\left(S(k,\theta),\hat{S}(k,\theta)\right) = \frac{1}{N}\sum_n |S_n(k,\theta) - \widehat{S_n}(k,\theta)|$, with N representing the number of colocalized SAR cross-spectra/directional wave spectra pairs in the training batch (here indexed by n).
- Importantly, this error metric was chosen given its enhanced robustness in the presence of outliers.
- A term designed to maintain structural information, namely the structural similarity (SSIM), typically used in image processing applications:

$$SSIM\left(S(k,\theta),\hat{S}(k,\theta)\right) = \frac{2.\mu_S(k,\theta)\,\mu_{\hat{S}}(k,\theta)+c_1}{\mu_S(k,\theta)+\mu_{\hat{S}}(k,\theta)+c_1} + \frac{2.\sigma_S(k,\theta)\,\sigma_{\hat{S}}(k,\theta)+c2}{\sigma_S(k,\theta)+\sigma_{\hat{S}}(k,\theta)+c_2} + \frac{2.\sigma_{S\hat{S}}(k,\theta)+c_2}{\sigma_S(k,\theta)\sigma_{\hat{S}}(k,\theta)+c_2} \quad \textbf{(3.4)}$$

where μ and σ represent, respectively, the means and variances/covariances of the groundruth and predicted wave spectra $S(k,\theta)$ and $\hat{S}(k,\theta)$.

This similarity measure is particularly relevant for the considered application, as the three terms represent, respectively, luminance, contrast and structure in image processing applications, and could be empirically associated to the energy level, variability and spectral structure when applied to directional wave spectra.

Importantly, $c_1 = (q_1 L)^2$, $c_2 = (q_2 L)^2$, with $q_1 = 0.01$ and $q_2 = 0.03$, and L being the dynamic range of the considered wave spectra, are variables that stabilize the division for cases where the denominator is small.

- A physics-informed term aiming at ensuring coherence for directional moments by minimizing the root-mean-square-error for the first 5 directional moments (obtained using a directional Fourier decomposition):

$$\sum_m \left(RMSE\big(a_m(S(k,\theta)), a_m(\hat{S}(k,\theta))\big) + RMSE\big(b_m(S(k,\theta)), b_m(\hat{S}(k,\theta))\big)\right) \quad \textbf{(3.5)}$$

where

$RMSE(y,\hat{y}) = \sqrt{\frac{1}{N}\sum_n (\hat{y}_n - y_n)^2}$ is the root root-mean-squared error,

$a_m(S(k,\theta)), b_m(S(k,\theta))$ and $a_m(\hat{S}(k,\theta))$,$b_m(\hat{S}(k,\theta))$ represent, respectively, the m-th order directional moments of the ground-truth two dimensional wave spectra $S(k,\theta)$ and the predicted two dimensional wave spectra $\hat{S}(k,\theta)$:

$$a_m(S(k,\theta))(k) = \int_\theta S(k,\theta)cos(m\theta)d,\ b_m(S(k,\theta))(k) = \int_\theta S(k,\theta)sin(m\theta)d\theta$$

Specially, this term is particularly relevant for seamless comparison and/or fine-tuning with in-situ measurements from buoys (from which directional moments can be derived with minimal computations)

Architecture details

Regarding other relevant architecture parameter and hyperparameter choices, multiple configurations are being tested. Currently, the best performing architecture relies on the following configuration:

- Inputs/Outputs:
  - The real and imaginary parts of the SAR cross spectra are fed as input into two independent channels in the first layer.
  - The target directional wave spectra occupy a single channel at the network output
  - Inputs/outputs are padded so that their dimensions match the closest power of two (128 in our preliminary tests) for efficiency.
    - The padding considered is a zero padding in wavenumber (which could be further improved by more advanced padding techniques such as mirroring or local averaging), and a circular padding in direction (accounting for periodicity in direction)

- Contractive path
  - The contractive path consists of N levels ($N = 4$ in our preliminary tests).
  - Each level consists of:

