



Global Flood Monitoring

Framework Service Contract No. 939866

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: 4.0

Issue Date: 14/10/2025

Issue By: Patrick Matgen (LIST)

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	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Document Release Sheet

Book captain:	Patrick Matgen		Date:
Approval:	Richard Kidd	Signature	Date:
Endorsement:	Peter Salamon (CEMS Manager)	Signature	Date:
Distribution:	Public (PU)		

Contributing Authors:

EODC GmbH:	▪ Christoph Reimer (CR), Tobias Stachl (TS), Dragana Milinkovic (DM)
GeoVille:	▪ Michaela Seewald (MS), Laurin Müller (LMü),
TU Wien:	▪ Bernhard Bauer-Marschallinger (BBM), Florian Roth (FRo), Felix Reuss (FRe), Mark Edwin Tupas (MT)
CIMA:	▪ Luca Molini (LM), Elisabetta Fiori (EF)
LIST:	▪ Patrick Matgen (PM), Marco Chini (MC), Renaud Hostache (RH), Ramona Pelich (RP), Yu Li (YL), Jefferson Wong (JW)
DLR:	▪ Marc Wieland (MW), Franziska Kraft (FK), Sandro Martinis (SM), Candace Chow (CC), Christian Krullikowski (CK)
JRC:	▪ Niall McCormick (NM), Peter Salamon (PS)

Change Log:

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	PM
0.02	15.07.2021	Inserted product and algorithm descriptions	All
0.04	01.09.2021	Reviewed the document, identified missing parts	RP, PM
1.2	26.10.2021	First major document revision	NM, PS
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
2.0	26.04.2024	Update of algorithm descriptions based on completed service evolutions	PM, MC, CK, SM, FK, FRo, BBM, NM
3.0	14.07.2025	Second major document revision, re-structuring, and updates based on completed service evolutions	PM, MC, FK, SM, FRo, BBM, NM
4.0	14.10.2025	Update of algorithm descriptions based on completed service evolutions and release of v4.0	PM, FRo, FK, NM, JW

For any queries, please contact Richard Kidd, GFM Service Manager (GFM_Service_Manager@eodc.eu).

Table of Contents

1	Introduction	1
2	Pre-processing and preparation of input Sentinel-1 and ancillary data.....	3
2.1	Sentinel-1 SAR image data pre-processing	3
2.1.1	Georeferenced, calibrated normalized radar backscatter coefficient (σ^0) data	3
2.1.2	Sentinel-1 temporal parameters.....	5
2.1.3	Sentinel-1 harmonic parameters	6
2.1.4	Sentinel-1 projected local incidence angle data	7
2.2	Ancillary data preparation	7
2.2.1	Satellite orbit state vectors	8
2.2.2	Copernicus DEM	8
2.2.3	Sentinel-1 Global Backscatter Model.....	9
2.2.4	Height above Nearest Drainage (HAND) information.....	9
2.2.5	Land cover data.....	10
3	GFM individual and Ensemble flood mapping algorithms.....	11
3.1	GFM flood mapping algorithm 1 (LIST)	12
3.1.1	Change detection by computing difference between the reference and new images .	12
3.1.2	Parameterizing the target and background distribution functions using HSBA	13
3.1.3	Mapping target class by HSBA-based histogram thresholding and region growing	18
3.1.4	Final post-processing of the LIST flood map	21
3.1.5	Computation of Likelihood Values by the LIST flood mapping algorithm	23
3.2	GFM flood mapping algorithm 2 (DLR)	24
3.2.1	Unsupervised initialization of water classification by tile-based thresholding	24
3.2.2	Fuzzy logic-based refinement of the initial water classification.....	30
3.2.3	Region growing of the defuzzified water classification	32
3.2.4	Final post-processing of the DLR flood map	33
3.2.5	Computation of Likelihood Values by the DLR flood mapping algorithm	34
3.3	GFM flood mapping algorithm 3 (TUW)	35
3.3.1	Computation of Sentinel-1 projected local incidence angle data	36
3.3.2	Computation of Sentinel-1 harmonic parameters.....	36
3.3.3	Estimation of backscatter distribution functions for water and land surfaces	39

3.3.4	Bayesian flood mapping and uncertainty estimation	41
3.3.5	Low sensitivity masking.....	42
3.3.6	Final post-processing of the TUW flood map	46
3.3.7	Computation of Likelihood Values by the TUW flood mapping algorithm.....	46
3.4	The GFM Ensemble flood mapping algorithm	47
4	Generation of the ten GFM output layers of flood information	53
4.1	GFM output layer: Observed Flood Extent	53
4.2	GFM output layer: Observed water extent.....	53
4.3	GFM output layer: Reference Water Mask.....	53
4.4	GFM output layer: Exclusion Mask	56
4.4.1	GFM Exclusion Mask sub-layer: No-sensitivity areas	57
4.4.2	GFM Exclusion Mask sub-layer: Water look-alikes.....	58
4.4.3	GFM Exclusion Mask sub-layer: Topographic distortions.....	59
4.4.4	GFM Exclusion Mask sub-layer: Radar shadow	60
4.4.5	GFM Exclusion Mask sub-layer: Low Sentinel-1 coverage.....	62
4.4.6	Global coverage of the input layers of the GFM Exclusion Mask	63
4.5	GFM output layer: Likelihood Values.....	64
4.6	GFM output layer: Advisory Flags	65
4.6.1	GFM Advisory Flags: Low regional backscatter	66
4.6.2	GFM Advisory Flags: Rough water surface	68
4.7	GFM output layer: Sentinel-1 Footprint and Metadata.....	69
4.8	GFM output layer: Sentinel-1 Schedule	70
4.9	GFM output layer: Affected Population.....	71
4.10	GFM output layer: Affected Land Cover	71
5	GFM product and service quality assessment	72
	Scientific references.....	74
	List of abbreviations and terminology	77

1 Introduction

The Global Flood Monitoring (GFM) product of the Copernicus Emergency Management System (CEMS) provides a continuous monitoring of flood events worldwide, by processing and analysing in near real-time all incoming Sentinel-1 Synthetic Aperture Radar (SAR) satellite imagery. The GFM product utilizes a data cube (or time-series) approach, enabling high product timeliness. Central to the GFM product is an ensemble flood mapping approach that integrates three independent, state-of-the-art SAR-based flood mapping algorithms, in order to optimize the robustness and accuracy of the flood and water extent maps, and to build a high degree of redundancy into the service (Salamon, et al., 2021), (Matgen, et al., 2020), (Wagner, et al., 2020). The ten output layers of flood-related information that are generated by the GFM product are described in Table 1.

Table 1: GFM output layers of global flood-related information, generated in near real-time based on Sentinel-1 SAR satellite imagery.

#	OUTPUT LAYER	DESCRIPTION	DATA FORMAT
1	Observed Flood Extent	Flooded areas mapped by applying the GFM ensemble flood mapping algorithm to the latest Sentinel-1 images of SAR backscatter intensity.	Raster (Geotiff); Vector (Shapefile)
2	Observed Water Extent	Open and calm water mapped as the union of the Observed Flood Extent and the Reference Water Mask.	Raster (Geotiff); Vector (Shapefile)
3	Reference Water Mask	Normal (i.e. permanent and seasonal) water mapped by applying the GFM ensemble water mapping algorithm to the most recent historical, five-year time-series ("data cube") of Sentinel-1 images of SAR backscatter intensity.	Raster (Geotiff); Vector (Shapefile)
4	Exclusion Mask	Areas where SAR-based water mapping is not technically feasible, due to no sensitivity (e.g. urban areas, dense vegetation), low backscatter (e.g. flat impervious areas, sandy surfaces), topographic distortions, radar shadows, or low coverage of Sentinel-1.	Raster (Geotiff)
5	Likelihood Values	Likelihood (0-100%) of correct flood classification for all pixels 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;	Raster (Geotiff)
7	Sentinel-1 Footprint and Metadata	Image boundaries of the Sentinel-1 data used, and in addition information on the "metadata", i.e. the acquisition parameters of the Sentinel-1 data used.	KML
8	Sentinel-1 Schedule	Next scheduled Sentinel-1 data acquisition.	KML
9	Affected Population	Number of people in flooded areas, mapped by a spatial overlay of Observed Flood Extent and gridded population, from the Copernicus GHSL project.	Raster (Geotiff)
10	Affected Land Cover	Land cover / use (e.g. artificial surfaces, agricultural areas) in flooded areas, mapped by a spatial overlay of Observed Flood Extent and the Copernicus GLS land cover.	Raster (Geotiff)

The GFM product and output layers of flood information are freely available through various dissemination channels, listed in Table 2, that are designed for flexible and easy access. All of the dissemination channels are routed to a central storage via a central access layer. The storage-access layer serves products in a desired format, by drilling down into the GFM Sentinel-1 SAR “data cubes” (time series), querying databases for meta-data or user-management, and receiving data for the WMS-T service. Further details on accessing the GFM product are available via the GFM wiki pages¹.

Table 2: Dissemination channels for accessing the CEMS Global Flood Monitoring (GFM) product.

GFM DISSEMINATION CHANNEL	URL
▪ Global Flood Awareness Sytem (GloFAS) Map Viewer	https://global-flood.emergency.copernicus.eu/
▪ European Flood Awareness Sytem (EFAS) Map Viewer	https://european-flood.emergency.copernicus.eu/en
▪ Full set of GFM-related Representational State Transfer (REST) APIs	https://api.gfm.eodc.eu/v2/
▪ GFM Web Map Service supporting temporal requests (WMS-T)	https://geoserver.gfm.eodc.eu/geoserver/gfm/wms
▪ GFM dedicated web portal	https://portal.gfm.eodc.eu/
▪ GFM SpatioTemporal Asset Catalogue (STAC)	https://services.eodc.eu/browser/#/v1/collections/GFM

This Product Definition Document (PDD) provides the main reference information required to understand all elements of the various data processing chains that are implemented in the GFM product, and explains the contents of the GFM output layers, the main assumptions underlying their generation, and the limitations of the data. Specifically, the PDD describes in detail the following main methods and input and output datasets that form the core of the GFM product:

- The pre-processing of the input Sentinel-1 SAR image datasets, and the preparation of the auxiliary input datasets, include the satellite orbit state vectors, the Copernicus DEM, the Sentinel-1 Global Backscatter Model, the Height above Nearest Drainage (HAND) information, and the land cover data.
- The three individual state-of-the-art SAR-based flood mapping algorithms (LIST, DLR and TUW), and the Ensemble algorithm used to combine the results of the three individual algorithms.
- The ten output layers of flood information generated by the GFM product, including Observed Flood Extent, Reference Water Mask, Likelihood Values, Exclusion Mask, and Advisory Flags.
- The procedures for the regular GFM product and service quality assessment.

¹ <https://extwiki.eodc.eu/GFM/PUM/DataAccess>

2 Pre-processing and preparation of input Sentinel-1 and ancillary data

In this Section, the methods that are used in the GFM product for the pre-processing and preparation of the main input Sentinel-1 SAR image datasets and auxiliary input datasets, which are listed in Table 3, are described in detail.

Table 3: The main input Sentinel-1 SAR and ancillary datasets used by the GFM product.

SENTINEL-1 SAR IMAGE DATASETS	ANCILLARY DATASETS
▪ <i>Georeferenced, calibrated Sentinel-1 SAR normalized radar backscatter coefficient (σ^0) data.</i>	▪ <i>Satellite orbit state vectors.</i>
▪ <i>Sentinel-1 SAR temporal parameters.</i>	▪ <i>Copernicus DEM.</i>
▪ <i>Sentinel-1 SAR harmonic parameters.</i>	▪ <i>Sentinel-1 Global Backscatter Model.</i>
▪ <i>Sentinel-1 SAR projected local incidence angle data.</i>	▪ <i>Height above Nearest Drainage (HAND) index.</i>
	▪ <i>Land cover data.</i>

2.1 Sentinel-1 SAR image data pre-processing

Sentinel-1 image data form the core dataset of the GFM product. The data are used not only as the primary information source on the extent of flooded areas, and of permanent (and seasonal) water bodies, but also for offline generation of parameters describing globally the C-band SAR backscatter signature at 20-metre pixel resolution, supporting the generation of the ten GFM output layers.

2.1.1 Georeferenced, calibrated normalized radar backscatter coefficient (σ^0) data

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 Sentinel-1 pre-processing workflow is triggered (see Figure 1). 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) commands² subsequently called from a Python package developed by TU Wien. This routine involves the steps described in Table 4 (where steps 2, 7, and 8 are done independently from SNAP). 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).

Table 4: Description of the steps in the Sentinel-1 pre-processing workflow, carried out using the SNAP toolbox via Graph Processing Tool (GPT) commands.

#	TASK	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 as σ^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.
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

² <https://senbox.atlassian.net/wiki/spaces/SNAP/pages/70503475/Bulk+Processing+with+GPT>

		(dB), scaled by a factor of 10, converted to Int16 data type, and written to disk as a Zstandard (ZSTD) compressed GeoTIFF file.
8	Resampling:	- 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. - 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 or σ^0) data as output.

2.1.2 Sentinel-1 temporal parameters

Based on the Sentinel-1 data, several temporal parameters are computed for a multiyear period 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 5.

Table 5: Overview of Sentinel-1 temporal parameters, and availability per orbit and for all orbits.

PARAMETER	DESCRIPTION	PER ORBIT	OVER ALL ORBITS
<i>Number of Observations (NOBS):</i>	Number of observations per pixels. This parameter is available per orbit and for all orbits combined.	X	X
<i>Low Backscatter Frequency (LTM15-FREQ):</i>	Frequency of observations below -15dB. This parameter is calculated per orbit and serves as input for low backscatter masking.	X	
<i>Total Aggregated Mean (TAG-MEAN):</i>	The TAG-MEAN is calculated per orbit and as a weighted average over all orbits. The latter is used as input for the DLR and LIST algorithms to extract the permanent water mask. (In a later version this will be replaced by the median over the entire time-series).	X	X
<i>5th Percentile (PERC-05):</i>	The 5th percentile is calculated per orbit and used as an input for the advisory flags.	X	
<i>Grouped Monthly Aggregated Median (GRMAG-MEDIAN):</i>	The GRMAG-MEDIAN is calculated for all months over two years, over all orbits. It is used as an input for the DLR and LIST algorithms to extract the seasonal water mask.		X

2.1.3 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.

The TUW flood mapping algorithm of the GFM product utilizes the harmonic model to simulate the backscatter seasonal variation and estimate normal, non-flooded conditions (see **Section 3.3** below for details). Figure 2 shows a modeled harmonic time series for an agricultural plot in Thessaly, Greece, superimposed on the VV backscatter time series. An example for an estimated backscatter image for one day in summer and one day in winter over Greece is given in Figure 3. In this example, one can notice the backscatter differences caused by varying seasonal vegetation states. This section describes 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 estimated backscatter of the harmonic model in summer and winter over Greece

2.1.4 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 Sentinel-1 backscatter (σ^0) image and the parameters of the harmonic model as an input of the TUW GFM flood mapping algorithm (see **Section 3.3** below for details).

2.2 Ancillary data preparation

In order to enable the calculations for the flood mapping and its by-products, the GFM product 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 GFM **Exclusion Mask** and **Advisory Flags** output layers.

Like the Sentinel-1 backscatter observations and the 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.

2.2.1 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 the SNAP toolbox from within the pre-processing pipeline.

2.2.2 Copernicus DEM

The Copernicus DEM³ 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 the SNAP toolbox 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 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:

$$h_{ellips} = H + N_{geoid} \quad \text{Eq. 2-1}$$

As the SNAP toolbox only works with a single external DEM file, both datasets 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 6.

³ <https://spacedata.copernicus.eu/web/cscda/dataset-details?articleId=394198#C4>

Table 6: 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
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 <code>gdal_warp</code> , <code>gdal_merge.py</code> or <code>gdal_translate</code> , 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 ellipsoidal heights.

2.2.3 Sentinel-1 Global Backscatter Model

The Sentinel-1 Global Backscatter Model (S1-GBM) was generated by the Remote Sensing Group of the TU Wien, 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 backscatter in VV- and VH-polarization at a 10 m sampling, giving a high-quality impression on surface structures and patterns.

The GFM product uses the temporal mean VV and VH backscatter (σ^0) 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.

2.2.4 Height above Nearest Drainage (HAND) information

The Height above Nearest Drainage (HAND) index, described by [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 globally based on elevation and drainage direction data provided by the Hydrosheds mapping product of [Lehner et al. \(2008\)](#) and its global extension by [Donchyts et al. \(2016\)](#).

The Hydrosheds product was derived based on the DEM of the Shuttle Radar Topography Mission (SRTM), which was hydrologically conditioned using a sequence of automated procedures (e.g. void-filling, filtering, stream burning, and upscaling) as well as manual corrections (Lehner et al., 2008).

For the GFM product, a HAND-derived Exclusion Mask (HAND-EM) is generated, which is used to reduce water-lookalike classes, depending on the hydrologic–topographic setting, and to separate flood-prone from not flood-prone areas.

2.2.5 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., \(2013\)](https://storage.googleapis.com/earthenginepartners-hansen/GFC-2022-v1.10/download.html); version 1.10⁴) characterises global forest extent and change from 2000 through 2022 and is based on the time-series analysis of Landsat images in characterizing global forest extent and change from 2000 through 2022. 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 <https://storage.googleapis.com/earthenginepartners-hansen/GFC-2022-v1.10/download.html>

3 GFM individual and Ensemble flood mapping algorithms

The CEMS Global Flood Monitoring (GFM) product provides a continuous, systematic monitoring of all major global flood events in near real-time, based on the latest Sentinel-1 SAR images, generating the ten output layers of flood information described in Table 1. The GFM product also provides a Sentinel-1 SAR-based processed archive of worldwide observed floods and water bodies, from 1 January 2015 until 2023. Central to the GFM product are three individual, state-of-the-art SAR-based flood mapping algorithms (i.e. LIST, DLR, and TUW), which are applied independently to the input Sentinel-1 SAR and ancillary datasets described in **Section 2** above. The final flood maps (and likelihood values) that are provided by the GFM product are generated by combining the results of the three individual GFM flood mapping algorithms, using an Ensemble flood mapping algorithm.

The LIST, DLR and TUW flood mapping algorithms were originally described by Chini et al. (2017), Martinis et al. (2015), and Bauer-Marschallinger et al. (2022), respectively. The technical characteristics and main steps of the three algorithms, as implemented in the GFM product, are summarized in Table 7. Detailed descriptions of the LIST, DLR, TUW, and Ensemble algorithms, and how they are implemented in the GFM product, are provided in the remainder of this Section.

Table 7: Summary of the technical characteristics and main steps of the LIST, DLR, and TUW flood mapping algorithms, as implemented in the GFM product.

STEP	LIST ALGORITHM	DLR ALGORITHM	TUW ALGORITHM
1.	Change detection by computing the difference between a (recent) reference and a new Sentinel-1 SAR image.	Unsupervised initialization of water classification by automatic tile-based thresholding of a Sentinel-1 SAR image.	- Sentinel-1 SAR image datacube: - Computation of Sentinel-1 projected local incidence angles. - Computation of Sentinel-1 harmonic parameters.
2.	HSBA-based parameterizing of the target (Water, Flood) and background (Non-water, Non-flood) distribution functions.	Fuzzy logic-based refinement of the initial water classification, based on backscatter (σ^0), terrain slope, and size of water bodies.	Estimation of backscatter distribution functions for water and land surfaces.
3.	Mapping target class by histogram thresholding and region growing.	Region growing of the defuzzified water classification.	- Bayesian flood mapping and uncertainty estimation. - Low sensitivity masking: - Exceeding incidence angles. - Conflicting distributions. - Outliers. - High uncertainties.
4.	Morphological post-processing.	Morphological post-processing.	Morphological post-processing.

