THE EXPERT FLOOD MONITORING ALLIANCE

Provision of an Automated, Global, Satellite-based Flood Monitoring Product for the Copernicus Emergency Management Service

Product Definition Document (PDD)

Document ID: GFM D6 Product Definition Document

Issue: 1

Version: 1.4

Issue Date: 20.04.2022

Issue By: Patrick Matgen (LIST)

THE EXPERT FLOOD MONITORING ALLIANCE

Provision of an Automated, Global, Satellite-based Flood Monitoring Product for the Copernicus Emergency Management Service

Dissemination Level		
PU	Public	x
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Document Release Sheet

Book captain:	Patrick Matgen		Date 23/09/2021
Approval:	Dragana Milinkovic	Signature	Date
Endorsement:	CEMS Manager Dr Peter Salamon	Signature	Date
Distribution:	Public (PU)		

Contributing Authors:

EODC GmbH
Christoph Reimer (CR), Dragana Milinkovic (DM)

GeoVille
Laurin Müller (LMü), Michaela Seewald (MS)

TU Vienna
Bernhard Bauer-Marschallinger (BBM), Florian Roth (FRo), Felix Reuss (FRe)

CIMA
Luca Molini (LM), Elisabetta Fiori (EF)

LIST
Patrick Matgen (PM), Renaud Hostache (RH), Marco Chini (MC), Ramona Pelich (RP), Yu Li (YL)

DLR
Candace Chow (CC), Marc Wieland (MW), Christian Krullikowski (CK), Sandro Martinis (SM)

Change Log:

Issue Version	Date	Description of Change	Author(s)
0.00	25.11.2020	Created template	RK
0.01	21.04.2021	Implemented structure of document and inserted first algorithm description (GFMS Product - Observed Flood event)	PM
0.02	15.07.2021	Inserted product and algorithm descriptions	all
0.03	01.08.2021	Started editing of document	RP, PM
0.04	01.09.2021	Reviewed the document, identified missing parts	RP, PM
1.0	23.09.2021	Integrated missing parts	all
1.1	23.09.2021	Document formatting	PM, DM
1.2	26.10.2021	JRC document revision	JRC
1.3	04.11.2021	Inserted references	RP, PM
1.4	20.04.2022	Update of algorithm descriptions in section 4.1 ("Output layer S-1 observed flood extent")	RP, CK, FRo

For any clarifications please contact Dragana Milinkovic, GFM Service Manager (GFM_Service_Manager@eodc.eu).

Table of Contents

Acronym List:

ARD	Analysis-Ready-Data
ASCAT	Advanced Scatterometer
CSCDA	Copernicus Space Component Data Access
dB	Decibel
DEM	Digital Elevation Model
DSM	Digital Surface Model
EC	European Commission
EO	Earth observation
ESA	European Space Agency
GDAL	Geospatial Data Abstraction Library
GFC	Global Forest Change
GFM	Global Flood Monitoring
GHSL	Global Human Settlement Layer
GPT	Graph Processing Tool
GRDH	Ground Range Detected at High resolution
HAND	Height Above Nearest Drainage index
HAND-EM	HAND-based Exclusion Map
HRL	High-Resolution Layer
IPF	Instrument Processing Facility
IW	Interferometric Wide-swath
JRC	Joint Research Centre of the European Commission
NRT	Near real-time
OS	Operating system
OSV	Orbit State Vector
PDF	Probability Density Function
PLIA	Projected Local Incidence Angle
POEORB	Precise Orbit Ephemerides files
RESORB	Restituted Orbit Ephemerides files
S1-GBM	Sentinel-1 Global Backscatter Model
SAR	Synthetic Aperture Radar
SNAP	Sentinel Application Platform (ESA's common software architecture for EO processing and analysis, supporting the Sentinel and other sensors).
SRTM	Shuttle Radar Topography Mission
SSE	Sum of Squared Estimate of Errors
WBM	Water Body Mask (one of the auxiliary information masks generated during the COPERNICUS DEM production process).
WSF2015	World Settlement Footprint 2015
Zstd	Zstandard (a lossless data compression algorithm).

1 Introduction

This Product Definition Document (PDD) provides detailed technical specifications for all of the product output layers of the new Global Food Monitoring (GFM) product of the Copernicus Emergency Management Service (CEMS). The PDD provides the reference information required to understand all elements of the various data processing chains, and explains the contents of the GFM product output layers, the main assumptions underlying their generation, and the limitations of the data.

Specifically, the PDD describes in detail the GFM products in terms of the physical and mathematical basis of the algorithms and systems for each step in the product generation. The remainder of this document is structured as follows:

- Section **Error! Reference source not found.** gives an overview of the eleven GFM product output layers and presents the overall product generation workflow.
- Section 3 provides a detailed description of the pre-processing steps that are applied to the Sentinel-1 satellite data, as well as the preparation of all ancillary data.
- Section 4 presents the input and output data, and describes in detail the algorithms used for generating the eleven GFM product output layers.
- Section 5 gives a general description of the GFM product access and dissemination system.
- Finally, Section 6 describes the procedure for the quality assurance of the GFM product.

2 Overview of data products

The eleven output layers of the new Global Flood Monitoring (GFM) product of the Copernicus Emergency Management Service (CEMS) are summarized in Table 1.

Table 1: The eleven output layers of the new Global Flood Monitoring (GFM) product of the Copernicus Emergency Management Service, generated in near real-time from Synthetic Aperture Radar imagery from ESA's Sentinel-1 (S-1) satellites.

#	OUTPUT LAYER	DESCRIPTION	DATA FORMAT
1	**S-1 observed flood extent**	Flooded areas mapped as the difference between S-1 observed water extent and the reference water mask.	Raster (Geotiff); vector (Shapefile).
2	**S-1 observed water extent**	Open and calm water mapped using Sentinel-1 SAR backscatter intensity.	Raster (Geotiff); vector (Shapefile).
3	**S-1 reference water mask**	Normal (i.e. permanent, seasonal) water mapped using Sentinel-1 SAR backscatter intensity.	Raster (Geotiff); vector (Shapefile).
4	**Exclusion mask**	Unclassified areas (e.g. urban areas; dense vegetation; flat and impervious areas; sandy surfaces; topographic effects).	Raster (Geotiff).
5	**Uncertainty values**	Estimated uncertainty of flood mapping, for all areas outside exclusion mask.	Raster (Geotiff).
6	**Advisory flags**	Flags indicating potential reduced quality of flood mapping, due to prevailing environmental conditions (e.g. wind, ice, snow, dry soil), or degraded input data quality due to signal interference from other SAR missions (e.g. Radarsat).	Raster (Geotiff).
7	**S-1 metadata**	Information on the acquisition parameters of the Sentinel-1 data used.	KML.
8	**S-1 footprint**	Image boundaries of the Sentinel-1 data used.	KML.
9	**S-1 schedule**	Next scheduled Sentinel-1 data acquisition.	KML.
10	**Affected population**	Number of people in flooded areas, mapped by a spatial overlay of observed flood extent and gridded population (from the Commission's GHSL project).	Raster (Geotiff).
11	**Affected land cover**	Land cover types (e.g. artificial surfaces, agricultural areas) in flooded areas, mapped by a spatial overlay of observed flood extent and Copernicus Global Land Service data.	Raster (Geotiff).

3 Satellite data pre-processing and ancillary data preparation

3.1 Sentinel-1 data

The Sentinel-1 data build the core dataset of the GFM service. We not only use them as primary information source on the extent of flooded areas and permanent and seasonal water bodies, but also for the (offline-) generation of parameters describing globally the C-band Synthetic Aperture Radar (SAR) backscatter signature at the 20m-pixel-level, supporting the generation of the different product layers.

3.1.1 Sentinel-1 pre-processing

The GFM service ingests observations from the Sentinel-1A/B satellites that are acquired in Interferometric Wide-swath mode and Ground Range Detected at High resolution (Sentinel-1 IW GRDH). The GRDH products consist of focused SAR data that has been detected, multi-looked and projected to ground range using an Earth ellipsoid model, and phase information is lost. The resulting product has approximately square spatial resolution pixels and square pixel spacing with reduced speckle at the cost of worse spatial resolution. In case of the here used high-resolution

product, the raw backscatter amplitude is sampled with a 10x10 m pixel size. For the GFM flood service, we process Sentinel-1 data in VV-polarisation (and neglect the VH-polarisation channel) due to its higher sensitivity in differentiating water from non-water surfaces.

The incoming Level-1 Sentinel-1 datasets must undergo the pre-processing routines before being forwarded to the algorithms for water and flood extent detection. The output of the pre-processing is Analysis-Ready-Data (ARD) that is formatted and gridded, and immediately forwarded to the GFM flood detection engine and added to the Sentinel-1 data cube.
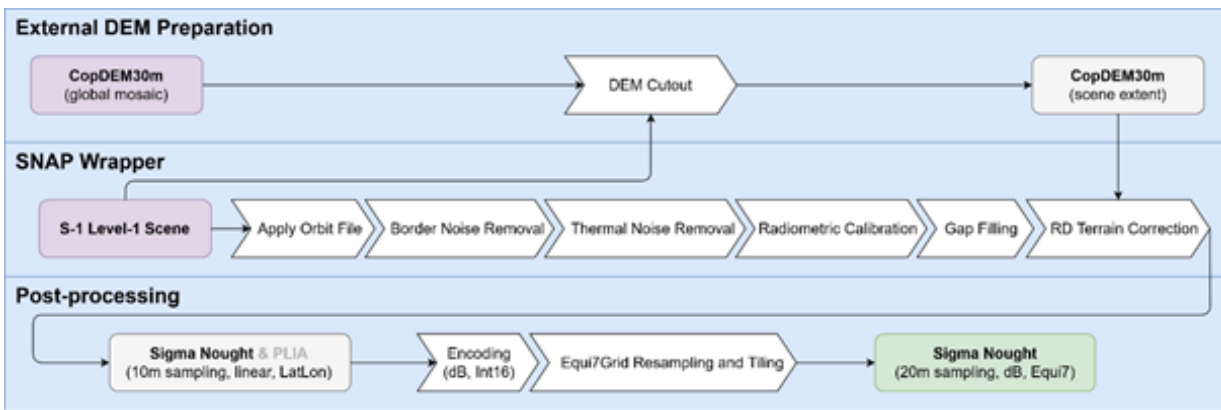
Based on single Sentinel-1 Level-1 GRDH scenes, the near real-time (NRT) Sentinel-1 pre-processing workflow (see Figure 1) is triggered. Before the actual SAR processing starts, the external, global 30m Copernicus DEM (Copernicus DEM) mosaic (see Section 3.2.1) is cropped to the projected extents of the scene to minimize the input file size in the Sentinel Application Platform (SNAP). Then, the SNAP processing is conducted via Graph Processing Tool (GPT) commands[1] subsequently called from a Python package developed by TU Vienna. This routine involves the following steps, where the step 2, 7, and 8 are done independently from SNAP:

Table 2: Description of the steps in the Sentinel-1 pre-processing workflow, carried out using the SNAP toolbox via Graph Processing Tool (GPT) commands.

#	STEP	DESCRIPTION
1	Apply Orbit File:	Applies precise orbit ephemerides (POEORB) if possible, otherwise, restituted orbit (RESORB) ephemerides are used, accessed from the Copernicus Hubs
2	Border Noise Removal:	Removes border noise pixels introduced by the conversion from the raw files to the Level-1 files conducted by the Instrument Processing Facility (IPF). In March 2018, IPF released a new processor version (IPF 2.90), which does not produce such artefacts at the image borders anymore. For older data we still apply our own, bidirectional-sampling algorithm to get rid of border noise (Ali et al., 2018).
3	Thermal Noise Removal:	Low backscatter from shallow or flat surfaces (e.g. lakes, rivers or streets) is often affected by a certain amount of additive noise, also known as thermal noise. This noise is reduced by applying the thermal noise calibration vectors stored in the Level-1 metadata provided by ESA.
4	Radiometric Calibration:	Converts and calibrates digital numbers (actually stored measurement) to backscatter quantities, in our case to Sigma Nought or the radar backscatter denoted here asσ^0.
5	Gap-Filling / Slice Assembly:	Level-1 products are disseminated by ESA by slicing the full swath into several smaller files along the movement direction (the individual S-1 scenes). These single scenes are consecutively aligned and are without any buffer. Since some processing steps (border noise removal, post-processing) remove a few boundary pixels to enhance the image quality, this leads to a small linear gap between adjacent scenes. To prevent that, we assemble neighbouring scenes and add a small buffer to the to-be-processed scene.
6	Terrain Correction:	The Range-Doppler terrain correction performs the geocoding of the scene from orbit to ground geometry. For this step, the Copernicus DEM-cutout is used as an auxiliary layer to provide the necessary information about terrain coordinates.

---

[1] https://senbox.atlassian.net/wiki/spaces/SNAP/pages/70503475/Bulk+Processing+with+GPT

| 7 | Encoding: | The output of the SNAP pipeline is a geocoded scene in the (native) LatLon projection system containing σ^0 backscatter values in the linear domain. This image is then converted to decibel (dB), scaled by a factor of 10, converted to Int16 data type, and written to disk as a Zstandard (ZSTD) compressed GeoTIFF file.

The logarithmic operator is applied in order to increase the dynamic of low intensity values and obtain more symmetrical distribution of classes of interest. This setup facilitates an efficient and storage-saving formation of the Sentinel-1 global datacube, providing a precision of 0.1dB that exceeds the radiometric accuracy of the C-band SAR sensor. |
| 8 | Resampling: | The last step resamples the data to the Equi7Grid system at 20m pixel sampling, using bilinear resampling from the Geospatial Data Abstraction Library (GDAL) libraries, and subsequently creates Equi7Grid-T3-tiles (squares covering 300x300km) to facilitate manageable file sizes in a stacked image structure. |

Figure 1: Illustration of the Sentinel-1 pre-processing workflow, with a Sentinel-1 Level-1 GRDH scene as input and georeferenced, calibrated backscatter (Sigma Nought) data as output.

Following the processing pipeline depicted in

Figure 1, we perform basic and fully automatic quality checks to ensure that the produced files and file contents fulfil a certain level of quality:

- Open and read file contents.
- Checking expected metadata attributes.
- Checking if the file only contains no data values.
- Validating the expected range of values.
- Checking if the spatial reference information is properly set in the GeoTIFF header.

An optional by-product of the SNAP geocoding workflow is the so-called projected local incidence angle (PLIA), which is the angle between the vector ground-satellite and the surface normal vector, projected into the range plane. PLIA is an essential information about the observation geometry of the satellite and varies across different orbits. Due to the Sentinel-1's stable orbit revisit over time,

we do not produce the PLIA files in the NRT service but only during archive reprocessing (hence it is depicted by grey text in the last section in

Figure 1).

3.1.2 Sentinel-1 temporal parameters

Based on the Sentinel-1 data, several temporal parameters are computed for the period 1.1.2019-31.12.2019 on a global scale. These serve as an input for the flood algorithms, exclusion layers and advisory flags. The parameters differ in their temporal aggregation. Parameter computed on monthly basis are in the following called GRMAG for grouped monthly aggregated, parameters aggregated over the total time period TAG. Moreover, parameters are either computed over all orbits together or per orbit. The temporal parameters are listed and described in Table 3.

Table 3: Overview of Sentinel-1 temporal parameters, and availability per orbit and over all orbits.

PARAMETER	DESCRIPTION	PER ORBIT	OVER ALL ORBITS
Number of observations (NOBS)	Contains the number of observations per pixels. This parameter is available per orbit and for all orbits combined.	x	x
Low backscatter frequency (LTM15-FREQ)	Contains the frequency of observations below -15dB. This parameter is calculated per orbit and serves as an input for the low backscatter masking.	x	
Mean (TAG-MEAN)	The total aggregated mean backscatter. The parameter is calculated per orbit and as a weighted average over all orbits. The latter is used as an input for the DLR and LIST algorithms to extract the permanent water mask and will in a later version be replaced by the median over the entire time series.	x	x
5th percentile (PERC-05)	The 5th percentile is calculated per orbit and used as an input for the advisory flags.	x	
Monthly grouped median (GRMAG-MEDIAN)	The GRMAG-MEDIAN is calculated for all month over the two years over all orbits. It is used as an input for the DLR and LIST algorithms to extract the seasonal water mask.		x

3.1.3 Sentinel-1 harmonic parameters

Radar signal interacts with the Earth's surface in many different ways. It can be absorbed, scattered, and reflected according to the surface states and characteristics of sensor. The surface state (including soil moisture content, vegetation, roughness, etc.) varies over time leading to variation of backscatter time series. Based on different periods of variation, time series of backscatter can be decomposed into trend, seasonality and short-term random variation.

GFM flood mapping algorithm 3 (see Section 4.1.3 for details) utilizes the harmonic model to simulate the backscatter seasonal variation and the estimation of normal, non-flooded conditions. Figure 2 shows a time series of VV backscatter and superimposed with the modeled harmonic time

series for an agricultural plot in Thessaly, Greece. An example for an estimated backscatter image for one day in summer and one day in winter over Greece is given in Figure 3. One can review the noticeable backscatter differences caused by varying seasonal vegetation state in this example. In this section, we describe the preparation performed to define the harmonic model on a global scale.
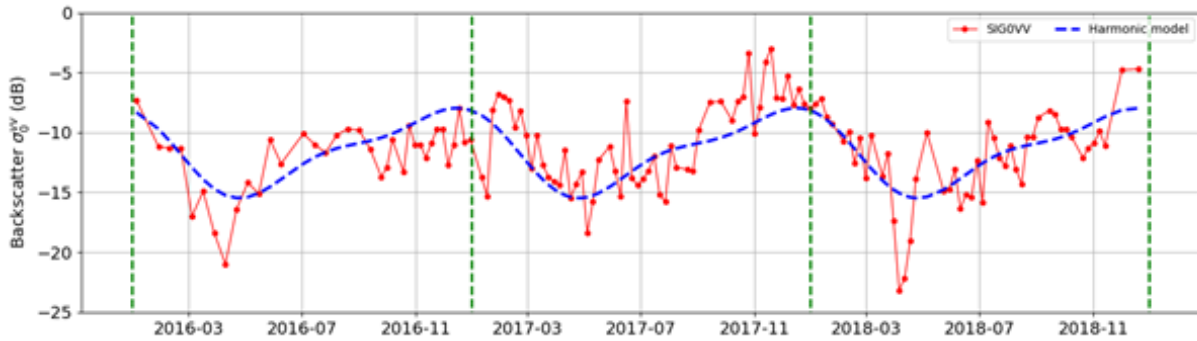
Figure 2: Comparison of actual backscatter time-series and its harmonic model
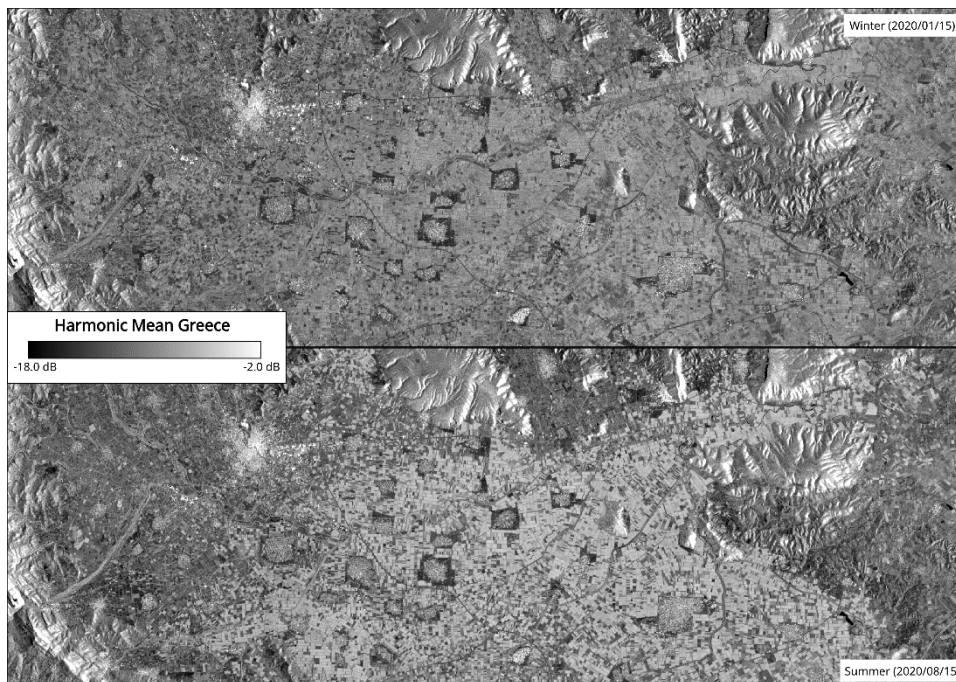
Figure 3: Example of the estimated backscatter of the harmonic model in summer and winter time over Greece

We introduce a harmonic model (see Eq. 3-1 below) following the approach of Schlaffer et al (2015), where $\hat{\sigma}^0_{t_{day}}$ the most probable radar backscatter at time t_{day} i.e. day of the year, is computed from $\overline{\sigma^0}$, the average backscatter for the time period, and C_i and S_i, the harmonic coefficients of the cosine and sine components. The harmonic coefficients and the average backscatter are further

referred to as harmonic parameters. As suggested by Schlaffer et al (2015), k is set to 3, representing processes of a time-scale of four months, sufficient to reduce the impact of long-lasting flood events.

$$\hat{\sigma}^0_{t_{day}} = \overline{\sigma^0} + \sum_{i=1}^{k} \left(C_i \cos \frac{2\pi i t_{day}}{n} + S_i \sin \frac{2\pi i t_{day}}{n} \right)$$ **Eq. 3-1**

The harmonic parameters are derived from a least-squares estimation based on the backscatter values and corresponding observation times of input Sentinel-1 time series. This approach differs from Schlaffer et al (2015) as the backscatter is used directly instead of using 10-day composites. However, backscatter coefficients are highly dependent on the acquisition geometry. While a normalisation approach could be utilised when sufficient incidence angle samples per pixel are present as with ENVISAT ASAR, this procedure is not feasible with Sentinel-1. Hence the parameter estimation is performed for each unique acquisition geometry corresponding to a relative Sentinel-1 orbit. Consequently, a unique harmonic parameter set per orbit per Equi7grid tile is required. Hence, to model the estimated backscatter $\hat{\sigma}^0_{t,r}$ for any given day t and relative orbit r, seven (at k=3) harmonic parameters are computed. In order to get a measure of how many samples support the estimation and to exclude ill-fitted harmonic parameters, the number of valid observations (NOBS) for each pixel is written as an additional layer. In addition, the standard deviation of the harmonic model is calculated as the square root of the sum of squared errors (SSE), divided by the number of data points (N_{points}) adjusted for the degrees of freedom of the model (see Eq. 3-2 below). The SSE is derived from the pixel's backscatter time-series $\sigma^0_{t,r}$ and its harmonic model $\hat{\sigma}^0_{t,r}$.

$$s = \sqrt{\frac{SSE(\sigma^0_{t,r}, \hat{\sigma}^0_{t,r})}{N_{points} - (2k + 1)}}$$ **Eq. 3-2**

3.1.4 Sentinel-1 projected local incidence angle data

The projected local incidence angle (PLIA) is the angle between the surface normal and the looking direction of the satellite (local incidence angle - LIA) which is further projected into the range plane. It provides essential information about the observation geometry of the satellite and varies across different orbits. It can be computed for each individual image within the terrain correction step (section 3.1.1 Sentinel-1 preprocessing), however, within the GFM project we exploit the stable nature of the Sentinel-1 orbit and precompute the mean PLIA values globally for all available orbits. As an input for this computation, we use the individual PLIA images aggregated to mean values from the all available data observed in the year 2020.

