

Reinforcing the AI4EU Platform by Advancing Earth Observation Intelligence, Innovation and Adoption

D5.1 Bootstrapping services and resources I

Grant Agreement ID	101016798	Acronym	AI4COPERNICUS
Project Title	Reinforcing the AI4EU Platform by Advancing Earth Observation Intelligence, Innovation and Adoption		
Start Date	01/01/2021	Duration	36 Months
Project URL	https://ai4copernicus-project.eu/		
Contractual due date	31/12/2021	Actual submission date	20/12/2021
Nature	DEM = Demonstrator	Dissemination Level	PU = Public
Author(s)	M. Lazzarini (SATCEN), O. Barrilero (SATCEN), M. Alexe (ECMWF), V. Baousis (ECMWF), G.Weikmann (UNITN), R. Hall (Equinor), A. Chietera (THA), A. Stoian (THA)		
Contributor(s)			
Reviewer(s)	T. Giannakopoulos (NCSR-D), P. Boza (INSEAD)		



This project has received funding from the *European Union's Horizon 2020 research* and innovation programme under grant agreement No 101016798.



Version	Date	Changes	Contributor(s)
V0.1	14/10/2021	First version of the document	M. Lazzarini, O. Barrilero
V0.2	29/11/2021	Added contributions of all authors and formatted	M. Lazzarini, O. Barrilero, M. Alexe, V. Baousis, G.Weikmann, R. Hall, A. Chietera, A. Stoian
V1	13/12/2021	Added summary table in Section 1 and addressed internal review comments	M. Lazzarini, O. Barrilero, M. Alexe, V. Baousis, G.Weikmann, R. Hall, A. Chietera, A. Stoian
		(pending review by EC)	
		•	(O)

Document Revision History (including peer reviewing & quality control)



Executive Summary

The deliverable will describe the set of services and resources made available from the Security, Agriculture, Energy and Health communities for the first batch of AI4Copernicus open calls winners.

The development of these bootstrapping services aimed to reduce the time and resources of the bidders in different tasks as data access (EO and ancillary data), pre-processing, labelling datasets, ML algorithm definition. The AI4Copernicus consortium support to the bidders allows to address open calls winner's effort on the development of innovative services based on AI.

Each service is documented in the section of the domain responsible of its deployment: following a cooperative approach, each service can be used by different domains, if relevant. The description of the services (e.g. the purpose, the input needed and the output produced) referred to the application deployed in the AI4Copernicus infrastructure.

This document is a demonstrator: it is meant to provide a valid support for the open calls winner to operate the services.



Т	able	of C	ontents
1	8		
	1.1	8	
	1.2	8	
	1.3	8	
	1.4	9	
2	12	2	
	2.1	12	
	2.2	12	
	2.	2.1	12
	2.	2.2	15
	2.	2.3	17
	2.	2.4	19
	2.	2.5	22
	2.	2.6	25
3	2	5	
	3.1	25	
	3.2	26	CX.
	3.	2.1	26
	3.	2.2	28
	3.	2.3	29
	3.	2.4	30
	3.	2.5	32
4	33	3	
	4.1	33	
	4.2	33	
5	34	1	
	5.1	34	
	5.2	34	
	5.	2.1	34
6	38	3	

D5.1 Bootstrapping services and resources I



List of Tables

Table 1. Summary table of services and resources	9
Table 2. Exposed parameters in Sentinel-1 GRD pre-processing	13
Table 3. Exposed parameters in Sentinel-1 SLC pre-processing	15
Table 4. Exposed parameters in Sentinel-2 pre-processing.	17
Table 5. Exposed parameters in S1 Change detection- Amplitude Change Detection and N	∕lulti-
temporal Coherence	19
Table 6. Exposed parameters in Sentinel-2 Change detection.	22
Table 7. Exposed parameters in Deep network for pixel-level classification of S2 patches.	27
Table 8. Dataset specifics for TimeSen2Crop	29
Table 9. Exposed parameters in Harmonization of pre-processed Time Series of Sentinel-2 data	29
Table 10. Exposed parameters in Long Short-Term Memory Neural Network for Sentinel-2.	31
Table 11. Labelled datasets for Energy domain.	33
Table 12. Input data for Probabilistic downscaling of CAMS air quality model data	34
Table 13. Output data for Probabilistic downscaling of CAMS air quality model data	36
Table 14. Commands to execute Probabilistic downscaling of CAMS air quality model data	36
Table 15. Exposed parameters in Probabilistic downscaling of CAMS air quality model data	37

List of Figures

List of Figures	
Figure 1. S1 GRD Processing graph	14
Figure 2. S1 SLC pre-processing graph for IW products.	16
Figure 3. S2 pre-processing graph.	18
Figure 4. Sentinel-1 IW processing.	20
Figure 5. Sentinel-1 SM processing.	21
Figure 6. S2 Change detection processing flow.	23
Figure 7. Example of classified changes in the polar representation using K-means.	24
Figure 8. Deep network for pixel-level classification of S2 patches.	26
Figure 9. Hierarchical folder structure of TimeSen2Crop.	28

List of Terms & Abbreviations

Abbreviation	Definition
AC	Atmospheric Composition
ACD	Amplitude Change Detection
Aol	Area of Interest
AQ	Air Quality
ARD	Analysis Ready Data
AI	Artificial Intelligence
AIS	Automatic Identification System
ВоА	Bottom of Atmosphere



C3S	Copernicus Climate Change Service
CAMS	Copernicus Atmosphere Monitoring Service
CD	Change Detection
CDS	Climate Data Store
CSV	Comma-Separated Values
CVA	Change Vector Analysis
DEM	Digital Elevation Model
DIAS	Data and Information Access Services
DL	Deep Learning
DTE	Digital Twin Earth
EC	European Commission
ECTL	Extract, Cleanse, Transform, Load
EO	Earth Observation
EU	European Union
GA	Grant Agreement
GAN	Generative Adversarial Network
GDAL	Geospatial Data Abstraction Library
GHG	Greenhouse gas
GPT	Graph Processing Tool
GRD	Ground Range Detected
IoT	Internet of Things
IW	Interferometric Wide
JRC	Joint Research Centre
JSON	JavaScript Object Notation
km	Kilometer
L1C	Level 1C
L2A	Level 2A
LAI	Leaf Area Index
LSTM	Long Short-Term Memory
ML	Machine Learning
MTC	Multi-Temporal Coherence
NN	Neural Network
OSM	Open Street Map
PM	Particulate Matter
S1	Sentinel-1
S2	Sentinel-2
SAR	Synthetic Aperture Radar
SLC	Single Look Complex
SM	Strip Map
SNAP	Sentinel Application Platform



SRGAN	Super-resolution GAN
SRTM	Shuttle Radar Topography Misiion
TOPSAR	Terrain Observation Progressive Scans SAR
UTM	Universal Transverse Mercator
WGS	World Geodetic System
WKT	Well-known text
WP	Work Package
XML	eXtensible Markup Language
YAML	YAML Ain't Markup Language

oral



1 Introduction

1.1 Purpose and Scope

The deliverable describes the set of services and resources made available from the Security, Agriculture, Energy and Health communities for the first batch of AI4Copernicus open calls winners.