3.1 GFM flood mapping algorithm 1 (LIST)

The LIST flood mapping algorithm, which was originally described by Chini et al. (2017), is a change detection approach that requires the following three main input datasets:

- The most recent Sentinel-1 SAR image scene to be processed (I_{t0}), i.e. **new image**.
- The previously recorded overlapping Sentinel-1 SAR image scene (I_{t0-1}), i.e. **reference image**.
- The corresponding previous GFM output layer **Observed Flood Extent**, generated by the GFM Ensemble flood mapping algorithm.

The LIST flood mapping algorithm consists of four main steps, which are described below:

- Change detection by computing the difference between the reference and new images.
- Parameterizing the target and background distribution functions using a hierarchical split-based approach (HSBA).
- Mapping the target class by HSBA-based histogram thresholding and region growing.
- Final post-processing of the flood map produced by the LIST algorithm.

3.1.1 *Change detection by computing difference between the reference and new images*

In order to compute the changes between the reference image (I_{t0-1}) and the new image (I_{t0}), a **difference image** (I_D) is computed as the difference between the reference image and the new image, i.e. $I_D = I_{t0-1} - I_{t0}$, with flooded pixels in the difference image expected to be positive because of the reduction of backscatter values caused by flooding.

The LIST flood mapping algorithm aims at detecting and mapping all increases and decreases of floodwater extent with respect to the reference image. As it is a change detection approach, this enables the differentiation of floodwater from permanent water bodies, as well as the filtering out of classes having water-like backscatter 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 to the new image is used. Sentinel-1 is well suited for this latter requirement, due to its repeat cycle of 6 days. Finally, in order to reduce false alarms further, and to speed-up the analysis, the approach uses as optional input data the Exclusion Mask and the HAND (Height above Nearest Drainage) terrain map.

Each time a new image scene is ingested and pre-processed, the algorithm processes the pair of scenes consisting of the new image (I_{t0}) and previous overlapping image from the same orbit (I_{t0-1}). A new flood map (FM_{t0}) is generated, updating the most recent ensemble flood map (FM_{t0-1}).

As in all statistical change detection or water mapping algorithms applied to SAR imagery, the parameterization of the distributions of the change (i.e. flood) and water classes depends on how easily identifiable the respective classes are in the histogram of backscatter values in the SAR imagery. More explicitly, classes such as flooded or changed areas typically represent only a small percentage of the total image, and so may not be easily identifiable on the histogram.

In order to address such a limitation, the LIST flood mapping algorithm uses a **hierarchical split-based approach (HSBA)** to locate specific regions (or tiles) of the difference image ($I_D = I_{t0-1} - I_{t0}$), and the corresponding new image (I_{t0}), where:

- The tile histogram is clearly **bimodal**, with two distributions representing the respective **target** (i.e. flood or water) and **background** (i.e. non-flood or non-water) classes.
- The two distributions in the tile histogram are **Gaussian**, and present with a **similar frequency**.

The logarithmic transformation that is applied to the backscatter values of Sentinel-1 SAR Ground Range Detected (GRD) images has the effect of converting the multiplicative noise (speckle) of SAR imagery to additive noise (which can be more easily removed), and increasing the dynamic range of low-intensity pixel values. Furthermore, the pixel values in a log-transformed image can be assumed to have a Gaussian distribution.

The “log-ratio” is a standard technique for change detection of SAR images, based on a pixel-by-pixel comparison of two images acquired at different times. The advantage of the log-ratio operator is that it takes account of the multiplicative model of speckle, and is also less affected by radiometric errors, so that the statistical distribution of the final image depends only on the changes between the two images. The log-ratio image is computed as the difference between the log-transformed reference and new images (i.e. $I_D = I_{t0-1} - I_{t0}$), based on the quotient rule for logarithms, namely:

$$\log(x / y) = \log(x) - \log(y)$$

Briefly, the distribution functions for the target (i.e. change and water) and background (i.e. non-change and non-water) classes in the difference and new images (I_D and I_{t0}) are parameterized, using HSBA. The target and background distributions are then used to map the target classes in I_D and I_{t0} , by histogram thresholding and region growing. In the region growing, the **posterior probabilities** of the target classes are used for the purposes of comparison with a **threshold value**, for selecting seed pixels for region growing, and a **tolerance criterion**, for stopping the region growing.

3.1.2 Parameterizing the target and background distribution functions using HSBA

Considering a SAR image, with backscatter measurements y , it is assumed that:

- Two classes are present, where $G_1(y)$ and $G_2(y)$ are their distribution functions.
- Both distributions can be approximated by Gaussian curves.
- The **prior probabilities** of the two classes are strongly imbalanced.

This last assumption is typical for SAR images covering large areas, where changes only affect a small part of the image. When this happens, the smaller 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). In order to cope with the problem of imbalanced populations, regions (or tiles) of the SAR image where the two classes are more balanced are identified using a **hierarchical split-based approach (HSBA)**, which consists of the following two steps, which are described below:

- Hierarchical tiling using quadtree decomposition.
- Selection of tiles that are statistically representative.

1. Hierarchical tiling using quadtree decomposition:

The first step in the HSBA is the hierarchical tiling of the SAR image using **quadtree decomposition**, which iteratively splits image regions into four quadrants, sub-quadrants, etc. Note that a quadtree is a hierarchical data structure that is used in image processing to partition recursively a two-dimensional image into four equally-sized quadrants or regions. Quadtrees are useful because of their ability to focus on the interesting subsets of the image.

Figure 4 illustrates the hierarchical tiling of a SAR image using a three-level quadtree decomposition. At each level of the quadtree, each quadrant or sub-quadrant corresponds with a specific **node**, representing a specific region (or **tile**) of the image, and has exactly four **children**, or no children at all. The latter is referred to as a **leaf node**. The first level of the quadtree (called the **root node**) represents the entire image region, while the last level of the quadtree (containing the leaf nodes) represents the smallest tiles (Laferte et al., 2000). For a quadtree with **N** levels, the **set of nodes** at the i^{th} level of the quadtree (i.e. S^i) contains 4^i nodes, where $i \in [0, N-1]$. For example, the set of nodes at the last level (i.e. S^2) of the quadtree shown in Figure 4 contains $4^2 = 16$ nodes.




Figure 4: Illustration of a three-level quadtree decomposition of a SAR image, where the set of nodes in the entire quadtree (S) is obtained from the union of the sets of nodes at the first, second, and last levels of the quadtree (S^0 , S^1 , and S^2 , respectively).

To summarize, for the hierarchical tiling of a SAR image (**I**) using an N-level quadtree decomposition:

- | |
|--|
| ▪ S is the set of nodes in the entire quadtree. |
| ▪ S^i is the set of nodes at the i^{th} level, where $i \in [0, N-1]$. |
| ▪ A given node at the i^{th} level is denoted as $S_j^i \in S$, where $j \in [1, 4^i]$. |
| ▪ The tile (i.e. region) of image I related to $S_j^i \in S$ is split into four tiles in the next level ($i + 1$), and is itself one of the four quadrants of a tile in the previous level ($i - 1$). |
| ▪ The tile of image I related to $S_j^i \in S$ is denoted as $I[S_j^i]$. |
| ▪ The lowest level of the quadtree decomposition is fixed based on the minimum tile size that still guarantees statistical representativeness. |

Thus, the resulting quadtree decomposition can be expressed as:

$$\mathbf{S} = \left\{ \mathbf{S}^0, \mathbf{S}^1, \dots, \mathbf{S}^i \mid \bigcup \mathbf{I}[\mathbf{S}_j^i] = \mathbf{I} = \mathbf{I}[\mathbf{S}^0], \text{ where } i \in [0, N-1] \text{ and } j \in [1, 4^i] \right\},$$

2. Selection of tiles that are statistically representative:

For a SAR image ($\mathbf{I}$), the histogram of any tile in the N -level quadtree (consisting of the set of nodes $\mathbf{S}_j^i$, where $i \in [0, N-1]$ and $j \in [1, 4^i]$, as described above) is expressed as $\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i])$.

An image tile $\mathbf{I}[\mathbf{S}_j^i]$ is selected for estimating the parameters of the Gaussian distributions of the target and background classes (to be used for the histogram thresholding and region growing of the SAR image), according to the following three conditions:

Condition 1:	▪ The tile histogram, $\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i])$, is clearly bimodal .
Condition 2:	▪ Both distributions represented in $\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i])$ are Gaussian .
Condition 3:	▪ Each distribution in $\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i])$ is present at a frequency of at least 10% of the other.

Recall that in mathematics, a Gaussian function of a normally distributed random variable $\mathbf{x}$, with mean μ and standard deviation σ , is of the form:

$$f(x) = \frac{1}{\sigma \sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)}$$

where $\frac{1}{\sigma \sqrt{2\pi}}$ is the function's maximum value (amplitude or height), which occurs at $x = \mu$.

For a SAR image ($\mathbf{I}$), we hypothesize that the histogram of each tile, $\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i])$, is a sum of two Gaussian distributions ($\mathbf{G}_1$ and $\mathbf{G}_2$), i.e.:

$$\mathbf{h}(\mathbf{I}[\mathbf{S}_j^i]) = \mathbf{h}(\mathbf{y}) \approx \mathbf{h}_f(\mathbf{y}) = \mathbf{G}_1(\mathbf{y}) + \mathbf{G}_2(\mathbf{y}) = \mathbf{A}_1 e^{\frac{-(\mathbf{y}-\mu_1)^2}{2\sigma_1^2}} + \mathbf{A}_2 e^{\frac{-(\mathbf{y}-\mu_2)^2}{2\sigma_2^2}}$$

where the abbreviations are defined as follows:

$\mathbf{y}$	The measurement (backscatter)
$\mathbf{h}(\mathbf{y})$	The tile histogram of image values.
$\mathbf{h}_f(\mathbf{y})$	The distribution fitted to the tile histogram.
$\mathbf{A}_1, \mu_1, \sigma_1$	Amplitude (height), mean, standard deviation of the target distribution.
$\mathbf{A}_2, \mu_2, \sigma_2$	Amplitude (height), mean, standard deviation of the background distribution.

The parameters of the two Gaussian distributions that together compose the histogram of each tile in the SAR image, are estimated using the **Levenberg-Marquardt algorithm**.

The Levenberg-Marquardt algorithm is a standard technique for estimating the parameters of the distribution functions that best fit a set of empirical data. It solves the non-linear least squares problems by combining the steepest descent and inverse-Hessian function fitting methods ([Marquardt et al., 1963](#)). Non-linear least squares methods involve an iterative optimization of parameter values to reduce the sum of squares errors between a fitting function and histogram values. Iterations are done until three consecutive repetitions fail to change the chi-squared value by more than a specified tolerance, or until a maximum number of iterations has been reached.

When the Levenberg-Marquardt algorithm is applied, the initial guess of the parameters should be as close as possible to the actual values, otherwise the solution may not converge. The **first-guess values** are retrieved using Otsu's method ([Otsu et al., 1979](#)), an automatic image thresholding algorithm which in its simplest form returns a single intensity threshold that separate pixels into two classes, foreground and background.

Given the Otsu-derived threshold y_{ot} for the tile histogram $h(y)$, the first-guess values of the amplitudes (A^0_1, A^0_2), means (μ^0_1, μ^0_2), and standard deviations (σ^0_1, σ^0_2) of the two distributions $G_1(y)$ and $G_2(y)$, can be derived as follows:

<i>First-guess parameters for $G_1(y)$:</i>	<i>First-guess parameters for $G_2(y)$:</i>
▪ $A^0_1 = h(y = \mu^0_1)$	▪ $A^0_2 = h(y = \mu^0_2)$
▪ $\mu^0_1 = \text{mean}(h(y < y_{ot}))$	▪ $\mu^0_2 = \text{mean}(h(y > y_{ot}))$
▪ $\sigma^0_1 = \text{standard deviation}(h(y < y_{ot}))$	▪ $\sigma^0_2 = \text{standard deviation}(h(y > y_{ot}))$

Once the distributions are fitted to all tile histograms in the SAR image, the statistical properties of each tile histogram are evaluated according to three conditions: only tiles with a **bimodal** histogram (Condition 1), and with two distributions that are **Gaussian** (Condition 2) and **balanced** (Condition 3), are selected for parameterizing the target and background distributions. The statistical properties of the tile histogram, and the indices used to evaluate them, are summarized in Table 8.

Table 8: The specific statistical properties, and corresponding statistical indices, used to select SAR image tiles that are suitable for parameterizing the target and background distributions.

STATISTICAL PROPERTY OF TILE HISTOGRAM	STATISTICAL INDEX USED	DEFINITION OF INDEX (See text below for abbreviations)
▪ Degree of separation of the two distributions (to evaluate bimodality).	Ashman's D coefficient (AD)	$AD(h_f(y)) = \sqrt{2} \frac{ \mu_1 - \mu_2 }{\sqrt{(\sigma_1^2 + \sigma_2^2)}}$
▪ Similarity of the fitted distribution and tile histogram (to evaluate Gaussianity).	Bhattacharyya Coefficient (BC)	$BC(h(y), h_f(y)) = \sum_k \sqrt{h(y_k)} \sqrt{h_f(y_k)}$
▪ Relative frequency of the two distributions (to evaluate class balance).	Surface Ratio (SR)	$SR(h_f(y)) = \frac{\min(A_1 \sigma_1, A_2 \sigma_2)}{\max(A_1 \sigma_1, A_2 \sigma_2)}$

where the abbreviations are defined as follows:

y	The measurement (backscatter)
$h(y)$	The tile histogram of image values.
$h_f(y)$	The distribution fitted to the tile histogram.

k :	The corresponding bin of the fitted and tile histograms.
A_1, μ_1, σ_1 :	Amplitude (i.e. height), mean, standard deviation of the target distribution.
A_2, μ_2, σ_2 :	Amplitude (i.e. height), mean, standard deviation of the background distribution.

The three statistical indices which are listed in Table 8, and which are used to evaluate the **bimodality** (Condition 1), **Gaussianity** (Condition 2) and **class balance** (Condition 3) of the distributions of the tile histograms, are briefly described as follows:

STATISTICAL INDEX	DESCRIPTION	REFERENCE
Ashman's D coefficient (AD):	Ashman's D coefficient quantifies how well two Gaussian distributions are separated, by considering the distance between their mean values and dispersions (i.e. standard deviations).	<i>Ashman et al. (1994)</i>
Bhattacharyya Coefficient (BC):	The Bhattacharyya Coefficient evaluates the similarity of two distributions. Here we use it to verify if the fitted distribution is a good approximation of the tile histogram, meaning that the hypothesis of having two Gaussian distributions is verified.	<i>Aherne et al. (1998)</i>
Surface Ratio (SR):	The Surface Ratio estimates the frequency of the smaller distribution relative to the larger one. It is computed as the ratio of areas under the curve of both distributions, where the product of a distribution's amplitude (representing its height) and standard deviation (representing its width) is a proxy estimate of its area under the curve.	<i>Chini et al. (2017)</i>

An overview of the procedure for selecting those tiles in the SAR image with a **bimodal** histogram (i.e. Condition 1), and two **Gaussian** and **balanced** distributions (i.e. Conditions 2 and 3), evaluated using the three statistical indices described above, is provided in Table 9 below.

In summary, for a given SAR image (I), the tiles of the entire N-Level quadtree are scanned from the set of nodes at level **0** (i.e. S^0) to the set of nodes at level **N-1** (i.e. S^{N-1}). Any node (and its associated image tile) that fulfils the following three conditions (C1, C2, and C3) is finally selected, and its "children" (i.e. tiles at subsequent levels in the quadtree) are not considered as candidate tiles:

$$\blacksquare \quad C1(I[S^i_j]) = C2(I[S^i_j]) = C3(I[S^i_j]) = 1 \text{ (see Table 9 below).}$$

These processing steps result in a binary map, called a **bimodal mask** (BM), which includes regions in image I where the distinctive populations, G_1 and G_2 , are present with a sufficient number of pixels, and where their distribution functions are clearly identifiable, and with more balanced prior probabilities. At this point, all tiles fulfilling the three conditions are selected, and the histogram of all pixels enclosed by these tiles (i.e. where $BM=1$) must be clearly bimodal.

Finally, the histogram $h(I[\text{where } BM=1])$ is used to estimate the parameters of the Gaussian distributions of the target and background classes (i.e. $G^{BM}_1(y)$ and $G^{BM}_2(y)$). To do so, the **Levenberg-Marquardt algorithm** is applied again, this time to fit two Gaussian curves on the histogram of pixel values included in all selected tiles.

Table 9: Overview of the procedure for selecting statistically representative tiles, based on the evaluation of bimodality (Condition 1), Gaussianity (Condition 2), and class balance (Condition 3).

CONDITION	DEFINITION	COMPUTATION (See text below for abbreviations)
Condition 1: (C1)	For a mixture of two Gaussian distributions, $AD > 2$ is required for a clear separation of the classes.	$C1(I[S_j^i]) = \begin{cases} 1 & \text{if } AD(h_f(I[S_j^i])) > 2 \\ 0 & \text{otherwise} \end{cases}$
Condition 2: (C2)	In the case of two almost identical histograms, BC approaches 1, so we assume that $BC > 0.99$ is required for a good fitting.	$C2(I[S_j^i]) = \begin{cases} 1 & \text{if } BC(h(I[S_j^i]), h_f(I[S_j^i])) > .99 \\ 0 & \text{otherwise} \end{cases}$
Condition 3: (C3)	The frequency of the smaller distribution should be at least 10% of the other, so $SR > 0.1$ is required for the target and background classes to be considered balanced.	$C3(I[S_j^i]) = \begin{cases} 1 & \text{if } SR(h_f(I[S_j^i])) > .1 \\ 0 & \text{otherwise} \end{cases}$

where the abbreviations are defined as follows:

AD :	Ashman's D coefficient (used to evaluate the bimodality of the tile histogram).
BC :	Bhattacharyya Coefficient (used to evaluate the Gaussianity of the tile histogram).
SR :	Surface Ratio (used to evaluate the class balance in the tile histogram).
S_j^i :	A given node at the i th level of the quadtree, where $j \in [1, 4^i]$.
$I[S_j^i]$:	The tile of image I related to $S_j^i \in S$ (the set of nodes in the entire quadtree).
$h(I[S_j^i])$:	The tile histogram of image values.
$h_f(I[S_j^i])$:	The distribution fitted to the tile histogram.

3.1.3 Mapping target class by HSBA-based histogram thresholding and region growing

The LIST flood mapping algorithm first parameterizes the target (i.e. water and change) and background (i.e. non-water and non-change) distributions in the new and difference SAR images (I_{t0} and I_D), using HSBA. The four distributions are then used to apply a histogram thresholding and region growing, in order to map the target classes (water and change) in I_{t0} and I_D .

The LIST flood mapping algorithm must handle different flood events, lasting from a few days to weeks (e.g. monsoon events). Thus, the algorithm addresses two different possible cases:

1. Increased and receded floodwater in the new SAR image.
2. Only receded floodwater in the new SAR image.

If neither of the cases is satisfied, meaning that the floodwater in the new SAR image (I_{t0}) has not changed relative to the previous SAR image (I_{t0-1}), the previous flood map (FM_{t0-1}) is not updated.

The way the LIST flood mapping algorithm is implemented for both cases, is explained below.

Case 1 - Increased and receded floodwater in the new SAR image:

In the first case, the hierarchical split-based approach (HSBA) described earlier is applied in parallel to the new and difference SAR images (I_{t0} and I_D), and the resulting **bimodal mask** is used to estimate the target and background distributions in I_{t0} and I_D . The estimated distributions are then used to apply a “two-input” region growing of both I_{t0} and I_D . A new flood map (FM_{t0}) is created by adding the new floodwater in I_{t0} to the previous flood map (FM_{t0-1}). In addition, any floodwater in FM_{t0-1} that has receded in I_{t0} , is removed from FM_{t0} .