The mean PLIA values of the corresponding orbit are used together with the σ^0 image and the parameters of the harmonic model as an input of the TUW GFM flood mapping algorithm (see Section 4.1.3 for details).

3.2 Ancillary data preparation

In order to enable the calculations for the flood mapping and its by-products, the GFM system relies on several auxiliary datasets. These datasets provide information on terrain, land cover, and Sentinel-1 orbit geometry and are necessary for the Sentinel-1 preprocessing, the flood- and water mapping itself, and the preparation of the Exclusion Mask and the Advisory Flags.

Like the Sentinel-1 backscatter observations and the thereof derived statistical- and local signature parameters, these auxiliary datasets are stored in the Equi7Grid-based data cube at their native sampling. Consequently, also the auxiliary datasets are more efficiently stored and are accessible via the same interfaces as the Sentinel-1 datasets.

3.2.1 Satellite orbit state vectors

The Restituted Orbit (RESORB) files for the Sentinel-1 satellites contain the restituted Orbit State Vectors (OSVs) based on the orbit determination performed by ESA's Precise Orbit Determination Service. The OSVs describe the geometry of the satellite's centre of gravity in an Earth-fixed coordinate frame and are a necessary input to the geocoding of the incoming Sentinel-1 IW GRDH scenes when geometric precision at the 20m scale is demanded. The RESORB files are selected as input in the pre-processing component of the GFM system, as the (more precise) Precise Orbit Ephemerides (POEORB) files are unsuitable in NRT operations, with a delivery of 20 days after acquisition.

The RESORB files are accessed from Copernicus POD Data Hub at a 5 minutes interval and accessed by SNAP from within the pre-processing pipeline.

3.2.2 Copernicus DEM

The Copernicus Digital Elevation Model (DEM)[2] is a digital surface model (DSM) which (unlike a digital terrain model) includes structures above ground, e.g. buildings, bridges, or vegetation. The Copernicus DEM is mainly based on data from the WorldDEM, which is a product originating from the TanDEM-X mission, but is further enhanced using auxiliary height data from ASTER, SRTM90, SRTM30, GMTED2010 ALOS World 3D-30m and many more. It is available in three resolutions at two spatial extent windows, EEA-10 (10m, Europe), GLO-30 (30m, global) and GLO-90 (90m, global). Within the scope of the GFM project, we used the 30m sampled global version of the Copernicus DEM to achieve a trade-off between a high spatial sampling and a global coverage.

Because SNAP 7/8 only supports SRTM height data internally for geocoding, every other DEM must be provided and prepared from outside. Since SAR geocoding routines work with ellipsoid based 3D coordinates, commonly used orthometric heights H, as is the case for the Copernicus DEM, must be

---

[2] https://spacedata.copernicus.eu/web/cscda/dataset-details?articleId=394198#C4

transformed to ellipsoidal heights h_{ellips}. To do so, geoid undulations (N_{geoid}), which are supplied by various Earth Gravitational Models (e.g. EGM96 or EGM2008), must be taken into account. These parameters are incorporated into the following formula to derive the desired ellipsoidal heights:

$$h_{ellips} = H + N_{geoid}$$
<div align="right">**Eq. 3-3**</div>

SNAP 7/8 only works with one single external DEM file, so both data sets need to be combined to a global mosaic. The workflow describing the data flow from the download of both tiled input data sets to the creation of the global mosaic in ellipsoidal heights, summarized in Table 4.

Table 4: Summary of the workflow from the download of the 30m Copernicus DEM and EGM2008 tiled input data sets to the creation of the global mosaic in ellipsoidal heights.

STEP	DESCRIPTION
1	Download of the tiled 30m Copernicus DEM layers (i.e. 29,310 data files) from the Copernicus Space Component Data Access (CSCDA). The zipped data is stored in a continent- and country-based folder structure, with unadjusted country names implicating file system issues and data gaps for countries, which have an OS-unfriendly name. Those issues could only be solved after manually downloading data of the affected countries and a back-filling of the global data set.
2	Several experiments were conducted to generate a mosaic with common GDAL tools like gdal_warp, gdal_merge.py or gdal_translate, but always resulted in artefacts due to a varying pixel sampling across different latitudes. Thus, the strategy was changed to apply a mosaicking in two steps, first for each latitude band and then merging all latitude bands together.
3	All Copernicus DEM files with the same pixel spacing were collected and combined to get one mosaic per latitude band.
4	Each mosaic was then resampled to the highest sampling of the Copernicus DEM, which is around 30m/0.00028° (measured along the equator), was adjusted to match the global extent of (-180°, -90°, 180°, 90°), converted to Int16 (1m vertical accuracy suffices SAR geolocation), tiled with 512x512 blocks and compressed with the ZSTD method.
5	The global 30m Copernicus DEM mosaic in orthometric heights was then generated by merging all resampled latitude bands to one file.
6	Download of the 2.5' EGM2008 tiles from http://earth-info.nga.mil/GandG/wgs84/gravitymod/egm2008/egm08_gis.html (original source, but deprecated) or https://www.agisoft.com/downloads/geoids/ .
7	All 45°x45° EGM2008 tiles were mosaicked to one global file.
8	The EGM2008 mosaic was then resampled to the Copernicus DEM mosaic sampling and extent.
9	The last step directly applied the aforementioned formula on the two mosaics and produced the desired output: the global 30m Copernicus DEM mosaic in ellipsoidic heights.

3.2.3 *Sentinel-1 Global Backscatter Model*

The Sentinel-1 Global Backscatter Model (S1-GBM) was generated by the Remote Sensing Group of the TU Vienna, within a dedicated project by the European Space Agency (ESA). The dataset describes Earth's complete land surface for the period 2016-17 by the temporal mean and standard deviation of the Sentinel-1 σ^0 backsctter in VV- and VH-polarization at a 10 m sampling, giving a high-quality impression on surface structures and patterns.

In GFM, we use the temporal mean VV and VH σ^0 backscatter and VH standard deviation (computed jointly from data from all orbits) for the identification of densely vegetated areas, where the Sentinel-1 SAR observations are insensitive to floods.

3.2.4 HAND drainage information

The Height Above Nearest drainage (HAND) index Rennó et al. (2008) contains the vertical distance between any location and the nearest location in the drainage network with respect to pre-identified flow directions and by minimizing hydrological distances. The HAND index is calculated near-globally (between 60°N and 56°S) based on elevation and drainage direction information provided by the Hydrosheds mapping product Lehner et al. (2008). This product was derived based on the digital elevation model of the Shuttle Radar Topography Mission (SRTM), which had to be hydrologically conditioned by using a sequence of automated procedures such as void-filling, filtering, stream burning, and upscaling techniques as well as manual corrections (Lehner et al., 2008).

In the GFM project, the HAND index was used to create the HAND Exclusion Mask (HAND-EM) which helps us to reduce water-lookalike classes depending on the hydrologic–topographic setting and separate the flood-prone from non-flood prone areas.

3.2.5 Land cover data

To support the masking of area where the Sentinel-1 sensors as C-band radar are in general insensitive to floods, we collect land cover data sets. Furthermore, the Affected Land Cover product layers provides such information to the users of the individual GFM products.The Global Human Settlement Layer (GHSL) is a static urban mask generated by JRC from Landsat-8 optical imagery at 30 m resolution. Information about the Affected Population (i.e., number of people living within flooded area) is derived from the GHSL-POP dataset. GHS-BUILD-S2 dataset contains the information about the build up area expressed in terms of probability grid at 10 m resolution. This information is derived from a Sentinel-2 global image composite.

The World Settlement Footprint 2015 (WSF2015) is a static urban mask generated by DLR from Sentinel-1 and Landsat-8 data at 10m resolution in 2015. The GHSL and the WSF2015 will be used to highlight urban areas where flood detection is probably not possible due to the side-looking viewing geometry of SAR-satellites and complex interactions of the SAR signal with urban structures.

The Global forest change dataset (Hansen et al., (2013); version 1.8: https://storage.googleapis.com/earthenginepartners-hansen/GFC-2020-v1.8/download.html) characterises global forest extent and change from 2000 through 2020 and is based on the time-series analysis of Landsat images in characterizing global forest extent and change from 2000 through 2020. The Global forest change dataset was used to highlight the densely vegetated areas, where their identification based on Sentinel-1 Global Backscatter Model (S1-GBM) was not possible because of the irregular temporal coverage of Sentinel-1.

4   Generation of the GFM product output layers

The new Global Flood Monitoring (GFM) product of the Copernicus Emergency Management Service provides a continuous monitoring of floods worldwide by processing and analysing in near real-time all incoming Sentinel-1 Synthetic Aperture Radar (SAR) images acquired in Interferometric Wide Swath (IW) mode, and generating the eleven output layers of flood information that are listed in Table 1. The remainder of this Section presents the key technical details underlying the generation of the eleven GFM product output layers.

4.1   Output layer "S-1 observed flood extent"

The GFM product output layer "S-1 observed flood extent" shows the flooded areas which are mapped in near real-time from Sentinel-1 (S-1) satellite imagery, as the difference between the GFM output layers "S-1 observed water extent" and "S-1 reference water mask". The four flood mapping algorithms which are used by the GFM product to generate the output layer "S-1 observed flood extent" are described below.

4.1.1   Flood mapping algorithm 1 (LIST)

GFM flood mapping algorithm 1 requires three main inputs: the most recent SAR scene to be processed, a previously recorded overlapping SAR scene acquired from the same orbit called reference SAR scene hereafter together with the corresponding "previous" ensemble flood extent map. As it is a change detection algorithm, it aims at detecting and mapping all increases and decreases of floodwater extent with respect to a reference one. A change detection approach is adopted because it allows to differentiate floodwater from permanent water bodies and, at the same time, filter out classes having water-like backscattering values such as shadows or smooth surfaces. Moreover, in order to reduce false alarms caused by different types of unrelated changes (e.g. vegetation growth), a reference image acquired close in time is used. Sentinel-1 fulfils well this latter requirement since it has a repeat cycle of 6 days. Moreover, to reduce false alarms and to speed-up the analysis, it uses as optional input data the exclusion mask and the HAND map. Each time a new Sentinel-1 scene is ingested and pre-processed, I_{t0}^{S1}, the algorithm processes the pair of S1 scenes consisting of the new image, I_{t0}^{S1}, and the previously recorded overlapping Sentinel-1 image acquired from the same orbit, I_{t0-1}^{S1}, and it provides the new floodwater map, FM_{t0}, updating the previous floodwater map, FM_{t0-1}. In the end, the algorithm makes use of the changes it mapped to update the most recently computed ensemble flood extent map (FM_{t-1}).

Since the preprocessed Sentinel-1 images are already log-transformed, the change image is computed as the backscatter difference between I_{t0}^{S1} and I_{t0-1}^{S1}. The algorithm 1 is based on the assumption that the distributions of the log-ratio image and the log-transformed single SAR image follow the Gaussian model. This assumption is based on evidence that the PDF of a random variable

affected by speckle approaches a Gaussian distribution Ulaby et al. (1989).

As in all statistical floodwater or change detection mapping algorithms, the parameterization of the distributions of "water" and "change" classes is conditioned by the identifiability of the respective classes. This depends on the shape of the histogram. More explicitly, some classes may not be easily identifiable on the histogram typically when flooded or changed areas represent only a small percentage of the total image. To circumvent such a limitation, algorithm 1 locates areas of the image where pixels belonging to the classes of interest are rather well balanced. To achieve this, the algorithm identifies regions in the flooded image, I_{t0-1}^{S1}, and the corresponding difference image, $I_{t0-1}^{S1} - I_{t0}^{S1}$, that exhibit bimodal histograms with two classes, i.e. target class and its background. Moreover, the two classes are i) well separated, ii) represented with a similar frequency and iii) that can be modelled as Gaussian distributions.

The algorithm is therefore composed of two main steps: 1) identification of bimodal areas hereafter termed Bimodal Mask (BM), with sufficient target class presence to enable a robust parameterization of the PDFs of both target and background classes; 2) use of the previously parameterized PDFs to map the target class in the entire image via a sequence of thresholding and region growing. Region growing assumes that pixels constituting the target class are clustered rather than randomly spread out over the entire image. The region growing algorithm needs to determine two threshold values, one for selecting its seeds and one for stopping the region growing process. The thresholds for the seeds and the tolerance criterion to stop the growing process are selected based on the distribution fitted to the class of interest. The steps to parametrize the PDFs of the target and background classes are described below.

Considering an image *I*, it is hypothesized that two classes are present, where $G_1(y)$ and $G_2(y)$ are their distribution functions and *y* is the measurement. It is further assumed that both distributions can be approximated by Gaussian curves. Finally, it is hypothesized that the prior probabilities of the two classes are strongly unbalanced. This last assumption is usually verified when processing images that cover large areas and where changes only impact a small part of the image. When this happens, the smallest class is dominated by the other class and its distribution is practically indistinguishable in the global histogram, thereby causing the classification problem to be ill-posed and the selection of the threshold highly uncertain Gong et al. (2016), Aach et al. (1995).

Hence, in order to cope with this ill-posed problem of unbalanced populations, areas where the two classes are more balanced are identified. The adopted hierarchical split based approach (HSBA) consists of two main steps. First, the image X is split into separate sub-regions (i.e. tiles) using a quad-three decomposition . Next, the tiles showing a clear bimodal behaviour with two Gaussian balanced populations are selected.
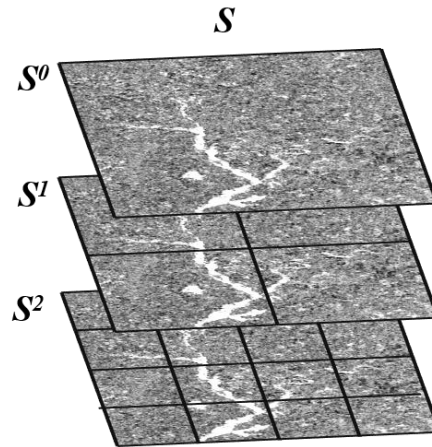
Figure 4: Example of a quad-tree tailing with three levels.

The hierarchical tiling consists of iteratively decomposing image regions into four equally sized quadrants, the so-called sub-quadrants. In the hierarchical approach - a schematic representation of which is shown in Figure 4 - each quadrant or sub-quadrant is related to a given node in the tree and has exactly four children or no children at all. The latter is referred to as a 'leaf node'. Each node is associated with a specific region of the image [Laferte et al (2000)]. Let I be an image and S_{nodes} the ensemble of tree nodes. A node in the i-th tree level is denoted $S_j^i \in S_{nodes}$, where S^0 is the root node, associated with the tile that completely encloses I, and S^N is the ensemble of nodes in the lowest level having the smallest tile size. The tile related to $S_j^i \in S_{nodes}$ is split into four tiles in the next lower level ($i+1$). It is itself one of the four quadrants of a tile in the level above ($i-1$). Thus, the resulting quad-tree decomposition can be expressed as:

$$S_{nodes} = \left\{ S^0, S^1, .., S^i \; s.t. \bigcup_{j=1}^{4^i} I[S_j^i] = I = I[S^0], i = 0, 1, .., N_{nodes} \right\},$$ **Eq. 4-1**

where $I[S_j^i]$ is the portion of the image related to node $S_j^i \in S_{nodes}$ and U represents the Union operator (see Figure 4).

The lowest level of the quad-tree decomposition is fixed based on the minimum tile size that still guarantees statistical representativeness. A tile is selected based on three conditions (C1, C2, C3):

- *C1*: The tile histogram, $h(I[S_j^i])$, is clearly bimodal;
- *C2*: The two populations represented in $h(I[S_j^i])$ are normally distributed;
- *C3*: Each population in $h(I[S_j^i])$ is present at a frequency of at least 10% of the frequency of the other class [Bazi et al. (2007)].

It is hypothesized that each $h(I[S_j^i])$ is a sum of two Gaussian curves, i.e.

$$h(I[S_j^i]) = h(y) \approx h_f(y) = G_1(y) + G_2(y) = A_1 e^{\frac{-(y-\mu_1)^2}{2\sigma_1^2}} + A_2 e^{\frac{-(y-\mu_2)^2}{2\sigma_2^2}} \qquad \textbf{Eq. 4-2}$$

where A_1 and A_2 are scale factors set to the maxima of the two curves; μ_1 and μ_2 are the distribution means; σ_1 and σ_2 are the distribution standard deviations; y is the measurement and $h_f(y)$ is the fitted histogram.

In order to estimate the parameters of the two Gaussian curves in (Eq. 4-2), $G_1(y)$ and $G_2(y)$, which in combination compose the histogram, the Levenberg-Marquardt algorithm is applied. It solves the nonlinear least squares problems by combining the steepest descent and inverse-Hessian function fitting methods Marquardt et al. (1963). Nonlinear least squares methods involve an iterative optimization of parameter values to reduce the sum of the squares of the errors between a fitting function and histogram values. Iterations are carried out until three consecutive repetitions fail to change the chi-squared value by more than a specified tolerance amount, or until a maximum number of iterations has been reached. The initial guess of the parameter values should be as close as possible to the actual values, otherwise the solution may not converge. The Otsu approach Otsu et al. (1979) is used for retrieving the first guess values. Given the Otsu-based thresholds, τ_{Otsu} for $h(y)$, the first guess values (denoted by the superscript 0) of the parameters of the two Gaussian distributions defined in (2) can be derived as:

$$\mu_1^0 = mean(h(y < \tau_{Otsu_{OT}})), \mu_2^0 = mean(h(y > \tau_{Otsu})), \qquad \textbf{Eq. 4-3}$$
$$\sigma_1^0 = std(h(y < \tau_{Otsu})), \sigma_2^0 = std(h(y > \tau_{Otsu})), \qquad \textbf{Eq. 4-4}$$
$$A_1^0 = h(y = \mu_1^0), A_2^0 = h(y = \mu_2^0). \qquad \textbf{Eq. 4-5}$$

When all distributions have been fitted, the Ashman D (*AD*) coefficient Ashman et al. (1994), the Bhattacharyya Coefficient (*BC*) Aherne et al. (1998) and the Surface Ratio (*SR*) of the two classes on image *I* are used for evaluating the bimodality of the histogram, the Gaussianity of each class distribution and the balance of each class in the image, respectively.

The *AD* coefficient quantifies how well two Gaussian distributions are separated, by considering the distance between their mean values and their dispersions, i.e. standard deviations, and it is expressed as:

$$AD(h_f(y)) = \sqrt{2} \frac{|\mu_1 - \mu_2|}{\sqrt{(\sigma_1^2 + \sigma_2^2)}} \qquad \textbf{Eq. 4-6}$$

For a mixture of two Gaussian distributions, *AD > 2* is required for a clear separation of the distributions Ashman et al. (1994). Consequently, the condition *C1* for selecting or not a tile is defined as:

$$C1(I[S_j^i]) = \begin{cases} 1 \ if \ AD(h_f I[S_j^i]) > 2 \\ 0 \ otherwise \end{cases}$$

<div align="right">**Eq. 4-7**</div>

The *BC* evaluates the similarity of two distributions. It is used here to verify if $h_f(y)$ is a good approximation of *h(y)*, meaning that the hypothesis of having two Gaussian distributions is verified and defined as:

$$BC\left(h(y), h_f(y)\right) = \sum_k \sqrt{h(y_k)} \sqrt{h_f(y_k)}$$

<div align="right">**Eq. 4-8**</div>

where *k* is the corresponding bin of the two discrete histograms. In the case of two almost identical histograms, the *BC* approaches 1 Aherne et al. (1998). As an approximation, we assume that *BC* has to be higher than *.99* for a good fitting. Consequently, the condition *C2* for selecting or not a tile is computed as:

$$C2(I[S_j^i]) = \begin{cases} 1 \ if \ BC\left(h(I[S_j^i]), h_f(I[S_j^i])\right) > .99 \\ 0 \ otherwise \end{cases}$$

<div align="right">**Eq. 4-9**</div>

Once all distributions have been parameterized, the verification of the remaining condition *C3* is straightforward. To do so, the *SR* between the smallest and the largest class is computed, i.e.

$$SR(h_f(y)) = \frac{\min(A_1 \sigma_1 \sqrt{2\pi}, \ A_2 \sigma_2 \sqrt{2\pi})}{\max(A_1 \sigma_1 \sqrt{2\pi}, \ A_2 \sigma_2 \sqrt{2\pi})}$$

<div align="right">**Eq. 4-10**</div>

Considering that *C3* stipulates that the smaller population should represent at least 10% of the other one, condition *C3* becomes:

$$C3(I[S_j^i]) = \begin{cases} 1 \ if \ SR(h_f(I[S_j^i])) > .1 \\ 0 \ otherwise \end{cases}$$

<div align="right">**Eq. 4-11**</div>

Technically speaking, the tiles of the entire tree are scanned starting from S^0 until S^N. Any node (and the associated tile) that fulfils the conditions $C1(I[S_j^i]) = 1 \ and \ C2(I[S_j^i]) = 1 \ and \ C3(I[S_j^i]) = 1$ (see Eq. 4-8 to Eq. 4-11) is finally selected and its children, tiles at levels with higher *i*, are no longer considered as possible tile candidates. These processing steps result in a binary map, henceforth named Bimodal Mask (*BM*). *BM* includes regions in *I* where the distinctive populations, G_1 and G_2, are present with a sufficient number of pixels and where their respective distribution functions are clearly identifiable with more balanced prior probabilities.

At this point, all subtiles fulfilling the three conditions *C1*, *C2* and *C3* are selected. Hence, the histogram of all the pixels enclosed by the selected tiles, i.e. *where BM=1*, has to be clearly bimodal. This histogram, $h(I[where \ BM = 1])$, is used to estimate the Gaussian distributions of the two

classes of interest, hereafter named $G_1^{BM}(y)$ and $G_2^{BM}(y)$. To do so, the Levenberg-Marquardt algorithm is re-applied to fit two Gaussian curves on the histogram pixel values obtained from the union of all selected tiles. Starting from the two derived Gaussian distributions, the objective is to select the threshold for separating the two classes.

The algorithm needs to handle different type of flood events, from those lasting few days up to those lasting weeks, as the case of monsoon events. As a consequence, the algorithm assumes that the floodwater extent in the newly acquired S1 scene, I_{t0}^{S1}, i) increases or ii) only decreases. Regarding the case (i), increase, the algorithm assumes that there are also areas where the water extent can recede. These two different floodwater evolutions are handled subsequentially: the condition (i) is checked firstly, if it is satisfied the newly floodwater is mapped, if not it is checked the condition (ii) and the receding water is removed from the previous floodwater map. If both conditions are not satisfied means that the floodwater states in I_{t0}^{S1} does not change with respect to the previous acquired scene, I_{t-1}^{S1}.

i) First, the image difference is computed between I_{t0}^{S1} and I_{t-1}^{S1}, i.e. $I_D^{S1} = I_{t0}^{S1} - I_{t-1}^{S1}$, then HSBA is applied in parallel to I_{t0}^{S1} and I_D^{S1} in order to select tiles that show a bimodality behaviour on both images, guaranteeing that I_{t0}^{S1} is affected by floodwater and I_D^{S1} is characterized by a class representative of a decrease of backscattering with respect to I_{t-1}^{S1}. The final BM is used to estimate PDFs of water and non-water, PDF(W) and PDF(NW), in I_{t0}^{S1} and change and non-change in I_D^{S1}, PDF(C) and PDF(NC). Before to estimate the four PDFs the permanent water layer (from Copernicus DEM or Permanent water from GFM) is removed in order to have PDFs more representative of floodwater pixels in both I_{t0}^{S1} and I_D^{S1}. All estimated PDFs are used to apply a 2D region growing, I_{t0}^{S1} and I_D^{S1}, where the thresholds for seeds and to stop the region growing are defined as:

Seeds: Pixels whose *posterior probability* of water class $p(W|\sigma^0) \geq 0.7$ and *posterior probability* of change class $p(C|\Delta\sigma^0) \geq 0.7$ are selected as seeds for region growing.

