



EXERCISE 2 – TUTORIAL

SAR for Land Applications 1 – SAR basics for Land monitoring using SNAP software

1 | Exercise outline

In this exercise, we will:

- Learn how to retrieve a land-use / land cover (LULC) classification based on Sentinel-1 image products
- Explore different types of analysis for the detection of discrete land cover classes.
- Use supervised and unsupervised classification algorithms to classify the satellite image into classes

2 | Background

Mapping land cover using Sentinel-1 data


Mapping land cover using Sentinel-1 data relies on the radar backscatter signals obtained from the satellite. The backscatter intensity, polarization, and texture of the radar imagery are analyzed to distinguish between different land cover types. Supervised classification methods involve training a classifier with labeled samples of land cover types, while unsupervised methods group pixels based on their spectral similarity without prior training. The accuracy of the land cover classification is validated using ground truth data, and the results can be refined iteratively to improve accuracy. The final land cover maps provide valuable information for various applications, including urban planning, agriculture, forestry, and environmental monitoring.



Source: <https://eo4society.esa.int/event/7th-advanced-training-course-radar-polarimetry/>

2.1 Study area and data used

For this exercise, we will use two Sentinel-1 images of the same area in northern Poland, downloaded for year 2023 from the Dataspace Copernicus Open Access Hub [[@https://dataspace.copernicus.eu/](https://dataspace.copernicus.eu/)].




However, the feature space of S1 data, referring to the number of variables available for predicting target classes, is restricted when compared to optical data. As a result, this tutorial utilizes two images captured in the same year but during different seasons. This approach enables the characterization of surfaces based on their temporal attributes to some extent. Consequently, Image 1 was obtained during the month with minimal precipitation (February), while Image 2 was acquired towards the beginning of July, a period associated with the highest precipitation.

2.2. SNAP - Open and explore product

Open SNAP Desktop, click Open Product and open 2 Sentinel-1 GRDH downloaded products with HH and HV polarizations acquired on 5 February and 11 July 2023 by double click on the zipped folders. The opened products will appear in Product Explorer window:

S1A_IW_GRDH_1SDV_20230205T050029_20230205T050054_047096_05A661_3C7A.zip
S1A_IW_GRDH_1SDV_20230711T050035_20230711T050100_049371_05EFD9_7392.zip

Click + to expand e.g. the content of product [1] from 5 February 2023, then expand Bands folder and double click on Intensity_VV band to visualize it. Depending on your computer's capabilities, this process could take some time. To minimize the data load, the next step involves creating a subset that encompasses only a portion of the dataset.



Preview of the S1 GRDH product

Note: The image appears inverted because it was captured during an ascending pass.

3.1 Pre-processing

We need to apply identical pre-processing steps to both of our scenes:

Create a subset

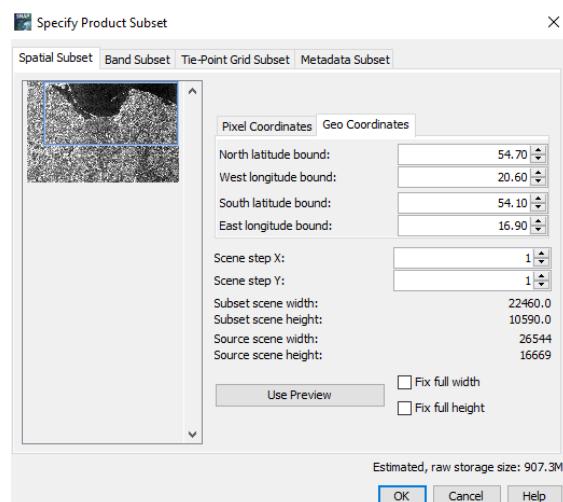
There is no need to process the whole image, instead, we can reduce the loaded data to a more manageable size – creating subset. This approach will decrease processing time in subsequent stages, especially when the analysis is concentrated on a specific area rather than the entire scene.

Navigate to Raster – Subset tab
and at Switch to “Geo Coordinates”

Enter the following numbers:

North latitude bound: 54.70
West longitude bound: 20.60
South latitude bound: 54.10
East longitude bound: 16.90

Confirm with Run and repeat the step
for the second S1 image.