The target and background distributions in I_{t0} and I_D are first estimated using HSBA, as follows:

▪	The difference image (I_D) between the new and previous SAR images (I_{t0} and I_{t0-1}) is computed as $I_D = I_{t0-1} - I_{t0}$.
▪	HSBA then is applied in parallel to I_{t0} and I_D , in order to select tiles that show a bimodality behaviour in both the new and difference SAR images. This guarantees that for these tiles: – I_{t0} is affected by floodwater, and – I_D is characterized by a class representing a decrease of backscattering with respect to I_{t0-1}.
▪	Permanent water is removed from the resulting bimodal mask, using the GFM Reference Water Mask and Copernicus DEM (Water Body Mask), to have distribution functions in I_{t0} and I_D more representative of flooding.
▪	The final bimodal mask is used to estimate the target and background distributions in the new image (i.e. water and non-water) and in the difference image (i.e. change and non-change).
▪	The four target and background distributions are used (see below) to apply a “two-input” region growing of the new and difference images (I_{t0} and I_D), to map areas that are both water in I_{t0} and change (i.e. flooded) in I_D .

Region growing assumes that the pixels of the target classes (i.e. water and change) in the new and difference SAR images are clustered, not randomly spread out over the entire image. The region growing algorithm that is used to map the target class requires two parameters: a **threshold value** for selecting seed pixels, and a **tolerance value** for stopping the region growing. In the region growing, the threshold and tolerance values are used (as described below) for the purposes of comparison with the **posterior probabilities** computed for the target class, which are inferred from the fitted distributions of the target and background classes, via Bayes’ theorem (see Table 10).

Table 10: Terms and definitions used to determine the posterior (or updated) probability of a hypothesis being true, given new evidence, using Bayes’ theorem.

TERM	DEFINITION
A, B :	Event A (the hypothesis), event B (the evidence).
$P(A / B)$:	The posterior (updated) probability of event A, given event B, which is inferred as: $P(A / B) = P(B / A) * P(A) / P(B)$
$P(B / A)$:	Conditional probability (or likelihood) of B, given A.
$P(A)$:	The prior (independent) probability of A.
$P(B)$:	The marginal (independent) probability of B.

How the target and background distributions are used to compute the posterior probabilities for the target classes in the new and difference SAR images, which are then compared with the threshold and tolerance values in region growing, is summarized in Table 11 and Table 12 below.

Table 11: Terms and definitions used to compute the posterior probabilities of the water class in the new SAR image (I_{t0}), for comparison with the threshold and tolerance values in region growing.

TERM	DEFINITION
σ^0	Backscatter coefficient value (in dB) in image I_{t0} .
$P(W \sigma^0)$	Posterior (updated) probability of water class, inferred (by Bayes' theorem) as: $P(W \sigma^0) = P(\sigma^0 W) * P(W) / P(\sigma^0)$
$P(\sigma^0 W)$	Conditional probability (or likelihood) of water class (from its distribution function).
$P(W)$	Prior (independent) probability of water class (set to 0.5).
$P(\sigma^0 NW)$	Conditional probability (or likelihood) of non-water class (from its distribution function).
$P(NW)$	Prior (independent) probability of non-water class (set to 0.5).
$P(\sigma^0)$	Marginal (independent) probability of backscatter coefficient value in image I_{t0} , calculated as: $P(\sigma^0) = [P(\sigma^0 W) * P(W)] + [P(\sigma^0 NW) * P(NW)]$

Table 12: Terms and definitions used to compute the posterior probabilities of the change class in the difference SAR image (I_b), for comparison with the threshold and tolerance values in region growing.

TERM	DEFINITION
$\Delta\sigma^0$	Backscatter difference value (in dB) in image I_b .
$P(C \Delta\sigma^0)$	Posterior (updated) probability of change class, inferred (by Bayes' Theorem) as: $P(C \Delta\sigma^0) = P(\Delta\sigma^0 C) * P(C) / P(\Delta\sigma^0)$
$P(\Delta\sigma^0 C)$	Conditional probability (or likelihood) of change class (from its distribution function).
$P(C)$	Prior (independent) probability of change class (set to 0.5).
$P(\Delta\sigma^0 NC)$	Conditional probability (or likelihood) of non-change class (from its distribution function).
$P(NC)$	Prior (independent) probability of non-change class (set to 0.5).
$P(\Delta\sigma^0)$	Marginal (independent) probability of backscatter difference value in image I_b , calculated as: $P(\Delta\sigma^0) = [P(\Delta\sigma^0 C) * P(C)] + [P(\Delta\sigma^0 NC) * P(NC)]$

The posterior probabilities for the target classes (water and change) in the new and difference SAR images are first used to select seed pixels for region growing, by a comparison with the threshold value. Selected seeds occurring in the height above nearest drainage (HAND) mask are removed, in order to avoid that false alarms (in areas not prone to flooding) could spread. Starting from the seed pixels, region growing is applied to neighbouring pixels with posterior probabilities within a specified range of tolerance values. The exact procedure is summarized as follows:

Threshold value (for seed pixels):	$P(W \sigma) \geq 0.7$ and $P(C \Delta\sigma) \geq 0.7$: ▪ Pixels with posterior probabilities for water class and change class ≥ 0.7 are selected as seeds for region growing.
Tolerance value (to stop region growing):	$P(W \sigma) > 0.3$ and $P(C \Delta\sigma) > 0.3$: ▪ Neighbouring pixels with posterior probabilities for water class and change class > 0.3 are subject to region growing. ▪ Otherwise, region growing stops.

Any new floodwater detected by the application of the region growing procedure to the new and difference SAR images (I_{t0} and I_D), is added to the previous flood map (FM_{t0-1}), producing a new flood map (FM_{t0}). The algorithm must also remove from FM_{t0-1} any floodwater that may have receded in the new SAR image (I_{t0}). To do this, the previously estimated water and non-water distributions are used to apply a “single-input” region growing of only the new SAR image (I_{t0}), in order to map the water in I_{t0} . Only pixels that are floodwater in the previous flood map (FM_{t0-1}) and water in the new SAR image (I_{t0}) are mapped as floodwater in the new flood map (FM_{t0}).

Case 2 - Only receded floodwater in the new SAR image:

If the bimodal map resulting from applying HSBA in parallel to I_{t0} and I_D (described above) is empty, this means that any floodwater in the previous SAR image (I_{t0-1}) has either only receded or has not changed in the new SAR image (I_{t0}). To remove any receded floodwater in I_{t0} from the previous flood map (FM_{t0-1}), HSBA is re-applied, but only to the new image (I_{t0}), not in parallel to I_{t0} and I_D . The estimated water and non-water distributions are then used to apply a “single-input” region growing, in order to map the water in I_{t0} . Again, only pixels that are floodwater in the previous flood map (FM_{t0-1}) and water in I_{t0} , are mapped as floodwater in the new flood map (FM_{t0}).

3.1.4 Final post-processing of the LIST flood map

Incidence angle-based processing and mask application:

The Observed Flood Extent output layer generated by the LIST flood mapping algorithm is processed using the GFM Exclusion Mask, the Height Above Nearest Drainage (HAND) mask, and the ocean pixels of the Copernicus DEM (Water Body Mask), to remove all pixels not part of flood-prone areas.

It should be noted that the LIST flood mapping algorithm is applied separately to strips that are a subset of the original input Sentinel-1 image, each strip having a narrower range of incidence angles. This reduces the impact of angle-dependent signal variations, which arise because backscatter can vary across an image, depending on the angle at which the radar signal hits the ground.

In general, backscatter decreases with increasing incidence angle. This dependency is observed for most land cover classes, but is most significant for smooth surface classes (e.g. water). For SAR wide-swath data, the backscatter distribution of a given class varies from near- to far-range. If the class of interest extends over the entire image, significant differences between incidence angles hamper the attribution of a unique backscatter distribution function.

Splitting the image into strips therefore improves the parameterization of the distribution functions of both the target (water and change) and background (non-water and non-change) classes in each strip. Each Sentinel-1 image is split into three strips, using incidence angle as a metric, as follows:

STRIP	INCIDENCE ANGLE
1	30° - 38°
2	36° - 44°
3	42° - 46°

In order to avoid creating gaps in the flood map where flooding covers multiple strips, the borders between strips overlap slightly. This ensures a smooth flood map, even for large flood zones. In short, splitting the image with overlap improves the accuracy of the parametrization of the target and background distribution functions, by considering the effect of incidence angle on the radar signal. As an added benefit, processing the Sentinel-1 image in separate strips enables them to be analysed simultaneously, making processing much faster.

Topography-informed object-based filtering:

The Observed Flood Extent output layer is further refined through a topography-informed, object-based classification and filtering approach. The purpose of this post-processing step is to preserve the most probable (“true”) flood extent while mitigating potential over-detection using HAND data.

First, the flood extent is segmented into discrete flood objects (F) based on D8 flow-direction connectivity. For each flood object, a two-pixel-wide buffer (B) is subsequently generated through two iterations of a 3x3 morphological dilation. Pixels within the buffer that overlap the water body mask are excluded from further analysis. The corresponding HAND values within both the flood objects and their buffers are then extracted for comparative evaluation.

This approach relies on the physical assumption that the HAND value distribution of a genuine flooded area should be shifted toward smaller values compared to that of its immediate surroundings, reflecting the gravitational tendency of water to accumulate within local topographic depressions. Thus, the statistical characteristics of the HAND value distributions within each flood object and its buffer are evaluated and classified according to five topography-based filters:

1	True filter:	This filter assesses if the HAND value distribution of the buffer exhibits a longer upper tail and higher HAND bins than that of the flood object, indicating that the flood object is topographically enclosed by its surroundings. In a physically plausible scenario, the buffer should typically display higher HAND values relative to the flood object. If fewer than 3 pixels populate the highest HAND bins within the buffer - even if those HAND values exceed those of the flood object – the comparison extends iteratively to the next-highest bins until at least three pixels are included. Flood objects satisfying this criterion are classified as Class 1 (True).
2	Bi-modality filter:	This filter examines if the HAND value distribution of the flood object exhibits a bi-modal pattern. Such a distribution may indicate possible overtopping of flood defences (e.g. levees, dikes, flood walls), resulting in a secondary peak in the HAND distribution implying distinct elevation zones within the flood object. Flood objects meeting this condition are classified as Class 2 (Bi-modal).
3	Overlap filter:	This filter evaluates the degree of similarity between the HAND value distributions of the flood object and its buffer using an overlap coefficient (threshold = 0.99). A high overlap implies that both areas locate on a flat topography, thereby questioning the likelihood of inundation. Flood objects meeting this condition are classified as Class 3 (Overlapping).
4	Mean difference filter:	This filter quantifies the difference between the mean HAND values of the flood object and its buffer. If the mean HAND of the flood object exceeds that of its buffer by more than 0.52m, the object is likely situated on higher ground than its surrounding, contradicting hydrological plausibility. These are classified as Class 4 (Elevated).
5	Mode filter:	This filter checks whether the mode of the flood objects’ HAND value distribution exceeds a predefined threshold (set to 14m). Exceedance of this threshold suggests the flood object occupies a relatively high terrain, inconsistent with typical flood behaviour. Such objects are classified as Class 5 (High mode).

Flood objects that do not meet the criteria of any of the five filters are classified as Class 6 (Unclassified) and are conservatively retained as potentially true flood objects. In cases where multiple filters assign conflicting classes to a single object (e.g. “True” and “High-mode”), a hierarchical overwriting sequence is applied in the order $5 > 4 > 3 > 2 > 1$, ensuring a consistent and physically meaningful classification outcome. This hierarchy reflects the relative confidence in identifying potential over-detections, whereby filters associated with stronger evidence of implausibility (e.g., a high-mode flood object) take precedence over those with weaker or more ambiguous conditions (e.g., overlapping distributions). For instance, if a flood object is simultaneously identified as both Class 1 and Class 4, it is finalized as Class 4 according to this rule.

Following the topography-informed filtering, flood objects classified as Classes 1, 2, and 6 are retained, representing plausible inundated areas. Conversely, flood objects classified as Classes 3, 4, and 5 are re-labelled as non-flood pixels, representing likely instances of over-detection.

3.1.5 Computation of Likelihood Values by the LIST flood mapping algorithm

The LIST flood mapping algorithm estimates the likelihood of flood classification using the posterior probabilities ($P(W | \sigma^0)$ and $P(\Delta\sigma^0 | C)$) computed for the water class in the new image and the change class in the difference image (see Table 11 and Table 12). For the two cases described in **Section 3.1.3**, the Likelihood Values are computed using two different methods, as described below.

For Case 1 (i.e. **increased and receded floodwater in the new SAR image**), both the new (I_{t0}) and difference (I_D) SAR images are used to estimate the Likelihood Values. In this case, those pixels with high posterior probabilities for both water and change classes (i.e. ($P(W | \sigma^0)$ and $P(\Delta\sigma^0 | C)$) are likely to be flooded pixels. Therefore, a grid-cell’s probability of being flooded ($P(F | \sigma^0)$) is computed as the minimum value of the posterior probabilities for the water and change classes, as follows:

$$P(F | \sigma^0) = \text{minimum}[P(W | \sigma^0), P(\Delta\sigma^0 | C)]$$

For Case 2 (i.e. **only receded floodwater in the new SAR image**), only the new SAR image (I_{t0}) is used to estimate the Likelihood Values. In this case, a grid-cell’s probability of being flooded ($P(F | \sigma^0)$) is computed as the posterior probability for the water class, as follows:

$$P(F | \sigma^0) = P(W | \sigma^0)$$

It should be noted that, for Case 2, because the likelihood is only calculated from the backscatter value in the new image (I_{t0}), false high flood probability can be caused by permanent water and other water look-alike dark areas. Such false alarms in the binary flood map have been removed by comparing the map with the previous flood map, as described above. To reduce these false high probabilities in the current likelihood map, for non-flood pixels in the new flood map, the flood probability ($P(F | \sigma^0)$) is the minimum value between the posterior probability for water ($P(W | \sigma^0)$) in the new SAR image and the probability of being flooded ($P(F | \sigma^0)$) in the previous likelihood map.

Finally, for both Cases 1 and 2, for grid-cells in the GFM Exclusion Mask, flood probability is set to 0.

3.2 GFM flood mapping algorithm 2 (DLR)

The DLR flood mapping algorithm makes use of an automatic hierarchical tile-based thresholding approach in combination with a fuzzy logic-based post-processing step for the unsupervised extraction of the flood extent in Sentinel-1 SAR data. The algorithm consists of three steps:

- Unsupervised initialization of the water classification by hierarchical automatic tile-based thresholding, to derive a global threshold backscatter value between water and non-water.
- Fuzzy logic-based refinement of the initial water classification, in order to exclude water look-alikes, using SAR backscatter values, topographic slope, and the size of detected water bodies.
- Region growing of the defuzzified water classification, in order to integrate the transient shallow water zone between flooded and non-flooded areas.

The DLR flood mapping algorithm requires the following four main input raster datasets:

- | |
|---|
| ▪ A full Sentinel-1 SAR image scene. |
| ▪ A digital elevation model (DEM), to determine topographic slope (see Section 2.2.2). |
| ▪ A height above nearest drainage index - exclusion mask (HAND-EM), to exclude non-flood-prone areas (see Sections 2.2.4 and 4.4.3). |
| ▪ A reference water mask, to exclude permanent (and seasonal) water bodies (see Section 4.3). |

Initial descriptions of the DLR flood mapping algorithm are provided by Martinis et al. (2015) and Twele et al. (2016). The main steps of the DLR flood mapping algorithm are described below.

3.2.1 Unsupervised initialization of water classification by tile-based thresholding

In the first stage of the DLR flood mapping algorithm an automatic, parametric tile-based thresholding procedure is applied to the Sentinel-1 SAR image scene, to derive a global threshold backscatter value between water and non-water classes. The global threshold is calculated for a selected limited number of representative image subsets (or tiles), and is then applied to the entire image scene by labelling as “water” all pixels with backscatter values lower than the threshold.

The selection of representative tiles for calculating the global threshold backscatter value is based on the assumption that tiles with low mean backscatter values and high standard deviations, have bimodal backscatter distributions and are likely to contain both water and non-water features.

This first step of the DLR flood mapping algorithm is carried out in two stages, as described below:

- Selection of tiles containing bimodal backscatter distributions and water-land boundaries.
- Calculation of the optimal global threshold backscatter value separating water and land.

1. Selection of tiles containing bimodal backscatter distributions and water-land boundaries:

In the first stage of step 1 of the DLR flood mapping algorithm, a two-level quadtree data structure is generated which, at the first or **parent** level (denoted as S^+), splits the SAR image scene into a discrete number of non-overlapping image subsets (called **parent tiles**), each of defined size C by C pixels (where $C = 200$). Each parent tile is represented at the second or **child** level (denoted as S^-) by four square **sub-tiles**, each of size $C/2$ by $C/2$ pixels. The two-level quadtree-based splitting of a SAR image by the DLR flood mapping algorithm, is illustrated in Figure 5.




Figure 5: Illustration of the two-level splitting of a Sentinel-1 SAR image. Each parent tile (level S^+) has a backscatter mean and standard deviation (μ^+ and σ_μ^+). Each parent tile is further split into four sub-tiles (level S^-) each with a backscatter mean (μ^-).

To determine the global threshold backscatter value between water and non-water, representative parent tiles are selected based on the probability (determined by the statistics of their backscatter values) that they contain a bimodal mixture of water and non-water pixels. The following situations (illustrated in Figure 6) can result in parent tiles that are not valid for threshold calculation:

- In rare cases, parent tiles intersect with the edge of a SAR scene, resulting in tiles smaller than the pre-defined size (i.e. C by C pixels). Such tiles are not valid for threshold computation.
- Due to the near-polar orbits of Sentinel-1, the along-track direction forms an angle of $\sim 10^\circ$ relative to north-south. This can result in parent tiles with variable amounts of **No Data** content. If over 50% of overall data content is **No Data**, that tile is not valid for threshold computation.
- Similarly, parent tiles with more than 20% of areas that are not flood-prone, based on the height above nearest drainage index - exclusion mask (HAND-EM), are also not valid.




Figure 6: In this Sentinel-1 SAR image scene, only green parent tiles are considered for threshold computation. Red parent tiles, which intersect with the edge of the SAR scene, and orange parent tiles, which have over 50% No Data content, are not considered.

The statistical parameters and procedure used to select representative parent tiles that most likely contain bimodal backscatter distributions and water-land boundaries, are described in Table 13.

Table 13: Summary of the statistical parameters and procedures used to select representative parent tiles likely to contain bimodal backscatter distributions and water-land boundaries.