Stop growing: Pixels with $0.3 < p(W|\sigma^0) < 0.7$ and $0.3 < p(C|\Delta\sigma^0) < 0.7$ are subject to region growing. The stop growing threshold pair $(\sigma_{stop}^0, \Delta\sigma_{stop}^0)$ is the one minimizing the root-mean-squared-error (RMSE) between $G_1^{BM}(y)$ and the histogram resulting from region growing $h_{rg}^{BM}(y_{seed}, y_{rg})$.

$p(W|\sigma^0)$ and $p(C|\Delta\sigma^0)$ are inferenced via the Bayes theorem:

$$p(W|\sigma^0) = \frac{p(\sigma^0|W)p(W)}{p(\sigma^0)} \qquad \text{Eq. 4-12}$$

$$p(C|\Delta\sigma^0) = \frac{p(\Delta\sigma^0|C)p(C)}{p(\Delta\sigma^0)} \qquad \text{Eq. 4-13}$$

$$p(\sigma^0) = p(\sigma^0|W)p(W) + p(\sigma^0|NW)p(NW) \qquad \text{Eq. 4-14}$$

$$p(\Delta\sigma^0) = p(\Delta\sigma^0|C)p(C) + p(\Delta\sigma^0|NC)p(NC) \qquad \text{Eq. 4-15}$$

Where $p(\sigma^0)$ is the *marginal distribution* of backscatter values in I_{t0}^{S1}, $p(\sigma^0|W)$ and $p(\sigma^0|NW)$ are *conditional probabilities* of water and non-water classes, respectively, $p(W)$ and $p(NW)$ are *prior probabilities* of water and non-water classes, both are set to 0.5. Similarly, $p(\Delta\sigma^0)$ is the *marginal distribution* of backscatter difference values in I_D^{S1}, $p(\Delta\sigma^0|C)$ and $p(\Delta\sigma^0|NC)$ are *conditional probabilities* of change and non-change classes, respectively, $p(C)$ and $p(NC)$ are *prior probabilities* of water and non-water classes, both are set to 0.5.

From the selected seeds are removed all those that occur in areas masked out by the HAND mask in order to avoid that false alarms in not prone flood areas could spread in plausible areas. The newly detected flooded areas are used to update the previous floodwater map.

The algorithm checks also areas where floodwater could have receded with respect to the previous acquisition. To do this, P(W) and P(NW) are used to apply a 1D region growing to map water on I_{t0}^{S1}. The resulting water map is compared with the previous floodwater map and only those pixels that belonging to water in both maps are mapped as floodwater in the new map.

The procedure above mentioned can have a dual role, to map new floodwater and at the same time to check the occurrence of a flood event which is assured by the selection or not of tiles of BM. If BM is an empty mask means that there is not a new flood event in the region or that the flood already ongoing at the time t-1 is only receding. In this latter case, the step ii is applied.

ii) In this case HSBA is applied only to I_{t0}^{S1} to fit P(W) and P(NW), which are subsequentially used for the 1D region growing on I_{t0}^{S1} to map water. The resulting water map is compared with the previous floodwater map and only those pixels that belonging to water in both maps are mapped as floodwater in the new map.

The exclusion layer, the HAND mask and the ocean mask from the Copernicus DEM are applied to mask out all pixels which are not part of flood prone areas.


4.1.2   Flood mapping algorithm 2 (DLR)

GFM flood mapping algorithm 2 requires a full Sentinel-1 scene as the main input and further exploits three ancillary raster datasets: a digital elevation model (DEM), areas not prone to flooding and reference water. Initial descriptions of the algorithm were published by Martinis et al. (2015) and Twele et al. (2016).

The algorithm applies a parametric tile-based thresholding procedure by labelling all pixels with a backscatter value lower than a threshold to the class *water*. The threshold is computed on a smaller, limited number of subsets of the SAR scene and applied to the entire scene. Subsets used to define the threshold are selected for based on their backscatter values, which should consist of both low backscatter and a strong variation in therein.

In a first step, a bi-level quadtree structure is generated with the division of a SAR scene into N_{tiles} quadratic non-overlapping subsets (*parent tiles*) of defined size c^2 on level S^+ (*parent tiles*). Each parent tile is represented by four quadratic child objects of size $(c/2)^2$ on level S^- (

Figure 5). The variable c is empirically set to 200 pixels. A limited number of tiles are selected out of N_{tiles}, based on the probability that the tiles contain a bi-modal mixture distribution of the *water* and *non-water* classes.
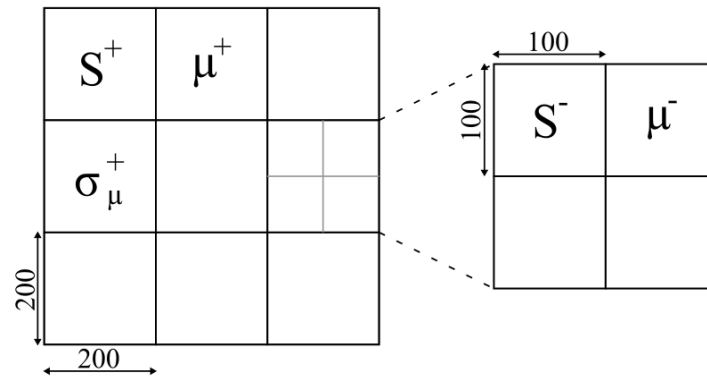

Figure 5: The SAR scene is tiled into N quadratic non-overlapping tiles S+. μ+ is the mean backscatter value of one tile S^+. Each tile S^+ is further divided into 4 sub tiles S^-. μ⁻ is the mean backscatter value of one sub tile S^-.

In rare cases, a parent tile intersects with the edge of a SAR scene, which results in dimensions that are smaller than the pre-defined size c^2. Consequently, these tiles are excluded from the threshold computation (

Figure 6). Furthermore, due to the near-cross polar orbits of Sentinel-1, the along track direction forms an angle of ~10° with respect to the north-south direction, which can result in parent tiles with variable amounts of no-data content. If no-data comprises of > 50 % of the overall data content, the tile is also excluded from the threshold computation (
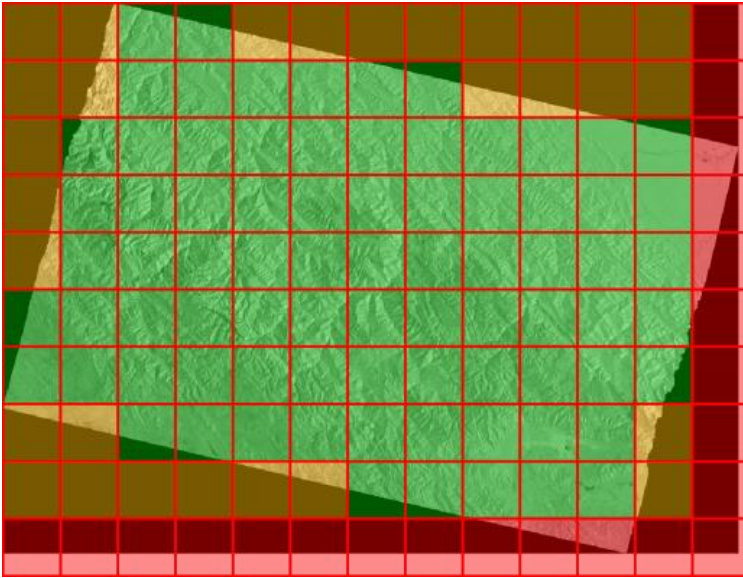
Figure 6).

Figure 6: Red tiles intersect with the edge of the SAR scene. Orange tiles have less than 50 % valid data content. Only green tiles are valid and considered for the threshold computation.

A parent tile S^+ must fulfil two conditions to be selected as a representative subset (Eq. 4-16):

- S^+ must have a lower mean than the global mean of the SAR scene; and
- S^+ must have a high standard deviation, which is computed on all S^-.

$$S^+ = \begin{cases} \mu^+ < \mu_{global} \\ \sigma_\mu^+ \geq \mu_{\sigma_\mu^+} + x \cdot \sigma_{\sigma_\mu^+} \end{cases}, \text{ with x = 2 by default} \qquad \textbf{Eq. 4-16}$$

Considering the first condition, for each parent tile S^+, the algorithm detects the mean backscatter value μ^+ of S^+, which must be lower than the mean backscatter value μ_{global} of the entire SAR scene (Eq. 4-17). We conduct this step to select for tiles with low backscatter content that are likely to contain water features.

$$\mu^+ < \mu_{global} \qquad \textbf{Eq. 4-17}$$

Considering the second condition, for each of the four child tiles S^- belonging to a parent tile S^+, the algorithm detects the mean backscatter value μ^- of S^-. In a further step, we compute the standard deviation σ_μ^+ of these four mean backscatter values $[\mu_1^-, \mu_2^-, \mu_3^-, \mu_4^-]$ (Eq. 4-18). We conduct this step to characterize the backscatter distribution of each parent tile S^+.

$$\sigma_\mu^+ = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n=4} (\mu_i^- - \overline{\mu^-})^2} \qquad \textbf{Eq. 4-18}$$

We assume that tiles with a low mean backscatter μ^+ and a high standard deviation σ_μ^+ show a bimodal backscatter distribution and are likely to contain both water and non-water features. In order to fulfil condition 2, we select a parent tile S^+ with a standard deviation σ_μ^+ that is as high as possible. When investigating σ_μ^+ of a specific tile and bringing it into relation to other tiles, we require σ_μ^+ to be greater or equal to the mean $\mu_{\sigma_\mu^+}$ of all standard deviations σ_μ^+ (Eq. 4-19), as well as two times the standard deviation $\sigma_{\sigma_\mu^+}$ of all standard deviations σ_μ^+ (Eq. 4-20). Figure 7 depicts a graphical representation of the concept. In Figure 7, plotting all standard deviations σ_μ^+ into a gaussian distribution, a standard deviation σμ+ that is greater or equal to the distribution's mean would still leave too many parent tiles S+, i.e. the majority of the tiles may not show a bimodal backscatter distribution. By restricting σ_μ^+ to even exceed the mean and one standard deviation we focus on a small number of parent tiles where the standard deviation σ_μ^+ is high enough to ensure bimodal backscatter distributions.
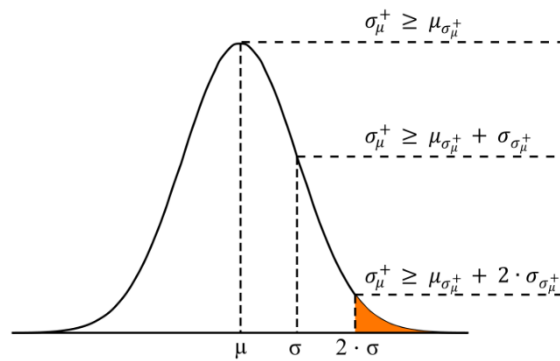
Figure 7: The standard deviation σ_μ^+ in relation to further statistical parameters. (See text for details).

$$\mu_{\sigma_\mu^+} = \frac{1}{n}\sum_{i=1}^{n}\sigma_{\mu}^+{}_i$$

Eq. 4-19

Eq. 4-20

$$\sigma_{\sigma_\mu^+} = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(\sigma_{\mu}^+{}_i - \overline{\sigma_\mu^+}\right)^2}$$

Applying Eq. 4-16 to all valid parent tiles selects suitable tiles for further threshold computation. If the count of suitable parent tiles S^+ is less than or equal to 10, Eq. 4-16 is re-applied with x = 1.28, thus the range of allowed standard deviations increases with the inclusion of more parent tiles into the selection. If the count of suitable parent tiles S^+ exceeds 10, we assume that there is a statistically sound basis of representative tiles and use the five tiles with the highest standard deviation σ_μ^+, i.e. the highest standard deviation of the means of the corresponding child tiles for further threshold computation.

Each of the five final parent tiles should have a bimodal backscatter distribution, which is likely to contain a valid water-land-boundary (

Figure 8). We apply the thresholding algorithm of Kittler & Illingworth (1986) to each of the five parent tiles. The algorithm is an iterative, cost-minimization approach, where the histogram is split into two classes with a threshold that identifies the class boundary. The optimal threshold τ separates both classes with minimum effort (

Figure 8). In Figure 8, if selecting a pure land tile (a1), the corresponding histogram (a2) is unimodal. If selecting a tile with both low and high backscatter values (b1), the corresponding histogram (b2) is bimodal and we assume depicting a water-land-boundary. If applying a threshold to the histogram b2, both classes can be separated giving the class water in the left part.
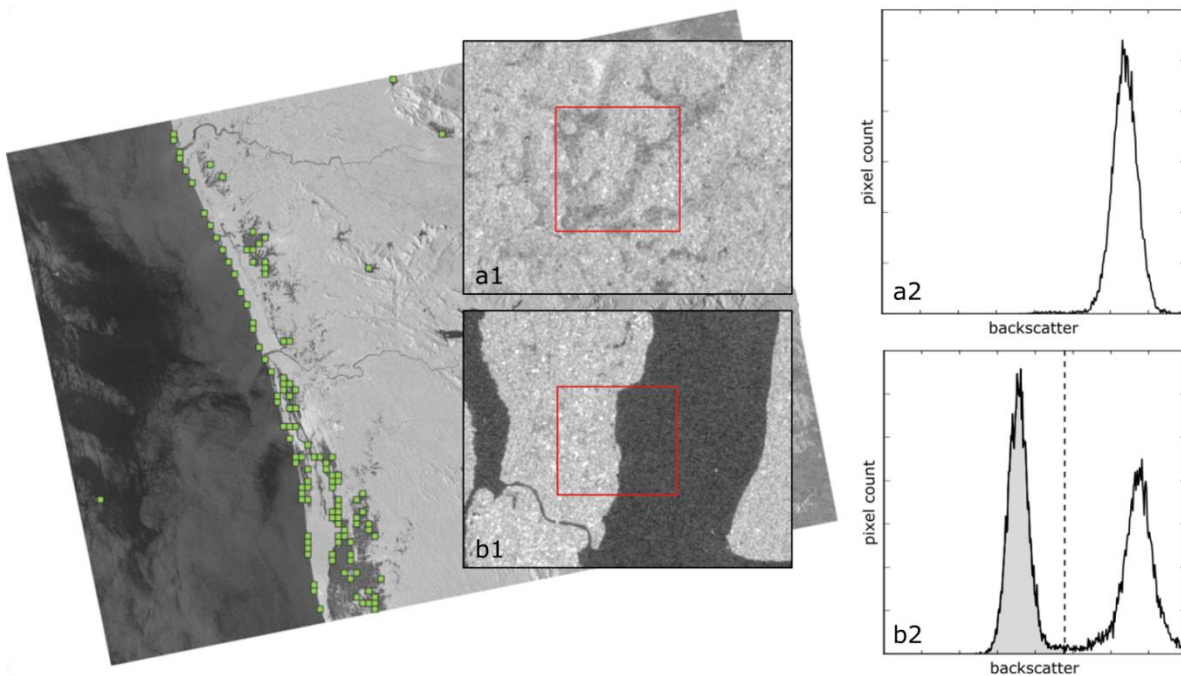
Figure 8: Selecting suitable subsets (green) from the SAR scene. (See text for details).

The final threshold separating the water and land classes is the mean of the thresholds of the parent tiles τ^+ (Eq. 4-21).

$$\tau = \frac{1}{n} \sum_{i=1}^{n=5} \tau_i^+$$

Eq. 4-21

We further compute the mean of the water class μ_{water} according to Kittler & Illingworth (1986), which describes the geometric centroid of the separated class (Eq. 4-22, with $\mu_i(\sigma^0)$ as a function

of the SAR backscatter value σ^0. h(g) is the SAR backscatter histogram, and g are the histogram bins g. $P_i(\sigma^0)$ represents the summed-up histogram h(g), up to the class-separating threshold (Eq. 4-23). As the thresholding algorithm is an iterative process, a and b denote updatable greyscale values (i.e. SAR backscatter) within the value range.

$$\mu_i(T) = \frac{\sum_{g=a}^{b} h(g)\, g}{P_i(\sigma^0)} \qquad \text{Eq. 4-22}$$

$$P_i(\sigma^0) = \sum_{g=a}^{b} h(g) \qquad \text{Eq. 4-23}$$

An average of all "water class centers" μ_{water} is also computed with the same method as the mean threshold τ for the entire SAR scene.

Applying this threshold to the entire SAR scene returns an initial basis for water detection.

4.1.3 Flood mapping algorithm 3 (TUW)

GFM flood mapping algorithm 3 Bauer-Marschallinger et al, (in review) requires three inputs: the SAR scene to be processed, a projected local incidence (PLIA) layer, and the corresponding parameters of the harmonic model. Based on these inputs, the probability of a pixel belonging to the flood or non-flood class is defined. For this purpose the a-priori flood and non-flood probability density function (PDF) is required. For the values of an incoming Sentinel-1 image, the most likely class is selected by the use of the Bayes decision rule. The algorithm's output consists of a flood extent map and the corresponding estimation of the classification's uncertainty.

Flood backscatter probability density function:

Due to specular reflections of the radar pulses by water surfaces, the received backscatter intensities by the sensor are significantly lower compared to most other land cover types. Consequently, a temporarily inundated surface is detectable by a significant drop in its backscatter relative to the time series. In order to ensure the drop results from inundation and no other effect, a detailed statistical knowledge of the backscatter behaviour over water surfaces is required. The proposed concept is the definition of the *flood backscatter probability density function* $p(\sigma^0|F)$ for any given incidence angle θ. It is assumed that the $p(\sigma^0|F)$ is a normal distribution and can be parametrized by the statistical value-pair of mean and a standard deviation.

Therefore, various backscatter observations (further referred to as $\sigma_{W,\theta}^0$) and their respective incidence angles were collected over oceans and inland waters from the Sentinel-1 datacube. By sorting the water backscatter collection $\sigma_{W,\theta}^0$ along the incidence angle (Figure 9), the expected indirect relation of water backscatter values and incidence angles can be confirmed.
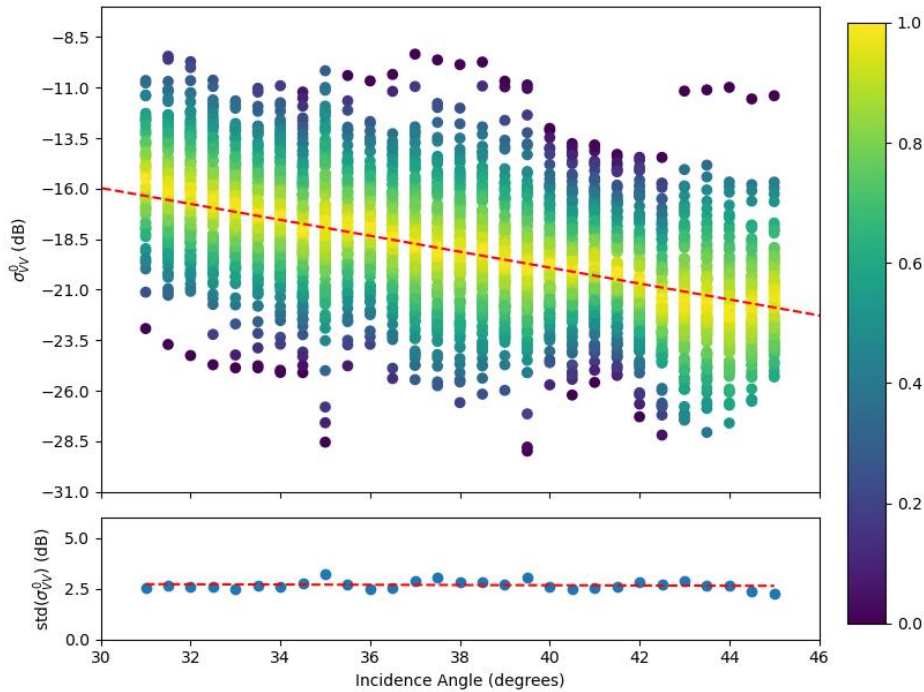
Figure 9: Scatterplot of water backscatter observations grouped by 0.5 degree bins (top). Standard deviation of these bins (bottom).

The mean backscatter value at a specific incidence angle can be derived by the use of a linear regression. The corresponding gradient was found to be 0.394 (red line in upper Figure 9). In order to analyse the standard deviation at a specific incidence angle, the water backscatter values were grouped into 0.5 degree bins. As shown in the bottom of Figure 9, the gradient of the standard deviation per bin is small with 0.008. Therefore, homoscedasticity is assumed and the total standard deviation is calculated by the square root of the sum of squared estimate of errors (SSE), normalised with respect to the number of data points (N_{points}). To summarize, the *flood backscatter probability density function* $p(\sigma^0|F)$ is defined by:

$$mean(\sigma_{W,\theta}^0) = -0.394\theta - 4.142 [dB] \qquad \text{Eq. 4-24}$$

$$std(\sigma_{W,\theta}^0) = \sqrt{\frac{SSE}{N_{points} - 2}} = 2.7 \ [dB] \qquad \text{Eq. 4-25}$$

Non-flood probability density function:

In order to model the normal, non-flooded circumstances, we propose the use of the harmonic model (details given in Section 3.1.3). The harmonic model allows retrieving an expected backscatter value $\widehat{\sigma_{t,r}^0}$ for any day of the year t_{day} and a specific relative orbit r. Therefore, the backscatter's seasonality is modelled by defining the pixel's harmonic parameters (C_i, S_i and $\overline{\sigma_0}$) based on its backscatter time-series σ_r^0 (see Eq. 3-1). Furthermore, this model is taken as base of

the *non-flood probability density function* $p(\sigma^0|NF)$, which is assumed to be a normal distribution. The mean backscatter value is set to the expected backscatter of the harmonic model $\widehat{\sigma_{t,r}^0}$. As mentioned in Section 3.1.3, in addition to the harmonic parameters the standard deviation s of the harmonic model is used as an input for the algorithm and standard deviation of the PDF.

$$mean\left(\sigma_{t,r}^0\right) = \hat{\sigma}_{t,r}^0 \qquad \text{Eq. 4-26}$$
$$std\left(\sigma_{t,r}^0\right) = s \qquad \text{Eq. 4-27}$$

A redundancy criterion is introduced to reduce the possibility of ill-fitting harmonic parameters (especially for sparse data sets). This criterion makes use of the NOBS layer (introduced in Section 3.1.3) which indicates the number of valid observations used for the estimation of the harmonic parameters. Given the requisite 2k + 1 samples for a unique solution to the harmonic equation, redundancy criteria is applied as a multiple of this number. Since the pixels which do not match the redundancy criteria are excluded from the flood mapping procedure, one needs to find a balance between number of excluded pixels and introduced noise. Based on initial tests, a redundancy value of 4x, e.g. 28 samples per pixel stack, can be set in order to minimizing noise **10**, specifically at edges where sparse samples might be present.