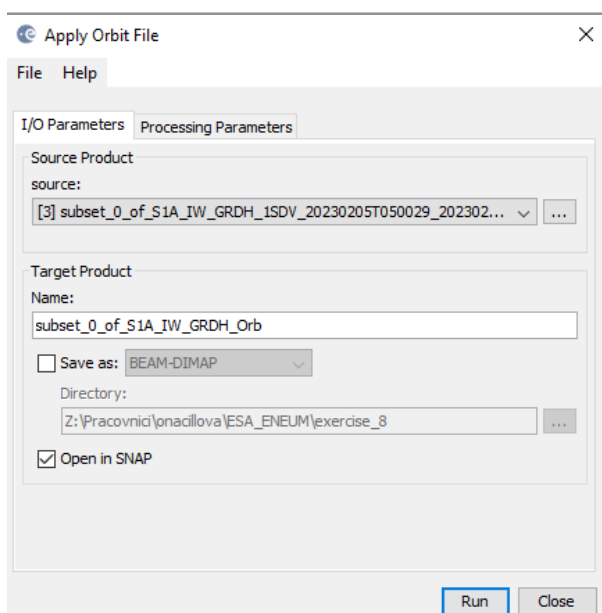
Apply orbit file

Navigate to Main Menu – Radar – Apply orbit file

In the I/O Parameters tab, select the first product and name the target product. There is no need to save the output as BEAM-DIMAP (we will save the time this way)

In the Processing Parameters accept the default settings and select the option „Do not fail if new orbit file not found”

Repeat for the second image.



Note: If precise orbits are not yet available for your product, restituted orbits can be selected which may not be as accurate as the Precise orbits but will be better than the predicted orbits available within the product

Radiometric calibration

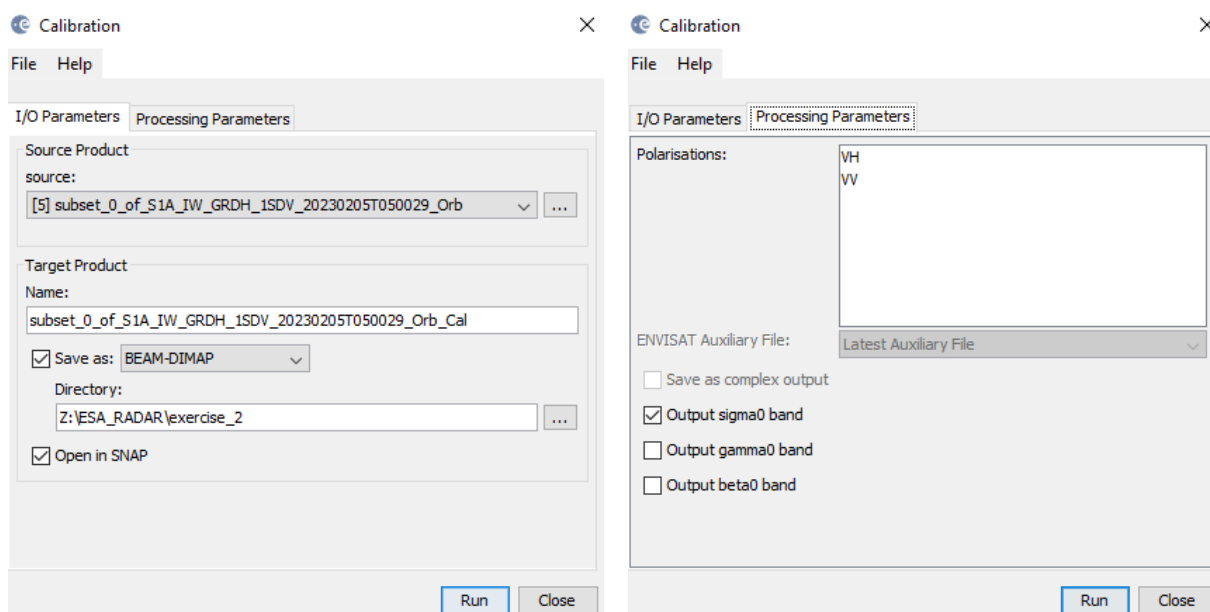
Navigate to Main Menu – Radar – Radiometric - Calibrate

In the I/O Parameters tab, select the product from the previous step (_Orb) as a source product. For the target product name, add _Cal at the end of the name as suggested In the case of this final product of preprocessing, please, save it to your folder for this exercise

In the Processing Parameters accept all default settings and then click Run.

Repeat for the second image. The saving might take some time.

Note: Thermal Noise Removal is not applied in this tutorial but should be considered when working with larger subsets or entire images.




3.2 Coregistration

After both images were radiometrically calibrated in the previous step, the coregistration brings both into one stack. Image coregistration is the process that involves aligning two or more images geometrically, ensuring that corresponding pixels depict the same area on the Earth's surface. While orbit state vectors alone can suffice to coregister images, for precise offset tracking, additional data from a digital elevation model (DEM) is necessary to enhance coregistration accuracy.