In the WP5 activities, each of the above domain oversaw and carried out the bootstrapping of AI4Copernicus with high-impact services and resources coming from the operational needs of users. Each domain task focused on the technical requirements, develop of the services and deployment of the resources in the AI4Copernicus infrastructure. First versions of services and resources for each area will have incremental updates and refinements until their final version, foreseen at M18.

D5.1 is a demonstrator: the description of the services referred to the application deployed in the Al4Copernicus infrastructure will constitute a valid support for the open calls winner to operate the services.

1.2 Approach for Work Package and Relation to other Work Packages and Deliverables

From an organisational point of view, the present D5.1 deliverable is a direct outcome of T5.1, T5.2, T5.3 and T5.4. These four tasks correspond to the bootstrapping services and resources of the targeted high-impact domains of Al4Copernicus and each task has a domain responsible: Security (SatCen), Agriculture (University of Trento), Energy (Equinor) and Health (ECMWF). All the tasks were performed with the support of other technical partners as Thales Six, University of Athens, NCSR "Demokritos" and Cloudferro.

The resources of WP5 have been initially described in the proposal and in the Grant Agreement (GA). During the first year of the project, these resources have been refined considering the users operational needs from WP2. The activities of WP5 have been constantly coordinated with the two other technical WPs: WP3 *Architecture and Technical Positioning*, to ensure the deployment of the services within the AI4Copernicus architecture; WP4 *Implementation, customisation, integration and testing*, to ensure the integration of the domain specific services with general services. For a comprehensive overview of the whole infrastructure please refer also to D3.1 and D4.1, available at M18.

As indicated before, the main objective of the WP is to provide the services for the open calls (WP6). The services will be refined and tailored based on the open calls requirements, when the first batch of the open calls selection process will be finalized (M12).

1.3 Methodology and Structure of the Deliverable

The deliverable is structured in accordance with the template and guidelines provided by the EC and is organised in an intuitive structure to directly refer the services to their domain. Section 2, 3, 4 and 5 are related to the services provided for Security (Section 2, SatCen), Agriculture (Section 3, University of Trento with the support of Thales Six), Energy (Section 4, Equinor), Health (Section 5, ECMWF). While the domain responsible was in charge to deploy the specific services, other domains can use the different services if considered useful for their own domain (e.g. pre-processing services).



from Security on Sentinel-1 and Sentinel-2 are useful also for other domains). This approach avoided the duplication of efforts to develop the same service and the simplification for the users. Each section briefly introduces the approach followed to deploy the services and listed all the services available for the first batch of open calls. The description of the services describes the purpose, the structure, the input needed, the output produced and explains how to run them.

1.4 Summary table of services and resources

The Table 1 summarizes the services and resources described in the next sections of this document.

Resource	Summary	Domain	Origin
Sentinel-1 GRD pre- processing	This pipeline processes a S1 GRD product in native format to generate a terrain corrected image representing the calibrated backscatter in GeoTiff format.	Security/ General	Al4Coper nicus
Sentinel-1 SLC pre-processing	This pipeline processes a S1 SLC product in native format to generate a terrain corrected image representing the calibrated backscatter in GeoTiff format.	Security/ General	Al4Coper nicus
Sentinel-2 pre- processing	This pipeline processes a S2 product in native format to generate a product with a common resolution for all the bands in GeoTiff format. The process allows to apply a land/sea mask and a cloud mask in order to have an output product ready for analysis.	Security/ General	Al4Coper nicus
Sentinel-1 Change detection- Amplitude Change Detection and Multi- temporal Coherence	This pipeline processes pairs of S1 SLC products in native format to generate a series of products to assess the changes between both images. These products include: coherence, ACD (Amplitude Change Detection), MTC (Multi-Temporal Coherence) and binary mask of changes.	Security/ General	Al4Coper nicus
Sentinel-2 Change Detection	This pipeline computes (and classifies) the changes using as input a pair of S2-L2A products by using the Change Vector Analysis approach.	Security/ General	Al4Coper nicus

Table 1. Summary table of services and resources



<u>Vector data of</u> <u>human features</u>	SatCen has pre-processed and ingested in a data model the OSM data and can provide the data as a service in the scope of the project	Security	Al4Coper nicus
Deep network for pixel-level classification of S2 patches	This service provides functionality for users to train a custom pixel-level classifier of Sentinel 2 patches. For example users can train a classifier for crop types (corn, sunflower, wheat, etc), land cover (urban vs. natural, water vs land), road extraction (road vs other).	Agriculture/ General	AI4Coper nicus
<u>TimeSen2Crop</u>	TimeSen2Crop is a pixel-based dataset made up of more than 1 million crop type samples of Sentinel-2 time series. The dataset includes atmospherically corrected images and reports the snow, shadows, and clouds information per labelled unit, as well as the spectral signature of the samples of nine Sentinel-2 spectral bands at 10m of spatial resolution.	Agriculture	Al4Coper nicus
Harmonizatio n of pre- processed Time Series of Sentinel-2 data	The harmonization of pre-processed time series of Sentinel-2 data considers a statistic-based approach that computes the median for each pixel in the different images acquired in a particular month. The pixel composite approach to mosaic generation provides consistent results at large scale, allowing the processing of harmonized acquisitions.	Agriculture	Al4Coper nicus
Long Short- Term Memory Neural Network for Sentinel-2	The Long Short Term-Memory architecture can be trained using samples selected by the user. The service exploits the data given by the user to train from scratch an LSTM and stores the resulting weights. Several parameters are exposed to allow the user to custom the model	Agriculture	Al4Coper nicus
Pre-Trained Long Short- Term Memory	The pre-trained Long Short Term-Memory architecture is already trained using the TimeSen2Crop database and is available in .h5 format. The service exploits a pre-trained architecture to classify the specified tile harmonized using the monthly composite approach	Agriculture	Al4Coper nicus
<u>Energy</u> <u>datasets</u>	Meteorological data: ERA5 Example: offshore wind farms are located (training data)	Energy	External referenc es



	JRC Open Power Plants Database (JRC-PPDB-OPEN) Open data from the floating offshore wind farm, Hywind Scotland		
Probabilistic downscaling of CAMS air quality model data	This service generates high-resolution (currently ~ 10km) air quality maps from low-resolution (~40 - 80km) CAMS model (re)analysis and/or forecast output	Health	AI4Coper nicus

oration



2 Security bootstrapping services and resources

2.1 Introduction

The Security Bootstrapping services and resources have been developed considering the objective of the open calls, which was "the development of EO applications leveraging on AI algorithms to detect, identify and/or predict features and events in response to current Security challenges. The applications are expected to exploit EO data, in conjunction with relevant collateral data sources as suitable (e.g. geolocalization, AIS, statistical data, climate/weather, in-situ sensors...) with the use of the latest technologies, also contributing to shape the development of a Digital Twin Earth (DTE) for Security".

The target of the open calls highlights the need of AI algorithms to detect changes using EO and collateral data. To build the above-mentioned services, data and resources coming from EO repositories (e.g. DIASes) needed to be pre-processed for further advanced processing. The AI4Copernicus consortium therefore decided to contribute with the provision of *Analysis Ready Data (ARD)* and possible training datasets for AI algorithms. The development of these bootstrapping services aimed to reduce the time and resources of the bidders in the data preparation and allow them to focus on the development of innovative services based on AI. Moreover, in-house change detection algorithms were also provided as benchmarks to compare the accuracy of possible change detection algorithms provided by bidders.