10

Bayesian Decision:

With the per-pixel knowledge of the relative orbit r, the day of year t_{day} and the incidence angle θ, one is able to define the *flood backscatter probability density function* $p(\sigma^0|F)$ and the *non-flood probability density function* $p(\sigma^0|NF)$ for each pixel. By using these PDF the belonging of the incoming Sentinel-1 image to either the flood (F) or non-flood (NF) class can be determined on a pixel-level. Therefore, the *Bayes decision rule* makes use of the class' *posterior probability*:

$$P(F|\sigma^0) = \frac{p(\sigma^0|F)P(F)}{p(\sigma^0)} \qquad \text{Eq. 4-28}$$
$$P(NF|\sigma^0) = \frac{p(\sigma^0|NF)P(NF)}{p(\sigma^0)} \qquad \text{Eq. 4-29}$$

Where $P(F)$ and $P(NF)$ are called *prior probabilities* and represent the prior knowledge of the pixel's belonging to a certain class. Since no prior information is available, it is set to equal weight (0.5). The *evidence* $p(\sigma^0)$ scales the *posterior probabilities* between 0 and 1 and is defined by:

$$p(\sigma^0) = p(\sigma^0|F)P(F) + p(\sigma^0|NF)P(NF) \qquad \text{Eq. 4-30}$$

By selecting the class (c), featuring the maximum *posterior probability* the *Bayes decision rule* selects the most likely class. Details can be seen within Eq. 4-31, where the class set $f_j = \{f_1, f_2\} = \{F, NF\}$ includes the flood (F) and non-flood (NF) classes. Figure 11 shows the decision step graphically. Finally, each pixel's belonging to the flood (F) class generates the preliminary binary flood extent map.

$$c = \arg max_j P(f_j|\sigma^0)$$ **Eq. 4-31**

Uncertainty Values:

The algorithm provides the ability to quantify the *uncertainty* of the classification by the use of the *conditional error* $P(error|\sigma^0)$. This measure is the posterior probability of the less probable class (i.e. the class not selected as flood), is defined between 0.0 and 0.5, and directly quantifies the uncertainty of the classification decision.
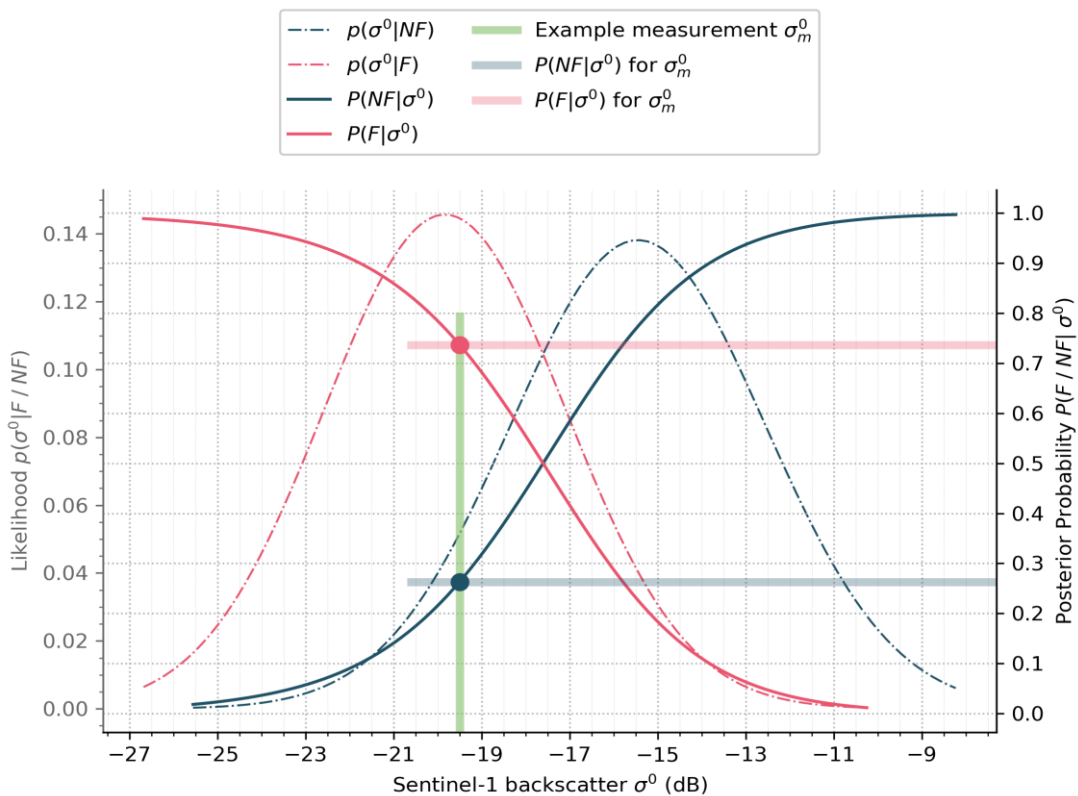

Figure 110: Exemplary Bayes flood mapping process for a single pixel. Including the probability density function and posterior probability of the flood and non-flood class. In this case the backscatter value of -19.5dB would be classified as flood.

Low-Sensitivity Masking:

For improving the reliability of Algorithm 3 in (non-supervised) operations, we introduce a set of routines to exclude locations where our statistical model does not allow a robust decision between flood and non-flood condition. Consequently, the following masks (PLIA, distributions conflicts, outlier, and high uncertainty) are applied on the preliminary flood extent map obtained from the

Bayes decision. Figure 11 shows an exemplary overview of the applied low sensitivity masks over Greece.
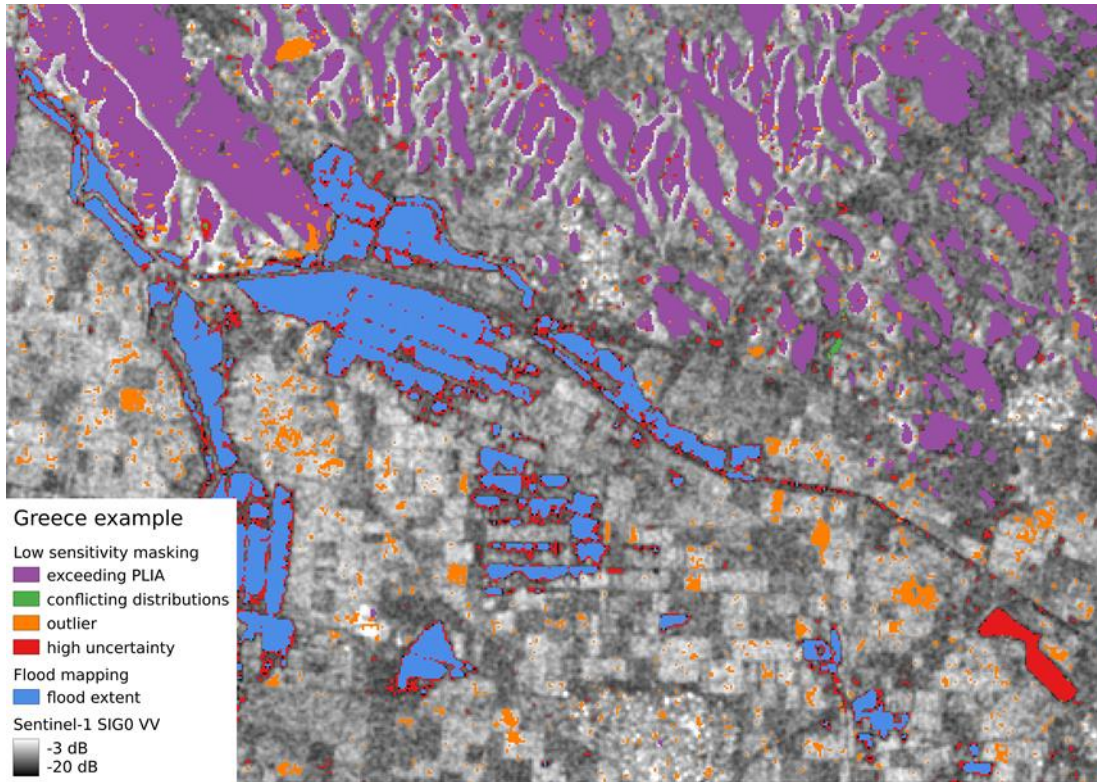
Figure 121: Example of low sensitivity masking in Algorithm 3 over Greece.

Flat areas are observed by Sentinel-1's IW mode within a viewing angle between 29° and 46°. Since (flat) water surfaces fall per definition into this PLIA-interval, the collection of water backscatter samples (described above) was limited accordingly. To extend the flood mapping capabilities to e.g. onsets of hills, the allowed range is relaxed to 27° to 48°. In order to pre-filter the pixels which feature an exceeding PLIA-value (θ), the following PLIA mask is introduced:

$$m_{PLIA} = \ \theta < 27° \ \& \ \theta > 48° \qquad \textbf{Eq. 4-32}$$

Our algorithm detects flood if a normally non-flooded pixel is temporary inundated which implies the pixel features a higher backscatter than a respective water surface. Putting it into context, the non-flooded PDF needs to be overall higher than the flood PDF. Typical locations where this requirement is not fulfilled are e.g. permanent water bodies, asphalt surfaces, salt panes, or very dry sand- and bedrock-areas. In order to exclude these cases of conflicting distributions, a pixel is masked if the mean of the non-flood PDF is lower than the mean of the flood PDF raised by half the standard deviation of the water PDF:

$$m_{dc} = mean(\sigma_{t,r}^0) < mean(\sigma_{W,\theta}^0) + \frac{1}{2} std(\sigma_{W,\theta}^0)$$

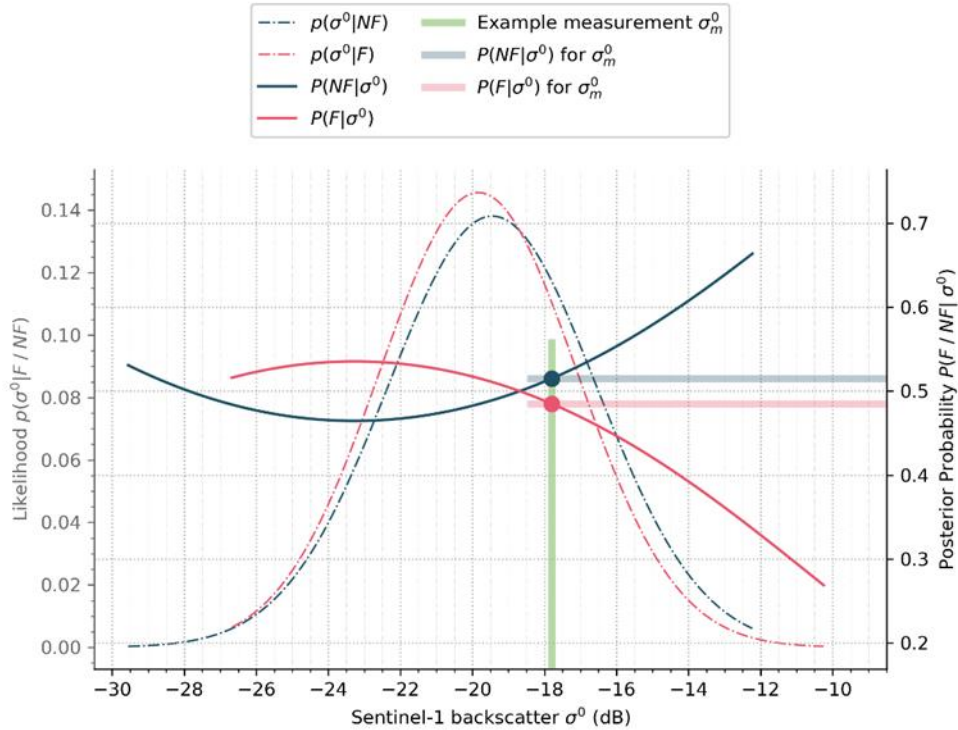<div align="right">**Eq. 4-33**</div>

Figure 12: Example of a conflicting distribution situation.

Besides the above mentioned all-season low backscatter situations (e.g. permanent water bodies, water look-alikes, etc.), the backscatter could be low for some seasons of the year only (e.g. rainy season, snow melting, etc.). An example of a conflicting constellation from the flood and non-flood PDFs is given in Figure 12. The consideration of such seasonal effects is one of the main features of the used harmonic model approach (details given in **Error! Reference source not found.**). In that perspective, if the used input observations show regularly low backscatter within a certain season, the corresponding expected DOY backscatter represents this for the flood mapping unfavourable behaviour. Figure 13 shows an example over France (near the city of Angers) in January 2022, where the expected DOY backscatter (a) shows a large area of low backscatter pixels, which may be related to seasonal inundation. Indeed, similar inundations can be observed in other recent years (b – e) as well. The harmonic model was trained in the current version by data from the years 2019 – 2020, and hence appear to reflect the local dynamics in that perspective.
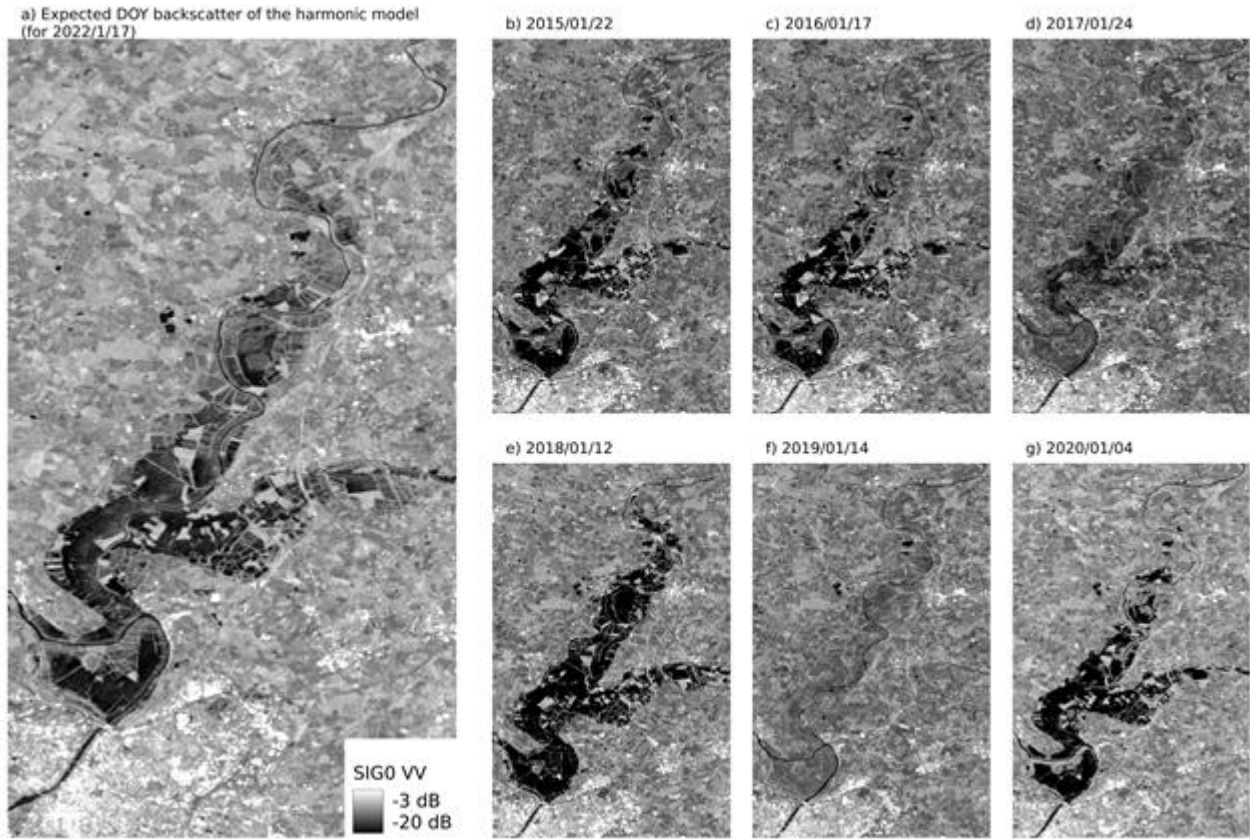
Figure 13: Example of the expected DOY backscatter compared to the Sentinel-1 observations from recent years (near the city of Angers, France).

Our statistical model is trained to provide robust decisions on normal flood conditions. Therefore, a meaningful Bayes decision cannot be provided in case of the input Sentinel-1 backscatter value (σ^0) not being represented by the flood or non-flood PDF. These extreme values could arise due to a statistical outlier for example and is masked by the following definition. Since it is assumed that the lower outliers of the flood PDF can still be classified as flooded, only higher outliers of the flooded PDF are considered within the masking. An example of an outlier observation is given in Figure 14.

$$m_{out}^{nf} = \sigma^0 < mean(\sigma_{t,r}^0) - 3std(\sigma_{t,r}^0) \; OR \; \sigma^0 > mean(\sigma_{t,r}^0) + 3std(\sigma_{t,r}^0) \quad \text{Eq. 4-34}$$

$$m_{out}^{f} = \sigma^0 > mean(\sigma_{W,\theta}^0) + 3std(\sigma_{W,\theta}^0) \quad \text{Eq. 4-35}$$

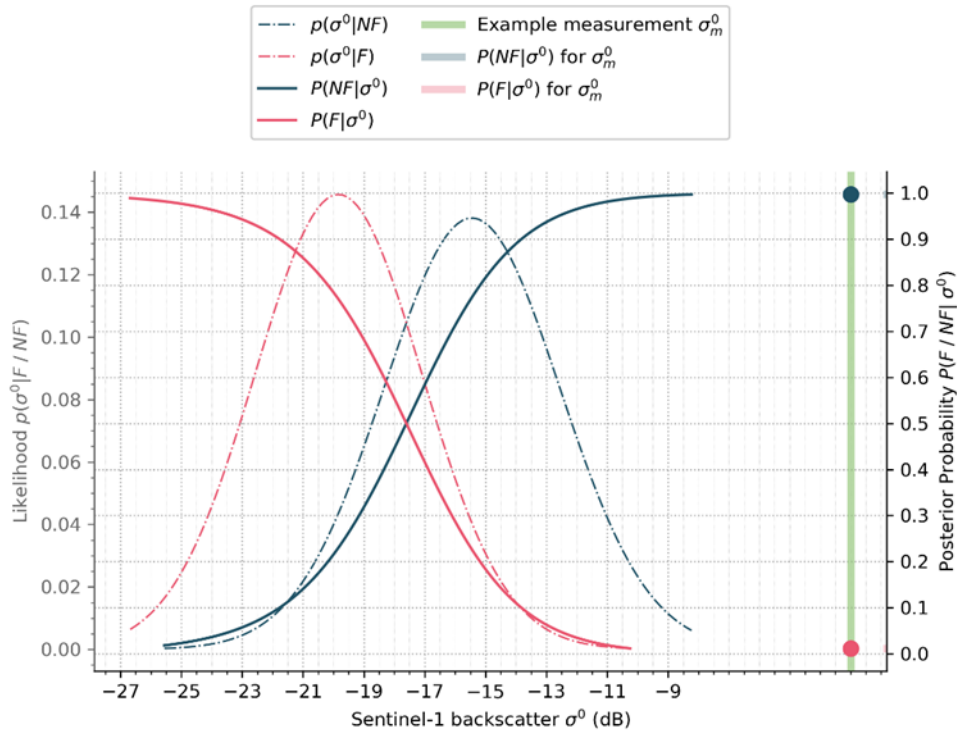$$m_{out} = m_{out}^{nf} \; AND \; m_{out}^{f} \quad \text{Eq. 4-36}$$

Figure 14: Example of a detected outlier measurement.

As the uncertainty (from the Bayes decision, details are given in Section 4.5.4) provides a measure for the confidence of the decision between flood and non-flood, it is used to limit the classification to certain decisions only. Therefore, we introduce a dedicated mask for high uncertainties which excludes pixels with uncertainties higher than 0.2:

$$m_{cert} = 0.2 < P(error|\sigma^0)$$ **Eq. 4-37**

An example of an uncertain decision is given in Figure 15. One can see that the incoming measurement is nearly equally probable to be assigned to one of the two classes (flooded and non-flooded).
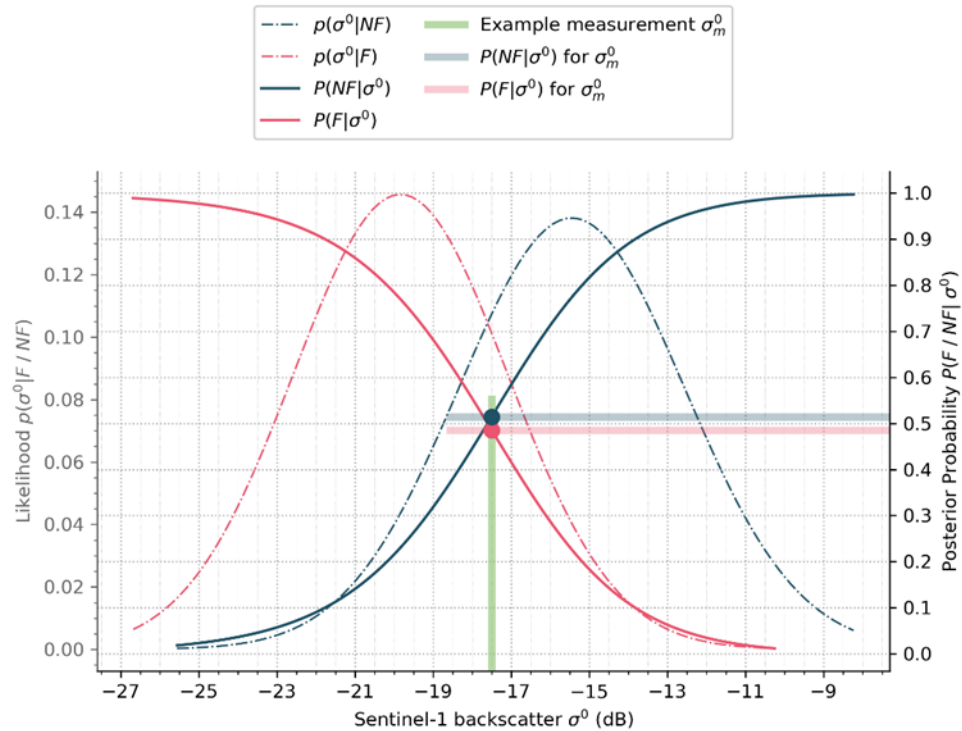
Figure 15: Example of pixel classification which shows a high uncertainty.

Morphological post-processing:

Due to the coherent nature of the radar signal, a single-time Sentinel-1 observation is in general affected by multiplicative noise (known as speckle). The random signal variation could lead to single pixels that feature a lower backscatter and could be confused with flood conditions. In order to reduce the influence of noise on the final flood extent map, a median filter of kernel size 5 is applied. This constitutes the final step and the result represents the flood extent and uncertainty output of algorithm 3.

4.1.4   Ensemble flood mapping algorithm

The GFM ensemble flood mapping algorithm combines the flood and likelihood results produced with all three individual flood mapping algorithms. The output is entirely pixel-based. A consensus decision of all algorithms, based on majority voting, determines if a pixel is marked as flooded or as non-flooded. To generate the combined product, each pixel is attributed with the ratio of the number of classifications as flooded to the number of algorithms that were applied. The value range is [0, 1] with:

Table 5: Degree of consensus to determine if a pixel is classified as flooded in the ensemble flood map

RATIO	DESCRIPTION	RESULT
0.0	No algorithm classifies given pixel as flooded	non-flooded
0.3	Only 1 of 3 algorithms classifies pixel as flooded	non-flooded
0.6	2 of 3 algorithms classifies pixel as flooded	flooded
1.0	3 of 3 algorithms classifies pixel as flooded	flooded

A value of 1.0 means that all three algorithms agreed on its classification as being flooded. A pixel is classified as flooded if it was classified as flooded by at least two algorithms.