PARAMETER	DESCRIPTION / DEFINITION	HOW PARAMETER IS USED				
S^+	First (parent) level of quadtree data structure.	Splits SAR image Y into N_{tiles} non-overlapping tiles, each of size C by C pixels (where $C = 200$).				
S^-	Second (child) level of quadtree data structure.	Splits each parent tile at level S^+ into four non-overlapping sub-tiles, each of size $C/2$ by $C/2$ pixels (i.e. 100 x 100 pixels).				
μ^-	For each of the four sub-tiles at level S^- , the mean backscatter value.	The backscatter distribution of a parent tile at level S^+ is characterized by the standard deviation (σ_{μ^+}) of mean backscatter values for the four sub-tiles ($\mu_1^-, \mu_2^-, \mu_3^-, \mu_4^-$) at level S^-.				
μ_{global}	For SAR image Y , the global mean backscatter value.	A parent tile with low mean backscatter (i.e. $\mu^+ < \mu_{global}$) is likely to contain water.				
μ^+	For a parent tile, the mean backscatter value.	- A parent tile that fulfils the two conditions of having a low mean backscatter (μ^+) and a high standard deviation (σ_{μ^+}), is likely to have a bimodal backscatter distribution, and to contain both water and non-water features. - Both conditions are evaluated as follows: <table><tr><td>Condition 1:</td><td>$\mu^+ < \mu_{global}$</td></tr><tr><td>Condition 2:</td><td>$\sigma_{\mu^+} \geq \mu_{\sigma_{\mu^+}} + x \cdot \sigma_{\sigma_{\mu^+}}$</td></tr></table> where $x = 2$.	Condition 1:	$\mu^+ < \mu_{global}$	Condition 2:	$\sigma_{\mu^+} \geq \mu_{\sigma_{\mu^+}} + x \cdot \sigma_{\sigma_{\mu^+}}$
Condition 1:	$\mu^+ < \mu_{global}$					
Condition 2:	$\sigma_{\mu^+} \geq \mu_{\sigma_{\mu^+}} + x \cdot \sigma_{\sigma_{\mu^+}}$					
σ_{μ^+}	For a parent tile, the standard deviation of the four mean backscatter values at level S^- (i.e. $\mu_1^-, \mu_2^-, \mu_3^-, \mu_4^-$), calculated as: $\sigma_{\mu^+} = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n=4} (\mu_i^- - \text{avg_}\mu^-)^2}$ where avg_ μ^- is the mean value of the four mean backscatter values (μ^-) at level S^- .					
$\mu_{\sigma_{\mu^+}}$	For all N parent tiles, the mean of all standard deviations (σ_{μ^+}), calculated as: $\mu_{\sigma_{\mu^+}} = \frac{1}{N} \sum_{i=1}^N \sigma_{\mu_i^+}$	- In Condition 2 above, the mean and standard deviation of all standard deviations (σ_{μ^+}) of parent tiles, are used to select parent tiles with standard deviations as high as possible. - Specifically, to be selected as a representative parent tile, the tile's standard deviation (σ_{μ^+}) must be greater than or equal to the mean of all standard deviations plus twice the standard deviation of all standard deviations.				
$\sigma_{\sigma_{\mu^+}}$	For all N parent tiles, the standard deviation of all standard deviations (σ_{μ^+}), calculated as: $\sigma_{\sigma_{\mu^+}} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (\sigma_{\mu_i^+} - \text{avg_}\sigma_{\mu^+})^2}$ where avg_ σ_{μ^+} is the mean value of all standard deviations at level S^+ .					

As is described in Table 13, in order to be selected as a representative image subset, a parent tile (i.e. at level S^+ in the quadtree) must fulfil two conditions:

Condition 1:	The mean backscatter value (μ^+) of the parent tile must be lower than the global mean backscatter of the entire Sentinel-1 SAR image scene. It is assumed that parent tiles with low mean backscatter values (μ^+) are likely to contain water.
Condition 2:	The standard deviation of backscatter values (σ_{μ}^+) for the parent tile, which is computed for all four of its sub-tiles (i.e. at level S^-), must be as high as possible. It is assumed that parent tiles with high standard deviations (σ_{μ}^+) are likely to have a bimodal backscatter distribution, and to contain both water and non-water features.

Specifically, in order to fulfil Condition 2, a parent tile's standard deviation (σ_{μ}^+) must be greater than or equal to the **mean of all standard deviations** plus twice the **standard deviation of all standard deviations**. This concept is illustrated in Figure 7, where all standard deviations (σ_{μ}^+) are plotted into a Gaussian distribution. As can be seen, a standard deviation (σ_{μ}^+) greater or equal to the distribution's mean, or even the mean plus one standard deviation, would still leave too many parent tiles, the majority of which may not show bimodal backscatter distributions.




Figure 7: Illustration of how selecting parent tiles with the highest possible backscatter standard deviation (σ_{μ}^+) results in a smaller sample of representative tiles (area coloured in orange) most likely to have bimodal backscatter distributions.

Once the sample of representative parent tiles most likely having bimodal backscatter distributions and water-land boundaries is selected (using Conditions 1 and 2), the following checks are made:

- If the number of selected representative parent tiles is insufficient (i.e. only 10 or less), a bigger sample is obtained by re-applying Conditions 1 and 2 (see Table 13), with an increased range of allowed standard deviations (i.e. set $x = 1.28$, instead of $x = 2$).
- If more than 10 representative parent tiles have been selected, a statistically sound sample of representative parent tiles is assumed. The five representative parent tiles with the highest backscatter standard deviations (σ_{μ}^+) are then used for further threshold computation.

2. Calculation of the optimal global threshold backscatter value separating water and land:

All of the representative parent tiles that are selected as described above are likely to contain bimodal backscatter distributions and valid water-land boundaries, as illustrated Figure 8 below.




Figure 8: Example of a Sentinel-1 SAR image showing (in green) the selected representative parent tiles likely to contain bimodal backscatter distributions and water-land boundaries. (See text for details).

In order to derive the optimal global threshold backscatter value between the water and non-water classes (indicated in panel **b2** of Figure 8), the automatic, histogram-based thresholding method of Kittler and Illingworth (1986) is applied to the five representative parent tiles with the highest backscatter standard deviations (σ_{μ}^+). This method (called the **minimum error thresholding** or **MET method**) is an iterative, cost-minimization approach that splits the histogram into two classes with a threshold that identifies the class boundary. (Note that Otsu's method of histogram thresholding, which is used in the LIST flood mapping algorithm, is a special case of the MET method).

The optimal global threshold backscatter value (τ) separates both classes with minimum effort. As can be seen in Figure 8, if selecting a pure land tile (panel **a1**), the corresponding histogram (panel **a2**) is unimodal. If selecting a tile with both low and high backscatter values (panel **b1**), the corresponding histogram (panel **b2**) is bimodal, and we assume a water-land-boundary. If applying a threshold to the latter histogram, both classes can be separated, giving the class water in the left part. The final global threshold backscatter value (τ) is obtained as the mean of the threshold values (τ^+) of the five individual representative parent tiles. Applying this threshold to the entire Sentinel-1 SAR image scene results in an initial water classification.

In addition to the optimal global threshold backscatter value (τ) separating water and non-water, we also compute the mean backscatter value of the water class (μ_{water}), as this is used in the fuzzy logic-based refinement of the initial water classification (as described in **Section 3.2.2** below).

In accordance with Kittler and Illingworth (1986), we compute the mean backscatter value of the water class (μ_{water}) as the geometric centroid of the separated class in the histogram, as follows:

PARAMETER	FORMULA
Total number of backscatter values in class "i" (i.e. up to threshold value τ in histogram):	$P_i(\tau) = \sum_{g=a}^b h(g)$
Mean backscatter value of class "i" (i.e. up to threshold value τ in histogram):	$\mu_i(\tau) = \frac{\sum_{g=a}^b h(g) g}{P_i(\tau)}$

where:

- Class "i" refers to the water class.
- τ is the global threshold backscatter value separating water and non-water.
- a is the minimum backscatter value in class "i".
- b is the maximum backscatter value in class "i".
- g is the backscatter value of a histogram bin.
- $h(g)$ is the histogram count of backscatter value "g".

As was done for calculating the global threshold backscatter value (τ), the mean backscatter value for water (μ_{water}) for the entire Sentinel-1 SAR image scene is calculated as the average of the mean backscatter values for water derived for the five representative parent tiles.

Finally, if the automatic tile-based histogram thresholding fails to compute a reliable global threshold backscatter value (τ), or returns an unusually high value (i.e. $\tau > -15$ dB) for a representative parent tile, a **fallback threshold mechanism** is activated. This fallback system determines thresholds individually for each Sentinel-1 scene by analysing backscatter values from known inland water bodies, identified in the Copernicus Water Body Mask (see **Section 2.2.3**).

Empirical testing has established that the 60th percentile of backscatter values from these inland water bodies provides a suitable fallback threshold when it falls between **-20 dB** and **-16 dB**. If this calculation fails, the system implements a graduated response, as follows:

- Default value of **-18 dB**, if calculation completely fails.
- Value of **-19 dB**, if calculated fallback threshold falls below **-20 dB**.
- Value of **-17 dB**, if calculated fallback threshold exceeds **-16 dB**.

The global threshold is reset to these fallback values when at least two parent tiles require fallback thresholds and the fallback value is lower than the mean threshold across the five parent tiles. To prevent underestimation, this reset mechanism is not triggered when all originally determined thresholds are consistently high (and therefore likely accurate).

3.2.2 Fuzzy logic-based refinement of the initial water classification

In the second stage of the DLR flood mapping algorithm, a fuzzy logic-based approach is used to refine the initial water classification that was derived by applying the global threshold backscatter value between water and non-water (τ_g) to the Sentinel-1 SAR image. The objective is to improve the thematic accuracy of the initial water classification, by removing potential water look-alikes.

In the fuzzy logic-based approach, the likelihood of a SAR image pixel being classified as water is determined by its **degree of membership** to the water class, ranging from 0 (not belonging) to 1 (completely belonging). For each pixel, the degree of membership to the water class will be low if:

- The pixel has a high backscatter value, close to the global threshold backscatter value (τ_g).
- The topographic slope at that pixel is high, since steeper surfaces are unlikely to retain water.
- The pixel has a low number of neighbouring water pixels, since dispersed small areas of low backscatter are commonly related to water look-alike areas.

Conversely, a pixel's degree of membership to the water class will be high if the pixel has a low backscatter value, low topographic slope and a high number of neighbouring water pixels.

To compute the degrees of membership to the water class (or **fuzzy values**) of SAR image pixels, the **Standard Z** or **S fuzzy membership functions** (see below) are applied to three input datasets: (a) the **backscatter values** (σ_0) in the Sentinel-1 SAR image; (b) **topographic slope**, computed from the DEM; and (c) the **size of individual water bodies** in the initial water classification.

The Standard Z and S fuzzy membership functions are illustrated in Figure 9. The Standard Z fuzzy membership function (or Negative S-function) assigns higher degrees of membership (fuzzy values) to lower input pixel values, and lower fuzzy values to higher input pixel values. The Standard S fuzzy membership function (or Positive S-function) does the opposite, assigning lower fuzzy values to lower input pixel values, and higher fuzzy values to higher input pixel values. As can be seen, for both fuzzy membership functions the range of input pixel values is specified by the lower and upper **threshold values** (a and c, respectively), with a cross-over point (b) defined as **b = (a + c) / 2**.




Figure 9: Illustration of the Standard S fuzzy membership function (called here the Positive S-function) and the Standard Z fuzzy membership function (called here the Negative S-function). Modified from Pal-Rosenfeld (1988). (See text for details).

In our case, since low degrees of membership to the water class are associated with both high backscatter values and high topographic slope, we apply the **Standard Z fuzzy membership function** to the input datasets of **backscatter values** and **topographic slope**, in order to compute the fuzzy values for both datasets. Conversely, since high degrees of membership to the water class are associated with high numbers of neighbouring water pixels, we apply the **Standard S fuzzy membership function** to the input dataset of **size of individual water bodies**.

Figure 10 illustrates the application of the Standard Z fuzzy membership function to discriminate between water and non-water, based on backscatter values (σ_0) in the Sentinel-1 SAR image. As can be seen in Figure 10, the range of input backscatter values for the fuzzy membership function is specified by the lower and upper threshold values (**X1** and **X2**), which are defined as follows:

Lower fuzzy threshold value (X1):	This is the mean backscatter value of the water class (μ_{water}), based on the initial flood classification. As the majority of water pixels share that backscatter value, the likelihood of correct classification as water is high. Therefore, pixels with backscatter values close to the mean (μ_{water}) are assigned high fuzzy values, approaching 1.
Upper fuzzy threshold value (X2):	This is the global threshold value (τ_g) separating water and non-water. This backscatter value represents the boundary of both classes, where uncertainty of correct classification is highest. Pixels with backscatter values close to the global threshold have a low likelihood of correct classification as water, as they may also belong to non-water (within the classification error range), and are assigned lower fuzzy values, approaching 0.




Figure 10: Illustration of the application of the Standard Z fuzzy membership function to discriminate between the water and non-water classes, based on backscatter values in the Sentinel-1 SAR image. (See text for details).

The upper and lower fuzzy threshold values that are used for applying the **Standard Z** and **S fuzzy membership functions** to compute the fuzzy values for the input datasets of **backscatter values**, **topographic slope**, and **size of individual water bodies**, are summarized in Table 14.

Table 14: Upper and lower fuzzy threshold values (X_1 and X_2) used to compute the degrees of membership to the water class, for the input datasets of backscatter values, topographic slope, and size of individual water bodies.

INPUT VARIABLE	FUZZY MEMBERSHIP FUNCTION	LOWER FUZZY THRESHOLD VALUE (X_1)	UPPER FUZZY THRESHOLD VALUE (X_2)
Backscatter values (σ_0)	Standard Z	Mean backscatter value (μ_{water}) of water pixels.	Global threshold value (τ_g) between water and non-water.
Topographic slope	Standard Z	0 degrees.	18 degrees.
Size of individual water bodies	Standard S	10 pixels.	500 pixels.

The result of applying the fuzzy logic-based approach to the three input datasets is a fuzzy set consisting of three layers (one for each variable) of **degrees of membership to the water class** (i.e. the fuzzy values), with floating point values in the range [0, 1]. For performance reasons, the fuzzy values are rescaled to the range [0, 100]. The three fuzzy layers are combined into a single layer of fuzzy values by calculating, for each pixel, the mean of the three degrees of membership to the water class. In the final step of the fuzzy logic-based approach, a **defuzzified water classification** is created, by labelling all pixels with a **mean fuzzy value ≥ 0.6** as correct water classifications.

3.2.3 Region growing of the defuzzified water classification

The transient shallow water zone between flooded and non-flooded areas is often characterized by successively increasing backscatter levels, mainly due to a higher signal return of emerging vegetation. The objective of the region growing step of the DLR flood mapping algorithm, is to integrate these areas in the flood classification, and to increase the spatial homogeneity of the detected flood plain. The following three input datasets are used as input for the region growing:

- The single layer of fuzzy values showing the mean degrees of membership to the water class.
- The defuzzified water classification, showing all image pixels with a mean fuzzy value ≥ 0.6 .
- A dataset showing the size (in pixels) of the water regions (i.e. sets of connected water pixels) in the defuzzified water classification.

The region growing is implemented in two stages, which are described as follows:

1	Selection of seed points:	▪ All water regions of at least 30 pixels with fuzzy values of ≥ 0.6 , are selected as seed regions for region growing.
2	Tolerance for region growing:	▪ Pixels directly connected to seed regions, with fuzzy values (FVAL) of $0.35 \leq \text{FVAL} < 0.6$, are potential water pixels. Pixels in this range but not connected to seed regions are excluded. Each seed region is allowed to grow into (merge with) adjacent potential water pixels.

The implementation of the region growing step for a 16x16 pixel sub-image is illustrated in Figure 11, where the contents of the six arrays shown are described below:

1	Fuzzy Array:	▪ The single layer of fuzzy values, computed as the mean of the fuzzy values for the three input datasets of backscatter values, topographic slope, and size of water bodies.
2	Defuzzy Array:	▪ The defuzzified water classification, created by labelling all pixels in the single layer of fuzzy values with a fuzzy value ≥ 0.6 as correct water classifications.
3	Labelled Array with region sizes:	▪ Each “water region” (i.e. set of connected water pixels), labelled as water and assigned with the region size in pixels.
4	Min Array:	▪ Seed pixels, defined as water regions with a minimum size of 30 pixels with fuzzy values of ≥ 0.6 . These regions are allowed to grow into the adjacent potential water pixels.
5	Max Array:	▪ Potential water pixels, defined as all pixels with fuzzy values (FVAL) of $0.35 \leq \text{FVAL} < 0.6$. The seed pixels will grow into adjacent potential water pixels.
6	Region Growing Array:	▪ Potential water pixels directly connected to the seed regions, are now merged. Any other potential water pixel, not connected to the seed region, is excluded from this operation.




Figure 11: Illustration of the implementation of the region growing step in the DLR flood mapping algorithm, for a 16x16 pixel image subset. See text for details and descriptions of the six arrays.

3.2.4 Final post-processing of the DLR flood map

Any small fragmented patches of the water and non-water classes in the region growing results are first eliminated using a blob removal method, which is implemented as follows:

- Water regions with a size less than the defined threshold (8 pixels) are considered too small to be certainly water, and so are re-classified as non-water, with a fuzzy value of 59% (the minimum likelihood for non-water).
- Non-water regions smaller than the defined threshold (31 pixels) and completely surrounded by water, are considered too small to be certainly non-water, and are re-classified as water, with a fuzzy value of 60% (the minimum likelihood for water).

An additional region growing operation is performed, based on the SAR backscatter values, to include directly connected pixels that are within a tolerance criterion of 1 decibel (dB). In contrast to the previous steps, the parameters for this region growing are determined by the algorithm, and cannot be static. The region growing is implemented in two stages, as described below.

1	Selection of seed points:	▪ The lower boundary for region growing is the global threshold backscatter value (τ_g) separating water and non-water, determined by automatic tile-based histogram thresholding.
2	Tolerance for region growing:	▪ The upper boundary for region growing is set to the global threshold backscatter value (τ_g) + 1 db, so that the water regions grow by 1 dB.

Finally, it should be noted that before the DLR flood mapping algorithm is implemented, the input Sentinel-1 SAR backscatter values are rescaled from the original range (**-40 - 0 db**) to **0 - 400**, to ensure positive input values for all processing steps. Thus, the additional region growing operation, for example, is implemented using rescaled backscatter values, as shown below:

Backscatter values (dB):	-40	-18	-16	-15	0
Re-scaled backscatter values:	0	220	240	250	400

3.2.5 Computation of Likelihood Values by the DLR flood mapping algorithm

The DLR flood mapping algorithm estimates the likelihood of flood classification for each grid-cell as the mean of the three **degrees of membership to the water class** (or fuzzy values), computed from the input datasets of backscatter values, topographic slope, and size of individual water bodies, as described earlier in **Section 3.2.2**.

3.3 GFM flood mapping algorithm 3 (TUW)

The TUW flood mapping algorithm, which was originally described by Bauer-Marschallinger et al. (2022), parameterizes a large fraction of the Sentinel-1 SAR backscatter signal history for each image pixel, which is stored within an extensive spatio-temporal **datacube** (or time-series) of SAR images that have been geocoded, gridded and stored as analysis ready data. As described in **Section 2.1** of this report, the datacube contains the complete historical archive of Sentinel-1 SAR IW GRDH products acquired in VV (vertical transmit and vertical receive) polarisation, which are preprocessed, co-registered and time-stacked over the Equi7Grid global spatial reference system. The Sentinel-1 SAR image data are stored as backscatter coefficient values (σ_0) in decibels, which have been resampled to a spatial resolution of 20 metres (from an original resolution of 10 metres), mainly for the purposes of noise reduction, as well as for optimal data storage and processing requirements.

The Sentinel-1 SAR image datacube serves as the source for two important datasets that are used by the TUW flood mapping algorithm, as described below: (a) the **projected local incidence angle** (PLIA) values, which describe the SAR observation geometry for each pixel; (b) the **harmonic parameters**, which describe the backscatter's seasonality for each pixel.