The following sections describe the services developed by SatCen in the frame of the Security domain.

2.2 Services

2.2.1 Sentinel-1 GRD pre-processing

2.2.1.1 Summary

The Sentinel-1 GRD pre-processing pipeline is available as a dockerized application that can be executed in any environment with a properly configured Docker client.

This pipeline processes a S1 GRD product in native format to generate a terrain corrected image representing the calibrated backscatter in GeoTiff format. Several parameters will be exposed to the users (e.g. final resolution, polarization and projection). Where possible, a standard value will be set, to facilitate the use by the less-expert users.

2.2.1.2 Input

The input of this pipeline is a Level-1 Ground Range Detected (GRD) product in its native SENTINEL-SAFE¹ format (zipped or unzipped product).

Regarding the acquisition mode, they are supported:

• Stripmap (SM): Stripmap (SM) mode acquires data with an 80 km swath at slightly better than 5 m by 5 m spatial resolution (single look). The ground swath is illuminated by a continuous sequence of pulses while the antenna beam is pointing to a fixed azimuth angle and an

¹ <u>https://sentinel.esa.int/web/sentinel/user-guides/sentinel-1-sar/data-formats/sar-formats</u>



approximately fixed off-nadir angle (this is subject to small variations because of roll steering). SM images have continuous along track image quality at an approximately constant incidence angle.²

• Interferometric Wide Swath (IW): the main acquisition mode over land and satisfies the majority of service requirements. It acquires data with a 250 km swath at 5 m by 20 m spatial resolution (single look). IW mode captures three sub-swaths using Terrain Observation with Progressive Scans SAR (TOPSAR). With the TOPSAR technique, in addition to steering the beam in range as in ScanSAR, the beam is also electronically steered from backward to forward in the azimuth direction for each burst, avoiding scalloping and resulting in homogeneous image quality throughout the swath³.

2.2.1.3 Exposed Parameters

Table 2. Exposed parameters in	Sentinel-1 GRD pre-processing	
Parameter	Valid values	Default Value
<u>CalibrationOutput:</u> backscatter convention selected for the output. Each value uses a different reference area.	Sigma, Gamma, Beta	Sigma
Polarization	VV,VH (valid values depend on the specific product)	VV
Speckle-filter (+ specific parameters if needed)	Supported by SNAP	Lee Sigma
Aol: Area of Interest	WKT polygon	None
<u>Resolution:</u> output resolution in meters. The minimum value recommended is the default value (10m for IW products and 5 m for SM products)	Any in meters	10m (IW), 5m (SM)
Projection: output projection	Any supported by SNAP	WGS84
Output format: output format	Supported by GDAL and SNAP	GeoTiff

2.2.1.4 Processing

The pipeline has been designed using SNAP⁴, the common software platform and host for the Sentinel Toolboxes.

² <u>https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-1-sar/acquisition-modes/stripmap</u>

³ <u>https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-1-sar/acquisition-modes/interferometric-wide-swath</u>

⁴ <u>https://step.esa.int/main/toolboxes/snap/</u>



One of the components of SNAP, the Graph Processing Tool (GPT), allows to execute SNAP in batchmode from command-line, with the possibility to create complex workflows using most of the SNAP operators (e.g. Readers, Subset, Reprojection, Band Math, Filters, Calibration...).

In the case of this pipeline, the SNAP graph executed is represented in Figure 1.



Where:

- Read: the operator in charge of reading a product to the SNAP internal data model.
- Apply-Orbit-File: search, download and apply the corrected information about the orbit that is provided some days after the S1 acquisition to improve the geolocation.
- ThermalNoiseRemoval: removes thermal noise
- Remove-GRD-Border-Noise: remove border noise
- Calibration: convert pixel values to calibrated sigma0, gamma0 or beta0
- Speckle-Filter: applies filter to reduce speckle
- Terrain-Correction: orthorectify the product
- Write: write the output product to the desired format.

(More information about the specific operators can be found in the SNAP help and documentation.)

2.2.1.5 Output

The output is a GeoTiff (by default) terrain-corrected image with one float32 band representing the calibrated backscatter (sigma0, gamma0 or beta0 depending on the selected parameters).

2.2.1.6 How to use

Inside the docker, the pipeline can be executed with the following command:

```
S1-GRD-preprocess --input "VALUE" [--calibration "VALUE"] [--polarization "VALUE"]
[--speckle "VALUE"] [--AoI "WKT"] [--resolution "VALUE"] [--projection "VALUE"] [-
output_format "VALUE"] --output_path "VALUE"
```



It can also be launched with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

In the case of any customization is needed in graph, it can be found in the source code in the docker and could be adapted by the users and executed directly using the SNAP's Graph Processing Tool (gpt).

2.2.2 Sentinel-1 SLC pre-processing

2.2.2.1 Summary

The Sentinel-1 SLC pre-processing pipeline is available as a dockerized application that can be executed in any environment with a properly configured Docker client.

This pipeline processes a S1 SLC product in native format to generate a terrain corrected image representing the calibrated backscatter in GeoTiff format. Several parameters are exposed (e.g. final resolution, polarization and projection), including when possible, a default value to facilitate the use by non-expert users.

2.2.2.2 Input

The input of this pipeline is a Level-1 Single Look Complex (SLC) product in its native SENTINEL-SAFE ⁵ format (zipped or unzipped product are both supported).

Regarding the acquisition mode, they are supported SM and IW (see 2.2.1.2 for a description of these modes).

2.2.2.3 Exposed Parameters

Parameter	Valid values	Default Value
<u>CalibrationOutput:</u> backscatter convention selected for the output. Each value uses a different reference area.	Sigma, Gamma, Beta	Sigma
Polarization	VV,VH (valid values depend on the specific product)	VV
Speckle-filter (+ specific parameters if needed)	Supported by SNAP	Lee Sigma
Aol: Area of Interest	WKT polygon	None
<u>Resolution</u> : output resolution in meters. The minimum value recommended is the default value (10m for IW products and 5 m for SM products)	Any in meters	10m (IW), 5m (SM)
Projection: output projection	Any supported by SNAP	WGS84
Output format: output format	Supported by GDAL and SNAP	GeoTiff

able 3. Exposed parameters in Sentinel-1 SLC pre-processing

⁵ <u>https://sentinel.esa.int/web/sentinel/user-guides/sentinel-1-sar/data-formats/sar-formats</u>



2.2.2.4 Processing

The main pipeline has been designed using SNAP. The SNAP graph executed for IW products is represented in Figure 2. For SM products, the graph varies a little since operators like TOPSAR-Deburst is not needed.



Figure 2. S1 SLC pre-processing graph for IW products.

Where:

- Read: the operator in charge of reading a product to the SNAP internal data model.
- Apply-Orbit-File: search, download and apply the corrected information about the orbit that is provided some days after the S1 acquisition to improve the geolocation.
- ThermalNoiseRemoval: removes thermal noise
- TOPSAR-Deburst: merge the bursts
- Calibration: convert pixel values to calibrated sigma0, gamma0 or beta0
- Speckle-Filter: applies filter to reduce speckle
- Terrain-Correction: orthorectify the product
- Write: write the output product to the desired format.

(More information about the specific operators can be found in the SNAP help and documentation.)