The ensemble algorithm accepts Geotiff files for both the classification and uncertainty layers of each algorithm. The algorithm checks if all layers can be loaded. If a flood layer from one of the algorithms cannot be loaded, it is likely that the particular algorithm failed to produce a result. The number of algorithms applied equals to the number of flood layers that were loaded successfully, such that if one layer failed to load, the number of algorithms applied is 2.

Each flood result from the algorithms is given between 0 and 1. If two algorithm results could be loaded and both algorithms agreed on a pixel as flooded, the final pixel would have a value 1.0. If two of the three algorithms that could be loaded classified a pixel as flooded, the resultant ratio is 0.6. In the current implementation, only degrees of consensus ≥ 0.6 are considered in the ensemble computation.

The final likelihood layer is generated as the arithmetic mean at each pixelwise location calculated with all likelihood layers provided by each of the algorithms. The term *likelihood* represents the classification accuracy of the respective binary class. Likelihood values lie in the interval [0, 100], where values toward 0 represents lower confidence and values toward 100 represents higher confidence of flood classification accuracy. A likelihood value of 50 separates both classes where non-flood pixels necessarily show likelihood values in the interval [0, 49] and flood pixels show likelihood values in the interval [50, 100]. The classification confidence of the respective class increases with likelihood values propagating to the lowest or highest class boundary. Non-flood likelihoods propagate towards 0 if the non-flood confidence is to increase. In contrast to that flood likelihoods propagate towards 100 if the flood confidence is to increase.

As the TUW algorithm produces uncertainty values, TUW uncertainties have to be flipped so that low uncertainty values propagating towards 0 are remapped to high likelihood values propagating towards 100 and vice versa.

A series of potential cases are presented to illustrate how the ensemble algorithm behaves.

*Case 1: all algorithms produce inputs with full consent*

When all three flood algorithms return the same classification, the ensemble reflects this complete agreement. The final likelihood layer is the pixelwise mean of inputs from all three algorithm likelihood layers.
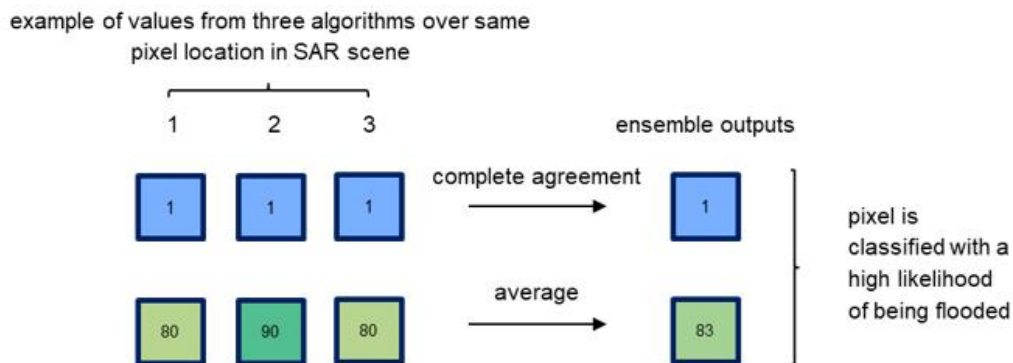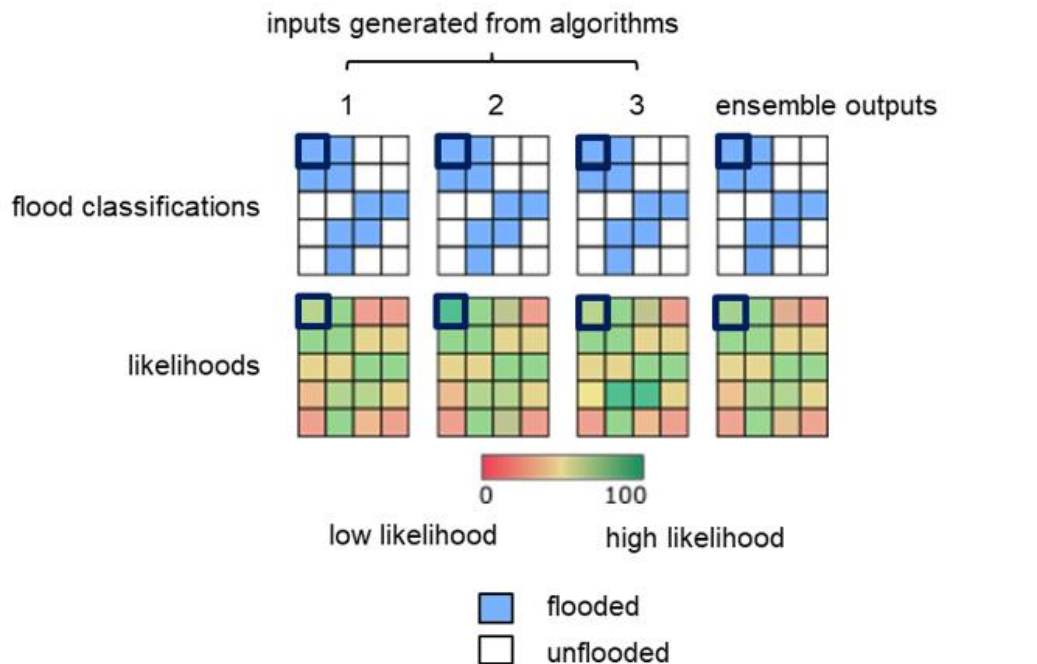
Figure 16: An example of full consent (complete agreement) among all three algorithms and also resulted in a given pixel being classified with a higher likelihood of being flooded.

Case 2: all algorithms produce inputs with majority consent

This is a variation of the usual case where all three algorithms produce inputs to the ensemble computation. If two of three algorithms agree that a pixel classification, the classification of the third algorithm is overruled. The likelihood result is the mean of all three likelihoods.
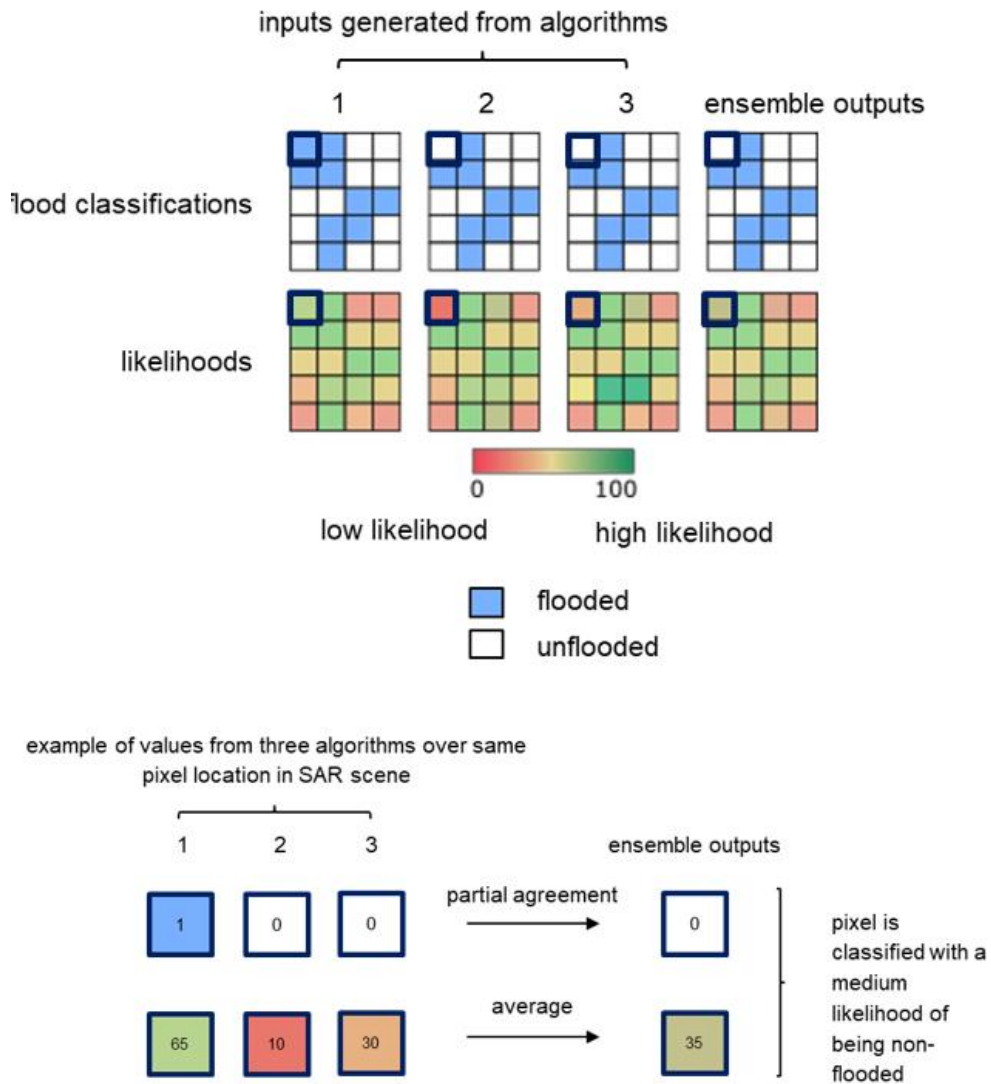

Figure 17: An example of majority consent (partial agreement) among the three algorithms and resulted in a given pixel being classified with a medium likelihood of being non-flooded.

Case 3: inputs from a single algorithm only

There may be cases where certain algorithms fail to produce inputs for the ensemble computation. However, since the ensemble algorithm requires the availability of inputs from at least 2 algorithms,

if only one provides results, an empty raster is generated and all pixels are assigned the lowest likelihood value of 0.
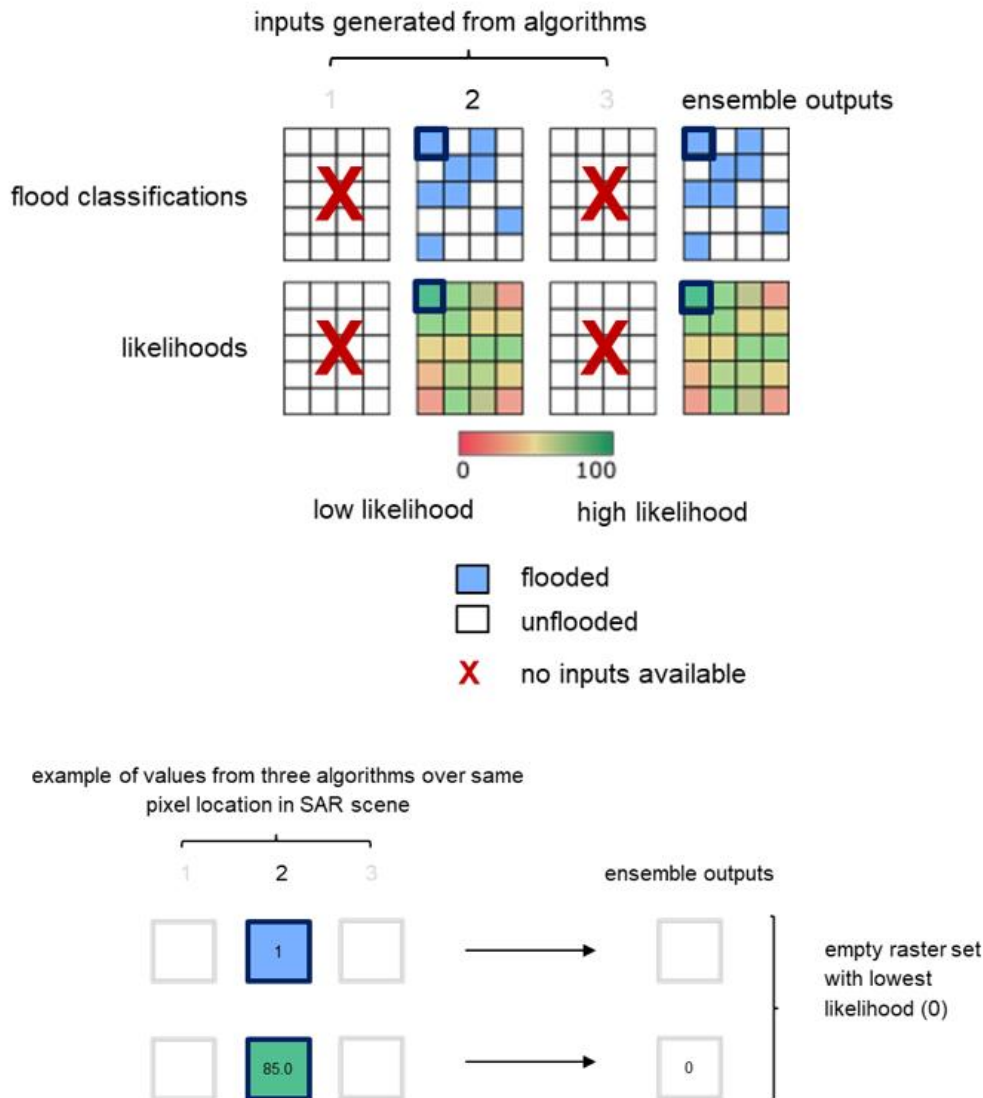

Figure 18: An example of classification and likelihood inputs being generated by a single algorithm only. Since the ensemble requires inputs from at least 2 algorithms, an empty flood classification raster is generated with the lowest likelihood of 0 assigned to each pixel.

Case 4: inputs from two algorithms only

In the unusual case that one algorithm fails to produce classification and likelihood inputs, the other two remaining algorithms decide on the flood classification For pixels where there is an agreement

between the two algorithms (Case 4A), the common classification is forwarded to the final result. This application is consistent with majority voting, which is the basis for the ensemble approach. For pixels where there is disagreement on the classification (Case 4B), i.e. one algorithm classified a given pixel as flooded while the other classified the same pixel as non-flooded, the pixel is assigned with the classification result that shows the highest confidence. The highest confidence is defined as the arithmetic distance to the basic likelihood interval boundaries [0, 100]. The likelihood value 50 has the greatest arithmetic distance.

In case of a split decision, the algorithm with the higher class confidence (greater distance from likelihood 50) overrules the detection from the algorithm with the lower class confidence (Case 4B). If both algorithms disagree on the flood classification but are equally confident in the respective class, the ensemble algorithm favours the flood detection over the non-flood detection (Case 4C).
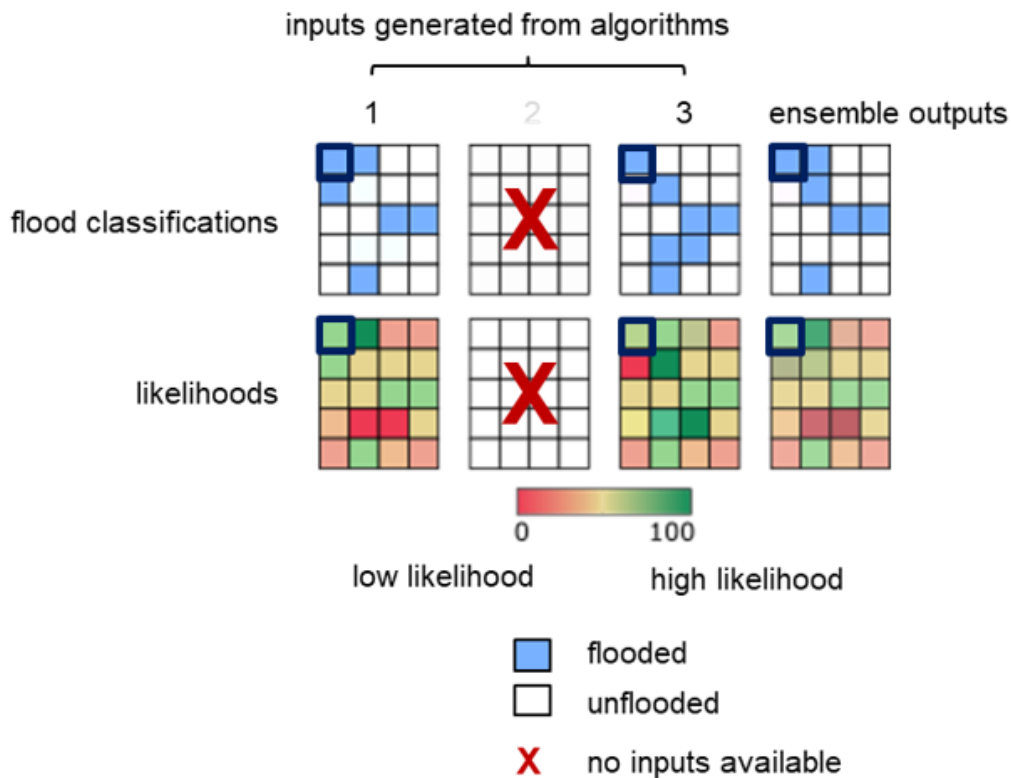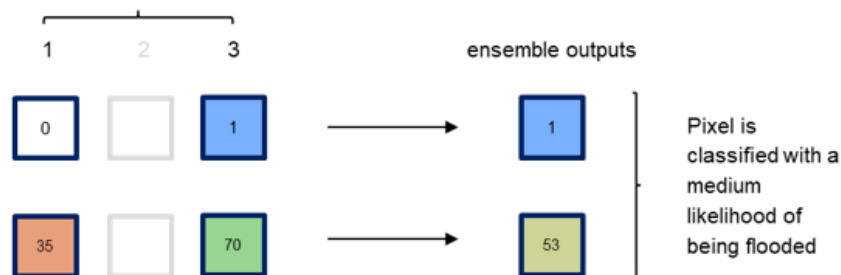

Case 4A: both algorithms agree on classification and have different likelihoods

example of values from three algorithms over same pixel location in SAR scene
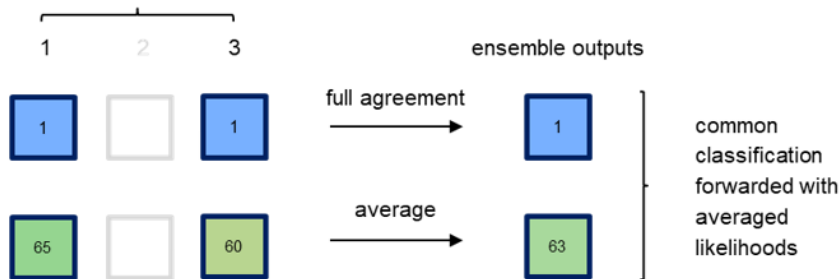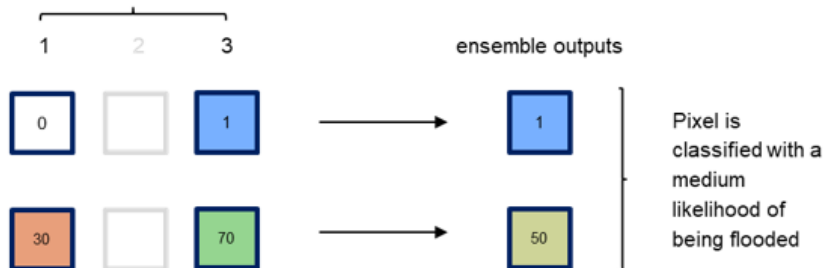
Case 4B: algorithms disagree on classification

example of values from three algorithms over same pixel location in SAR scene

example of values from three algorithms over same pixel location in SAR scene

Case 4C: algorithms disagree on classification and have equally strong likelihoods

Figure 19: Examples of three cases when only two algorithms produce inputs, where there is agreement (Case 4A), disagreement with different likelihoods (Case 4B) and equally strong likelihoods (Case 4C) on the classification, respectively.

We apply a region size dependant correction of the ensemble flood pixels and remove any flood pixel that forms a coherent region with its 8 neighbours and is less than a pre-defined blob size of 60. The indices of the removed flood pixels are applied to the likelihood layer so both layers show a consistent handling of valid pixels. As non-flood pixels can have a maximum likelihood of only 49,

we correct the likelihood value of pixels that were changed due to their blob size. Pixels that were corrected had a likelihood value of at least 50 (denoting flood) and are now assigned with a likelihood value of 49, i.e. the highest likelihood value that is allowed for non-flood pixels.

For all cases, we apply the reference water layer to the final flood and likelihood layers. This layer is comprised of both permanent and seasonal water extents. At any pixel location where there is the flood layer declares flood and the permanent or seasonal reference water layer also declare the presence of water, the pixel value of the flood layer is set to 0, i.e. no flood.

In a final step, we apply an exclusion layer to the flood and likelihood layer, correcting pixels that were misclassified due to radar shadow, permanent low backscatter and urban areas. We also apply the Copernicus Water Body Mask masking out any oceanic area from the flood ensemble classification.

4.2 Output layer "S-1 observed water extent"

The GFM product output layer "S-1 observed water extent" identifies the pixels classified as open and calm water based on Sentinel-1 SAR backscatter intensity, and is derived by combining with a logical OR the ensemble flood map obtained with the algorithm described in section 4.1.4 and the seasonal water mask as described in section 4.3. While the floodwater extent is produced using an ensemble approach which use all three algorithms (i.e. LIST, DLR and TUW), the ensemble approach to provide permanent and seasonal water uses only two algorithms (i.e. LIST and DLR). The "S-1 observed water extent" output layer values are described in Table 6.

Table 6: Description of "S-1 observed water extent" output layer values.

Value	Description	User Interpretation
0	No Water	No water detected on the Sentinel-1 image
1	Water	Water detected on the Sentinel-1 image
NaN	No data available	No data available at location (i.e. not covered by the satellite).