The TUW flood mapping algorithm takes advantage of Sentinel-1 orbit repetition and *a priori* generated probability parameters for flood and non-flood conditions. A globally applicable **flood signature** is obtained from manually collected wind- and frost-free images, and provides expected backscatter Sentinel-1 values over water, depending on incidence angle. Through harmonic analysis of each pixel's full time-series, a local seasonal **non-flood signature** is derived, comprising the expected backscatter values for each day-of-year over land pixels. From these predefined probability distributions, all incoming Sentinel-1 images are classified in near real-time by simple Bayes inference, also providing the uncertainty values that are ingested in the GFM likelihood layer.

The TUW flood mapping algorithm requires the following three main inputs datasets:

- The latest Sentinel-1 SAR image scene to be processed.
- Sentinel-1 projected local incidence angle data (used to derive the Flood backscatter signature).
- Sentinel-1 harmonic parameters (used to derive the Non-flood backscatter signature).

In the following sub-sections, details are presented on the main datasets and analysis methods used by the TUW flood mapping algorithm, namely:

- Computation of Sentinel-1 projected local incidence angle data.
- Computation of Sentinel-1 harmonic parameters.
- Estimation of backscatter distribution functions for water and land surfaces.
- Bayesian flood mapping and uncertainty estimation.
- Low sensitivity masking.
- Morphological post-processing.

3.3.1 Computation of Sentinel-1 projected local incidence angle data

Under the conditions that water bodies are open, calm, and non-frozen, the Sentinel-1 SAR backscatter signal can be assumed to be of universal character, and primarily dependent on the (local) incidence angle of the incoming Sentinel-1 radar wave.

Given a point target on a scattering surface, characterized by a certain local slope, the local incidence angle (θ_l) is the angle between the incoming radar wave direction and the normal direction to the scattering surface. As can be seen in Figure 12, the local incidence angle (θ_l) differs from the incidence angle (θ), which is the angle between the incoming wave direction and the vertical direction to the ellipsoid. The projected local incidence angle (PLIA) is simply the local incidence angle projected into the range (i.e. across-track) plane. The PLIA provides essential information about the observation geometry of the satellite, and varies across different orbits.




Figure 12: Illustration of SAR observation geometry. From Rizzoli and Bräutigam (2014).

The PLIA values are available as a by-product of the terrain correction step of the SAR preprocessing chain. Because of the self-repeating orbit geometries of the Sentinel-1 mission, almost identical observation angles are established at each overpass. Globally, the Sentinel-1 mission has 175 (repeating) relative orbits, with locally up to 9 orbits. Consequently, when working with Sentinel-1 data separately per relative orbit (r), a pixel's PLIA value (θ_r) is assumed to be constant.

We capitalise on this, and use as input to the TUW flood mapping algorithm a set of constant θ_r values, which are computed a priori and per-orbit as average θ of the Sentinel-1A and 1B observations for the year 2020. As is described in **Section 3.3.3** below, these mean PLIA values of the corresponding orbit are used together with the backscatter (σ^0) images and the parameters of the harmonic model, as input for the TUW flood mapping algorithm.

3.3.2 Computation of Sentinel-1 harmonic parameters

The 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 the sensor. The surface state (e.g. soil moisture content, vegetation, roughness) 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.

The TUW flood mapping algorithm utilizes a harmonic model to simulate the backscatter seasonal variation and to estimate normal, non-flooded conditions (see **Section 3.3.3** below). This section describes the preparation performed to define the harmonic model on a global scale.

Following the approach of Bauer-Marschallinger et al. (2022), the TUW algorithm uses a harmonic model to compute the most probable radar backscatter (σ^0) at time t_{day} (i.e. day of the year), based on the harmonic parameters ($C_i, S_i, \overline{\sigma^0}$). The terms and definitions used to compute the harmonic parameters (i.e. the harmonic coefficients and the average backscatter) are described in Table 15.

Table 15: Terms and definitions used to derive the parameters of the harmonic model describing the backscatter's seasonality for Non-flood conditions. For details, see Bauer-Marschallinger et al. (2022) and Schlaffer et al (2015).

TERM	DEFINITION
t_{day}	Day of the year, derived from the actual acquisition time (t) of the radar measurement.
$\hat{\sigma}_{t_{day}}^0$	The estimated (i.e. most probable) radar backscatter at time t_{day} (day of the year), calculated as: $\hat{\sigma}_{t_{day}}^0 = \overline{\sigma^0} + \sum_{i=1}^k (C_i * \cos(i * v) + S_i * \sin(i * v))$
$\overline{\sigma^0}$	The average backscatter for the period.
k	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.
v	$v = \frac{2\pi * t_{day}}{365}$
C_i and S_i	The i th harmonic cosine and sine coefficients (i.e. the harmonic parameters).
r	The relative Sentinel-1 orbit (i.e. 1 to 175).
s	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 (Npoints) adjusted for the degrees of freedom of the model, as follows: $s = \sqrt{\frac{SSE(\sigma_{t,r}^0, \hat{\sigma}_{t,r}^0)}{N_{points} - (2k + 1)}}$
$\sigma_{t,r}^0$	The pixel's backscatter time-series, corresponding to a relative Sentinel-1 orbit (r).
$\hat{\sigma}_{t,r}^0$	The pixel's harmonic model, corresponding to a relative Sentinel-1 orbit (r)
SSE	The sum of squared errors, derived from the pixel's backscatter time-series $\sigma_{t,r}^0$ and harmonic model $\hat{\sigma}_{t,r}^0$.

Figure 13 (panels a - d) show examples from the local non-flood distribution dataset for four selected days of the year, i.e. the expected backscatter for **relative orbit D080** resulting from the harmonic analysis defined in row 2 of Table 15. Figure 13 (panels e - f) show, for two relative Sentinel-1 orbits (descending D080, and ascending A175), example temporal plots of the expected local distribution over observed backscatter (in green), and the water distribution for the local incidence angle (in red), which is static over time, with two low spikes corresponding to flood events.

Note that for **Version 3.x** of the GFM product, the harmonic parameters were updated, to improve accuracy over agriculture, for example. After evaluating impacts for different periods, the 3-year baseline period of 2019-2021 was selected as an optimal dataset, providing a best parameterization of the TUW algorithm in terms of data availability, orbit stability, and robust performance.




Figure 13: Panels a - d: Examples of the local non-flood distribution dataset for four selected days of the year (day), over Thessaly, Greece; Panels e - f: Examples of temporal plots for two relative Sentinel-1 orbits. See text for details. (From Bauer-Marschallinger et al., 2022).

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. It should be noted that the derivation of the harmonic parameters within the GFM product differs from the method of Schlaffer et al (2015), as the backscatter values are used directly, instead of 10-day composites.

Backscatter coefficients are highly dependent on acquisition geometry. While a normalization approach could be utilized when sufficient incidence angle samples per pixel are present (as with ENVISAT ASAR), this is not feasible with Sentinel-1. Hence, parameter estimation is performed for each unique acquisition geometry, corresponding to a relative Sentinel-1 orbit. A unique set of harmonic parameters is required per orbit per Equi7grid tile. To model the estimated backscatter for any given day (t) and relative orbit (r), seven harmonic parameters are computed (for $k = 3$).

In order to measure 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. The standard deviation of the harmonic model is also calculated (see Table 15), based on the sum of squared errors derived from the pixel's backscatter time-series and its harmonic model.

3.3.3 Estimation of backscatter distribution functions for water and land surfaces

The backscatter probability distribution functions for the Flood and Non-flood classes are derived based on dedicated statistical parameters generated from the Sentinel-1 multi-year data archive (i.e. the datacube), as described below. These distribution functions are subsequently used as input for the Bayesian flood mapping and uncertainty estimation, to compute the posterior probabilities for the Flood and Non-flood classes, as described in **Section 3.3.4** below.

1. Flood backscatter probability density function:

Due to specular (or mirror-like) reflection of the radar pulses by water surfaces, the backscatter intensities received by the sensor are significantly lower compared with most other land cover types. A temporarily flooded surface is thus detectable by a significant decrease in its backscatter relative to the time-series. In order to ensure that the decrease is due to flooding, and no other effect, a detailed statistical knowledge of the backscatter behaviour over water surfaces is required.

The proposed concept is to define the **Flood backscatter probability density function** (PDF), for any given (projected local) incidence angle (θ). It is assumed that the Flood backscatter PDF is a normal distribution, and can be parametrized by a **mean** and a **standard deviation**. Therefore, various observed backscatter values for water bodies and their respective incidence angles (further referred to as $\sigma^0_{w,\theta}$) were collected from the Sentinel-1 datacube over oceans and inland waters. In Figure 14, the collected backscatter values ($\sigma^0_{w,\theta}$) are plotted against incidence angle, which confirms the expected indirect dependence of water backscatter values and incidence angles.

For each specific incidence angle, the mean backscatter value was derived by linear regression. The slope of the regression of mean backscatter values (red line in top part of Figure 14) is **0.394**. The slope of the regression of standard deviations (bottom part of Figure 14) is small (**0.008**), so "homoscedasticity" (i.e. same variance) is assumed. The total standard deviation is the square root of the sum of squares error (**SSE**), normalised with respect to the number of data points (**N_{points}**).




Figure 14: Top: Scatterplot of observed water backscatter values, grouped by 0.5 degree bins of incidence angles. Bottom: Standard deviation of these bins.

Therefore, the **Flood backscatter PDF** (referred to as $P(\sigma^0 \mid F)$) is defined by its mean backscatter value and standard deviation, which are calculated per incidence angle as follows:

$$\begin{aligned} \text{mean}(\sigma_{w,\theta}^0) &= (-0.394 * \theta) - 4.142 \text{ [dB]} \\ \text{std}(\sigma_{w,\theta}^0) &= \sqrt{\frac{\text{SSE}}{N_{\text{points}} - 2}} = 2.75 \text{ [dB]} \end{aligned}$$

2. Non-flood backscatter probability density function:

In order to model the normal, non-flooded conditions of each pixel in the Sentinel-1 SAR image datacube, we use the harmonic model described in **Section 3.3.2** above. The harmonic model can be used to retrieve the expected backscatter value ($\hat{\sigma}_{t,r}^0$) for any day of the year (t_{day}) and specific relative Sentinel-1 orbit (r) (see Table 15). Therefore, the backscatter's seasonality for each pixel in the datacube can be modelled by defining the harmonic parameters ($C_i, S_i, \bar{\sigma}_0$), based on the pixel's backscatter time-series (σ_r^0) (see Table 15). This harmonic model is used to define the **Non-flood probability density function** (PDF) for each pixel, which is assumed to be a normal distribution.

The mean backscatter value is set to the **expected backscatter of the harmonic model** ($\hat{\sigma}_{t,r}^0$). As well as the harmonic parameters, the **standard deviation** (s) of the harmonic model is also computed (see Table 15), and is used as the standard deviation of the Non-flood PDF. Therefore, the **Non-flood backscatter PDF** (i.e. $P(\sigma^0 \mid \text{NF})$), is defined by its mean backscatter value and standard deviation, which are calculated for each pixel in the Sentinel-1 SAR image datacube as follows (see Table 15):

$$\begin{aligned} \text{mean}(\sigma_{t,r}^0) &= \hat{\sigma}_{t,r}^0 \\ \text{std}(\sigma_{t,r}^0) &= s \end{aligned}$$

Finally, a redundancy criterion is used to reduce the possibility of ill-fitting harmonic parameters (especially for locations with sparse data). This criterion is derived from the NOBS layer (see **Section 3.3.2**) which indicates the number of valid observations used to estimate the harmonic parameters.

Given the requisite $[2k + 1]$ samples for a unique solution to the harmonic equation, the redundancy criterion is set as a multiple of this number. Since pixels not matching the redundancy criterion are excluded from the flood mapping, a balance must be found between number of excluded pixels and introduced noise. Based on initial tests, a redundancy value of 4 (e.g. 28 samples per pixel stack), can minimize noise, specifically at edges where sparse samples might be present.

3.3.4 Bayesian flood mapping and uncertainty estimation

Given the per-pixel knowledge of the relative orbit (r), day of year (t_{day}), and incidence angle (θ), for each SAR image pixel one can define the parameters (mean and standard deviation) of the backscatter **probability density functions** (PDFs) for the Flood (F) and Non-flood (NF) classes (i.e. $P(\sigma^0 | F)$ and $P(\sigma^0 | NF)$, respectively). For each pixel in the new Sentinel-1 SAR image, the **posterior** (i.e. updated) **probabilities** of belonging to the Flood and Non-flood classes can then be computed, using Bayes theorem, based on the pixel's backscatter value (σ^0) and the Flood and Non-flood PDFs.

Recall that Bayes' theorem is used to determine the posterior probability of a hypothesis (in our case Flood or Non-flood) being true, given new evidence (in our case backscatter value). The computation of the posterior probabilities is summarized in Table 16.

Table 16: Terms and definitions used to compute the posterior (i.e. updated) probabilities of pixels in a new Sentinel-1 SAR image belonging to the flood (F) and non-flood (NF) classes.

TERM	DEFINITION
σ^0	Backscatter coefficient value (in dB) in SAR image.
$P(F \sigma^0)$	Posterior (updated) probability of flood class , inferred (by Bayes' theorem) as: $P(F \sigma^0) = P(F) * P(\sigma^0 F) / P(\sigma^0)$
$P(\sigma^0 F)$	Conditional probability (or likelihood) of flood class (from its distribution function).
$P(F)$	Prior (independent) probability of flood class (set to 0.5).
$P(NF \sigma^0)$	Posterior (updated) probability of non-flood class , inferred (by Bayes' theorem) as: $P(NF \sigma^0) = P(\sigma^0 NF) * P(NF) / P(\sigma^0)$
$P(\sigma^0 NF)$	Conditional probability (or likelihood) of non-flood class (from its distribution function).
$P(NF)$	Prior (independent) probability of non-flood class (set to 0.5).
$P(\sigma^0)$	Marginal (independent) probability of backscatter coefficient value in image I_{t0} , calculated as: $P(\sigma^0) = [P(\sigma^0 F) * P(F)] + [P(\sigma^0 NF) * P(NF)]$
$P(error \sigma^0)$	Conditional error (uncertainty) defined by lower posterior probability (0.0 - 0.5), calculated as: $P(error \sigma^0) = \min[P(F \sigma^0), P(NF \sigma^0)]$

Note that the **prior probabilities** of the Flood and Non-flood classes (i.e. $P(F)$ and $P(NF)$) shown in Table 16 represent the prior knowledge of the pixel belonging to a certain class. As no prior information is available, both probabilities are set to equal weight (i.e. **0.5**). Furthermore, the **marginal** (or independent) **probability** of the backscatter value (i.e. $P(\sigma^0)$) has the effect of scaling the values of the posterior probabilities to the range 0 to 1.

Bayes decision rule is then used to assign the pixel to the class (Flood or Non-flood) with the higher computed posterior probability. An example of the Bayesian flood mapping process for a single SAR image pixel, is illustrated in Figure 15, which includes both the probability distribution functions (i.e. $P(\sigma^0 | F)$ and $P(\sigma^0 | NF)$) and the posterior probabilities (i.e. $P(F | \sigma^0)$ and $P(NF | \sigma^0)$) for the Flood and Non-flood classes. As can be seen, the given backscatter value (-19.5dB) would be classified as flooded, since it has a higher computed posterior probability for the Flood class ($P(F | \sigma^0)$). Finally, a preliminary binary flood extent map is generated containing all pixels classified as flooded.




Figure 15: Illustration of the Bayesian flood mapping for a SAR image pixel (with backscatter value = -19.5dB), showing the probability distribution functions and posterior probabilities for Flood (F) and Non-flood (NF). The pixel is classified as flooded. (See text for details).

As can be seen in Table 16, the Bayesian flood mapping algorithm also provides the conditional error (i.e. $P(\text{error} | \sigma^0)$), which directly quantifies the uncertainty of the classification. It is the lower posterior probability of the two classes (i.e. of the Non-flood class), and consequently its values are defined between 0.0 and 0.5, which are interpreted as follows:

- An uncertainty close to zero implies a very confident decision, since the probabilities for both classes (flood and non-flood) support a clear decision.
- An uncertainty of close to 0.5 implies the opposite, i.e. the probabilities of the input Sentinel-1 SAR backscatter value belonging to the two classes are nearly identical.
- In the case of high conditional errors (i.e. close to 0.5), no certain decision can be made. In order to exclude such dubious decisions (i.e. with low reliability), a dedicated uncertainty mask is applied in the TUW flood mapping algorithm (see Section 3.3.5 below).

3.3.5 Low sensitivity masking

To improve the reliability of the preliminary flood extent map generated automatically in near real-time by the TUW algorithm, a set of routines is applied to exclude those “low sensitivity” locations where the Bayes flood mapping does not allow a robust decision between flood and non-flood conditions. Therefore, four masks are generated that exclude pixels with (1) **exceeding incidence angles**, (2) **conflicting backscatter distributions**, (3) **outlier values in backscatter distributions**, and (4) **high uncertainties of Bayesian flood mapping**. Figure 16 shows an example of the four low sensitivity masks. The generation of the masks is summarized in Table 17, and described below.




Figure 16: Example over Greece showing the four low sensitivity masks applied to the preliminary flood extent map produced by the TUV flood mapping algorithm.

Table 17: Terms and definitions used for the low sensitivity masking of the preliminary flood extent map produced by the TUV flood mapping algorithm.

LOW SENSITIVITY MASK	DEFINITION
Mask for exceeding incidence angles (M_{PLIA})	$M_{PLIA} = \theta < 27^\circ \text{ and } > 48^\circ$ Where θ is the (projected local) incidence angle for a pixel.
Mask for conflicting distributions (M_{cd})	$M_{cd} = \text{mean}(\sigma_{t,r}^0) < [\text{mean}(\sigma_{w,\theta}^0) + \frac{1}{2} \text{std}(\sigma_{w,\theta}^0)]$ Where $\sigma_{t,r}^0$ is a pixel's backscatter time-series for a given relative orbit (see Error! Reference source not found.), and $\sigma_{w,\theta}^0$ is a pixel's mean backscatter value for water surfaces at a given θ.
Mask for outliers for Flood class (M_{out}^F)	$M_{out}^F = \sigma^0 > [\text{mean}(\sigma_{w,\theta}^0) + 3\text{std}(\sigma_{w,\theta}^0)]$
Mask for outliers for Non-flood class (M_{out}^{NF})	$M_{out}^{NF} = \sigma^0 < [\text{mean}(\sigma_{t,r}^0) - 3\text{std}(\sigma_{t,r}^0)] \text{ or } [> \text{mean}(\sigma_{t,r}^0) + 3\text{std}(\sigma_{t,r}^0)]$ Where $\text{mean}(\sigma_{t,r}^0)$ and $\text{std}(\sigma_{t,r}^0)$ are, respectively, the mean value and standard deviation of a pixel's backscatter values in the datacube.
Mask for outliers (M_{out})	$M_{out} = M_{out}^F \text{ and } M_{out}^{NF}$
Mask for high uncertainties (M_{cert})	$M_{cert} = P(\text{error} \sigma^0) > 0.2$ Where $P(\text{error} \sigma^0)$ is the uncertainty of a pixel's backscatter value (see Table 11).

1. Mask for exceeding incidence angles:

Flat areas are observed by Sentinel-1's IW mode at (projected local) incidence angles within the range $29^\circ - 46^\circ$. By definition, this range includes (flat) water surfaces, and therefore the collection of water backscatter samples (described in Section 3.1.3.4 above) was limited to this range. In order to extend the flood mapping to gentle slopes, the range of incidence angles was relaxed to 27° to 48° , by applying a mask of exceeding incidence angles (see Table 17).