2.2.2.5 Output

The output is a GeoTiff (by default) terrain-corrected image with one float32 band representing the calibrated backscatter (sigma0, gamma0 or beta0 depending on the selected parameters).

2.2.2.6 How to use

Inside the docker, the pipeline can be executed with the following command:



```
S1-SLC-preprocess --input "VALUE" [--calibration "VALUE"] [--polarization "VALUE"]
[--speckle "VALUE"] [--AoI "WKT"] [--resolution "VALUE"] [--projection "VALUE"] [-
-output_format "VALUE"] --output_path "VALUE"
```

It can be also executed with "*docker run*" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

In the case of any customization is needed in graph, it can be found in the docker and could be adapted by the users and executed directly using *gpt*.

2.2.3 Sentinel-2 pre-processing

2.2.3.1 Summary

The Sentinel-2 pre-processing pipeline is available as a dockerized application that can be executed in any environment with a properly configured Docker client.

This pipeline processes a S2 product in native format to generate a product with a common resolution for all the bands in GeoTiff format. The process allows to apply a land/sea mask and a cloud mask in order to have an output product ready for analysis.

Several parameters are exposed (e.g. DEM, cloud mask type), including when possible, a default value to facilitate the use by non-expert users.

2.2.3.2 Input

The input of this pipeline is a Sentinel-2 L2A product in its native SENTINEL-SAFE format (zipped or unzipped products are both supported).

2.2.3.3 Exposed Parameters

Parameter	Valid values	Default Value
<u>Resolution</u> : output resolution in meters. The minimum value recommended is the default value (10m)	Any in meters	10
Bands	Any combination of S2 L2A bands: B1,B2, B3, B4, B5, B6, B7, B8, B8A, B9, B10, B11, B12	B2, B3, B4, B8
Land/Sea mask: type of pixels to be removed considering SRTM3Sec. If "Sea" is selected, al values with elevation=0 in SRTM 3Sec will be set to NoData.	Land/Sea/None	None
Aol: Area of Interest	WKT polygon	None
<u>CloudMask</u> : assign NoData to cloudy pixels according to the cloud mask type selected	L1C/L2A/other?/None	None
Projection: output projection	Any supported by SNAP	UTM(Auto)
Output format: output format	Supported by GDAL and SNAP	GeoTiff

Table 4. Exposed parameters in Sentinel-2 pre-processing.



2.2.3.4 Processing

The main workflow has been designed using SNAP. The SNAP graph executed is represented in Figure 3.



Where:

- Read: the operator in charge of reading a product to the SNAP internal data model.
- S2Resampling: this operator resamples the product to a common resolution.
- Land/Sea Mask: It applies land/sea mask based on srtm 3sec.
- BandMaths: it is used to compute the cloud mask when needed.
- Subset: it filters out the non-desired bands and crop to the AoI.
- Reproject: reprojects to the selected projection.
- Write: write the output product to the desired format.

(More information about the specific operators can be found in the SNAP help and documentation.)

2.2.3.5 Output

The output is a GeoTiff (by default) terrain-corrected image containing the selected bands. Depending on the selected parameters, pixels in sea/land or/and cloudy have been set to NoData.

2.2.3.6 How to use

Inside the docker, the pipeline can be executed with the following command:

```
S2-preprocess --input "VALUE" [--bands "XX,XX,XX"] [--landseamask "VALUE"] [--
cloudmask "VALUE"] [--AoI "WKT"] [--resolution "VALUE"] [--projection "VALUE"] [-
-output_format "VALUE"] --output_path "VALUE"
```

It can be also executed with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.



In the case of any customization is needed in graph, it can be found in the docker and could be adapted by the users and executed directly using *gpt*.

2.2.4 Sentinel-1 Change detection – Amplitude Change Detection and Multi-temporal Coherence

2.2.4.1 Summary

The Sentinel-1 Change Detection resource is available as a dockerized application that can be executed in any environment with a properly configured Docker client.

This pipeline processes pairs of S1 SLC products in native format to generate a series of products to assess the changes between both images. These products include:

- the coherence (the amplitude of correlation between the images),
- the ACD (Amplitude Change Detection), which is a RGB composite of the backscatter of the input images
- the MTC (Multi-Temporal Coherence), which is a RGB composite of the backscatters and the coherence
- binary mask of changes

Several parameters are exposed (e.g. resolution, speckle), including when possible, a default value to facilitate the use by non-expert users.

2.2.4.2 Input

The input of this pipeline is a pair of Sentinel-1 SLC product in their native SENTINEL-SAFE format (zipped or unzipped product are both supported). The inputs shall be acquired with the same acquisition geometry (i.e. same relative orbit), which means in practice that they are acquired in a multiple of six days (6, 12, 18 days) of difference (considering both satellites, S1A and S1B). The validity of the inputs will be checked by the pipeline and when the requirements are not fulfilled, an error will be raised.

2.2.4.3 Exposed Parameters

Table 5. Exposed parameters in S1 Change detection – Amplitude Change Detection and Multi-temporal Coherence

Parameter	Valid values	Default Value
<u>Resolution</u> : output resolution in meters. The minimum value recommended is the default value (10m for IW products and 5 m for SM products)	Any in meters	10m (IW), 5m (SM)
<u>CalibrationOutput</u> : backscatter convention selected for the output. Each value uses a different reference area.	Sigma, Gamma, Beta	Sigma
Polarization	VV,VH (valid values depend on the specific product)	VV

D5.1 Bootstrapping services and resources I



<u>Land/Sea mask:</u> type of pixels to be removed considering SRTM3Sec. If "Sea" is selected, al values with elevation=0 in SRTM 3Sec will be set to NoData.	Land/Sea/None	None
Aol: Area of Interest	WKT polygon	None
Speckle-filter (+ specific parameters if needed)	Supported by SNAP	Lee Sigma
Projection: output projection	Any supported by SNAP	WGS84
Output format: output format	Supported by GDAL and SNAP	GeoTiff

2.2.4.4 Processing

The pipeline has been designed mostly using SNAP and GDAL. The SNAP graphs executed to generate the coherence are represented in Figure 4 and Figure 5.



Figure 4. Sentinel-1 IW processing.

AI4 copernicus



Figure 5. Sentinel-1 SM processing.

Where:

- Read: the operator in charge of reading a product to the SNAP internal data model.
- Apply-Orbit-File: search, download and apply the corrected information about the orbit that is provided some days after the S1 acquisition to improve the geolocation.
- TOPSAR-Split: split the product using the indicated subswath.
- Back-Geocoding: co-register the input products.
- ESD: Enhance the coregistration when more than one burst have been selected.
- Calibration: convert pixel values to calibrated sigma0, gamma0 or beta0
- Speckle-Filter: applies filter to reduce speckle
- Coherence: estimates the coherence
- Terrain-Correction: orthorectify the product
- CreateStack + Cross-Correlation + Warp: coregister the input products.
- Write: write the output product to the desired format.

(More information about the specific operators can be found in the SNAP help and documentation.)