4.3 Output layer "S-1 reference water mask"

The GFM product output layer "S-1 reference water mask" outlines pixels that are classified as water, both permanent and seasonal, by an ensemble water mapping algorithm and using Sentinel-1 SAR mean backscatter intensity over a certain time period. The permanent water extent mapping is based on the mean backscatter of all Sentinel-1 data from a reference time period of two years. The seasonal reference water mapping uses as input the median backscatter of all Sentinel-1 data from a given month over the same reference time period. As a result, twelve masks are available, one per month, which include information on the permanent and seasonal reference water extent. The application of a permanent water mask is reliable in environments with stable hydrological

conditions over the year. In geographic regions with strong hydrological dynamics related to temporal change and variation in space of the water extent (e.g., in areas affected by monsoonal effects) the additional use of a seasonal reference water mask, which considers the usual water extent of the respective time period, is preferable. Figure 20 shows an example of reference water masks computed over a test area in Myanmar for the time-period 2019-01-01 to 2020-12-31. Seasonality, by means of monthly variations of water body areas, can be depicted and separated from permanent water bodies, which do not show alterations in their extent during the two years reference time-period. In Figure 20, each mask separates permanent water (pixels classified as water during the whole reference time-period) and seasonal water (pixels classified as water during a specific month within the reference time-period).
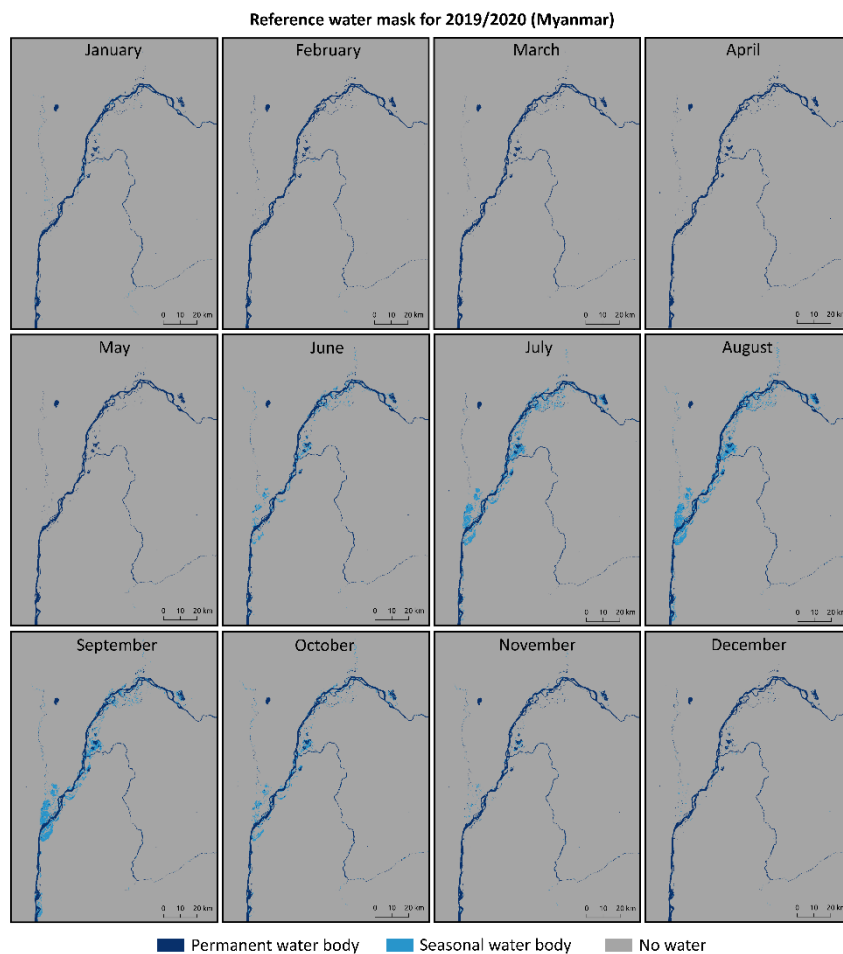
Figure 20: Exemplary reference water masks computed over Myanmar for the time-period 2019-01-01 to 2020-12-31. (See text for details).

The reference water mask product is updated once a year considering the previous two years of observation. This means that for example, the flood monitoring system running in 2021 will rely on the Sentinel-1 reference water mask being extracted from the Sentinel-1 mean backscatter over the reference time period from 2019 and 2020. The reference period of two years is chosen as

compromise in order to base the computation on a sufficient number of Sentinel-1 data and to remove long-term hydrologic changes which would be integrated in the reference waters masks if longer time periods would be considered. The Sentinel-1 reference water mask product values are described in the following table.

Table 7: Description of "S-1 reference water mask" output layer values.

Value	Description	User Interpretation
0	No Water	Users can assume this location normally not affected by the presence of water.
1	Permanent Water Body	Users can assume this location permanently covered by water.
2	Seasonal Water Body (for the current month)	Users can assume this location seasonally covered by water.
NaN	No data available	No data are available at this location (i.e., not covered by the satellite).

Algorithm for generating the S-1 reference water mask:

Observed permanent and seasonal water extent mapping is done by an ensemble water mapping algorithm that combines the water and uncertainty results produced with the LIST (section 3.1.3) and DLR (section 3.1.3) flood algorithms and having as input the mean backscattering value of S1 images available over a two-year time period (for permanent water) or the mean backscattering value of S1 images available over a two-year time period per month (for seasonal water). The results of the single water extent maps are combined into twelve masks, one per month, which include information on the permanent and seasonal reference water extent.

Both LIST and DLR flood algorithms produce a water extent map and associated likelihoods as by-product of the flood detection. In particular, the step ii of the LIST algorithm is used, i.e. the one that map the water extent using a single image and a 1-D region growing as described in the section 4.1.1. The TUW algorithm (section 3.1.3) is specifically designed to detect temporary flooded water bodies, does not produce a water extent map and hence is excluded for the more generic task of observed water extent mapping.

Therefore, the ensemble water mapping algorithm follows the same logic as the flood ensemble algorithm (section 3.1.3) but is specifically targeted towards the task of water mapping from two (not three) input sources. This means that a consensus decision of the two algorithms decides if a pixel is marked as water or as no water. To generate the combined product, each pixel is attributed with the ratio of the number of water classifications to the number of algorithms that were applied. A number of 1 would mean that both algorithms agreed on its classification as water. For pixels where they disagree on the classification (one algorithm classified that pixel as water and the other algorithm classified that pixel as no water) the algorithm with the highest likelihood for that classification dictates the result. That means, if the first algorithm is certain that a pixel is water (its likelihood value is high) and the second algorithm declares the pixel as no water with a low classification accuracy (its likelihood value is low), the first algorithm classifying that pixel as water dictates the final result. The likelihood result is the mean of all providing likelihood values for that

pixel. If only one algorithm provides results, the final water and uncertainty values equal the results from that particular algorithm.

Finally , for selected tiles where false negatives could be observed (e.g., large lakes with roughened surface that has falsely been classified as land) the Copernicus Water Body Mask is used to correct these. To enforce a consistent land-sea border, a land-sea mask derived from the Copernicus Water Body Mask is applied to all results.

4.4   Output layer "Exclusion mask"

The identification of conditions under which the GFM algorithms are deemed reliable is considered essential for an automated global flood monitoring system and is hence integral part during NRT flood detection and final product preparation. The key to a successful SAR-based identification and delineation of water bodies, and in subsequence of flood extents, is the low backscatter signature of water surfaces. Under common conditions, the water surfaces scanned by the SAR show significantly lower backscatter values than the surrounding area, the desired discrimination into water and non-water surfaces can be achieved with high accuracy and reliability. With the inclusion of information generated from time series stored in the datacube, additional discrimination criterions are established, further improving the reliability of the GFM flood extent. Nevertheless, there are several situations when the flood extent delineation is hampered, or even completely impeded, and where an **Exclusion Mask** is deemed necessary to prevent misclassification.

First, the flood detection is challenged by non-water surfaces that feature low backscatter themselves and would yield false alarms when not treated properly. For instance, very dry or sandy soils, frozen ground, wet snow and flat impervious areas (e.g. smooth tarmac covers as airfields or roads) feature very low backscatter signatures and appear thus as water-look-alikes in SAR imagery. Another common effect in SAR remote sensing is radar shadowing, which appears over strong terrain (especially at the far-range section of the SAR image) as well as in the vicinity of high objects above the ground, like high buildings and along forest borders. Whereas the first can be recognised well with terrain- and observation geometry analysis, the second can only be identified by backscatter time series analysis, considering the global scope of the service (as real Digital Surface Models are only available for some areas).

Second, floods occurring in urban or vegetated areas bear the danger of missed alarms, as co-located flood extents potentially feature high instead of low backscatter. The formation of corner reflectors for the SAR's microwaves through perpendicular buildings (urban) or plant stems (vegetation) over standing water surfaces leads generally to high backscatter and thus common SAR water mapping approaches fail in detecting water surfaces. Moreover, the detection of water-bodies under densely vegetated canopies is complicated by the absorption and diffuse scattering in the canopy itself, reducing the sensitivity to processes on the ground.

The GFM product output layer "Exclusion Mask" addresses static topography and land cover effects, where the interaction of C-band microwaves with the land surface is in general complex. In several situations, flooded areas cannot be detected by Sentinel-1 C-band SAR over certain land cover types and terrain conditions for physical reasons. We **grouped** these reasons and **defined four types**:

1. Sentinel-1 does not receive sufficiently strong signals from the ground surface to distinguish a flooded from a non-flooded surface. In such a case, we encounter **No-Sensitivity** to detect water surfaces.

2. Sentinel-1 senses the ground, but backscatter from the non-flooded surface is in general so low as to be indistinguishable from backscatter from smooth open water. Here we encounter water-look-alikes due to a **Low Backscatter** signature of the ground.

3. The Sentinel-1 signals are heavily distorted by terrain effects, effectively enhancing the noise and signal disturbances to such a degree to that it becomes larger than the change in backscatter due to potential flooding. As this problem generally occurs over areas with strong topography, we refer to this with **Topographic Distortions**.

4. Sentinel-1 receives no signals from certain regions of the land surface due to **Radar Shadows** casted by mountains, high vegetation canopies or anthropogenic structures.

Reflecting this grouping of physical and geometrical effects, we generate the Exclusion Mask by combining layers that address the four individual effect groups. With the addition of the no-data values from the expert flood detection algorithm output, it indicates on a binary map the pixel locations where the SAR data could not deliver the necessary information for robust flood delineation. With respect to the complexity of the identified effects, and to potential spatial overlaps and interactions, a combined, single Exclusion Mask that is also easy for the user to interpret, is provided.

4.4.1 Algorithm for generating the Exclusion Mask

The Exclusion Mask is based on offline-generated Sentinel-1 SAR parameters and on auxiliary thematic datasets. It is accessed and subset in NRT and added to the other product layers. The necessary operations are done at the 20m-sampling of the Sentinel-1 preprocessed data cube. In the following, we describe the methods to address the four types of disturbing effects.

4.4.1.1 Masking of no-sensitivity areas

The no-sensitivity layer delineates all land cover types and areas where Sentinel-1 CSAR is not sensitive to flooding – and in fact any other type of change – of the ground surface. These non-sensitive areas encompass both densely vegetated- and urban areas. In both cases, the loss in

sensitivity reflects the fact that the Sentinel-1 signal is dominated by backscatter echoes from other parts of the observed scene.

In the case of vegetation, backscatter from the vegetation canopy increasingly dominates the signals coming from the ground for increasing biomass levels. Given the limited penetration capability of C-band waves, Sentinel-1 becomes essentially insensitive to flooding over high biomass areas such as forests and shrubland with biomass levels larger than 30-50 t/ha Quegan et al., (2000). Backscatter from these high biomass land cover classes is in general rather stable and high, a feature that is often exploited in land cover classification schemes.

To identify the densely vegetated areas, parameters from the Sentinel-1 Global Backscatter Model (S1-GBM) were used, namely the mean backscatter in VV ($\mu_{VV}^{\sigma^0}$) and VH ($\mu_{VH}^{\sigma^0}$) polarization, mean cross polarization ratio ($\mu_{CPR}^{\sigma^0}$, computed as VV/VH backscatter) and standard deviation in VH polarization ($std(\sigma_{VH}^0)$). Locally variable thresholds (τ_i) were defined for each parameter to discriminate between the vegetated and not vegetated areas setting the rule that the pixel is classified as dense vegetation in case that the following rules apply:
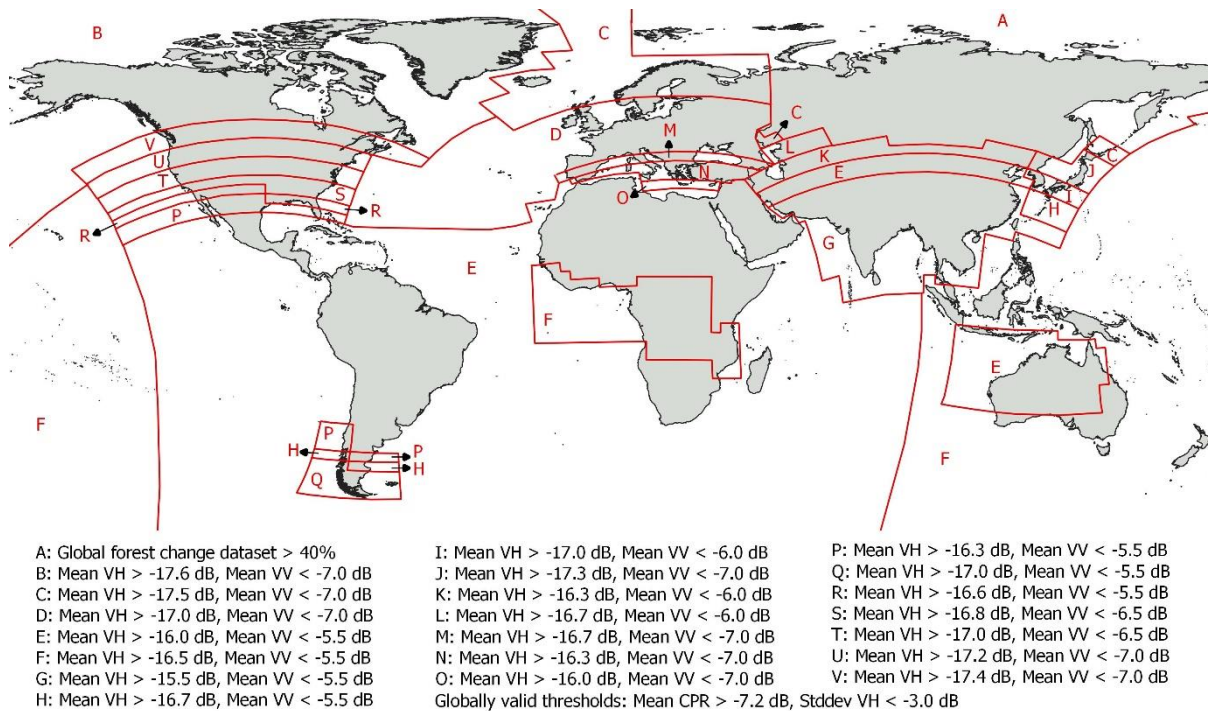
$$mean(\sigma_{VV}^0) > \tau_{\mu-VV} \qquad \text{Eq. 4-38}$$
$$mean(\sigma_{VH}^0) < \tau_{\mu-VH} \qquad \text{Eq. 4-39}$$
$$mean(\sigma_{CPR}^0) > \tau_{\mu-CPR} \qquad \text{Eq. 4-40}$$
$$std(\sigma_{VH}^0) > \tau_{\sigma-VH} \qquad \text{Eq. 4-41}$$

The thresholds were optimized separately for each continent and varied with northing and vegetation type. The Global Forest Change dataset (Hansen et al., (2013); version 1.8: https://storage.googleapis.com/earthenginepartners-hansen/GFC-2020-v1.8/download.html) was used as a reference dataset for the threshold selection and is used directly as a dense vegetation mask in regions, where the irregular coverage of Sentinel-1 data caused artefacts (e.g. stripes of higher and lower backscatter) in the S1-GBM layers. The global distribution of the thresholds is shown in Figure 21.

A: Global forest change dataset > 40%
B: Mean VH > -17.6 dB, Mean VV < -7.0 dB
C: Mean VH > -17.5 dB, Mean VV < -7.0 dB
D: Mean VH > -17.0 dB, Mean VV < -7.0 dB
E: Mean VH > -16.0 dB, Mean VV < -5.5 dB
F: Mean VH > -16.5 dB, Mean VV < -5.5 dB
G: Mean VH > -15.5 dB, Mean VV < -5.5 dB
H: Mean VH > -16.7 dB, Mean VV < -5.5 dB

I: Mean VH > -17.0 dB, Mean VV < -6.0 dB
J: Mean VH > -17.3 dB, Mean VV < -7.0 dB
K: Mean VH > -16.3 dB, Mean VV < -6.0 dB
L: Mean VH > -16.7 dB, Mean VV < -6.0 dB
M: Mean VH > -16.7 dB, Mean VV < -5.5 dB
N: Mean VH > -16.3 dB, Mean VV < -7.0 dB
O: Mean VH > -16.0 dB, Mean VV < -7.0 dB
Globally valid thresholds: Mean CPR > -7.2 dB, Stddev VH < -3.0 dB

P: Mean VH > -16.3 dB, Mean VV < -5.5 dB
Q: Mean VH > -17.0 dB, Mean VV < -5.5 dB
R: Mean VH > -16.6 dB, Mean VV < -5.5 dB
S: Mean VH > -16.8 dB, Mean VV < -6.5 dB
T: Mean VH > -17.0 dB, Mean VV < -6.5 dB
U: Mean VH > -17.2 dB, Mean VV < -7.0 dB
V: Mean VH > -17.4 dB, Mean VV < -7.0 dB

Figure 21: Global distribution of thresholds used to identify densely vegetated areas using S1-GBM and Global Forest Change dataset

Over urban areas backscatter as measured by C-band SAR sensors is often dominated by comparably few, but very strong echoes received from buildings and other artificial objects (Sauer et al., 2011). The presence of large corner reflectors in form of buildings perpendicular to smooth horizontal ground create an extremely high backscatter signature, often close to saturation. Therefore, in the absence of sufficiently large open spaces between the buildings and urban vegetation, inundation in urban areas is difficult to detect.

To identify the urban areas, GHSL and WSF2015 static urban masks were used. The urban areas were defined as pixels, where (i) the GHS-BUILD value exceeds 30% probability of built up area and (ii) are identified as settlements in WSF2015 urban mask.

Lastly, the dense vegetation mask and urban areas mask were combined into the binary No-sensitivity mask.

4.4.1.2 Masking of non-water low backscatter areas

There is another situation where Sentinel-1 is not sensitive to the flooding of the ground, but for quite a different reason as for vegetation and urban areas. In this second case, the lack of sensitivity is related to the fact that backscatter from rather smooth surfaces (smooth compared to the 5.5 cm

wavelength of Sentinel-1) can be as low as backscatter from a water surface. Smooth non-water surfaces may act similar like water as a specular reflector scattering most energy of the SAR signal in forward direction, leading to very low backscatter values. Examples of such permanently low-backscatter areas are tarmacs (e.g. airports, motorways, etc.) or sandy surfaces (e.g. in riverbanks and deserts). The separation between water and smooth land surfaces based on a single SAR acquisition is hardly feasible. Therefore, the generation of a water look-alike layer with permanent low backscatter is performed based on statistical information from Sentinel-1 time-series data; pixels are masked in case that the occurrence of low backscatter values (below –15 dB) exceeds 70% of all values in the respective time-series.

4.4.1.3 Masking of topographic distortions

The consideration of information on local topography is essential to flood detection in two ways. There are areas that must be excluded during flood detection due to the geometry of the SAR observation – this is dealt with measures discussed in next subsection.

Then there are areas that shall be excluded from flood detection because of the nature of the phenomenon itself, as floods appear commonly over non-sloped surfaces, in vicinity to streams, rivers, and other permanent water bodies. The latter is addressed by the Height Above Nearest Drainage (HAND) index. The consideration of the HAND index is of vital importance to the reliability of the flood detection algorithms, as SAR signal disturbances over strong topography can be substantial and dominate the surface backscatter signature, and its change over time, leading to otherwise uncontrollable over- and underestimation of flood indications.

Based on the HAND index, a binary exclusion mask (termed "HAND-EM") has been calculated to separate flood-prone from non-flood prone areas based on the elevation of each pixel to the nearest water-covered pixel Twele et al., (2016); Chow et al. (2016), which we use as input the Exclusion Mask.

Both binary classes are determined using an appropriate threshold value. Choosing the threshold value too high may lead to misclassifications (i.e. the inclusion of flood-lookalikes in areas much higher than the actual flood surface and drainage network) while a threshold value set too low would eliminate valid parts of the flood surface. The choice of an appropriate threshold is thus critical, and was derived through a series of empirical tests including more than 400 Sentinel-1 and TerraSAR-X datasets of different hydrological and topographical settings Twele et al., (2016).

Due to the global application scope of the GFM Sentinel-1 flood processing chain, a rather conservative threshold value of ≥15 m was finally chosen for defining non-flood prone areas. The HAND-EM has been further shrunk by one pixel using an 8 neighbour function to account for potential geometric inaccuracies between the exclude layer and the radar data.

It must be noted that the Copernicus DEM is not hydrologically conditioned, hence the HAND index cannot be computed on this version of the elevation model.

4.4.1.4 Masking of Sentinel-1 radar shadows

Radar shadow is a term for unreachable areas by a radar signal, where no information can be gained. Such areas are commonly located in mountainous regions, around high anthropogenic structures, and along forest borders (forest lines). The extent of radar shadows increases with radar range, thus the far range section of a SAR scene is more affected by it. After preprocessing, radar shadow areas have extreme low backscatter values, and are therefore very likely to be confused with water surfaces leading to false alarms in flood detection.

Radar shadows caused by mountains can be rather easily detected by analysing the SAR sensing geometry with respect to a Digital Elevation Model (DEM). This can be computed offline with a geometric observation model of the Sentinel-1 orbits, without using the Sentinel-1 SAR archive.

However, a DEM may only describe the terrain elevation not the surface elevation or may have a too coarse resolution to catch finer details, thus it does not necessarily include anthropogenic structures or forests. The radar signal (at C-band) is mainly reflected at the surface leading to discrepancies with respect to a terrain model. Therefore, a geometric observation model is not applicable for such smaller structures causing radar shadows. For example, shadows on the boundary of a forest as shown in Figure 22 would be missed. Fortunately, we can make use of the fact that radar shadow depends also on the azimuthal looking angle and can be extracted by analyzing backscatter from multiple orbits in different orbit pass directions (ascending and descending). Radar shadow areas are characterized by huge differences in backscatter for opposite orbit directions, where one direction is prone to shadowing and the other to double bounce scattering effects causing high backscatter values.

We analyze these differences for orbits of opposite directions, and generate a radar shadow mask for each Sentinel 1 relative orbit number (which feature a very stable viewing geometry): By comparing mean of backscatter at one orbit with mean of backscatter from orbits in the opposite direction, the extent of the orbit specific shadow areas can be estimated.

For a given orbit, pixels will be masked as shadow, when the mean of backscatter at current orbit is less than -15 dB and mean of backscatter from orbits in the opposite direction is higher than -10 dB. (see also Figure 22 below for illustration).
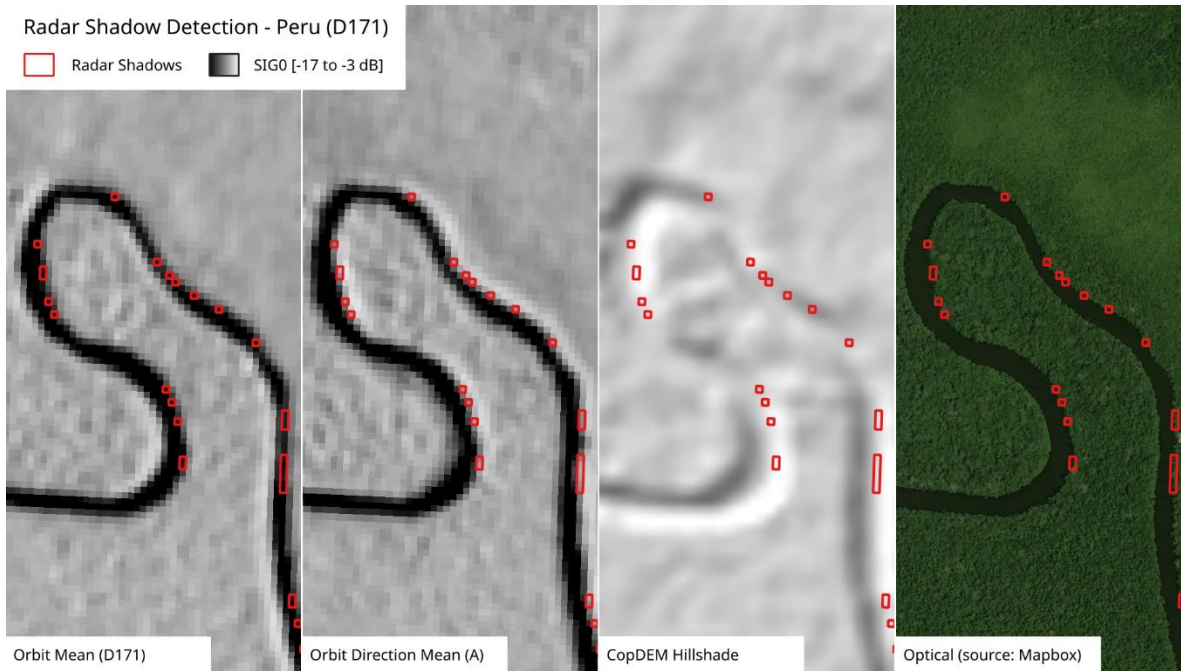
Figure 22: Radar shadow effect on the boundary of forest (columns from left to right): mean backscatter from decending orbit 171, mean backscatter from ascending pass, Copernicus DEM, optical image from Mapbox satellite map

Because of the new availability of the Copernicus DEM as a digital surface model with an appropriate spatial resolution it could be considered to derive a shadow mask based on a geometric observation model. However, a benchmarking of both outputs is missing, but due to its simplicity, computation time and independency on auxiliary data, preference is given to backscatter differences of opposite orbit directions.