2. Mask for conflicting distributions:

The TUW flood mapping algorithm detects if a pixel that is normally non-flooded is temporarily flooded, implying that the pixel normally has a higher backscatter than a respective water surface. In other words, the Non-flood probability distribution function (PDF) must be higher overall than the Flood PDF. Typical locations where this condition is NOT fulfilled are permanent water, asphalt surfaces, salt flats, and very dry sand or bedrock areas. In addition to these “all-season” low backscatter conditions (permanent water, water look-alikes, etc.), a pixel’s backscatter may be low only for some seasons of the year, for example due to the rainy season, snow melting, etc. In order to exclude such areas showing conflicting distributions, a pixel is masked if the mean of the Non-flood PDF is lower than that of the Flood PDF plus half the standard deviation of the Water PDF (see Table 17). Figure 17 shows an example of a pixel that is masked due to conflicting distributions.




Figure 17: Illustration of the Bayesian flood mapping for a SAR image pixel (backscatter value ≈ -17.9 dB), showing the probability distribution functions and posterior probabilities for Flood (F) and Non-flood (NF). This pixel is masked due to conflicting distributions. (See text for details).

The consideration of such seasonal effects is an important feature of the harmonic model approach used by TUW algorithm (see Section 3.3.2). If the input observations for the harmonic model regularly show low backscatter for a certain season, the corresponding expected backscatter for a given day of the year should represent this. This is illustrated in Figure 18 for an area (near the city of Angers, France), where the expected backscatter for a given day of the year in 2022 shows a large area of low backscatter, which may be related to seasonal inundation. As can be seen in Figure 18, a similar pattern of inundations is observed also for other recent years. In this case, the harmonic model was trained within the current version (v3.x) of the GFM product, using data from 2019-2021, and appears to reflect the local dynamics in that perspective.




Figure 18: Example (near Angers, France) showing the Non-flood backscatter values, for a given day of the year, predicted by the harmonic model for 2022 (a), compared with those observed for 2015 (b), 2016 (c), 2017 (d), 2018 €, 2019 (f) and 2020 (g). (See text for details).

3. Mask for outliers:

The statistical model used by the TUW flood mapping algorithm is trained to provide robust decisions for normal flood conditions. Therefore, if the input backscatter value (σ^0) is NOT represented by the Flood or Non-flood PDFs, the Bayes decision will not be reliable. In order to mask extreme values, only high outliers of the Flood PDF are considered, as it is assumed that low outliers of the Flood PDF can still be classified as flooded. The rule used for masking outliers is defined in Table 17. Figure 19 shows an example of such a situation with an outlier.




Figure 19: Illustration of the Bayesian flood mapping for a SAR image pixel (backscatter value ≈ -6 dB), showing the probability distribution functions and posterior probabilities for Flood (F) and Non-flood (NF). This pixel is masked as an outlier. (See text for details).

4. Mask for high uncertainties:

The conditional error produced by the Bayesian flood mapping (see Table 16) provides a measure of the confidence in the decision between Flood and Non-flood, and is used to limit the classification to certain decisions. Masking for high uncertainties is done by excluding pixels with $P(\text{error} \mid \sigma^0) > 0.2$ (see Table 17). Figure 20 shows an example of the masking of an uncertain decision. As can be seen, the backscatter value is almost equally likely to be classified as Flood or Non-flood.




Figure 20: Illustration of the Bayesian flood mapping for a SAR image pixel (backscatter value ≈ -17.5 dB), showing the probability distribution functions and posterior probabilities for Flood (F) and Non-flood (NF). This pixel is masked due to a high uncertainty. (See text for details).

3.3.6 Final post-processing of the TUV flood map

Due to the coherent nature of the radar signal, a single Sentinel-1 image is affected by multiplicative noise (speckle). This random signal variation can lead to a noise-like pattern of pixels with lower or higher backscatter compared to their surroundings. Lower backscatter might be confused with flooded conditions, while higher backscatter could prevent a correct flood classification. With GFM v4.0, to reduce the influence of speckle on the flood extent map, a minimum mapping unit (MMU) of 17 pixels is applied and holes smaller than 7 pixels are filled. The result of this step represents the final flood extent and estimated uncertainty of the TUV flood mapping algorithm.

3.3.7 Computation of Likelihood Values by the TUV flood mapping algorithm

The TUV flood mapping algorithm estimates the **uncertainty** of flood classification ($P(\text{error} \mid \sigma^0)$) for each grid-cell, as the minimum of the two posterior probabilities of the flood and non-flood classes ($P(F \mid \sigma^0)$ and $P(NF \mid \sigma^0)$) (see Table 16). In order to be used by the GFM Ensemble flood mapping algorithm (see **Section 3.4** below), these uncertainty values are “flipped”, so that low uncertainty values propagating towards 0 (indicating high likelihood of non-flood) are remapped to high likelihood values propagating towards 100 (indicating high likelihood of flood), and vice versa.

3.4 The GFM Ensemble flood mapping algorithm

The GFM Ensemble flood mapping algorithm combines the flood classification results and likelihood values generated by the three individual flood mapping algorithms (LIST, DLR and TUW) for each input Sentinel-1 image grid-cell. The GFM Ensemble algorithm first assigns to each grid-cell a ratio equal to the proportion of individual algorithms providing inputs that classified the grid-cell as **Flooded**. A majority-based consensus rule is then used to classify the grid-cell as **Flooded** (i.e. ratio > 0.5), **Unflooded** (i.e. ratio ≤ 0.5), or **Unclassified** (i.e. no individual algorithm provided an input). The ten possible cases for classifying a grid-cell with the GFM Ensemble algorithm are shown in Table 18. The GFM Ensemble algorithm also computes the Likelihood Value for each grid-cell, as the average Likelihood Value computed by the individual algorithms providing inputs for that grid-cell.

Table 18: The ten possible cases of classifying a grid-cell with the GFM Ensemble flood mapping algorithm, using a majority-based consensus (i.e. ratio > 0.5).

CASE	RATIO	INDIVIDUAL ALGORITHMS PROVIDING INPUTS	INDIVIDUAL ALGORITHMS CLASSIFYING GRID-CELL AS FLOODED	ENSEMBLE RESULT
1	1.00	3	3	FLOODED
2	1.00	2	2	FLOODED
3	1.00	1	1	FLOODED
4	0.67	3	2	FLOODED
5	0.50	2	1	UNFLOODED
6	0.33	3	1	UNFLOODED
7	0.00	3	0	UNFLOODED
8	0.00	2	0	UNFLOODED
9	0.00	1	0	UNFLOODED
10	-	0	0	NOT CLASSIFIED

The GFM Ensemble flood mapping algorithm accepts Geotiff files containing the flood classification results and likelihood values from each individual algorithm that provided inputs. The Ensemble algorithm checks if all layers can be loaded. If a layer from one algorithm cannot be loaded, it is likely that that algorithm failed to produce a result. Therefore, the number of flood layers loaded successfully gives the number of algorithms that provided inputs. For example, if two layers load successfully, and one layer fails to load, the number of individual algorithms applied is two.

Five examples of the behaviour of the GFM Ensemble algorithm in computing flood classification results and Likelihood Values, are illustrated in Figure 21 to Figure 25, and are summarized below:

#	EXAMPLE	COMMENTS
1	Inputs from all three individual algorithms, and full agreement: (See Figure 21).	- All three individual algorithms return the same result (Flooded). - Ensemble result: Flooded. - Likelihood Value: Average likelihood from three individual algorithms.
2	Inputs from all three individual algorithms, and partial agreement: (See Figure 22).	- All three individual algorithms provide inputs. Two algorithms classify the grid-cell as Unflooded. - Ensemble result: Unflooded. - Likelihood Value: Average likelihood from three individual algorithms.

3	Inputs from only two individual algorithms, and agreement: (See Figure 23).	▪ Exceptional case where one individual algorithm fails to provide inputs. ▪ The other two individual algorithms agree that the grid-cell is Flooded. ▪ Ensemble result: Flooded. ▪ Likelihood Value: Average likelihood from two individual algorithms.
4	Inputs from only two individual algorithms, and disagreement: (See Figure 24).	▪ Exceptional case where one individual algorithm fails to provide inputs. ▪ The other two individual algorithms disagree. ▪ Ensemble result: Flooded. ▪ Likelihood Value: Average likelihood from two individual algorithms. ▪ Note: This may result in the unusual event (depicted here) where an Unflooded grid-cell has a Likelihood Value ≥ 50.
5	Inputs from only one individual algorithm: (See Figure 25).	▪ Exceptional case where two individual algorithms fail to provide inputs. ▪ The other algorithm classifies the grid-cell as Flooded. ▪ Ensemble result: Flooded. ▪ Likelihood Value: Likelihood from the individual algorithm.




Figure 21: Flood classification and likelihood value computed by the GFM Ensemble algorithm, with inputs from all three individual algorithms (LIST, DLR, TUW): Ensemble result = Flooded (i.e. ratio = 1.00).

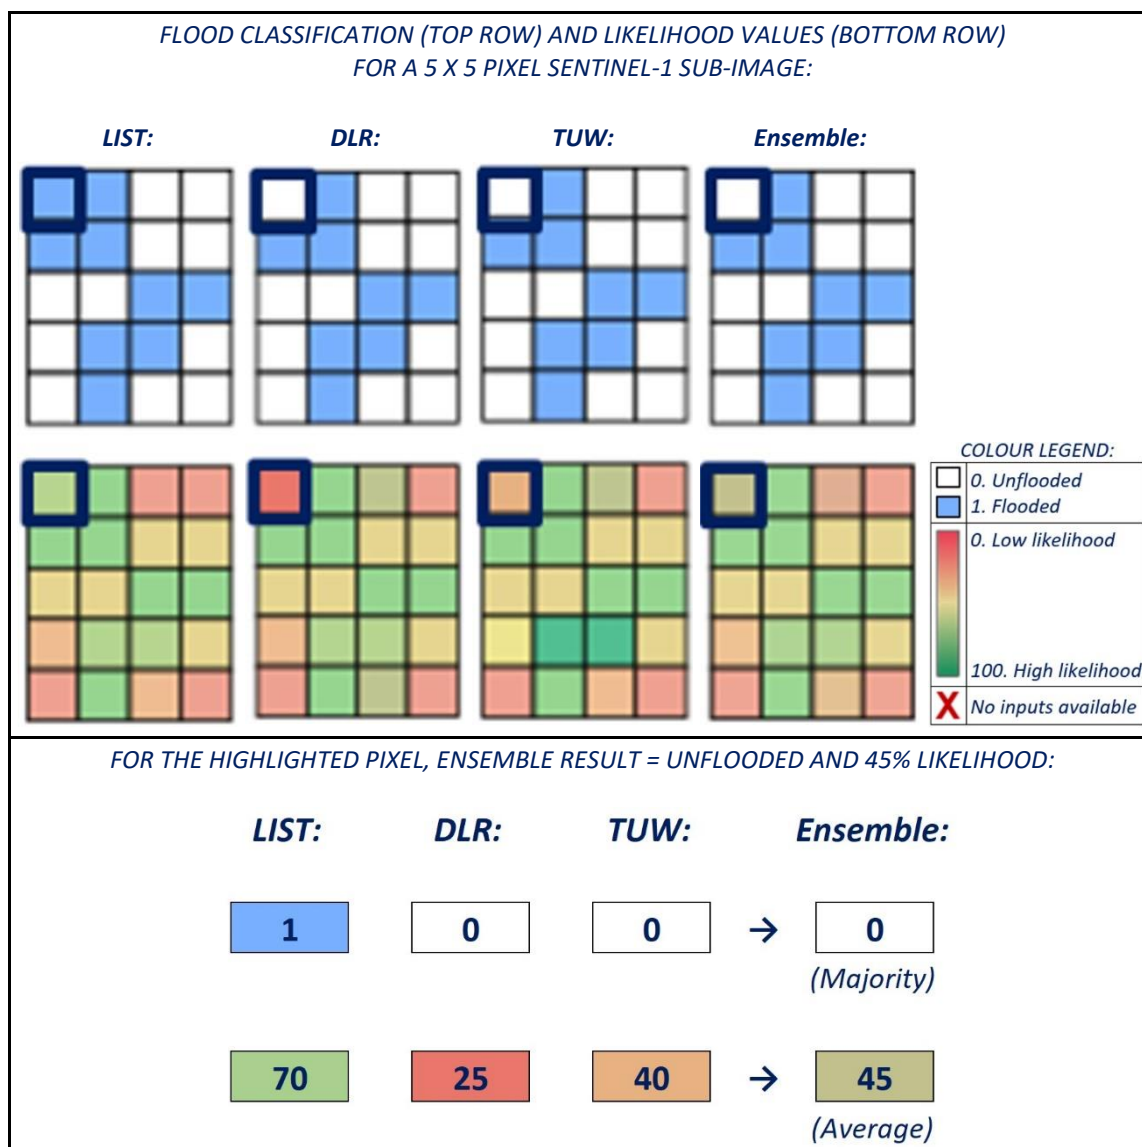


Figure 22: Flood classification and likelihood value computed by the GFM Ensemble algorithm, with inputs from all three individual algorithms (LIST, DLR, TUW): Ensemble result = Unflooded (i.e. ratio = 0.33).

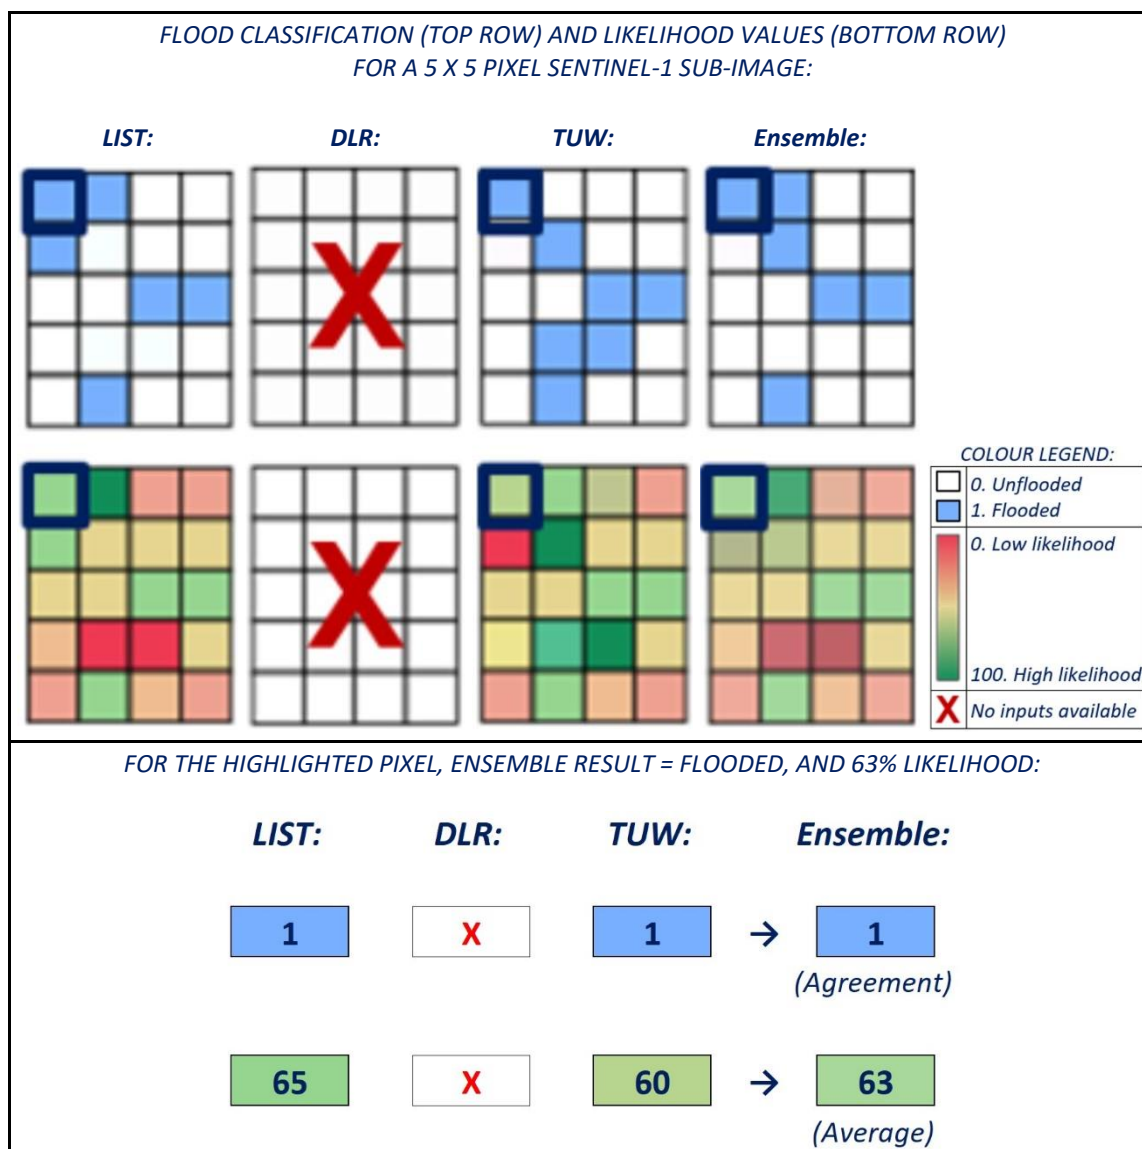


Figure 23: Flood classification and likelihood value computed by the GFM Ensemble algorithm, with inputs from only two individual algorithms (LIST, TUW): Ensemble result = Flooded (i.e. ratio = 1.00).




Figure 24: Flood classification and likelihood value computed by the GFM Ensemble algorithm, with inputs from only two individual algorithms (LIST, TUW): Ensemble result = Unflooded (i.e. ratio = 0.50).




Figure 25: Flood classification and likelihood value computed by the GFM Ensemble algorithm, with inputs from only one individual algorithm (DLR): Ensemble result = Flooded (i.e. ratio = 1.00).

4 Generation of the ten GFM output layers of flood information

The Global Flood Monitoring (GFM) product of CEMS provides a continuous monitoring of floods worldwide, by processing and analysing in near real-time all incoming Sentinel-1 SAR images acquired in Interferometric Wide Swath (IW) mode, and generating the ten output layers of flood-related information (see Table 1). In this Section, the key technical features underlying the generation of the ten GFM output layers of flood information, are described in detail.

4.1 GFM output layer: Observed Flood Extent

The GFM output layer Observed Flood Extent shows the flooded areas which are mapped in near real-time from Sentinel-1 satellite imagery, using three individual, state-of-the-art flood mapping algorithms (LIST, DLR, and TUW). Briefly, for each input Sentinel-1 SAR image scene, the three independent flood maps that are generated by the three individual flood mapping algorithms are combined into a single Observed Flood Extent output layer, using the GFM Ensemble flood mapping algorithm. Detailed technical descriptions of the LIST, DLR, TUW, and Ensemble flood mapping algorithms, which form the core of the GFM product, are provided in Section 3 of this PDD.

The results of the individual flood mapping algorithms are harmonized using the GFM Reference Water Mask and Exclusion Mask, respectively, to remove permanent (and seasonal) water bodies, and to exclude surface types where SAR-based water mapping is unreliable. The Observed Flood Extent and Likelihood Values are corrected by changing **Flooded** grid-cells included in the Reference Water Mask, to **Unflooded**. The Copernicus Water Body Mask is used to mask out ocean areas. The GFM Reference Water Mask and Exclusion Mask are described in **Sections 4.3** and **4.4** below.

4.2 GFM output layer: Observed water extent

The GFM output layer Observed Water Extent identifies grid-cells classified as open and calm water based on Sentinel-1 SAR backscatter intensity, and is derived by combining (using a logical OR) the flood map generated by the GFM ensemble algorithm (see **Section 3.4** above) with the seasonal and permanent water of the GFM Reference Water Mask (see **Section 4.3** below). The values of the GFM output layer Observed Water Extent are defined in Table 19.