2.2.4.5 Output

The outputs are:

- Coherence: GeoTiff image that represents the amplitude of correlation between the images. The pixel type is float32.
- Backscatters: 2 GeoTiff products (one for each of the input images) with one float32 band representing the calibrated backscatter.
- ACD: a RGB composite of the backscatter of the input images.
 - GeoTiff with three bands. R: backscatter of image 1; G : backscatter of image 2; B: backscatter of image 2
 - Pixel type is Byte, where byte value is computed by data conversion of float values using a linear interpolation taking as min and max values the percentiles 2.5 and 7.5.
- MTC (Multi-Temporal Coherence): a RGB composite of the backscatters and the coherence



- GeoTiff with three bands. R: backscatter of image 1; G : backscatter of image 2; B: coherence
- Pixel type is Byte, where byte value is computed by data conversion of float values using a linear interpolation taking as min and max values the percentiles 2.5 and 7.5 for backscatter band. For the coherence band the linear conversion is (0,1)->(0,255)
- Binary mask of changes. The pixel type is byte.

2.2.4.6 How to use

Inside the docker, the pipeline can be executed with the following command:

S1-CD --input1 "VALUE" --input2 "VALUE" [--calibration "VALUE"] [--polarization
"VALUE"] [--landseamask "VALUE"] [--speckle "VALUE"] [--AoI "WKT"] [--resolution
"VALUE"] [--projection "VALUE"] [--output format "VALUE"] --output path "VALUE"

It can be also executed with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

If any customization is needed in the processing graphs, they can be found in the docker and could be adapted by the users and executed directly using *gpt*.

2.2.5 Sentinel-2 Change detection

2.2.5.1 Summary

The Sentinel-2 Change Detection pipeline is available as a dockerized application that can be executed in any environment with a properly configured Docker client.

This pipeline computes (and classifies) the changes using as input a pair of S2-L2A products by using the Change Vector Analysis approach.

Several parameters are exposed (e.g. resolution, bands, number of change classes), including when possible, a default value to facilitate the use by non-expert users.

2.2.5.2 Input

The input of this pipeline is a pair of Sentinel-2 L2A products in their native SENTINEL-SAFE format (zipped or unzipped products are both supported). The inputs shall correspond to the same tile (e.g. same relative orbit).

2.2.5.3 Exposed Parameters

Parameter	Valid values	Default Value
<u>Resolution</u> : output resolution in meters. The minimum value recommended is the default value (10m for IW products and 5 m for SM products)	Any in meters	10m (IW), 5m (SM)
<u>Bands:</u> list of S2 bands that are going to be used for computing the changes.	Any combination of S2 L2A bands	B2,3,B4,B8

Table 6. Exposed parameters in Sentinel-2 Change detection.

D5.1 Bootstrapping services and resources I



Aol: Area of Interest	WKT polygon	None
Projection: output projection	Any supported by SNAP	WGS84
<u>NumberOfClasses</u> : number of classes in which changes will be automatically classified.	Any integer	4
<u>ReferenceVector</u> : Reference vector to be used for computing the angle in CVA methodology.		1,0,0,
Level of confidence:	Any float value lower than 100.	99,99
Output format: output format	Supported by GDAL and SNAP	GeoTiff

2.2.5.4 Processing

The standard approach when computing changes is simplified in Figure 6.



Figure 6. S2 Change detection processing flow.

The S2 change detection pipeline developed implements the following steps:

- 1) Pre-processing:
 - a. Resampling: Bands that are needed for the processing are resampled to the common selected resolution. These bands include the bands selected by the user and the scene classification bands. This step is performed using SNAP.
 - b. Radiometric correction/histogram matching: in order to minimize errors caused by not accurate radiometric corrections (including atmospheric), a relative radiometric



correction is applied to one of the inputs. For this, it is computed a linear regression using as references the pixels with less changes after removing the ones affected by clouds or where the land cover is different.

- c. Crop the image to the Aol.
- d. Generate cloud masks.
- 2) Computation of normalized vector of differences in pre-processed images
- 3) Compute the module of the vector and angle with respect the reference vector.
- 4) Compute binary mask of changes by assuming that difference between bands follow gaussian distributions and the module of the change follows a chi-squared distribution:
 - After having applied the radiometric correction/histogram matching in 1.b, it is assumed that the difference of the same band in two images is following a Gaussian distribution. The values that cannot be explained with this distribution are classified as changes.
 - When taking into account the full set of bands selected for the processing, the amplitude of the change is computed with the normalized differences and it is assumed that it follows a chi-squared distribution.
- 5) Classify changes using K-means algorithm in the polar representation of the vector of differences (amplitude and angle).



Figure 7. Example of classified changes in the polar representation using K-means.

2.2.5.5 Output The outputs are:

- CVA: GeoTiff image with two bands. The first band is the amplitude of the change and the second band is the angle with respect the reference vector.

- S2-CD: GeoTiff image with one band with pixel type Byte. It represents the classes of the detected changes.

2.2.5.6 How to use

Inside the docker, the pipeline can be executed with the following command:

```
S2-CD --input1 "VALUE" --input2 "VALUE" [--bands "XX,XX,XX"] [--AoI "WKT"] [--
resolution "VALUE"] [--projection "VALUE"] [--numberClasses "VALUE"] [--
referenceVector "VALUE"] [--levelConfidence "VALUE] [--output_format "VALUE"] --
output_path "VALUE"
```



It can be also executed with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

If any customization is needed in the processing graph, it can be found in the docker and could be adapted by the users and executed directly using *gpt*.

2.2.6 Vector data of human features

2.2.6.1 Summary

The vector data of human features service is based in OpenStreetMap (OSM). OSM is a collaborative project to create a free editable map of the world with an open-content license. The OSM License allows free (or almost free) access to the map images and all of the underlying map data. But this access is not always easy since there are limitations in the servers and APIs. Moreover, the data structure is not the preferred by Security domain.

SatCen has pre-processed and ingested in their own data model (SatCen Data Dictionary) the OSM data and can provide the data as a service in the scope of the project.

2.2.6.2 Input and parameters

The input data in which the service is based is the OSM data, but it is offered to the users in SatCen's data model.

For obtaining the data, the user has to indicate the features of interest (e.g. airport) and the area of interest.

2.2.6.3 Processing

The process can be considered a ECTL process (Extract, Cleanse, Transform and Load workflow). It was developed in FME and based in schema mapping documents.

2.2.6.4 Output

The outputs are the desired features in the AoI in geojson format.

2.2.6.5 How to use

This service is available as a web service where user can select the desired feature and define the AoI.

3 Agriculture bootstrapping services and resources

3.1 Introduction

The Agriculture Bootstrapping service and resources have been developed to support the development of EO applications leveraging on AI algorithms for the food security and agriculture fields. The services aim to facilitate the integration of the proposals of the bidders in the AI4Copernicus environment.



To be effective at large scale, data and resources need to be properly pre-processed to accurately perform crop dynamic monitoring. The bootstrapping services described in this section allow the definition of a generic processing pipeline to perform crop type mapping, and to support the testing and implementation of innovative AI algorithms.

3.2 Services

The services described are available as a dockerized application that can be executed in any environment with a properly configured Docker client.

3.2.1 Deep network for pixel-level classification of S2 patches

3.2.1.1 Summary

This service provides functionality for users to train a custom pixel-level classifier of Sentinel 2 patches. For example users can train a classifier for crop types (corn, sunflower, wheat, etc), land cover (urban vs. natural, water vs land), road extraction (road vs other).

The service will be available as a dockerized application that can be executed in any environment with a properly configured Docker client.