4.5 Output layer "Uncertainty values"

The GFM product output layer "Likelihood values" provides the estimated likelihood that is associated with the applied GFM flood mapping algorithms, for all areas outside exclusion mask.

4.5.1 Estimation of uncertainty for flood mapping algorithm 1 (LIST)

The likelihood of floodwater classification is characterized by flood probability. In the case (i), both I_{t0}^{S1} and I_{D}^{S1} are used for likelihood estimation, the pixels those have high posterior probability of both water class and change class are likely to be real flooded pixels. The probability of been flooded for a given pixel is defined as the minimum value between $p(W|\sigma^0)$ and $p(C|\Delta\sigma^0)$:

$$p(F|\sigma^0, \Delta\sigma^0) = min(p(W|\sigma^0), p(C|\Delta\sigma^0))$$
Eq. 4-42

In the case (ii), only I_{t0}^{S1} is considered for likelihood estimation of flood classification:

$$p(F|\sigma^0) = p(W|\sigma^0) \qquad \text{Eq. 4-43}$$

As in this case the likelihood is only calculated from the backscatter value in I_{t0}^{S1}, false high flood probability can be caused by permanent water and other water look-alike dark areas, these false alarms in binary map has been removed by comparing the resulting flood map with the previous flood map as described above. To reduce these false high probabilities in current likelihood map, for non-flood pixels in the new flood map, their flood probability is the minimum value between $p(W|\sigma^0)$ and the value in the latest previous likelihood map.

In both cases of (i) and (ii), the flood probability of pixels located in exclusion layers is set to 0.

4.5.2 Estimation of uncertainty for flood mapping algorithm 2 (DLR)

A fuzzy logic-based approach is applied to measure and to reduce the uncertainty associated with the water classification. Three cases influence the classification likelihood:

1. The likelihood of a SAR backscatter pixel being classified as water is low if that pixel value is close to the threshold τ.

2. The likelihood of a SAR backscatter pixel being classified as water is high if the slope at that location is low, since steeper surfaces are unlikely to retain water.

3. The likelihood of a SAR backscatter pixel being classified as water is high if that pixel forms a coherent area with other neighbouring water pixels and the area is relatively large.

Figure 23 demonstrates the fuzzy logic approach in detail for the first case considering the SAR backscatter likelihood. In Figure 23, the threshold τ of the parent tile is the upper fuzzy value x_2 and separates the water and non-water (land) classes. This value represents the boundary of both classes, where the uncertainty of a correct classification is highest. The mean backscatter value of the class water μ_{water} is associated with a lower fuzzy value. As the majority of water pixels share that backscatter value, the likelihood of a correct classification is high. In fuzzy logic terminology, high likelihoods map to high degrees of membership to a particular class, i.e. a high likelihood of a correct classification to the water class corresponds to a high degree of membership to the water

class and is assigned a higher fuzzy value. The converse is also true, where a low likelihood corresponds to a low degree of membership to the water class and is assigned a low fuzzy value.

In Figure 23, , SAR backscatter pixel values close to the threshold τ have a low likelihood of being correctly classified as water, since they may also belong to the non-water class, within the classification error range. Following the fuzzy logic terminology, SAR backscatter pixel values propagating to the threshold τ have a lower degree of membership to the class and are assigned a lower fuzzy value approaching 0. SAR backscatter values that are assigned to the water class and are also close to the class mean μ_{water} have a higher certainty of being correctly classified as water. Therefore, they are assigned with a higher degree of membership to the water class and are assigned with higher fuzzy values approaching 1.
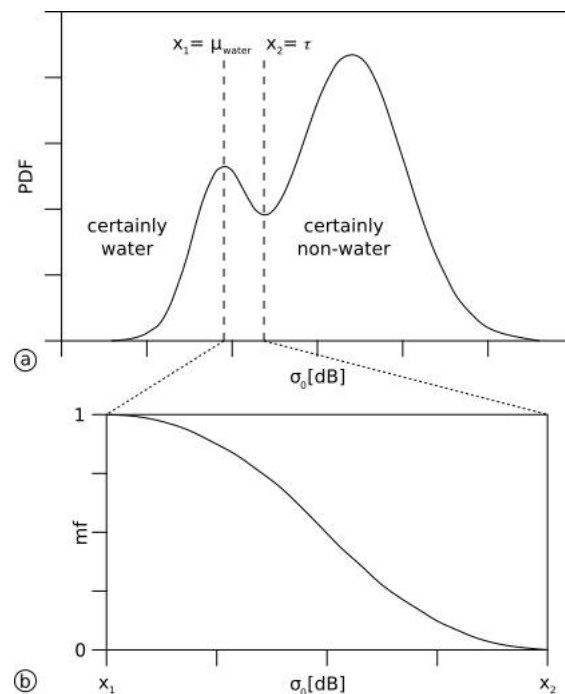
Figure 23: Fuzzy logic approach for discriminating water and non-water classes based on SAR backscatter values. (See text for details).

The diagram in Figure 23b depicts the membership function as the standard S-function, described in Pal et al. (1988). As high SAR backscatter values propagate to the threshold τ (low membership degree) and low SAR backscatter values propagate to the mean of the water class μ_{water} (high membership degree), we negate the standard S-function (

Figure 24).

In Figure 24, the positive S-function applies when assigning low pixel values with a low degree of membership and high pixel values with a high degree of membership. The negative S-function applies vice versa. Any SAR backscatter value that is less than or equal to μ_{water} is assigned a high degree of membership, thus a fuzzy value of 1. Any SAR backscatter value that is greater than or equal to τ is assigned a low degree of membership, thus a fuzzy value of 0.

Slope values equal to 0 have a high degree of membership, thus a fuzzy value of 1. Any slope that is greater than or equal to an empirically determined value of 18° has a low degree of membership, thus a fuzzy value of 0. For the slope, we also apply the negative S-function.

Considering the size of an area detected as water, we define a minimal, empirically determined mapping unit range of [10, 500], i.e. any water area with a size less than or equal to 10 pixels has a low degree of membership, thus a fuzzy value of 0. Any water area with a size greater than or equal to 500 has a high degree of membership, thus a fuzzy value of 1. For the minimal mapping unit, we apply the positive S-function.
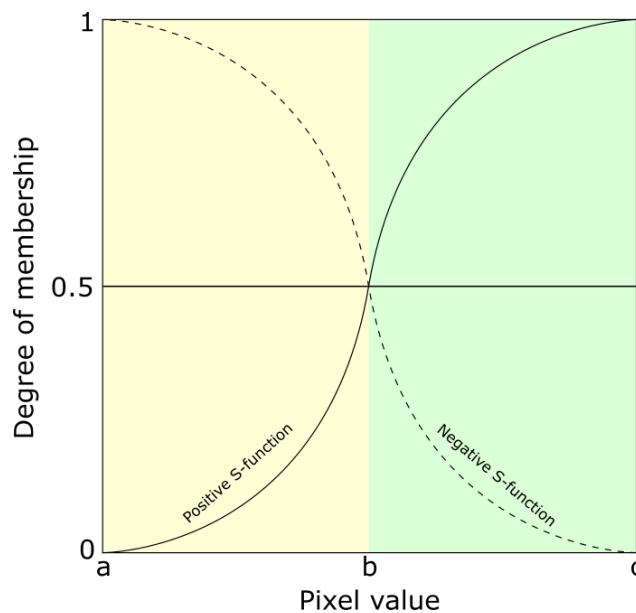
Figure 24: Standard S-function modified from Pal-Rosenfeld (1988). (See text for details).

The result of the fuzzy logic approach are three fuzzy layers in the range [0, 1]. For performance reasons, the float values were rescaled to the range [0, 100]. The resulting fuzzy layer is the mean of all three individual fuzzy layers.

A region-growing algorithm is applied as post-processing step, given the fuzzy layer and the initial water detection as inputs.

Water pixels with a fuzzy value of ≥ 0.6 are treated as correct water classifications. Water pixels with a fuzzy value of ≥ 0.45 and < 0.6 are treated as potential water pixels that are subject to region-growing.

Water pixels with a high fuzzy value of ≥ 0.7 are treated as seed pixels. Any potential water pixel, i.e. with a fuzzy value of ≥ 0.45 and < 0.6, that is also directly adjacent to a seed pixel is then remapped as a certain water pixel. As the original fuzzy value of that pixel was below the critical certainty mark of 0.6, we also remap the fuzzy value to the lowest uncertainty corresponding to the water class, i.e. 0.6.

As a final step, we further modify water detections based on the size of the region in which they are located. A region detected as water must have a minimum size of 30 pixels. Otherwise, that region is treated as a misclassification and remapped as land. The converse is also true, such that a region detected as land must have a minimal size of 10 pixels. Otherwise, the region is treated as a misclassification and remapped as water.

Any original water region that is remapped is assigned a fuzzy value representing land with a low classification accuracy, i.e. 0.45. Any original land region that is remapped as water is assigned a fuzzy value indicating water with a low classification accuracy, i.e. 0.6.

4.5.3   Estimation of uncertainty for flood mapping algorithm 3 (TUW)

As mentioned in Section 4.1.4, Algorithm 3 Bauer-Marschallinger et al, (2021) yields posterior probabilities of a backscatter measurement for belonging to either the flood or non-flood class. Based on the Bayes decision rule the higher posterior probability defines the class allocation. Additionally, the *conditional error* $P(error|\sigma^0)$ can be defined by the lower posterior probability.

$$P(error|\sigma^0) = \min\left[P(F|\sigma^0), P(NF|\sigma^0)\right] \qquad \text{Eq. 4-44}$$

The conditional error as measure for the *uncertainty* allows directly to quantify the (non-) confidence of the provided decision. Since per definition the posterior probabilities sum up to 1, a higher posterior probability for one class results in a lower posterior probability for the other class. Consequently, the uncertainty is defined between 0.0 and 0.5. An uncertainty close to zero implies a very confident decision, since the probabilities for both classes (flood and non-flood) support a clear decision. The opposite condition would be an uncertainty of close to 0.5, which implies that the probabilities of the input Sentinel-1 value belonging to the two classes is nearly identical.

For all pixels of the incoming Sentinel-1 image, the conditional errors ($P(error|\sigma^0)$) are forwarded to the ensemble algorithm, representing the uncertainties of Algorithm 3's flood mapping, per individual pixel.

In case of high conditional errors (i.e. close to 0.5), no certain decision can be made. In order to exclude such poorly based decisions (with low reliability), a dedicated uncertainty mask is already applied in Algorithm 3 (see Section 4.1.4 for more details).

4.6 Output layer "Advisory flags"

Various meteorological factors may hamper or even prevent the detection of flooded areas. As discussed by Matgen et al. (2019) strong winds, rainfall, as well as the presence of wet snow, frost and dry soils are of particular concern. These factors are all important because they reduce the contrast between backscatter from open water surfaces and backscatter from the surrounding areas, thereby reducing the separability of flooded areas and surrounding land. While wind, frozen water surface and rainfall might reduce the backscatter contrast by increasing backscatter over water due to wind- and rainfall-induced roughening of the water surface, wet snow, frost and dry soils reduce the contrast by decreasing backscatter over the surrounding areas.

Unfortunately, all these factors are difficult to capture because of their extremely dynamic nature and high spatial heterogeneity. Therefore, wind, rainfall, temperature, and snow data coming from sparsely distributed meteorological stations and/or numerical weather prediction with its much coarser resolution are hardly suited to capture the exact situation at the time of the S-1 acquisition and are not ideal for providing the required advisory flags. It should be noted that optical remote sensing satellites such as MODIS or Sentinel-3 would provide sufficient spatial details, but neither do optical data depict the environmental conditions as seen by Sentinel-1, nor do they provide timely observations at all times due to cloud cover and poor lightning conditions. Therefore, in the scope of the GFM service, the only robust and applicable source for the environmental advisory flags is microwave remote sensing data. Our approach builds on using the S-1 NRT data and Sentinel temporal parameters (see section 3.1.2) that were derived from time series within the Sentinel-1 data cube archive.

The GFM product output layer "Advisory flags" aims for raising awareness that meteorological processes comprising wind or frozen conditions might impair the water body detection. As the here defined Advisory Flags can only be retrieved at a coarser resolution, we do not forward the information of the flags to the masking or to the Exclusion Mask. As coming in the form of the additional layer, it should guide the users when interpreting the product, allowing additional insight on its local reliability at the time of Sentinel-1 acquisition. The Advisory Flags indicate locations where the SAR data might be disturbed by such processes during the acquisition, but leaves the flood and water extent layers unmasked. In summary, pixels marked by the Advisory Flag are not excluded by the Exclusion Mask, but users are advised caution when using the observed water and flood extents over flagged areas. As outlined in the following section, the Advisory Flag layer consists of 2 separate flags, defining 4 potential flag values:

VALUE	DESCRIPTION	USER INTERPRETATION
0	**No Advisory Flag** was set	Users can assume high quality
1	**Low Regional Backscatter** (Snow, Ice, or Dryness)	Users are advised to caution, due to frozen or dry or snow-covered soil affecting the Flood mapping reliability.
2	**Rough Water Surface** (Wind, rainfall, or frozen water surface)	Users are advised to caution, due to local wind, rainfall, or frozen water surface affecting the Flood mapping reliability.
3	**Both**: Low Regional Backscatter **&** Rough Water Surface	Users are advised to caution, due to frozen or dry or snow-covered soil & local wind, rainfall, or frozen water surface affecting the Flood mapping reliability.

Figure 25 shows an example of the output layer "Advisory flags" for a Sentinel-1 scene over Greece from 2020-08-10 at 16:32 (from orbit A175), a point in time when windy conditions where prevalent. In Figure 25, the Advisory flags are overlaid on the backscatter image with transparent colours, with red for Low Regional Backscatter (Snow, Ice, or Dryness), green for Rough Water Surface (Wind), and blue for both.
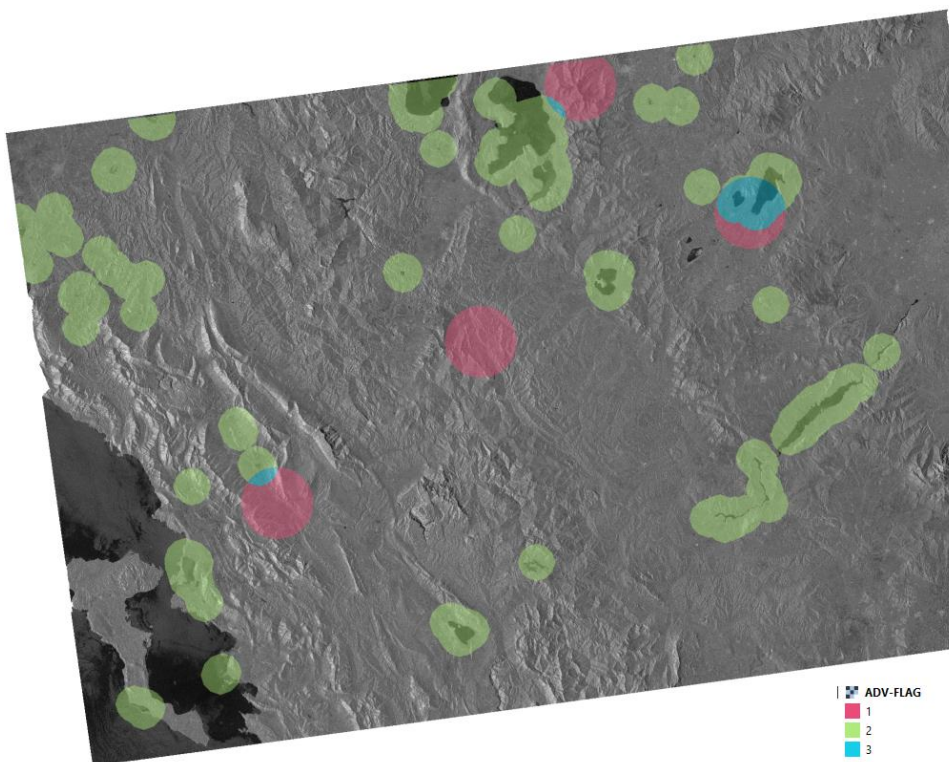
Figure 25: Example of the GFM product output layer "Advisory flags" for a Sentinel-1 scene over Greece from 2020-08-10 at 16:32 (from orbit A175), a point in time when windy conditions where prevalent. (See text for details).

4.6.1 Algorithm for generating the Advisory flags

For the Advisory Flag layer production, we need to combine the near real-time input data flow, as well as read out the corresponding statistical parameters and auxiliary layers. After processing the flood extent for the incoming Sentinel-1 scene, the involved Equi7Grid-tiles are identified, and the collocated parameters and auxiliary layers are read from the database. Figure 26 illustrates the data flow and how the flag values are set logically.
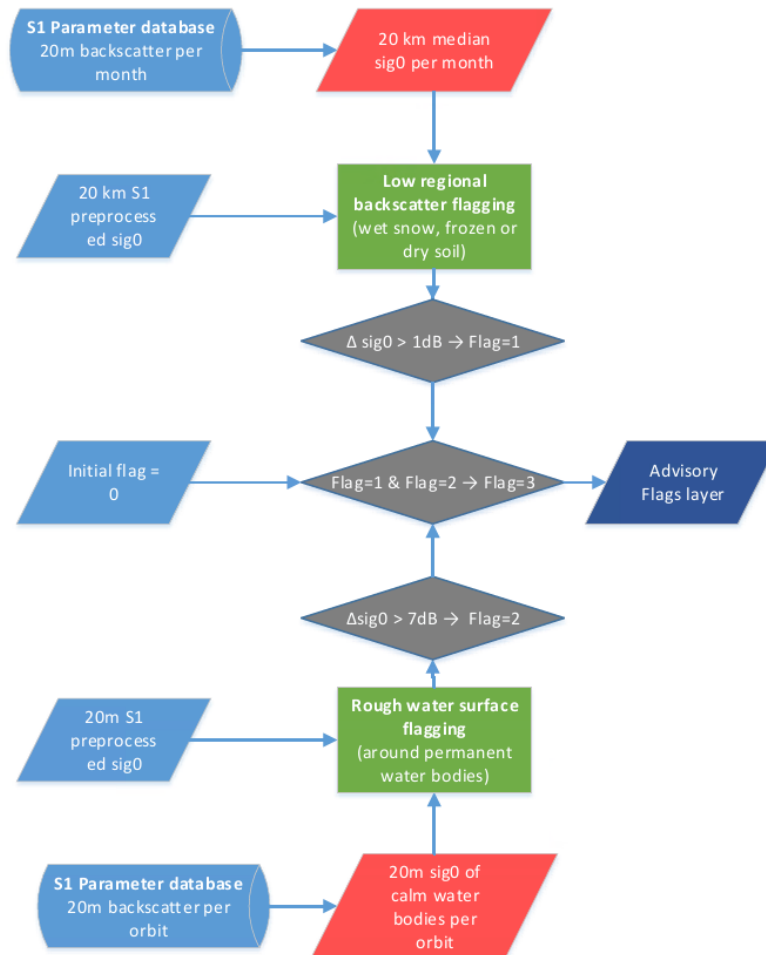
Figure 26: Advisory Flag workflow combining the low regional backscatter flag (Flag 1) and the rough water surface flag (Flag 2). The final Advisory Flags layer includes four values, 0 if no dynamic influence is present, 1 for Flag 1, 2 for Flag 2, and 3 if both flags are true.

4.6.1.1 Low Regional Backscatter Flag

During the snow accumulation period, dry snow and ice-packs are almost transparent to microwaves. As a result, the SAR signal penetrates the snow / ice-pack up to several meters and the main contribution to the backscattering is from the snow–ground interface Rott et al. (1987). During the melting period, however, the increase of the amount of free liquid water inside the snow and ice bodies causes high dielectric losses, thereby increasing the absorption coefficient, featuring very

low backscatter. In addition, the occurrence of meltwater puddles might change the backscattering behaviour of the surface, leading to components with specular microwave reflection, further decreasing the received amplitude at the sensor. Such patches of very low backscatter from those combined effects act easily as water-look-alikes and are source of false alarms.

Similarly, frozen soils with no free liquid water components in the upper soil layers show very low backscatter signatures, appearing from the radar perspective as quasi-dry soils. Dried out soils are a globally more common issue for water and flood body detection, they often show backscatter values at the same low level as calm water bodies. This ambiguity is hard to come by even with exploitation of time series analysis and statistical parametrisation, as the occurrence of dried soils is as erratic and a-priori unpredictable as flood events.

Low backscatter over large regions characterizes all the above-described effects. Thus, to derive an Advisory Flag 1 called "Low Regional Backscatter", we resample the Sentinel-1 acquisitions to 20 km resolution and compare the regional backscatter to the regional monthly grouped median value (see section 3.1.2) which represents seasonal backscatter signature of the corresponding region. Prior to resampling, each acquisition as well as the median image is masked using the reference water mask (see section 4.3), radar shadow mask (see section 4.4.1.3) and build-up area mask (see section 3.2.5). The difference between regional backscatter and regional monthly grouped median backscatter reveals regionally low backscatter values that are likely to be caused by snow, frost or dry conditions while taking the seasonal variations due to e.g., annual changes in vegetation cover into account. On the condition, that this difference exceeds 1 dB, the low backscatter flag is set to value "1":

$$median(sig_{S1}^0) > sig_{S1}^0 + 1dB \rightarrow Flag = 1 \qquad \text{Eq. 4-45}$$

These values are subsequently resampled to the 20m pixel sampling and a simple radial buffer of 14 km radius is applied to the identified pixels with regionally low backscatter.

Known limitation of the used approach is the low number of flagged data in areas, where frozen or dry conditions are typical for full months (e.g., permanently frozen soil during winter months in northern regions).

4.6.1.2 Rough Water Surface Flag

Wind, strong rainfall or frozen water surface over flood surfaces can lead to missed alarms, as they undermine the initial assumption of low backscatter due to specular reflection on smooth water surfaces. As a response to wind stress, short waves on the order of centimetres to decimetres are formed at the water surface causing an enhanced backscatter signal that occurs due to a constructive interference that reinforces the backscatter signal. This effect is called Bragg resonance effect and is dependent on wind speed and direction as well as on the radar wavelength. For C-

Band, the minimum threshold wind speed that can cause this effect and thus enhance the backscatter from water surface is estimated to be 3.3 m/s. Similarly, strong rainfall events roughen the water surface and thus lead to the enhanced backscatter. In case of the frozen lake surface, the enhanced backscatter values are caused by the scattering from water/ice transitions. For the detection of these effects, we make use of the backscatter signature of permanent water bodies inferred from time series within the Sentinel-1 data cube archive.