        - A convolutional layer of C channels with kernels of size (K_k, K_θ) and stride (s_k, s_θ) ($C = 16$, $K_k = K_\theta = 8$ and $s_k = s_\theta = 2$ in our preliminary tests)
        - A dropout layer with probability p ($p = 0.1$ in our preliminary tests),
        - A second convolutional layer, identical to the first one
        - A batch normalization layer
        - A ReLU activation function to account for non-linearity
        - A max-pooling operator of size (M_k, M_θ) ($M_k = M_\theta = 2$ in our preliminary tests)
    - The number of channels C is duplicated at each subsequent level (yielding 4 levels of 16, 32, 64 and 128 channels in our preliminary tests)

- Expansive path:

    - The expansive path mirrors the contractive path (excluding dropout) and also consists of N levels ($N = 4$ in our preliminary tests).
    - Each level consists of:
        - A transposed convolutional layer (convolution+upsampling) of C channels with kernels of size (K_k, K_θ) and stride (s_k, s_θ) ($C = 128$, $K_k = K_\theta = 8$ and $s_k = s_\theta = 2$ in our preliminary tests)
        - A concatenation between the output of the transposed convolutional layer and the output of the skip connection at the corresponding level in the contractive path
        - A second convolutional layer, identical to the first one
        - A batch normalization layer
        - A ReLU activation function to account for non-linearity
            - The number of channels C is divided in half at each subsequent level (yielding 4 levels of 128, 64, 32 and 16 channels in our preliminary tests)

- Skip connections:

    - Skip connections connect corresponding levels between the contractive and expansive paths to maintain structural information
    - At each level, the skip connection consists of:
        - A convolutional layer with the appropriate number of channels for the level (16, 32, 64 and 128 channels in our preliminary tests), and kernels of size (K_k, K_θ) and stride (s_k, s_θ) ($K_k = K_\theta = 1$ and $s_k = s_\theta = 2$ in our preliminary tests)
        - A softmax activation function to select, at each level, the most relevant parts of the spectra

- Output layer:

    - The output layer combines the 2 channels of the last level of the expansive path into a single channel. while also smoothing the output to prevent artefacts and counteract checkerboard effects (to which convolutional networks are prone to). It consists of:
        - A single channel convolutional layer with kernels of size (K_k, K_θ) and stride (s_k, s_θ) ($K_k = K_\theta = 1$ and $s_k = s_\theta = 1$ in our preliminary tests)
        - A linear activation function

- Output cropping:
  - A cropping operator, needed to counteract the input/output padding, is applied to the network output to retrieve its central part and obtain the final prediction of the directional wave spectra

- Training:
  - The colocalized SAR cross-spectra and directions wave spectra dataset is split into training, test and validation datasets using a random 64/20/16 split
  - The network is trained using the Adam optimizer (Kingma et al. 2014)
  - The initial learning rate considered is $lr = 10^{-5}$
  - The learning rate is updated iteratively following an exponential decay scheme with a decay rate of $d = 10^{-6}$
  - The training uses batches of size $b = 256$ is performed for a maximum of $E_{max} = 10000$ epochs, or when the validation loss does not decrease after 25 epochs.

3.2.4.2. Partitioning and partition parameter estimation

To obtain partition parameters, the directional wave spectra obtained using the proposed deep-learning based approach are subsequently partitioned using a classic watershed approach, and partition parameters are computed at the partition level for each identified partition.

For each retrieved directional wave spectrum $\hat{S}_n(k,\theta)$, and for each partition p identified by the watershed partitioning scheme,, denoted hereafter as $\widehat{S_n}^p(k,\theta)$, the following parameters are computed:

- Partition significant wave height H_s^p

$$H_s{}^p = 4\sqrt{\int_k \int_\theta \widehat{S_n}^p(k,\theta) k \, dk \, d\theta} \quad \textbf{(3.6)}$$

- Partition peak wavenumber k_{peak}^p

$$k_{peak}{}^p = argmax_k\left(\widehat{S_n}^p(k,\theta)\right)$$

(3.7)

- Partition peak azimuth θ_{peak}^p

$$\theta_{peak}{}^p = argmax_\theta\left(\widehat{S_n}^p(k,\theta)\right) \quad \textbf{(3.8)}$$