The service will implement code to train a U-NET model on Sentinel2 hyperspectral images (patches of 64x64 pixels). The weights of a part of the U-NET model will be pre-trained with a self- supervised approach, increasing final classification performance.



Figure 8. Deep network for pixel-level classification of S2 patches.

3.2.1.2 Input/output

The docker container will take as input a directory containing patches (64x64 patches) and their labels (patches with pixels = 1..N labels). The output will be a U-NET model which makes the classification by pixel with these N classes.

D5.1 Bootstrapping services and resources I



3.2.1.3 Pipeline

The service design is based on a network architecture with a backbone, a ResNet-50, which is pretrained separately on SEN12MS. By using self supervised learning, an emerging unsupervised training procedure, we will learn good features on Sentinel-2 images, without requiring labelling, to initialize our network's backbone (Ref. Citation: Ciocarlan, A.; Stoian, A. Ship Detection in Sentinel 2 Multi-Spectral Images with Self-Supervised Learning. Remote Sens. 2021, 13, 4255. https://doi.org/10.3390/rs13214255)

3.2.1.4 Exposed parameters

Table 7. Exposed parameters in Deep network for pixel-level classification of S2 patches.		
Parameter	Valid values	Default Value
<u>Predict/train</u> : toggle whether the dockerized application will train the model or use a trained model to predict	Predict, Train	Train
Epochs: number of training epochs	1-1000	100
Learning rate	Any non-zero positive real value	1E-4
Number of classes	2-255	2
<u>Bands</u>	Any combination of S2 L2A bands: B1,B2, B3, B4, B5, B6, B7, B8, B8A, B9, B10, B11, B12	6 bands: B2, B3, B4, B6, B10, B11
Patch mode:	full tile or patch	Full tile
Loss:	Focal loss or Cross Entropy	Focal Loss
Output format: ONNX format or PYTORCH pth for the a UNET model	Onnx/Pytorch	Pytorch

3.2.1.5 How to use

First of all the user should obtain a data set of S2 image tiles with any combination of bands. Furthermore, the user should then provide pixel level annotation of these image tiles. These tiles can be of any size, with the lowest size being 64x64 pixels. Larger image tiles will be cut into 64x64 patches by the service, automatically.

The annotations do not need to cover all image pixels, some pixels that have unknown/irrelevant types will have class 0. This can be considered a background class. For example if the user wants to segment roads, she/he will annotate all road pixels with class 1 and leave the rest of the pixels as class 0.

To train a model the user will invoke the dockerized applications through the command line or eventually through the cloud interface. He/she will point the application to a directory containing



Sentinel 2 images and their annotations. After training, the application will produce a "trained model" that can produce predictions on new image data. Calling the application in "Predict" mode while supplying such a saved trained model will allow the user to predict pixel level classification on new data.

3.2.2 TimeSen2Crop

3.2.2.1 Summary

TimeSen2Crop is a pixel-based dataset made up of more than 1 million crop type samples of Sentinel-2 time series. The dataset includes atmospherically corrected images and reports the snow, shadows, and clouds information per labelled unit, as well as the spectral signature of the samples of nine Sentinel-2 spectral bands at 10m of spatial resolution.

3.2.2.2 Dataset Description

The dataset is organized hierarchically, as shown in Figure 8. The data are organized per Sentinel-2 tiles, i.e. 16 folders. Each folder contains 16 sub-folders, each associated to a particular crop type. The multispectral temporal signature is stored in a .csv file, which provides a matrix where each row indicates the acquisition date, and each column indicates the spectral band associated. The last column in the file shows the condition of the pixel. A csv containing the list of the dates in which the samples have been acquired is stored inside the tile folder.







Dataset Specifics	Values
Spatial Resolution	10 [m]
Spectral Resolution	B2-490 [nm], B3-560 [nm], B4-665 [nm], B5-705 [nm], B6-740 [nm], B7-783 [nm], B8A-865 [nm], B11- 1610 [nm], B12-2190 [nm]
Pixel condition	clear [0], cloud [1], shadow [2], snow [3]
Agronomic Year	September 2017 - August 2018
Crop Types – Classification Scheme	Legumes [0], Grassland [1], Maize [2], Potato [3], Sunflower [4], Soy [5], Winter Barley [6], Winter Caraway [7], Rye [8], Rapeseed [9], Beet [10], Spring Cereals [11], Winter Wheat [12], Winter Triticale [13], Permanent Plantations [14], Other Crops [15]
Sentinel-2 Tiles	32TNT, 32TPT, 32TQT, 33TUM, 33TUN, 33TVM, 33TVN, 33TWM, 33TWN, 33TXN, 33UUP, 33UVP, 33UWP, 33UWQ, 33UXP

Table 8. Dataset specifics for TimeSen2Crop

3.2.3 Harmonization of pre-processed Time Series of Sentinel-2 data

3.2.3.1 Summary

One of the challenges of crop type mapping at large scale is the definition of a regular temporal sampling grid to perform analysis in a standardized manner. This optical pre-processing harmonizes the irregular time series of images and mitigates the cloud coverage problem. The harmonization of pre-processed time series of Sentinel-2 data considers a statistic-based approach that computes the median for each pixel in the different images acquired in a particular month. The pixel composite approach to mosaic generation provides consistent results at large scale, allowing the processing of harmonized acquisitions.

3.2.3.2 Input

The inputs of this service are a Sentinel-2 Bottom of Atmosphere time series and the LAI images associated. The input to this service can be obtained using the pre-processing pipeline defined in the previous section. The input_dir specified to execute the service must contain two folders, "REFBOA" and "LAI", each containing the .tif images.

3.2.3.3 Exposed Parameters

Table 9. Exposed parameters in Harmonization of pre-processed Time Series of Sentinel-2 data		
Parameter	Valid Values	Default Value



tile: UTM code associated with the selected tile (i.e. "33UVP")	string	None
<u>season_start:</u> Year at which the agronomic year starts defined from September (A) to August (B).	number	2017

3.2.3.4 Processing

The integrated algorithm has been designed using python. The script reads the LAIs of the images in a particular month, sets to 0 the corresponding pixels in the BOAs, and performs the median operation masking the cloudy samples.

3.2.3.5 Output

The outputs of the Harmonization step are 12 monthly composite GeoTiff images in uint16 representing the agronomic year specified by the user.

3.2.3.6 How to use

Inside the docker, the pipeline can be executed with the following command:

```
python3 monthlycomposite.py --input_dir="VALUE" --output_dir="VALUE" --
season_start="VALUE" --tile="VALUE"
```

It can be also executed with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

```
docker run -u $(id -u):$(id -g) -v $(pwd):$(pwd) -w $(pwd) docker_name python3
script_name.py --input_dir="VALUE" --output_dir="VALUE --season_start="VALUE" --
tile="VALUE"
```

The user can check the help page to have an overview of the variables that can be defined and a short description of each parameter.

monthlycomposite --help

3.2.4 Long Short-Term Memory Neural Network for Sentinel-2

3.2.4.1 Summary

The Long Short Term-Memory architecture can be trained using samples selected by the user. The service exploits the data given by the user to train from scratch an LSTM and stores the resulting weights. Several parameters are exposed to allow the user to custom the model.