Table 19: Data values of the GFM output layer Observed Water Extent.

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 GFM output layer: Reference Water Mask

The GFM output layer Reference Water Mask outlines grid-cells that are classified as both permanent and seasonal water, using an ensemble water mapping algorithm (described below), and based on Sentinel-1 SAR median backscatter intensity over a certain time period, defined as follows:

- Permanent water extent is mapped based on each grid-cell's median backscatter of all Sentinel-1 data over a reference period of five years.
- Seasonal water extent is based on each grid-cell's median backscatter of all Sentinel-1 data for a given month over the same reference period.

Therefore, 12 monthly reference water masks are created, providing information on permanent and seasonal reference water extent.

The use of a permanent water mask is reliable in environments with stable hydrological conditions over the year. In regions with strong hydrological dynamics and spatio-temporal variations in observed water extent (e.g. areas affected by monsoons), the additional use of a seasonal reference water mask is desirable, as it also considers the seasonality of the observed water extent.

Unlike the GFM Ensemble flood mapping algorithm, the Ensemble water mapping algorithm uses only the LIST and DLR algorithms. The TUW algorithm is specifically designed to detect temporary flooded areas. It does not produce a water extent map, and is unsuited for the generic task of water mapping. The key elements of the Ensemble water mapping algorithm are summarized below:

- | |
|---|
| <ul style="list-style-type: none"> ▪ The GFM Ensemble reference water algorithm combines the water and uncertainty results produced by the LIST and DLR flood mapping algorithms (see Sections 3.1 and 3.2), and having as input the median backscattering value of Sentinel-1 images over (a) a five-year period (for permanent water) and (b) the same time period per month (for seasonal water). The results of the single water maps are combined into twelve masks, one per month, which include on the mapped permanent and seasonal water extent. |
| <ul style="list-style-type: none"> ▪ Both the LIST and DLR flood mapping algorithms produce water maps and associated likelihoods, as a by-product of the flood detection. In particular, the implementation of the LIST flood mapping algorithm for Case 2 (see Section 3.1.3 above) is used, which maps the water extent using a single image and a region-growing. |
| <ul style="list-style-type: none"> ▪ The GFM Ensemble reference water algorithm follows the same logic as the GFM Ensemble flood mapping algorithm (see Section 3.4 above), but is specifically targeted for mapping water from two (versus three) input sources. A consensus decision of the two algorithms decides if a grid-cell is marked as water or as no water, as explained below: <ul style="list-style-type: none"> – To generate the combined product, each grid-cell is attributed with the ratio of the number of water classifications to the number of algorithms that were applied. A ratio of 1 means that both algorithms agreed on the water classification. – For each grid-cell where the two algorithms disagree (i.e. one classified it as water and the other classified it as no water), the algorithm with the highest likelihood value for that classification dictates the result. – Therefore, if one algorithm is certain that a pixel is water (high likelihood) and the other algorithm classified the pixel as no water with a low likelihood value, the algorithm classifying that grid-cell as water dictates the final result. The likelihood value is the average of the provided likelihood values for that grid-cell. – If only one of the two algorithms provides a result for a grid-cell, the final water classification and likelihood values are those from that particular algorithm. |
| <ul style="list-style-type: none"> ▪ Finally, for selected Equi7Grid-tiles where misclassifications are observed, the following post-processing is applied: <ul style="list-style-type: none"> – An exclusion layer is applied that corrects grid-cells that were misclassified due to radar shadow and water-look-alikes. To remove such false positives, the results of the water algorithm are masked with the inverse of a buffered (100 m) version of the “Maximum Extent Layer” of the Global Surface Water product (Pekel et al. 2016). This removes false positives outside the maximum water extent over a 30-year time period observed by Landsat. – The Copernicus Water Body Mask is used to fill larger patches of false negatives (e.g. large lakes with roughened surface, mis-classified as land). The Water Body Mask is also applied to enforce a consistent land-sea border. |

Figure 26 shows an example (in Myanmar) of the GFM monthly reference water masks. Seasonal (i.e. monthly) variations of water bodies can be separated from permanent water bodies, which do not change their extent during the reference time-period. In Figure 26, each mask separates permanent water (i.e. grid-cells classified as water during the whole reference period) and seasonal water (i.e. grid-cells classified as water during a specific month within the reference period).




Figure 26: Examples of GFM Reference Water Mask output layers computed over Myanmar, for the period 2019-01-01 to 2020-12-31. (See text for details).

The GFM Reference Water Mask is a static dataset that is computed once for a reference period of five years, covering the years 2017 to 2021. A reference period of five years is chosen as a compromise, to have enough Sentinel-1 data for computation, and to remove long-term hydrologic changes that might be integrated in the Reference Water Mask if longer time periods were used. Ending the reference period in 2021 is crucial to also include data from Sentinel-1 B.

The data values of the GFM Reference Water Mask are defined in Table 20.

Table 20: Data values of the GFM output layer Reference Water Mask.

VALUE	DESCRIPTION	USER INTERPRETATION
0	No water	Location normally not affected by the presence of water.
1	Permanent water body	Location permanently covered by water.
2	Seasonal water body (for the current month)	Location seasonally covered by water.
NaN	No data available	Location not covered by the Sentinel-1 satellite.

4.4 GFM output layer: Exclusion Mask

The identification of conditions under which the SAR-based flood mapping algorithms are deemed reliable, is an essential element of the GFM product, or indeed of any automated global flood monitoring system. This consideration is therefore an integral part of the near real-time flood detection and final preparation of the results of the GFM product.

The central assumption underlying the successful SAR-based identification and delineation of water bodies, and subsequently of flood extent, is the low backscatter signature of water surfaces. Under normal conditions, water surfaces observed by the Sentinel-1 SAR sensor show significantly lower backscatter values than surrounding areas, and the desired discrimination between water and non-water surfaces can be achieved with high accuracy and reliability. The inclusion of historical information stored in the backscatter “datacube” (time-series), which is a key element of the GFM product, provides additional discrimination criteria, further improving the GFM product’s reliability.

Nonetheless there are several surface types where SAR-based delineation of flood extent is hampered, or completely impeded. In order to prevent misclassification, such surface types are excluded from the flood mapping results, using the GFM Exclusion Mask. The GFM 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 SAR, over certain land cover types and terrain conditions, for physical reasons. The main surface types that are excluded from the GFM flood mapping results, are defined in Table 21.

Table 21: Surface types that are excluded from the GFM results using the Exclusion Mask.

#	SURFACE TYPE	DESCRIPTION
1.	<i>No-sensitivity:</i>	For certain surfaces, such as urban areas and dense vegetation, 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, there is no sensitivity for the SAR-based detection of water surfaces.
2.	<i>Water look-alikes:</i>	For certain surfaces, such as flat impervious areas and sand surfaces, Sentinel-1 senses the ground, but the backscatter of the non-flooded surface is in general so low as to be indistinguishable from backscatter of smooth open water.
3.	<i>Topographic distortions (and low probability of flood occurrence):</i>	In areas of strong topography, 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. Flooding anyway commonly appears over non-sloped surfaces (e.g. in the vicinity of streams, rivers, and other permanent water bodies).
4.	<i>Radar shadow:</i>	In some areas, Sentinel-1 receives no signals from certain surfaces due to radar shadows cast by mountains, high vegetation canopies or anthropogenic structures.
5.	<i>Low Sentinel-1 coverage:</i>	In some areas, mainly the far North of Siberia and Canada, there is low (or no) coverage by Sentinel-1, and so there is an inadequate historical SAR time-series. Areas with a lack of Height above Nearest Drainage (HAND) information represent an additional problem.

The GFM Exclusion Mask is generated by combining sub-layers that address distinct physical and geometric effects, while ignoring permanent or seasonal water pixels from the reference water mask. In addition, no-data values from the Ensemble flood mapping algorithm indicate on a binary map grid-cells where SAR data could not deliver the information needed for robust flood detection.

Considering the complexity of the physical and geometrical effects described in Table 21, and their potential spatial overlaps and interactions, one combined Exclusion Mask is provided that is easy for users of the GFM product to interpret. The Exclusion Mask is based on offline-generated Sentinel-1 parameters and auxiliary thematic datasets. It is accessed in near real-time, and provided in addition to the other GFM product output layers. The necessary operations are done at the 20m-sampling of the Sentinel-1 preprocessed data cube. The methods used to generate the sub-layers of the GFM Exclusion Mask, corresponding to the surface types in Table 21, are described below.

4.4.1 GFM Exclusion Mask sub-layer: No-sensitivity areas

Floods in urban or vegetated areas are at risk of being missed alarms in terms of flood detection, as co-located flood extents potentially feature high instead of low backscatter. The formation of corner reflectors for the SAR 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 no-sensitivity sub-layer of the GFM Exclusion Mask delineates all land cover types and areas where Sentinel-1 C-band SAR is not sensitive to flooding - or any other type of change - of the ground surface. These no-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 densely vegetated areas, four parameters from the Sentinel-1 Global Backscatter Model (S1-GBM) were used: mean backscatter in VV ($\mu_{VV}^{\sigma^0}$) and VH ($\mu_{VH}^{\sigma^0}$) polarisation, mean cross polarisation ratio ($\mu_{CPR}^{\sigma^0}$, computed as VV/VH backscatter) and standard deviation in VH polarisation ($std(\sigma_{VH}^0)$). Locally variable thresholds (τ_i) were defined for each parameter to discriminate vegetated from not vegetated areas, with grid-cells classified as dense vegetation as follows:

PARAMETER	RULE WITH LOCAL THRESHOLD (τ_i)
▪ Mean backscatter (σ^0) in VV polarisation:	$mean(\sigma_{VV}^0) > \tau_{\mu-VV}$
▪ Mean backscatter (σ^0) in VH polarisation:	$mean(\sigma_{VH}^0) < \tau_{\mu-VH}$
▪ Mean cross polarization ratio:	$mean(\sigma_{CPR}^0) > \tau_{\mu-CPR}$
▪ Standard deviation in VH polarization:	$std(\sigma_{VH}^0) > \tau_{\sigma-VH}$

Thresholds were optimized separately per continent, and varied with northing and vegetation type. For the updated GFM product v3.x, the Global Forest Change dataset (Hansen et al., 2013; v1.10)⁵ was used as a reference dataset for threshold selection, and used directly as a dense vegetation mask in regions where irregular Sentinel-1 coverage caused artefacts (e.g. stripes of high and low backscatter) in the S1GBM layers. The global distribution of thresholds is shown in Figure 27.




Figure 27: Global distribution of thresholds used to identify densely vegetated areas using S1-GBM and Global Forest Change dataset.

Over urban areas backscatter measured by SAR C-band 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, in GFM v3.x products the GHSL R2023A E2018 data (at 10m)⁶ was used. The urban areas were defined as grid-cells where the GHS-BUILD value exceeds 35% probability of built-up area.

Finally, the dense vegetation mask and urban areas masks are combined with a logical “or” operation into one binary No-sensitivity mask.

4.4.2 GFM Exclusion Mask sub-layer: Water look-alikes

There is another situation where Sentinel-1 is not sensitive to the flooding of the ground, but for quite a different reason from that for vegetation and urban areas. Here the flood detection is challenged by non-water surfaces that have low backscatter themselves and would yield false alarms if not addressed. For example, very dry or sandy soils (e.g. on riverbanks or in deserts), frozen ground, wet snow and flat impervious areas (e.g. smooth tarmac such as airports, motorways, etc.) feature very low backscatter signatures and appear as water-look-alikes in SAR imagery.

⁵ <https://storage.googleapis.com/earthenginepartners-hansen/GFC-2022-v1.10/download.html>

⁶ https://ghsl.jrc.ec.europa.eu/ghs_buS2023.php

In this case, the lack of sensitivity is due to the fact that backscatter from surfaces that are smooth (compared to the 5.5 cm wavelength of Sentinel-1's C-band SAR sensor) can be as low as backscatter from a water surface. Smooth non-water surfaces may act like water as a specular reflector, scattering most energy of the SAR signal in forward direction, leading to very low backscatter values.

Separating water and smooth land surfaces based on a single SAR acquisition is hardly feasible. Therefore, the generation of a water look-alike sub-layer with permanent low backscatter is based on statistical information from Sentinel-1 time-series data. Grid-cells are masked if the occurrence of low backscatter values (below -15 dB) exceeds 70% of all values in the respective time-series. For the GFM product v3.x, the corresponding parameter baseline period for the determination of non-water low backscatter areas was set to 2020-2022, to reflect most recent conditions and to take into account recent land cover changes, e.g. ongoing desertification in arid regions.

4.4.3 GFM Exclusion Mask sub-layer: Topographic distortions

The consideration of local topography is essential for SAR-based flood detection in two ways. Firstly, there are areas that must be excluded during flood detection due to the geometry of the SAR observation. This is dealt using measures discussed in **sub-section 4.4.4** below. Secondly, SAR signal disturbances over areas of strong topography can be substantial, and can dominate the surface backscatter signature, and its change over time, leading to otherwise uncontrollable over- and under-estimation of flood extent. For this second case, consideration of the Height above Nearest Drainage (HAND) index is of vital importance to the reliability of the flood detection algorithms.

A binary Exclusion Mask, based on the HAND index (HAND-EM), has been calculated to separate flood- from non-flood prone areas, based on the elevation of each grid-cell to the nearest water-covered grid-cell ([Twele et al., \(2016\)](#); [Chow et al. \(2016\)](#)). This is used as an input sub-layer of the GFM Exclusion Mask. Both binary classes are determined using an appropriate threshold value. A threshold value that is too high may lead to misclassifications (i.e. inclusion of flood look-alikes in areas much higher than the actual flood surface and drainage network). A threshold value that is too low may eliminate valid parts of the flood surface. Selecting an appropriate threshold is thus critical, and was carried out 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 scope of the GFM product's Sentinel-1 flood detection algorithms, a threshold value of ≥ 10 m was finally chosen to define 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. On a final note, since the Copernicus DEM is not hydrologically conditioned, the HAND index cannot be computed using this version of the DEM.

4.4.4 GFM Exclusion Mask sub-layer: Radar shadow

Radar shadow refers to areas that are unreachable by a radar signal, so no information can be gained. Radar shadow is a common effect in SAR imaging, appearing over strong terrain (especially at the far-range section of SAR images), and near high objects (e.g. mountains, high anthropogenic structures, forest borders). While the former are easily recognised with terrain and observation geometry analysis, the latter can only be identified using backscatter time-series, due to the global scope of the GFM product, and as real Digital Surface Models are only available for some areas.

The extent of radar shadow increases with radar range, so the far-range section of a SAR scene is more affected. After preprocessing, areas of radar shadow have very low backscatter values, so are likely to be confused with water surfaces, potentially leading to false alarms in flood detection. Two methods are used to provide radar shadow as one sub-layer of the GFM Exclusion Mask:

1	Statistical orbit-based approach:	This method makes use of the fact that a 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. By comparing mean backscatter at one orbit with mean backscatter from orbits in the opposite direction, the extent of the orbit-specific shadow areas can be estimated. For a given orbit (which feature a very stable viewing geometry), pixels will be masked as shadow, when the mean of backscatter at current orbit < -15 dB and mean of backscatter from orbits in the opposite direction > -10 dB (see also Figure 28).
2	Geometric DEM-based approach:	This method detects shadows by analyzing the SAR imaging geometry with respect to a DEM, i.e. the 30m Copernicus DEM (CopDEM). This can be computed offline with a geometric observation model of the Sentinel-1 orbits, without using the Sentinel-1 SAR backscatter archive. Based on the actual location of the satellite, the known incidence angle and the DEM, unreachable areas can be detected geometrically, by describing the viewing angle of the sensor from its orbit position relative to surface topography. This is applicable at the full extent of the CopDEM, and so is used also in areas covered by only one Sentinel-1 orbit direction. Due to the of the global CopDEM's spatial resolution (30m), this is not as accurate as the orbit-based method (Figure 28), but it captures strong shadows, making the global radar shadow mask more reliable.

The orbit-based approach is well suited for masking radar shadow for the GFM product, as it models directly the backscatter signature of the Sentinel-1 observations. However it can only be used over areas with overpasses in both directions (i.e. ascending and descending). If one orbit direction is missing, this radar shadow mask cannot be produced. Figure 29 maps the global areas for which both directions are available, and where the orbit-based approach is used.

On the other hand, a DEM may only describe terrain elevation not surface elevation, or may have a too coarse a 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 a terrain model. Therefore, a geometric observation model is not applicable for such smaller structures causing radar shadows, but delivers shadows caused by larger topographic features (e.g. mountains). The resulting differences between the two approaches can be seen in Figure 30. The combination of both methods in a single sub-layer allows for full coverage of all Sentinel-1 measurements, and the consideration of topography-related shadows, while shadows caused by finer details (e.g. forests) are considered only in areas of better Sentinel-1 coverage.

Consequently, both masking methods, the orbit-approach and the DEM-approach, are applied independently, and in a last step are joined to the final radar shadow masking. The join is done by logical “or” so that it is masked when at least one approach indicates a shadow configuration.




Figure 28: Radar shadow effects on the boundary of forest next to the Danube river in Austria, illustrating typical differences from opposite orbits. Left to right: mean backscatter from descending orbit; mean backscatter from ascending pass; CopDEM hillshade; optical image from Mapbox satellite map.




Figure 29: World map showing the areas where Sentinel-1 overpasses from both orbit directions (ascending and descending), and where the orbital data is used for radar shadow masking.




Figure 30: Radar shadows detected based on relative orbits (left) or a DEM (right).

4.4.5 GFM Exclusion Mask sub-layer: Low Sentinel-1 coverage

In order to ensure high reliability of the results of the GFM product, areas where input data is insufficient and parameters and masks could not be produced, are also masked. This affects mainly the far North of Siberia and Canada due to poor or zero coverage by Sentinel-1, and lack of HAND-EM data. Furthermore, in response to erroneous Sentinel-1 images in the GFM data cube, a few tiles among the global land extent are masked, as no parameters or Exclusion Mask could be generated there. This is temporary safety measure in the initial project phase, and the low coverage layers are updated (i.e. reduced) on a regular basis to account for changes in the Sentinel-1 mission coverage, and the health of the GFM data cube. In particular, two low coverage input layers are used internally:

1	Low coverage on a pixel-level (LOWCOVPX):	▪ The LOWCOVPX sub-layer of the GFM Exclusion Mask is regularly updated, and masks individual pixels with insufficient Sentinel-1 coverage (data density), or where no HAND data are available. The layer is calculated on a pixel level, and masking pixels with < 50 observations for years 2020-2021 - or with no HAND data. The layer is updated regularly to account for changes in the Sentinel-1 mission coverage. A mask file is only created if for the specific tile one of the above reasons apply. Thus, if no LOWCOVPX-file exists, no masking is done for this tile.
2	Low coverage on the T3-tile-level of the Equi7Grid (LOWCOVTILE):	▪ The LOWCOVTILE sub-layer of the GFM Exclusion Mask is also updated often, and excludes entire tiles where input parameters for the flood algorithms or exclusion layers are missing. The layer accounts for the following factors: – Erroneous images (broken data files): Due to broken/corrupted files in the Sentinel-1 data cube on the GFM hard disk, some parameters required as input to the algorithms could not be computed for this tile. – No reference water mask (RWM): No RWM available for tile, which would lead to false flood detection. – No low backscatter mask: No low backscatter mask available for tile, which would lead to false flood detection. ▪ If one of these factors apply, that entire tile is masked. A mask file is only created if one of the factors applies, so if no LOWCOVTILE-file exists, no masking is applied. In the regular updating of the layer, a check is made of whether the masking factor(s) for a tile are still valid. If not, the LOWCOVTILE-file is removed for the specific tile.