- Partition spectral spread K_{sp}^p

$$k_{sp}{}^p = \hat{\sigma}_k \ni [\widehat{\mu_k}, \hat{\sigma}_k] = argmin_{\mu_k, \sigma_k}\left(RMSE\left(N(\mu_k, \sigma_k) - \left(\int_\theta \widehat{S_n}^p(k,\theta)\, k \, d\theta\right)\right)\right) \quad \textbf{(3.9)}$$

- Partition directional spread θ_{sp}^p

$$\theta_{sp}{}^p = \hat{\sigma}_\theta \ni [\widehat{\mu_\theta}, \hat{\sigma}_\theta] = argmin_{\mu_\theta, \sigma_\theta}\left(RMSE\left(N(\mu_\theta, \sigma_\theta) - \left(\int_k \widehat{S_n}^p(k,\theta)\, dk\right)\right)\right) \quad \textbf{(3.10)}$$

where $N(\mu,\sigma)$ denotes the normal distribution of mean μ and variance σ.

In particular, it should be noted that partition spectral and directional spreads ${k_{sp}}^p$ and ${\theta_{sp}}^p$ are estimated as the variances of two independent Gaussian functions fitted, respectively, onto the partial integrals of partition wave directional spectra over wavenumber k and direction θ.

3.2.5. Processing details and workflow

As previously stated, ODL is currently testing and comparing different network architectures to tackle the retrieval of wave partition parameters from Sentinel 1-A and Sentinel 1-B Wave Mode (S1-A WV & S1-B WV) data, with the target wave partition parameters being computed from colocalised reference wave numerical model data (Alday and Ardhuin 2023). The resulting S1-A/S1-B derived wave partition parameters dataset is to be included into the v4 release of the CCI+ Sea State Climate Record.

The Sentinel-1 derived model developed during the first stage will be used as a base model that will be adapted and applied to ERS-1, ERS-2 and Envisat WV data, with the target wave partition parameters being computed from colocalised reference wave numerical model data.

Moreover, the base model may also be adapted and applied to Sentinel-1A and Sentinel-1B IW and EW data, depending on the availability of SAR cross-spectra computed from these datasets, to produce a demonstration dataset that illustrates the relevance of the proposed approach for Sentinel-1A and Sentinel1-B IW and EW modes. In this regard, the demonstration dataset, not expected to be made part of the public release, will not necessarily span the whole Sentinel-1A and Sentinel-1B mission lifespan, with no reprocessing effort for the whole dataset being envisaged.

Regarding model adaptation to other missions, the preferred approach is based on exploiting the initial 'base' model, trained on Sentinel-1 data, and adapting its characteristics and parametrizations to the subsequent SAR missions considered on a mission-per-mission basis. This approach would also require retraining the model, at least partially, on the associated datasets. To this end, fine-tuning (Fu et al. 2023) and transfer learning (Thrun & Pratt 2012) approaches will provide powerful tools for this task. Alternatively, a second approach could rely on the joint standardisation of all missions to obtain a consistent multi-mission dataset. The standardisation across datasets will be achieved by means of classical resampling and interpolation techniques. We note that some degradation and information loss is to be expected following this adaptation. The standardised multi-mission dataset can then be exploited to retrain/fine-tune the base model and produce a single multi-mission model. In this approach, the base model characteristics and parametrizations might need to be modified to better suit the multi-mission dataset (as is the case for the alternative single-mission adapted multi-model approach).

Regarding training and verification/validation, as previously stated, the deep learning based inversion of wave partition parameters from SAR cross-spectra will be initially trained on colocalized Sentinel-1 and WaveWatch3 data, with results for the whole Sentinel-1 mission span being integrated into the CCI Sea State V4 dataset release. The training dataset will be a global dataset spanning all oceans. The training dataset period will span at least a whole year, which should be representative enough and include enough variability to produce a robust neural network with adequate generalization capabilities. Internal verification of the proposed methods will partially exploit the training dataset. To this end, the whole year will be split into a training dataset, an in-training verification dataset (to monitor and prevent model overfit) and an internal verification dataset.