Navigate to Main Menu – Radar – Coregistration – Coregistration

In the ProductSet-Reader tab, load the last two preprocessed products (calibrated and with orbit files applied)

In the Write tab name the target product and set the directory to save it
Leave all other settings as predefined

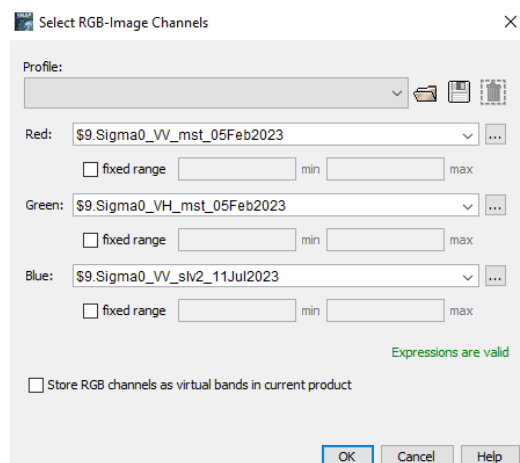


Note: If the coregistration fails it is advisable to increase the number of GCPs and also to have the initial coarse offset estimated by the operator.

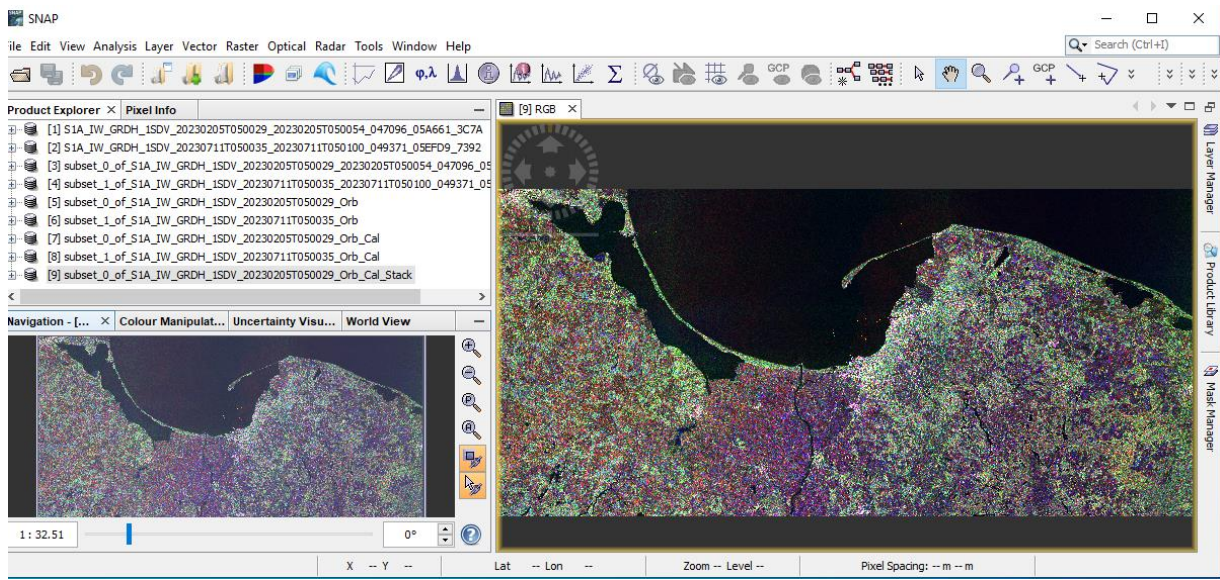
In addition to verifying the accuracy of the residuals, it is advisable to visually inspect the stack's quality. One way to do this is by creating an RGB representation of both the master and slave products, which helps determine if the images are properly aligned.

Right click on the output stack product
– Open RGB Image Window:

Choose February VV for the red channel,
February VH for the green channel,
and July VV for the blue channel.




Magnify an area featuring distinct surfaces, such as the boundary between land and water. The RGB image should be vivid and well-defined, except for alterations in land cover or scattering mechanisms that may have transpired between the first and second image acquisitions. In the figure, rivers, roads, and agricultural fields are distinctly depicted in all colors. White pixels represent high backscatter in all three bands (urban areas), while black pixels denote low backscatter values in all three bands (water). Blue pixels signify a slightly shifted shoreline in July compared to February.




Speckle Filter

Despeckling removes thermal noise introduced by the sensor from the image to remove potential sources of error in analysis.

Navigate to Radar > Speckle Filtering Single Product Speckle Filter. In the I/O Parameters select the last stack product and in the Speckle-Filter tab, choose the simple Lee Sigma filter with default window sizes.