3.2.4.2 Input

The service requires two different .npy files stored inside the input_dir specified, one containing the multitemporal spectral signature of the samples used to train the LSTM ("train.npy"), and a .npy file containing the labels associated with the samples ("y_train.npy"). The dimension of the array containing the samples used to train the model must be [N,ts,b], where N is the number of samples, ts the number of time steps in the Time Series, and b the number of bands.

3.2.4.3 Exposed Parameters

Table 10. Exposed parameters in Long Short-Term Memory Neural Network for Sentinel-2.

Parameter	Valid Values	Default Value
<u>n_epochs</u> : number of epochs to update the internal model parameters	Any number >1	50
batch size: number of samples used to update the internal models	Any number >2 and power of 2	32
<u>class</u> weights: flag used in the training to give different weights based on the a-priori probability of each class.	0 (not used) - 1 (used)	0
learning rate: number describing the step size.	Any number > 0 and < 1	1e-3
<u>val</u> : string pointing to the path where a validation set is stored.	path to folder containing the .npy data	None

3.2.4.4 Processing

The service has been designed using python. The script performs a normalization on the input data and trains the neural network.

3.2.4.5 Output

The output of the Long Short-Term Memory network is a .h5 file containing the weights and the model configuration.

3.2.4.6 How to use

Inside the docker, the pipeline can be executed with the following command:

```
python3 train_lstm.py --input_dir="VALUE" --output_dir="VALUE" --n_epochs="VALUE"
--batch_size="VALUE" --class_weights="VALUE" --learning_rate="VALUE" --
val="VALUE"
```

It can be also executed with "*docker run*" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

```
docker run -u $(id -u):$(id -g) -v $(pwd):$(pwd) -w $(pwd) docker_name python3
train_lstm.py --input_dir="VALUE" --output_dir="VALUE" --n_epochs="VALUE" --
batch_size="VALUE" --class_weights="VALUE" --learning_rate="VALUE" --val="VALUE"
```

The user can check the help page to have an overview of the variables that can be defined and a short description of each parameter.

train_lstm --help



3.2.5 Pre-Trained Long Short-Term Memory

3.2.5.1 Summary

The pre-trained Long Short Term-Memory architecture is already trained using the TimeSen2Crop database and is available in .h5 format. The service exploits a pre-trained architecture to classify the specified tile harmonized using the monthly composite approach.

3.2.5.2 Input

The input of this service is a GeoTiff image containing the information regarding the agricultural samples that the LSTM must classify (if no mask is provided, the service will classify the entire image), and a monthly composite time series. The .tif image and the .h5 pre-trained network must be in the --input_dir folder.

3.2.5.3 Processing

The service has been designed using python. The monthly composite time series is classified according to the TimeSen2Crop classification scheme. The GeoTiff image defined by the user is used to mask the samples not related to agriculture and the samples where classification is not needed. If the binary crop mask is specified, pixels with 0 value will be classified, while pixels with other values will be ignored. The script performs a block-wise classification of the monthly composites and applies morphological operators to the result.

3.2.5.4 Output

The output is the crop type map associated to the monthly composite time series as a GeoTiff image in uint8.

3.2.5.5 How to use

Inside the docker, the pipeline can be executed with the following command:

```
python3 tile_classifier.py --input_dir="VALUE" --output_dir="VALUE"
mc_dir="VALUE"
```

It can be also executed with "docker run" taking into account that a volume has to be mounted in order to write on it the output file so it is accessible at the end of the processing.

```
docker run -u $(id -u):$(id -g) -v $(pwd):$(pwd) -w $(pwd) docker_name python3
tile classifier.py --input dir="VALUE" --output dir="VALUE"
```

The user can check the help page to have an overview of the variables that can be defined and a short description of each parameter.

tile classifier --help



4 Energy bootstrapping services and resources

4.1 Introduction

The Al4Copernicus Energy Bootstrapping services have been identified to support the application of Al technology towards satellite data to produce models or forecasts that address one of the three broad energy questions.

The three broad questions are:

- As a user I want to know where I can and cannot build low carbon and renewable energy infrastructure
- As a user I want a better understanding of energy consumption and energy needs of a society
- As a user I want to know where to carry out precision, pre-emptive, maintenance in my energy infrastructure

The main function of the services are labelled data sets used to train machine learning and artificial intelligence services. In addition to these domain specific resources, the energy domain will also exploit the bootstrapping services from other domains (e.g. pre-processing services of Sentinel-1 and -2 from the security) if requested by open calls winners.

4.2 Resources

Table 11. Labelleu datasets for Energy domain.	
Meteorological data:	CREODIAS - ERA5 product home page
ERA5	How to order CDS ERA5 products - FAQ Answer - CREODIAS
Example: offshore wind farms are located (training data)	EMODNET Data Catalogue EMODnet Product Catalogue - EMODnet (europa.eu)
JRC Open Power Plants Database (JRC-PPDB- OPEN)	WMS service identified (https://data.europa.eu/data/datasets/9810feeb-f062-49cd-8e76- 8d8cfd488a05?locale=en)
Open data from the floating offshore wind farm, Hywind Scotland	Platform for Operational Data (POD) POD (catapult.org.uk)

able 11. Labelled datasets for Energy domain.



5 Health bootstrapping services and resources

5.1 Introduction

The Al4Copernicus Health Bootstrapping services and resources are being developed to address current public health and air pollution (or air quality) challenges using Earth observation and in-situ measurement data. The services are focused on probabilistic downscaling (super-resolution) of air quality (AQ) and atmospheric composition (AC) model output. Current AC and AQ models output forecasts at relatively low-resolution (e.g., ca. 10 to 40 km for the CAMS EU and global domains). Previous research has demonstrated that it is possible to make use of Earth observation (EO) and insitu measurement data to downscale (i.e., increase the spatio-temporal resolution) of dynamical model output, allowing the identification of pollution or greenhouse gas (GHG) emission hotspots. This service is well aligned with all proposals submitted to the AI4Copernicus first open call under the health domain. These proposals have identified the need and market potential for hyper-local short to medium-term air pollution forecasts, with up to street-scale resolution in densely populated areas that are subject to heavy pollution episodes. Real-time knowledge of AQ plus high-resolution forecasts up to a few days ahead allow policymakers to make science-informed decisions (e.g. limit traffic, issue health advisories, deploy additional health personnel, etc.). Citizens can also benefit by reducing their exposure to air pollutants - staying indoors during aerosol peaks, for example, or using AQ-aware routing when navigating through traffic-congested cities.

The health services also provide an estimate of the systematic uncertainties present in the downscaled product.

It is important to note that the technical specification of the service described below is subject to changes. The main objectives are to (1) provide a health bootstrap service that is useful for the winning bidders of the first open call and (2) ensure interoperability with other services described in this document if advantageous for end users.

5.2 Services

5.2.1 Probabilistic downscaling of CAMS air quality model data

This service generates high-resolution (currently ~ 10km) air quality maps from low-resolution (~40 - 80km) CAMS model (re)analysis and/or forecast output.

5.2.1.1 Input data

The input data sources for the probabilistic downscaling service are summarized in the table below.