We examine the NRT 20m Sentinel-1 backscatter data over the areas identified by the reference water mask (see section 4.3) and additional criteria based on backscatter. Concretely, for the pixel to be identified as water and used for the rough water surface flag decision, it needs to be identified as permanent or seasonal water body by the reference water mask of the given month and it has to have backscatter value below −10 dB and reference backscatter value (represented by the 5th percentile of Sentinel-1 sig0 backscatter) below −20 dB. The criteria based on backscatter are needed, as the water bodies shape might deviate from the seasonally defined reference water mask. For this reason, masking of very high backscatter values helps to avoid the situation, that the rough water surface mask would be issued due to the change of the water body shape instead of due to the roughening of its surface.

Over the identified water pixels, we examine, whether the backscatter deviates from a calm water signature. In the case of a roughened water surface, it is most likely to observe increased backscatter that is indicative of strong winds, rainfall, or frozen water surface, which in turn are likely to increase the backscatter from the nearby water surfaces over the flooded areas too.

For the calm water signature, we chose the 5th percentile of Sentinel-1 sig0 backscatter (see section 3.1.2) as baseline. An analysis of backscatter images at windy conditions led to the selection of 7dB difference as a threshold to distinguish rough from non-rough water surfaces, with respect to the local calm water signature. On the condition that this limit is exceeded over water surface, we set, on the 20m pixel scale, the rough water surface flag to value "2":

$$p_{5\%}\left(sig^0_{S1\_water}\right) + 7dB < sig^0_{S1\_water} \rightarrow Flag = 2 \qquad \text{Eq. 4-45}$$

As a final step, all pixels identified as roughened permanent water bodies are spatially filtered by using morphological operations. On the assumption that wind, rainfall or frozen water surface induced patterns are highly correlated within the neighbouring area, a simple L2/radial buffer of 5km radius is applied to the identified rough-water-surface-pixels, effectively enlarging the identified area. All pixels within the buffered area are labelled with the Advisory Flag 2 called "Rough Water Surface". This should also reflect the initial expectation that the majority of floods appear in the vicinity of rivers and permanent or seasonal water bodies, and the wind flag is needed by users in those areas.

It should be noted that this realisation of the wind flag constitutes the initial version here in GFM. We foresee tweaking and optimisation (in respect to the 7dB threshold and the 5km radius) towards the alert accuracy, once more experience within the operational service is gained.

Finally, the 20m pixel array for the Advisory Flag is obtained from the combination of both intermediate flag outputs. All pixel locations that have values in the incoming Sentinel-1 image are initiated with value 0, and are updated with the following operation:

$$0 + ind_{LowBSFlag} * 1 + ind_{WindFlag} * 2 \qquad \text{Eq. 4-46}$$

Effectively, all pixels with no flag set get the value 0, and all pixels where both flags are set get the value 3.

Known limitation of the used approach are flags that are issued due to the land cover change (e.g. water body changed its shape when compared to the reference water mask and the enhanced backscatter caused by this change is misinterpreted as roughened water surface). Generally, the quality of this flag is strongly dependent on the quality and topicality of the reference water mask.

4.7   Output layer "S-1 metadata"

The GFM product output layer "S-1 metadata" consists of all available metadata attributes provided with each Sentinel-1 GRD data product used in the generation of the global flood monitoring product. Metadata of each Sentinel-1 GRD scene is encompassed in the distributed Sentinel "Standard Archive Format for Europe (SAFE)" format included in the "manifest.safe" file. The manifest file is an XML file containing the mandatory product metadata. Attributes contained in the manifest file are classified into four categories: summary, product, platform, and instrument. Platform and instrument related attributes are considered as static for the different Sentinel-1 satellites. A total number of 29 attributes are contained in the manifest file such as information about the absolute orbit number, pass direction, polarisation, sensing start and end date and the product timeliness category. An abstract of the included attributes is given in hereafter as an example.

Table 8: Example of metadata fields available in the corresponding layer.

Acquisition Type: NOMINAL
Cycle number: 201
Footprint:<gml:Polygon srsName="http://www.opengis.net/gml/srs/epsg.xml#4326" xmlns:gml="http://www.opengis.net/gml">
<gml:outerBoundaryIs> <gml:LinearRing> <gml:coordinates>43.920486,20.600445 44.328014,23.858505 42.829479,24.188070 42.421612,21.009506 43.920486,20.600445</gml:coordinates> </gml:LinearRing> </gml:outerBoundaryIs> </gml:Polygon>
Format: SAFE
Ingestion Date: 2020-05-25T20:19:46.329Z
JTS footprint: MULTIPOLYGON (((21.009506 42.421612, 24.18807 42.829479, 23.858505 44.328014, 20.600445 43.920486, 21.009506 42.421612)))
Mission datatake id: 248418
Orbit number (start): 32724
Orbit number (stop): 32724
Pass direction: ASCENDING
Phase identifier: 1
Polarisation: VV VH
Product class: S

Product class description: SAR Standard L1 Product
Product composition: Slice
Product level: L1
Product type: GRD
Relative orbit (start): 102
Relative orbit (stop): 102
Resolution: High
Sensing start: 2020-05-25T16:25:04.241Z
Sensing stop: 2020-05-25T16:25:29.239Z
Slice number: 9
Start relative orbit number: 102
Status: ARCHIVED
Stop relative orbit number: 102
Timeliness Category: Fast-24h

4.8   Output layer "S-1 footprint"

The footprint of a Sentinel-1 GRD scene as provided in each data product is included in the before discussed manifest file. The footprint is represented as human readable Java Topology Suite (JTS) object named "JTS footprint". The JTS footprint is converted into Well-known Text (WKT) and Well-known Binary (WKB) in the process of parsing and ingesting the acquired manifest files into the operated metadata database. WKT and WKB are originally defined by the Open Geospatial Consortium (OGC) to describe simple features. On-the-fly conversion to GeoJson is provided to aggregate the requested footprint features into the corresponding layer request by the user.

4.9   Output layer "S-1 schedule"

Sentinel-1 observations follow a strict acquisition planning often referred to as acquisition segments. Information on the planned future acquisition is provided by ESA in form of Keyhole Markup Language (KML) files. A single file usually covers an acquisition period of about 12 days, with the start and stop time of the future planned acquisitions already given in the file name. KML files are publish by ESA on a regular base, well before activation, with potential last-minute changes due to requests from the Copernicus Emergency Management Service. Information provided by the KML files are organized based on the planned data takes. Parameters listed in following table are included in the KML. KML files are regularly checked, downloaded and ingested into the described metadata database for further analysis. All parameters are exposed in the layer to extract the requested schedule information indicating the next planned Sentinel-1 GRD acquisition for a given location.

Table 9: Information provided with the next planned Sentinel-1 GRD acquisition.

PARAMETER	DESCRIPTION
Satellite ID	Satellite identifier.
DatatakeId	Corresponds to a unique product identifier (Hexadecimal value). This identifier is also reported in the filename of any product generated from this acquisition.

Mode	Instrument acquisition mode (IW, EW, SM)
Swath	Instrument swath (from 1 to 6 for SM, not applicable for IW and EW)
Polarisation	Instrument polarisation for the acquired data take.
ObservationTimeStart	UTC start date and time of the planned data take
ObservationTimeStop	UTC stop date and time of the planned data take
ObservationDuration	Duration of the planned data take in seconds
OrbitAbsolute	Absolute orbit number at the start time of the data take
OrbitRelative	Relative orbit number at the start time of the data take

4.10   Output layer "Affected population"

The information about the affected population is extracted from the Global Human Settlement Layer (GHSL), in particular the GHS-POP dataset. This data contains a raster representation of the population's distribution and density as the number of people living within each grid cell. The information is available at various spatial resolutions and for different epochs. For the GFM service, we will consider the dataset at 250m resolution (highest possible resolution) and for the latest available time-step, 2015.

The dataset was re-projected to the same grid system as the flood map itself, which is the Equi7Grid with a 20m pixel-spacing and a 300km gridding (T3 level). Furthermore, the data was resampled from 250m to 20m resolution.

4.11 Output layer "Affected land cover"

Land cover information is extracted from the Copernicus Land Monitoring Service making use of the Global Land Cover dataset which is available at a global level with a resolution of 100m. The Copernicus Global Land Cover3 includes 23 classes and provides annual updates. The dataset was reprojected to the same grid system as the flood map itself, which is the Equi7Grid with a 20m pixel-spacing and a 300km gridding (T3 level). The data was then resampled from 100m to 20m resolution.

This information provides a first assessment of affected land cover or land use types, e.g. how much agricultural area is affected by the flood within the flood extent area or the area of interest.

5   Product access and dissemination system for the GFM product

The Product and Dissemination API is documented by using Swagger UI and can be accessed at the following web-link: https://api.gfm.geoville.com/v1/

Swagger UI allows anyone - be it your development team or your end consumers - to visualize and interact with the API's resources without having any of the implementation logic in place. It's

automatically generated from your OpenAPI (formerly known as Swagger) Specification, with the visual documentation making it easy for back end implementation and client side consumption.

6   Quality control of the GFM product

The various quality control (QC) methods which are used to measure, on a regular basis, the product and service quality of the Global Flood Monitoring (GFM) product, are summarized in Table 10: Summary of quality control (QC) methods applied to the Global Flood Monitoring (GFM) product.

Table 10: Summary of quality control (QC) methods applied to the Global Flood Monitoring (GFM) product.

QC OBJECTIVE	QC METHOD
Regular assessment of thematic accuracy:	▪ Implementation of automated, continuous QC procedures to ensure that all input and output data undergoes a quality check before releasing the data to the dissemination system. ▪ Implementation of systematic quality control review to check product consistency and accordance with product specifications. ▪ Systematic thematic quality assessment based on worldwide distributed reference samples and representative set of flood cases considering systematic and stratified sampling approaches.
Regular checking and monitoring of product timeliness:	▪ Implementation of automated and continuous monitoring of the data flow throughout the entire service from data acquisition to the product delivery to the user. ▪ Systematic use and monitoring of metadata inherent to data products and the processing platform to check product timeliness after each task in the product generation workflow.
Regular checking and monitoring of service availability:	▪ Establishing an active and continuous health check monitoring system scraping health check metrics over time enabling an instant alerting mechanism based on a strict threshold-based detection of negative service performances. ▪ Anonymous usage statistics are collected in a time-series database to analyse user uptake on a regular interval to deduce measures to enhance the user experience. ▪ Realisation of fully automated test-bots to simulate defined user service interactions to check and monitor the service performance experience in view of the user. ▪ Regular online surveys to collect user feedback to assess and act on changing user requirements.
Detailed overall quality assessment of the GFM product:	▪ Pre-operational product and service quality assessment report. ▪ Annual product and service quality assessment reports.

7   Scientific references

▪ Aach T. and Kaup A., "Bayesian algorithms for adaptive change detection in image sequences using Markov random fields", Image Commun., vol. 7, no. 2, pp. 147–160, 1995. https://doi.org/10.1016/0923-5965%2895%2900003-F

- Aherne F. J., Thacker N. A., and Rockett P. I., "The Bhattacharyya metric as an absolute similarity measure for frequency coded data", Kybernetika, vol. 34, pp. 363–368, Jun. 1998. http://www.kybernetika.cz/content/1998/4/363/paper.pdf

- Ali, I., Cao, S., Naeimi, V., Paulik, C., & Wagner, W. (2018). Methods to remove the border noise from Sentinel-1 synthetic aperture radar data: implications and importance for time-series analysis. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 11(3), 777-786. https://doi-org.proxy.bnl.lu/10.1109/JSTARS.2017.2787650

- Ashman K. M., Bird C. M. and Zepf S. E., "Detecting bimodality in astronomical datasets", Astron. J., vol. 108, pp. 2348–2361, Aug. 1994. https://doi.org/10.1086/117248

- Bauer-Marschallinger, B., Cao, S., Wagner, W., Navacchi, C., Tupas, M., Roth, F., Melzer, T., Freeman, V., in review. Satellite-based flood mapping through Bayesian inference from Sentinel-1 SAR datacube. Remote Sensing of Environment.

- Bazi Y., Bruzzone L. and Melgani F., "Image thresholding based on the EM algorithm and the generalized Gaussian distribution", Pattern Recognit., vol. 40, no. 2, pp. 619–634, 2007. https://doi.org/10.1016/j.patcog.2006.05.006

- Chow, C., Twele, A., Martinis, S. (2016). An assessment of the 'Height Above Nearest Drainage' terrain descriptor for the thematic enhancement of automatic SAR-based flood monitoring services. Proc. of SPIE, 9998, 1-11. https://doi.org/10.1117/12.2240766

- European Commission. 2020. Technical Specifications for Call for tenders JRC/IPR/2020/OP/0551 - Provision of an Automated, Global, Satellite-based Flood Monitoring Product for the Copernicus Emergency Management Service. https://etendering.ted.europa.eu/document/document-file-download.html?docFileId=77006

- Gong M., Li H. and Jiang X., "A multi-objective optimization framework for ill-posed inverse problems", CAAI Trans. Intell. Technol., vol. 1, no. 3, pp. 225–240, 2016 https://doi.org/10.1016/J.TRIT.2016.10.007

- Hansen, M.C.; Potapov, P.V.; Moore, R.; Hancher, M.; Turubanova, S.A.; Tyukavina, A.; Thau, D.; Stehman, S.V.; Goetz, S.J.; Loveland, T.R.; et al. 2013. High-Resolution Global Maps of 21st-Century Forest Cover Change. Science, 342, 850–853. https://doi.org/10.1126/science.1244693

- Kittler, J., & Illingworth, J. (1986). Minimum error thresholding. *Pattern Recognit., 19*, 41-47. https://doi.org/10.1016/0031-3203%2886%2990030-0

- Laferte J. M., Perez P. and Heitz F., "Discrete Markov image modeling and inference on the quadtree", IEEE Trans. Image Process., vol. 9, no. 3, pp. 390–404, Mar. 2000 https://doi.org/10.1109/83.826777

- Lehner, B., Verdin, K., & Jarvis, A. (2008). New global hydrography derived from spaceborne elevation data. Eos, Transactions American Geophysical Union, 89(10), 93-94. https://doi.org/10.1029/2008EO100001

- Marquardt D. W., "An algorithm for least-squares estimation of nonlinear parameters", J. Soc. Ind. Appl. Math., vol. 11, no. 2, pp. 431–441, 1963. https://doi.org/10.1137/0111030

- Matgen, P., S. Martinis, W. Wagner, V. Freeman, P. Zeil, and N. McCormick. 2020. Feasibility assessment of an automated, global, satellite-based flood-monitoring product for the Copernicus Emergency Management Service. JRC Technical Report. EUR 30073 EN. Publications Office of the European Union, Luxembourg. ISBN 978-92-76-10254-0. 47p. https://doi.org/10.2760/653891

- Otsu N., "A threshold selection method from gray-level histograms", IEEE Trans. Syst., Man, Cybern., vol. SMC-9, no. 1, pp. 62–66, Jan. 1979. https://doi.org/10.1109/TSMC.1979.4310076

- Pal S. K. and Rosenfeld A., "Image Enhancement and Thresholding by Optimization of Fuzzy Compactness", Patt. Recog. Lett., vol. 7, pp. 77-86, 1988. https://doi.org/10.1016/0167-8655%2888%2990122-5

- Pekel J.F., Cottam A., Gorelick N., Belward A.S., "High-resolution mapping of global surface water and its long-term changes", Nature, vol. 540, pp. 418-422, 2016. https://doi.org/10.1038/nature20584

- Quegan, S.; Le Toan T.; Yu J.J.; Ribbes F.; Floury N. 2000. Multitemporal ERS SAR Analysis Applied to Forest Mapping. IEEE Transactions on Geoscience and Remote Sensing 38 (2): 741–753. https://doi.org/10.1109/36.842003

- Rennó, C.D., Nobre, A.D., Cuartas, L.A., Soares, J.V., Hodnett, M.G., Tomasella, J., Waterloo, M.J., 2008. HAND, a new terrain descriptor using SRTM-DEM: Mapping terrafirme rainforest environments in Amazonia. Remote Sensing of Environment, 112, 3469-3481. https://doi.org/10.1016/j.rse.2008.03.018

- Rott H., Mätzler C. Possibilities and limits of synthetic aperture radar for snow and glacier surveying, Ann. Glaciol., 9, 195–199, 1987.
  https://doi.org/10.1017/S0260305500000604

- Salamon, P., N. McCormick, C. Reimer, T. Clarke, B. Bauer-Marschallinger, W. Wagner, S. Martinis, C. Chow, C. Böhnke, P. Matgen, M. Chini, R. Hostache , L. Molini, E. Fiori, A. Walli. 2021. The new, systematic global flood monitoring product of the Copernicus Emergency Management Service. International Geoscience and Remote Sensing Symposium (IGARRS) 2021. Paper WE4.O-6.4.
  https://igarss2021.com/view_paper.php?PaperNum=3884

- Sauer, S., Ferro-Famil, L., Reigber, A., Pottier, E., 2011. Three-dimensional imaging and scattering mechanism estimation over urban scenes using dual-baseline polarimetric InSAR observations at L-band. IEEE Transactions on Geoscience and Remote Sensing, 49(11), 4616-4629.
  https://doi.org/10.1109/TGRS.2011.2147321

- Schlaffer, S.; Matgen, P.; Hollaus, M.; Wagner, W. Flood detection from multi-temporal SAR data using harmonic analysis and change detection.International Journal of Applied Earth Observation and Geoinformation 2015,38, 15 – 24.
  https://doi.org/10.1016/j.jag.2014.12.001

- Twele, A., Cao, W., Plank, S., Martinis, S., 2016. Sentinel-1 based flood mapping: a fully automated processing chain. International Journal of Remote Sensing, 37 (13), 2990-3004.
  https://doi.org/10.1080/01431161.2016.1192304

- Ulaby F. and Dobson M. C., Handbook of Radar Scattering Statistics for Terrain. Norwood, MA, USA: Artech House, 1989
  https://www.amazon.com/Handbook-Scattering-Statistics-Terrain-Sensing/dp/0890063362

- Wagner, W., V. Freeman, S. Cao, P. Matgen, M. Chini, P. Salamon, N. McCormick, S. Martinis, B. Bauer-Marschallinger, C. Navacchi, M. Schramm, C. Reimer, and C. Briese. 2020. Data processing architectures for monitoring floods using Sentinel-1. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, V-3-2020, 641–648.
  https://doi.org/10.5194/isprs-annals-V-3-2020-641-2020

8   Annexes:

Table 11: Glossary of mathematical symbols.

SYMBOL	MEANING
Ad	Ashman d coefficient

A_1, a_2	Scale factors set to the maxima of the image distribution functions
A, b	Updatable greyscale values
Bc	Bhattacharyya coefficient
Bm	Bimodal mask
C^2	Tile size (c= 200 pixels)
C_i	Harmonic coefficients of the cosine components
C	Class??
$C1, C2, C3$	Bimodal conditions
FM_{t0}	New floodwater map
FM_{t-1}	Previous floodwater map
G	Histogram bin
$G_1(y), G_2(y)$	Gaussian distribution functions
$G_1^{BM}(y), G_2^{BM}(y)$	Gaussian distributions of the two classes of interest
h_{ellips}	Ellipsoidal height
H	Orthometric height
$h_f(y)$	Image fitted histogram
$h(I[S_j^i])$	Image tile histogram
H(g)	Sar backscatter histogram
I	Sar image
I_D^{S1}	Current and previous sentinel-1 image difference
I_{t0}^{S1}	Current sentinel-1 image
I_{t-1}^{S1}	Previous sentinel-1 image
K	Number of sine and cosine terms in the harmonic model
m_{out}^f	Flood mask of statistical outliers
m_{PLIA}	Mask of exceeding plia
m_{DC}	Mask of conflicting distributions
m_{out}^{nf}	Non-flood mask of statistical outliers
m_{out}	Mask of statistical outliers
m_{cert}	Mask for high uncertainties
N_{geoid}	Geoid undulations
Nw	Non-water
N_{points}	Number of data points
N_{tiles}	Number of overlapping subsets (tiles)
N_{nodes}	Number of tree level
$p(\sigma^0)$	Marginal distribution of backscatter values
$P(W), P(NW)$	Water and non-water prior probabilities
$P(C), P(NC)$	Change and non-change prior probabilities
$p(\sigma^0|NW)$	Conditional pdf of the non-water backscatter
$p(\sigma^0|W)$	Conditional pdf of the water backscatter
$p(\sigma^0|F)$	Conditional pdf of the water's flood backscatter
$p(\sigma^0|NF)$	Conditional pdf of the non-flood backscatter
$p(W|\sigma^0)$	Posterior probability of water
$p(\Delta\sigma^0)$	Marginal distribution of backscatter difference values
$p(C|\Delta\sigma^0)$	Posterior probability of change
$p(\Delta\sigma^0|C), p(\Delta\sigma^0|NC)$	Conditional probabilities of change and non-change classes
$P(F|\sigma^0), P(NF|\sigma^0)$	Flood and non-flood posterior probability
$p(F|\sigma^0, \Delta\sigma^0), p(F|\sigma^0)$	Conditional probability of a given pixel being flooded given its backscatter and backscatter difference values
$P_i(T)$	Summed-up histogram h(g)

$P(error\|\sigma^0)$	Conditional error p
r	Relative orbit
s^-	Child tile (sub-tile)
S^n	Ensemble of nodes in the lowest level tiles
S_{nodes}	Ensemble of tree nodes
S_i	Harmonic coefficients of the sine components
S	Image tile
S^+	Parent tile
S^0	Root node
s	Standard deviation of the harmonic model
S-function	Standard s-function
Sr	Surface ratio
t	Day
t_{day}	Day of the year
W	Water
Y	Image measurement
θ	Incidence angle
$\mu_i(T)$	Geometric centroid of the separated class
M_1, μ_2	Mean distribution values for $G_1(y)$, $G_2(y)$
$\mu_{VH}^{\sigma^0}$	Mean backscatter in vh
$\mu_{VV}^{\sigma^0}$	Mean backscatter in vv
μ_{global}	Mean backscatter value
μ^-	Mean backscatter value of one tile s^-
μ^+	Mean backscatter value of one tile s^+
$\mu_{CPR}^{\sigma^0}$	Mean cross polarization ratio
$\mu_{\sigma_\mu^+}$	Mean of all standard deviations σ_μ^+
μ_{water}	Mean of the water class
σ^0	Radar (sar) backscatter value (sigma nought)
$\overline{\sigma^0}$	Average sar backscatter for the time period
σ_r^0	Backscatter time-series of a specific pixel and relative orbit r
σ_1, σ_2	Distribution standard deviation values
$\sigma_{t,r}^0$	Backscatter from a specific pixel, day of year t and relative orbit r
$\hat{\sigma}_{t,r}^0$	Estimated backscatter from the harmonic model for any given day of year t and relative orbit r
$\hat{\sigma}_{t_{day}}^0$	Most probable radar backscatter
$\sigma_{VH}^{\sigma^0}$	Standard deviation in vh polarization
$\sigma_{\sigma_\mu^+}$	Standard deviation of all standard deviations σ_μ^+
Σ_μ^+	Standard deviation of the child tiles mean backscatter (μ^-)
$\widehat{\sigma_{t,r}^0}$	The expected backscatter of the harmonic model
$\sigma_{W,\theta}^0$	Various backscatter observations along the incidence angle
τ	Final threshold
τ_i	Locally variable thresholds
τ_{Otsu}	Otsu threshold
T^+	Thresholds of the parent tiles

List of Figures.