Following a similar approach, the developed model will be adapted, retrained and verified/validated on colocalized Envisat/ERS-1/ERS-2 and WaveWatch3 data, with results for the whole span of these missions being integrated into the CCI Sea State V5 dataset release. Similarly, the training datasets for these subsequent missions will be global datasets spanning all oceans. The training dataset periods for each mission will span at least a whole year, which should be representative enough and include enough variability to produce a robust neural network with adequate generalization capabilities. Internal verification of the proposed methods will partially exploit these training datasets. To this end, each whole year dataset will be split into a training dataset, an in-training verification dataset (to monitor and prevent model overfit) and an internal verification dataset.

Internal verification may also exploit colocalised data from in-situ buoys (when and where available), which could provide estimates of two-dimensional wave spectra that can be partitioned to produce additional wave partition parameters. These in-situ derived wave partition parameters could then be exploited as a ground-truth reference to further verify the performance and accuracy of the proposed methodology. If sufficient in-situ data is available, fine-tuning (Fu et al., 2023) may be explored as a means to further increase the generalization and inversion capabilities of the proposed approach.

Both internal and external validation performed during the first proposed development stage, as well as feedback from the V4 release of the CCI+ Sea State Data Product, will serve as inputs for the second development stage to tackle any caveats or unforeseen issues, and further refine and improve the algorithms developed.

4. References

References altimetry

Brown, G., 1977. The average impulse response of a rough surface and its applications. IEEE Transactions on Antennas and Propagation 25, 67–74. https://doi.org/10.1109/TAP.1977.1141536

Dodet, G., Piolle, J.-F., Quilfen, Y., Abdalla, S., Accensi, M., Ardhuin, F., Ash, E., Bidlot, J.-R., Gommenginger, C., Marechal, G., Passaro, M., Quartly, G., Stopa, J., Timmermans, B., Young, I., Cipollini, P., Donlon, C., 2020. The Sea State CCI dataset v1: towards a sea state climate data record based on satellite observations. Earth System Science Data 12, 1929–1951. https://doi.org/10.5194/essd-12-1929-2020

Hayne G.S., 1980. Radar altimeter mean return waveforms from near-normal-incidence ocean surface scattering IEEE Transactions on Antennas and Propagation, 28 (1980), pp. 687-692

Nelder, J.A., & Mead, R. (1965). A Simplex Method for Function Minimization. Comput. J., 7, 308-313.

Passaro, M., Cipollini, P., Vignudelli, S., Quartly, G.D., Snaith, H.M., 2014. ALES: A multi-mission adaptive subwaveform retracker for coastal and open ocean altimetry. Remote Sensing of Environment 145, 173–189. https://doi.org/10.1016/j.rse.2014.02.008

Queffeulou, P., 2016. Validation of Jason-3 altimeter wave height measurements. Presented at the OSTST.

Sepulveda, H., Queffeulou, P., Ardhuin, F., 2015. Assessment of SARAL/AltiKa Wave Height Measurements Relative to Buoy, Jason-2, and Cryosat-2 Data. Marine Geodesy 38, 449–465. https://doi.org/10.1080/01490419.2014.1000470

Wolberg, A., 2006. Data analysis using the method of least squares: extracting the most information from experiments. Springer Science & Business Media

References SAR

Alday, M., & Ardhuin, F. (2023). On consistent parameterizations for both dominant wind-waves and spectral tail directionality. Journal of Geophysical Research: Oceans, e2022JC019581.

Camps-Valls, G., Tuia, D., Zhu, X. X., & Reichstein, M., 2021. Deep learning for the Earth Sciences: A comprehensive approach to remote sensing, climate science, and geosciences (p. 432). Wiley. doi: 10.1002/9781119646181

Dietterich, Thomas G., 2020. Ensemble methods in machine learning. International workshop on multiple classifier systems. Berlin, Heidelberg: Springer Berlin Heidelberg.

Fu, Z., Yang, H., So, A. M.-C., Lam, W., Bing, L., & Collier, N., 2023. On the effectiveness of parameter-efficient fine-tuning. Proc. AAAI Conf. Artificial Intelligence, 37(11), 12799-12807.