Table 12. Input data for Probabilistic downscaling of CAWS all quality model data		
Dataset	Summary	Format(s)
CAMS European air quality analysis and forecasts	This dataset provides daily air quality analyses and forecasts for Europe (east boundary=25.0° W, west=45.0° E,	netCDF
	south=30.0° N, north=72.0° N) at a spatial resolution of 0.1	

Table 12. Input data for Probabilistic downscaling of CAMS air quality model data



	 degrees (ca. 10km). The set of atmospheric variables includes particulate matter, pollen, ozone, dust, CO, NO, SO2, NO2 and ammonia. The downscaling service currently uses only surface data. Other vertical levels may be incorporated depending on user need. 	
CAMS global (re-)analysis	The EAC4 reanalysis. The downscaling service currently uses only surface data. Other vertical levels may be incorporated depending on user need and/or impact on downscaled field quality.	netCDF
Sentinel-2	The use of Level-2A Sentinel-2 image data for identification of pollution hotspots is currently being investigated. Satellite-retrieved aerosol optical depth (AOD) ⁶ can be correlated to aerosol conditions over urban environments at a fine scale. Furthermore, it has been demonstrated that data from the Sentinel-2 MultiSpectral Instrument can be used to detect and quantify anomalously large methane point sources. ⁷	GeoTIFF (can be the output of the Sentinel-2 pipeline described for the "Security" domain above)
Sentinel-5P ⁸	The Copernicus Sentinel-5 Precursor mission provides atmospheric measurements of air quality and atmospheric composition (aerosol, ozone, NO2, CH4, CO, HCHO and SO2) with high temporal and spatial resolution (7 x 7 km at nadir resolution). Sentinel-5P NO2 total columns are assimilated operationally into the CAMS-IFS 4D-Var model. The data are available in near-real-time (within three hours of sensing, except for total ozone which has a 24hrs delay).	xml (header) and netCDF (data)
CDS weather data	The Climate Data Store (CDS, <u>http://cds.climate.copernicus.eu</u>) is the operational data access portal of the Copernicus Climate Change Service (C3S), which is implemented by ECMWF on behalf of the European Commission. It provides access to the C3S portfolio of products through a web interface and API. The CDS includes high-resolution weather and climate reanalysis datasets such as the ERA5 and UERRA (EU-	netCDF, GRIB

⁶ <u>https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-3-slstr/product-types/level-2-aod</u> (product under development)

⁷ <u>https://amt.copernicus.org/articles/14/2771/2021/</u> and

https://www.esa.int/Applications/Observing the Earth/Copernicus/Sentinel-5P/Mapping highresolution methane emissions from space

⁸ <u>https://sentinel.esa.int/web/sentinel/sentinel-data-access</u>



	regional) multi-decadal re-analyses. Air quality and atmospheric composition are affected by weather, hence the addition of weather information often leads to an improvement in the accuracy of the downscaled data product.	
In-situ data: e.g. low- cost sensors, IoT devices, weather stations,	These data will be made available and incorporated into the service at a later date, subject to licensing agreements.	XML, JSON, CSV,
Population density and built area maps	Regridded to 0.1 degrees. Original data available at 1km global resolution.	netCDF
Orography (surface elevation)	Regridded to 0.1 degrees. High-resolution data available e.g. from EU-DEM ⁹ for the European domain (25m pixels)	netCDF, Geotiff

5.2.1.2 Algorithm

The probabilistic downscaling engine is based on a generative adversarial network (GAN). Recent GAN-based models have achieved impressive performance in artificial high-resolution "image" generation for the Earth sciences¹⁰. Crucially, GANs allow the generation of "image" (i.e., 2D dataset) ensembles that quantify the uncertainty in the high-resolution output fields.

5.2.1.3 Output Data

Table 13. Output data for Probabilistic downscaling of CAMS air quality model data

High- resolution air quality data	10km (or finer) spatial resolution for the initial deployment of the service on WEkEO and CREODIAS. This is envisaged to go up to street-level after hyper-local data sources are incorporated into the machine learning algorithms. PM2.5 is currently supported. Other variables to be added in subsequent iterations.	netCDF, CSV (for point locations of interest)
High- resolution atmospheric composition data	10km (or finer resolution). NO2	netCDF, CSV

5.2.1.4 How to use

The downscaling pipeline is available as a dockerized application that can be executed in any environment with a properly configured Docker client. Inside the docker, the pipeline can be executed with the following commands:

⁹ <u>https://land.copernicus.eu/user-corner/publications/eu-dem-flyer/view</u>

¹⁰ e.g., Ledig et al., 2016: <u>https://arxiv.org/abs/1609.04802</u>; Stengel et al., 2020:

https://www.pnas.org/content/117/29/16805; Leinonen et al., 2020: https://arxiv.org/abs/2005.10374;



Pre-process	aqgan-preprocstart-date <yyyymmddhh>end-date</yyyymmddhh>	The YAML configuration file
	config.yaml	contains
		settings, e.g. the
		size of a data "tile" (bi-res: 64
		x 64, low-res: 8 x
		8). An example
		provided in the
		source tree
	\sim	aqgan/src/config
Train the GAN model on a	aqgan-trainmodel srganstart-date <yyyymmddhh></yyyymmddhh>	Start and end dates for training
predefined	cha date (IIIIIII) config config. Jami	must be
data interval		provided. Currently only
		the srgan (super-
		resolution GAN) model is
		supported.
Generate new high-	aqgan-generatemodel srganstart-date	
resolution	config.yaml	
frames with a pre-trained		
SR-GAN model		

Table 14. Commands to execute Probabilistic downscaling of CAMS air quality model data

5.2.1.5 Exposed parameters

The user can change the model settings through the YAML configuration file. This way the model architecture can be adapted to the task at hand, new input data sources can be ingested, etc.

Table 15. Exposed parameters in Probabilistic downscaling of CAMS air quality model data

Parameter	Default Value
Input and output frame size	8 x 8 and 64 x 64
Type of super-resolution model	srgan ¹¹
Area of Interest (AoI)	lat-lon bounding box
Output path	netCDF

¹¹ inspired by <u>https://arxiv.org/abs/1609.04802</u>



Number of filters in the generator and discriminator	64 and 32
Learning rate	10^-5
Batch size	128
Type of adversarial loss	Wasserstein ¹²
Low-resolution inputs	"pm2p5", "u10", "v10", "t2m", "lsm", "sp", "z", "pm10"
High-resolution (static) inputs	"urban_frac" (urban fraction), "orog_scal" (min-max scaled orography)
High-resolution output(s)	"pm2p5" (PM2.5)
Input, output data paths	See the config file.
Output "ensemble" size	10

6 Conclusions

The present deliverable offers an overview of the current status of services and resources made available from the AI4Copernicus 4 domains (Security, Agriculture, Energy and Health): each specific section described the services and resources to support the AI4Copernicus open calls winners.

This document will be used, together with other technical documents provided by WP3 and WP4, as technical documentation to access the whole suite of resources and services provided by AI4Copernicus (general services for all the domains, geospatial tools, user manual for the EO clusters, domain specific resources).

The present document will be updated at M18: together with the selection of the open calls winners, the domain bootstrapping services and resources will be fine-tuned according to the possible requirements coming from the different selected consortia.

A deeper integration with AI4EU platform is expected in the second part of the project: an initial set of requirements from AI4Copernicus has been already submitted to AI4EU consortium, in order to take into account the needs from the EO sector in the future development of AI4EU architecture.

¹² Arjovsky et al., 2017: <u>https://arxiv.org/abs/1701.07875</u>



brati