4.4.6 Global coverage of the input layers of the GFM Exclusion Mask

The Sentinel-1 and auxiliary datasets required for running the GFM product are not available everywhere. Figure 31 shows the global coverage of the GFM Exclusion Mask sub-layers. As can be seen, the global coverage of the GFM Exclusion Mask sub-layers can be summarized as follows:

- **No-sensitivity** masking (green) is globally applied.
- **Non-water low-backscatter** masking (dark blue) is applied quasi-globally (not in Greenland).
- **Strong topography** masking (light blue) is applied globally.
- **Radar shadow** masking (orange) is applied in most tiles.
- **Pixel-level low Sentinel-1 coverage** masking (violet) is done over scattered land areas in Asia, Africa, and North America.
- **Tile-level low Sentinel-1 coverage** masking (purple) is done over scattered areas in Asia, Africa, and North America over Arctic lands and islands.




Figure 31: Global coverage maps of the technical input layers to the GFM Exclusion Mask. GFM products over the colored T3-tiles make use of the dedicated input layer masking.

4.5 GFM output layer: Likelihood Values

The GFM output layer Likelihood Values provides the estimated likelihood that is computed by the GFM Ensemble flood mapping algorithm, for all areas outside the GFM Exclusion Mask. In this context, the term likelihood represents flood classification accuracy for a given pixel. The methods used by the LIST, DLR, and TUW flood mapping algorithms to compute the Likelihood Values for a Sentinel-1 grid-cell, are described in **Sections 3.1.5, 3.2.5 and 3.3.7**, respectively. The method used by the GFM Ensemble algorithm to combine the Likelihood Values computed by the LIST, DLR and TUW algorithms, into a final Likelihood Value for each grid-cell, is described in **Section 3.4**.

Likelihood values lie in the interval $[0, 100]$, where:

▪ Likelihood values toward 0 represents lower confidence of flood classification accuracy.
▪ Likelihood values toward 100 represents higher confidence of flood classification accuracy.
▪ A likelihood value of 50 separates both classes (Unflooded and Flooded).
▪ Unflooded pixels necessarily show likelihood values in the interval $[0, 49]$.
▪ Flooded pixels show likelihood values in the interval $[50, 100]$.
▪ Unflooded likelihoods propagate towards 0, as Unflooded confidence increases.
▪ Flooded likelihoods propagate towards 100 as Flooded confidence increases.

Finally, as mentioned in Section 3.7 above, the TUW flood mapping algorithm produces uncertainty values. In order to be used by the GFM Ensemble flood mapping algorithm, these uncertainty values have to be “flipped” so that low uncertainty values propagating towards 0 are remapped to high likelihood values propagating towards 100, and vice versa.

4.6 GFM 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.

All of these deleterious effects are difficult to capture, due to their very dynamic nature and high spatial heterogeneity. Therefore, wind, rainfall, temperature, and snow data 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 Sentinel-1 acquisition, and are not ideal for providing the required advisory flags. While optical remote sensing satellites (e.g. MODIS, Sentinel-3) would provide sufficient spatial detail, optical data do not depict the environmental conditions shown by Sentinel-1, nor do they provide timely observations at all times, due to cloud cover and poor illumination conditions. Therefore, within the scope of the GFM product, microwave remote sensing data are the only robust and applicable source for environmental advisory flags. The applied approach is based on Sentinel-1 near real-time data and temporal parameters (see section 3.1.2) derived from the Sentinel-1 data cube (time series) archive.

The purpose of the GFM output layer Advisory Flags is to raise awareness that meteorological conditions (wind, frozen conditions, etc.) may impair the detection of water bodies. As the Advisory Flags can only be retrieved at a coarser spatial resolution than Sentinel-1, this information is not forwarded to the Exclusion Mask. The Advisory Flags are provided as an additional layer, to guide users when interpreting the GFM results, allowing additional insight on local reliability at the time of Sentinel-1 acquisition. The Advisory Flags highlight grid-cells where SAR data may be disturbed by such processes during image acquisition, but the observed flood and water extents remain unmasked. In summary, grid-cells marked by Advisory Flag are not included in the Exclusion Mask, but users are advised to use with caution the observed water and flood extents for these areas.

The GFM Advisory Flags consist of two separate flaggings (i.e. for low regional backscatter and rough water surface), and are provided as four possible data values, as shown below:

VALUE	DEFINITION	USER INTERPRETATION
0	▪ <i>No Advisory Flag set</i>	High quality can be assumed.
1	▪ Low regional backscatter	Caution advised, due to snow-covered, frozen, or dry soil affecting flood mapping reliability.
2	▪ Rough water surface	Caution advised, due to local wind, rainfall, or frozen water surface affecting flood mapping reliability.
3	▪ Both low regional backscatter and rough water surface	Caution advised, due to both (a) snow-covered, frozen, or dry soil, and (b) local wind, rainfall, or frozen water surface affecting flood mapping reliability.

Figure 32 shows an example of the GFM Advisory Flags for a Sentinel-1 scene over Greece, acquired at a time when windy conditions were prevalent. In Figure 32, the Advisory Flags are overlaid on the backscatter image, with RED indicating **low regional backscatter** (i.e. snow, ice, or dryness), GREEN indicating **rough water surface** (i.e. wind, etc.), and BLUE indicating both conditions.




Figure 32: Example of the GFM output layer Advisory Flags for a Sentinel-1 scene over Greece from 2020-08-10 at 16:32 (from orbit A175), at a time when windy conditions were prevalent. (See text for details).

To generate the GFM Advisory Flags, it is necessary to combine the near real-time input data flow, with extracting the corresponding statistical parameters and auxiliary datasets. After processing the flood extent for the incoming Sentinel-1 scene, the Equi7Grid-tiles are identified, and the collocated parameters and auxiliary layers are read from the database. Figure 33 illustrates the data workflow, and how the flag values are set logically.

4.6.1 GFM Advisory Flags: Low regional backscatter

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 in the upper soil layers, show very low backscatter, appearing from the radar perspective as quasi-dry soils. Dried out soils are globally a more common issue for water and flood body detection, often showing backscatter values as low as calm water bodies. This ambiguity is hard to address, even with exploitation of time series analysis and statistical parameters, as the occurrence of dry soils is as erratic and unpredictable as flood events.




Figure 33: Workflow for generating GFM Advisory Flags, combining low regional backscatter (Flag 1) and rough water surface flag (Flag 2). The final Advisory Flags include four values: 0 (no dynamic influence present), 1 (i.e. Flag 1), 2 (i.e. Flag 2), and 3 (i.e. Flags 1 and 2).

Low backscatter over large regions characterizes all the effects described above. Thus, to derive Advisory Flag 1 (**low regional backscatter**), the Sentinel-1 acquisitions are resampled to 20 km resolution, and the regional backscatter is compared to the regional monthly grouped median value (see **Section 2.1.2**) 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 4.4.1.4**) and the built-up area mask (see **Section 2.2.5**). 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 account of seasonal variations, such as annual changes in vegetation cover. If this difference exceeds 1 dB, the low backscatter flag is set to value 1, as follows:

$$\text{median}(\sigma_{S1}^0) > \sigma_{S1}^0 + 1\text{dB} \rightarrow \text{Flag} = 1$$

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. A 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).

4.6.2 GFM Advisory Flags: Rough water surface

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 (of 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 so-called “Bragg resonance effect” is dependent on wind speed and direction as well as on the radar wavelength.

For C-band SAR, 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.

The near real-time 20m Sentinel-1 backscatter data over the areas identified by the GFM Reference Water Mask (see **Section 4.3**) and additional criteria based on backscatter, are analysed. For the grid-cell to be identified as water and used for setting the **rough water surface** flag, it must to be identified as permanent or seasonal water by the Reference Water Mask for the given month, has to have backscatter value **below –10 dB** and deviate from the expected calm water signature. 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 setting the **rough water surface** flag due to the changed water body shape instead of roughening of its surface. Additionally, an erosion step is applied to the selected water bodies, to remove false positives in shallow waters at the water borders.

Over the identified water pixels, a check is made of whether the backscatter deviates from a calm water signature. In the case of a **rough water surface**, there is most likely an increased backscatter indicative of strong winds, rainfall, or frozen water surface, which in turn are likely to increase the backscatter from the nearby water surfaces also over the flooded areas.

For the calm water signature, we chose the **linear water backscatter model** ($\sigma_{w,\theta}^0$), dependent on the project local incidence angle θ (see **Section 3.3.1**) and also used in the TU Wien algorithm (see **Section 3.3**) as baseline. An analysis of backscatter images at windy conditions showed the suitability of the water backscatter model’s standard deviation ($\text{std}(\sigma_{w,\theta}^0)$) 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, as follows:

$$\sigma_{w,\theta}^0 + 0.8 * \text{std}(\sigma_{w,\theta}^0) < \sigma_{s1_water}^0 \rightarrow \text{Flag} = 2$$

Finally, all grid-cells identified as roughened permanent water bodies are spatially filtered using morphological operators. Assuming that wind, rainfall or frozen water surface effects are highly correlated within the neighbourhood, a simple L2/radial buffer of 5km radius is applied to the identified rough water surface grid-cells, effectively enlarging the identified area.

All pixels within the buffered area are labelled with Advisory Flag 2 (**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 current version in the GFM product. Some tweaking and optimisation towards the alert accuracy, is foreseen, as more experience with 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 (which effectively sets to value 0 all grid-cells with no flag set, and sets to value 3 all grid-cells where both flags are set):

$$0 + ind_{LowBSFlag} * 1 + ind_{WindFlag} * 2$$

Note that a known limitation of the applied approach is that flags issued due to land cover change (e.g. changed shape of water body compared with the Reference Water Mask, with the resulting enhanced backscatter misinterpreted as roughened water surface). Generally, the quality of this flag is strongly dependent on the quality and timeliness of the Reference Water Mask.

4.7 GFM output layer: Sentinel-1 Footprint and Metadata

The GFM output layer Sentinel-1 Footprint and Metadata includes all available metadata (i.e. acquisition parameters) provided with each Sentinel-1 GRD image processed by the GFM product. The metadata for each Sentinel-1 GRD scene is provided in the distributed Sentinel SAFE (Standard Archive Format for Europe) format included in the “manifest.safe” file. This is an XML file containing the mandatory product metadata. Attributes in the manifest file are classified into four categories:

<i>Summary.</i>	<i>Product.</i>	<i>Platform.</i>	<i>Instrument.</i>
-----------------	-----------------	------------------	--------------------

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, including information about the absolute orbit number, pass direction, polarisation, sensing start and end date and the product timeliness category. A summary of the included attributes is shown in Table 22.

The footprint of a Sentinel-1 GRD scene, as provided in each data product, is included in the aforementioned manifest file. The footprint is represented as a human readable JTS (Java Topology Suite) 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 OGC (Open Geospatial Consortium) 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.

Table 22: Example of metadata fields in the GFM output layer Sentinel-1 Footprint and Metadata.

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 GFM output layer: Sentinel-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 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 published by ESA on a regular base, well before activation, with potential last-minute changes due to requests from CEMS. Information in the KML files is organized based on planned data takes.

The parameters included in the KML file are listed in Table 23. KML files are regularly checked, downloaded and ingested into the metadata database for further analysis. All parameters are available for requesting the schedule information indicating the next planned Sentinel-1 GRD acquisition for a given location.

Table 23: Information in the GFM output layer Sentinel-1 Schedule, indicating the next planned Sentinel-1 GRD acquisition.

PARAMETER	DESCRIPTION
<i>Satellite ID:</i>	▪ Satellite identifier.
<i>Datatakeld:</i>	▪ Corresponds to a unique product identifier (Hexadecimal value). This identifier is also reported in the filename of any product generated from this acquisition.
<i>Mode:</i>	▪ Instrument acquisition mode (IW, EW, SM).
<i>Swath:</i>	▪ Instrument swath (from 1 to 6 for SM, not applicable for IW and EW).
<i>Polarisation:</i>	▪ Instrument polarisation for the acquired data take.
<i>ObservationTimeStart:</i>	▪ UTC start date and time of the planned data take
<i>ObservationTimeStop:</i>	▪ UTC stop date and time of the planned data take
<i>ObservationDuration:</i>	▪ Duration of the planned data take in seconds
<i>OrbitAbsolute:</i>	▪ Absolute orbit number at the start time of the data take
<i>OrbitRelative:</i>	▪ Relative orbit number at the start time of the data take

4.9 GFM output layer: Affected Population

The GFM output layer Affected Population provides information about the flood-affected population, which is extracted from the Global Human Settlement Layer (GHSL) of CEMS, 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. The GFM product uses the dataset at the highest possible (250m) resolution and for the latest available time-step, (2015). The dataset has been re-projected to the same grid system as the flood map, i.e. the Equi7Grid with a 20m pixel spacing and 300km gridding (T3 level), and resampled to 20m resolution.

4.10 GFM output layer: Affected Land Cover

The GFM output layer Affected Land Cover provides the land cover information which is extracted from the Copernicus Land Monitoring Service⁷, making use of the Global Land Cover dataset, available at a global level at 100m resolution. The Copernicus Global Land Cover (3) 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, such as how much agricultural area is affected by flooding within the GFM Observed Flood Extent or the area of interest.

⁷ <https://land.copernicus.eu/en>

5 GFM product and service quality assessment

Implementing and operating the GFM product requires a set of procedures to validate the technical and scientific quality of the GFM output layers of flood information, as well as that of the generating service, with the overall aim to deliver the main GFM output layers (i.e. Observed Flood Extent, Reference Water Mask, Likelihood Values, and other supplementary information) with the best possible quality. The various methods that are used for the GFM product and service quality assessment, which is performed on a quarterly basis each year, are summarized in Table 24.

Central to the GFM product and service quality assessment is a set of Key Performance Indicators (KPIs), listed in Table 25, which are used for the quarterly monitoring and reporting of all aspects of the GFM service and product delivery performance. The 2022, 2023, and 2024 annual product and service quality assessment (QA) reports of the GFM product, as well as the GFM pre-operational QA report, have been published as Technical Reports of the Joint Research Centre (JRC) of European Commission, and are freely accessible via the GFM wiki pages⁸.

Table 24: Overview of the various GFM product and service quality assessment (QA) tasks.

TASK	DESCRIPTION
<i>Regular assessment of thematic accuracy:</i>	- Implementation of automated, continuous QA procedures to ensure that all input and output data are quality-checked before being released via the dissemination system. - Implementation of systematic QA review to check product consistency and accordance with product specifications. - Systematic thematic quality assessment based on worldwide distributed reference samples and representative Use Cases of worldwide flood events, considering systematic and stratified sampling approaches.
<i>Regular checking and monitoring of product timeliness:</i>	- Implementation of automated and continuous monitoring of the data flow throughout the entire workflow of the GFM product, from data acquisition to the product delivery to users. - 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.
<i>Regular checking and monitoring of service availability:</i>	- 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.

⁸ https://extwiki.eodc.eu/GFM/GFM_QA_Reports

Table 25: Key Performance Indicators (KPIs) used for GFM product and service quality assessment.

KPI	NAME	DESCRIPTION	TARGET VALUE
KPI-1	Service Availability	▪ Percentage the service was available to users per quarter of a year.	>=99 %
KPI-2	Product Timeliness	▪ Percentage of products delivered within 8 hours.	>=95 %
KPI-3	Thematic Accuracy	▪ The Critical Success Index (CSI), and other accuracy metrics, computed by comparing the GFM Observed Flood Extent and Reference Water Mask with independent reference datasets.	>=70 % (CSI)
KPI-4	Unique Visitors	▪ Number of unique users visiting via API / WMS-T (front-end application).	-
KPI-5	Total Visitors	▪ Total number of user visits via API / WMS-T (front-end application).	-
KPI-6	Total Downloads	▪ Number and volume of data downloads via API / WMS-T / web download (front-end application).	-
KPI-7	Service Performance Experience	▪ Percentage change in service performance experience to the user over a quarter of a year.	< 20 %

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.; Ca1o, S.; Tupas, M.E.; Roth, F.; Navacchi, C.; Melzer, T.; Freeman, V.; Wagner, W. Satellite-Based Flood Mapping through Bayesian Inference from a Sentinel-1 SAR Databcube. Remote Sens. 2022, 14, 3673. <https://doi.org/10.3390/rs14153673>
- 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>
- Donchyts, G.; Schellekens, J.; Winsemius, H.; Eisemann, E.; Van de Giesen, N. A 30 m Resolution Surface Water Mask Including Estimation of Positional and Thematic Differences Using Landsat 8, SRTM and OpenStreetMap: A Case Study in the Murray-Darling Basin, Australia. Remote Sens. 2016, 8, 386. <https://doi.org/10.3390/rs8050386>
- 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 terra firme rainforest environments in Amazonia. *Remote Sensing of Environment*, 112, 3469-3481.
<https://doi.org/10.1016/j.rse.2008.03.018>

- Rizzoli P. and B. Bräutigam. 2014. Radar Backscatter Modeling Based on Global TanDEM-X Mission Data. IEEE Transactions on Geoscience and Remote Sensing, Vol. 52, No. 9, September 2014. <https://doi.org/10.1109/TGRS.2013.2294352>
- 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., 2015. Flood detection from multi-temporal SAR data using harmonic analysis and change detection. International Journal of Applied Earth Observation and Geoinformation, 38, 15-24. <https://doi.org/10.1016/j.jag.2014.12.001>
- Tuele, 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., 1989. Handbook of Radar Scattering Statistics for Terrain. Norwood, MA, USA: Artech House Publishers. Pages: 372. ISBN: 978-0890063361.
- Ulaby, F. and Long, D.G., 2014. Microwave Radar and Radiometric Remote Sensing, University of Michigan Press, Ann Arbor, Michigan. Pages: 1116. ISBN: 978-0472119356.
- 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>

List of abbreviations and terminology

TERM	DEFINITION	TERM	DEFINITION
ARD:	Analysis-Ready-Data	IPF:	Instrument Processing Facility
ASCAT:	Advanced Scatterometer	IW:	Interferometric Wide-swath
CEMS:	Copernicus Emergency Management Service (a service of Copernicus, the EO component of the EU's space programme)	JRC:	Joint Research Centre of the EC
CSCDA:	Copernicus Space Component Data Access	NRT:	Near real-time
dB:	Decibel	OS:	Operating system
DEM:	Digital Elevation Model	OSV:	Orbit State Vector
DSM:	Digital Surface Model	PDF:	Probability Density Function
EC:	European Commission	PLIA:	Projected Local Incidence Angle
EO:	Earth observation	POEORB:	Precise Orbit Ephemerides files
ESA:	European Space Agency	Polarisation:	Orientation of the plane in which the SAR signal oscillates, in either the transmit or receive paths.
GDAL:	Geospatial Data Abstraction Library	RESORB	Restituted Orbit Ephemerides files
GFC:	Global Forest Change	S1-GBM:	Sentinel-1 Global Backscatter Model
GFM:	Global Flood Monitoring of CEMS	SAR:	Synthetic Aperture Radar
GHSL:	Global Human Settlement Layer of CEMS	SNAP:	Sentinel Applications Platform (ESA's common architecture for processing and analysing data from many EO missions)
GPT:	Graph Processing Tool	SRTM:	Shuttle Radar Topography Mission
GRDH:	Ground Range Detected, High resolution	SSE:	Sum of Squared Estimate of Errors
HAND:	Height Above Nearest Drainage (a DEM normalized using nearest drainage)	WBM:	Water Body Mask (an auxiliary information mask generated during the Copernicus DEM production process)
HAND-EM:	HAND-based Exclusion Mask	WSF2015:	World Settlement Footprint 2015
HRL:	High-Resolution Layer	Zstd:	Zstandard (a lossless data compression algorithm).