Gennatas, E. D., Friedman, J. H., Ungar, L. H., Pirracchio, R., Eaton, E., Reichmann, L. G., Interian, Y., Luna, J. M., Simone, C. B., Auerbach, A., Delgado, E., van der Laan, M. J., Solberg, T. D., & Gilmer Valdes, 2020. Expert-augmented machine learning. Proc. Nat. Acad. Sci., 117(9), 4571–4577. doi: 10.1073/pnas.1906831117

Goan, E., and C. Fookes, 2020.Bayesian neural networks: An introduction and survey. Case Studies in Applied Bayesian Data Science: CIRM Jean-Morlet Chair, Fall 2018: 45-87.

Goodfellow, I, Bengio, Y., Courville, A., 2016., Deep Learning. MIT Press.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. 2020. Generative adversarial networks. Communications of the ACM, 63(11), 139-144

Hao., Z., Wang, Z., Bai., D., Tong, X., 2022. Surface Defect Segmentation Algorithm of Steel Plate Based on Geometric Median Filter Pruning. Sec. Bionics and Biomimetics. Vol. 10 2022. https://doi.org/10.3389/fbioe.2022.945248

Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L., 2021. Physics-informed machine learning. Nature Reviews Physics, 3(6), 422-440.

Kingma, Diederik P. "Adam: A method for stochastic optimization." arXiv preprint arXiv:1412.6980 (2014).

Li, X., Xiong, H., Li, X., Wu, X., Zhang, X., Liu, J., ... & Dou, D., 2022. Interpretable deep learning: Interpretation, interpretability, trustworthiness, and beyond. Knowledge and Information Systems, 64(12), 3197-3234.

Molnar, C., 2020. Interpretable machine learning. Lulu.com

Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., & Bin Yu. (2019). Definitions, methods, and applications in interpretable machine learning. Proc. Nat. Acad. Sci., 116(44), 22071–22080. doi: 10.1073/pnas.1900654116

Niu, Z., G. Zhong, and H. Yu. A review on the attention mechanism of deep learning. Neurocomputing 452 (2021): 48-62.

Pleskachevsky, A., Jacobsen, S., Tings, B., and E. Schwarz, 2019. Estimation of sea state from Sentinel-1 Synthetic aperture radar imagery for maritime situation awareness. IJRS, Vol. 40-11, pp. 4104-4142.

Pleskachevsky, A., Tings, B., and S, Jacobsen, 2022: Multiparametric Sea State Fields from Synthetic Aperture Radar for Maritime Situational Awareness, RSE, vol. 280, 22 pp.

Pleskachevsky, A., Tings, B., Jacobsen, S., Wiehle, S., Schwarz, E., and D. Krause, 2024: A System for Near Real Time Monitoring of the Sea State using SAR Satellites. IEEE Transactions on Geoscience and Remote Sensing, VOL. 62, 2024

Pleskachevsky, A., W. Rosenthal, and S. Lehner, 2016: Meteo-Marine Parameters for Highly Variable Environment in Coastal Regions from Satellite Radar Images." JPRS, Vol. 119, pp. 464-484., 2016.

Quach, B., Y Glaser, JE Stopa, AA Mouche, P Sadowski, 2020. Deep learning for predicting significant wave height from synthetic aperture radar, IEEE Trans. Geosci. Rem. Sens. 59 (3), 1859-1867.

Raissi, M., Perdikaris, P., & Karniadakis, G. E., 2017. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. arXiv preprint arXiv:1711.10561.

Reichstein, M., Camps-Valls, G., Stevens, B. et al., 2019. Deep learning and process understanding for data-driven Earth system science. Nature 566, 195–204. doi: 10.1038/s41586-019-0912-1

Ronneberger, O., Fischer, P., Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015. MICCAI 2015. Lecture Notes in Computer Science(), vol 9351. Springer, Cham. https://doi.org/10.1007/978-3-319-24574-4_28

Shorten, C., and Taghi M. Khoshgoftaar, 2019. A survey on image data augmentation for deep learning. Journal of Big Data 6.1, 1-48.

Thrun S., & Pratt, L., 2012. Learning to Learn. Springer US

Thuerey, N., Holl, P., Mueller, M., Schnell, P., Trost, F., & Um, K., 2021. Physics-based Deep Learning. WWW. https://physicsbaseddeeplearning.org

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I., 2017. Attention is all you need. Advances in neural information processing systems, 30.