You can compare the image pre and post-filtering by double-click on the VV polarizations and utilizing the Split Window feature found in the toolbar. As usual, different types of filters will yield somewhat varied outcomes.



Sigma0 VV before (left) and after (right) application of a speckle filter

Terrain-Correction

Navigate to Radar – Geometric – Terrain Correction – Range Doppler Terrain Correction. The last step in our graph for image pre-processing is to apply terrain correction to the product, ensuring that all the pixels are moved to the right locations (eg if the nadir angle of the image is off, it will align the pixels correctly so it is closer to a top-down view of the imagery).

Our data are still in radar geometry, moreover due to topographical variations of a scene and the tilt of the satellite sensor, the distances can be distorted in the SAR images. We need to apply Terrain Correction to compensate for the distortions and reproject the scene to geographic projection.

Navigate to Main Menu: Radar - Geometric - Terrain Correction - Range-Doppler Terrain Correction

In the I/O Parameters tab set as "Source Product" the velocity product.

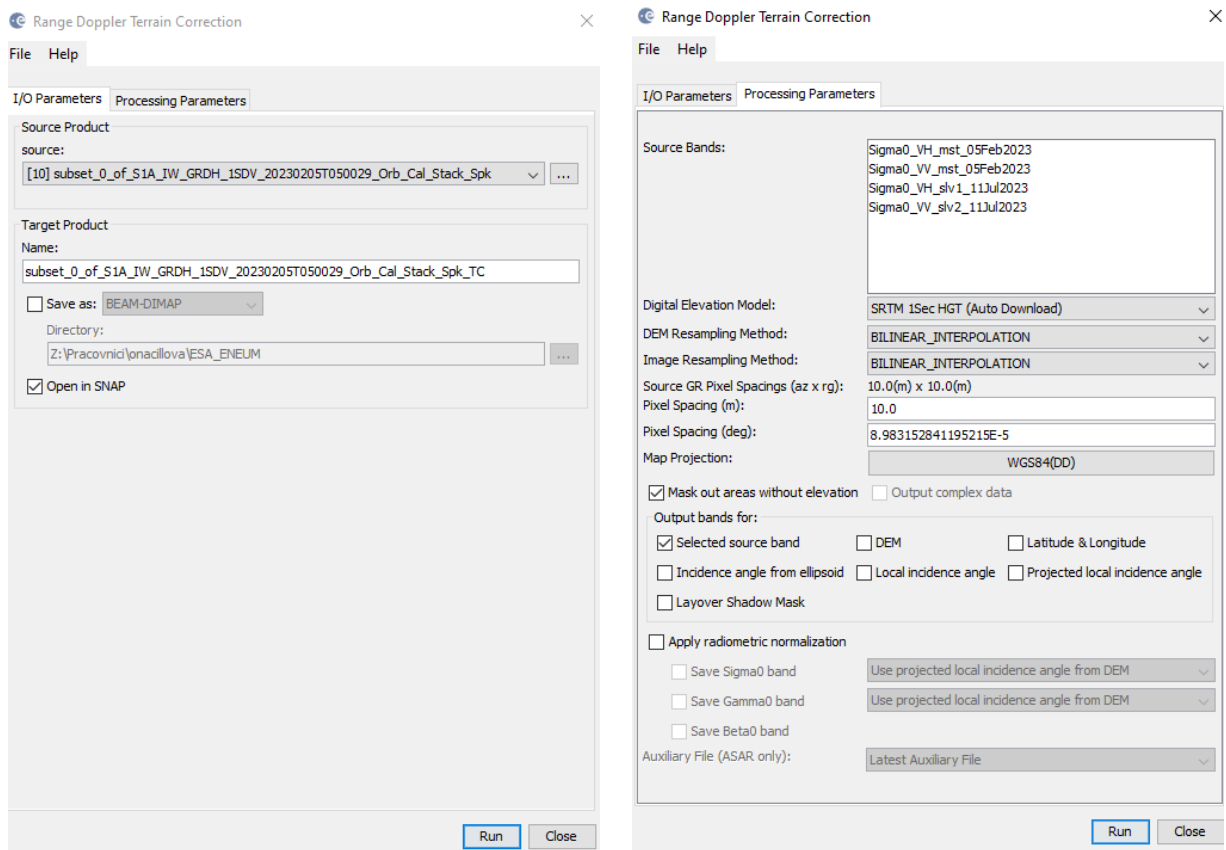
In "Target Product", keep the default name and set the Directory

In the Processing Parameters tab set:

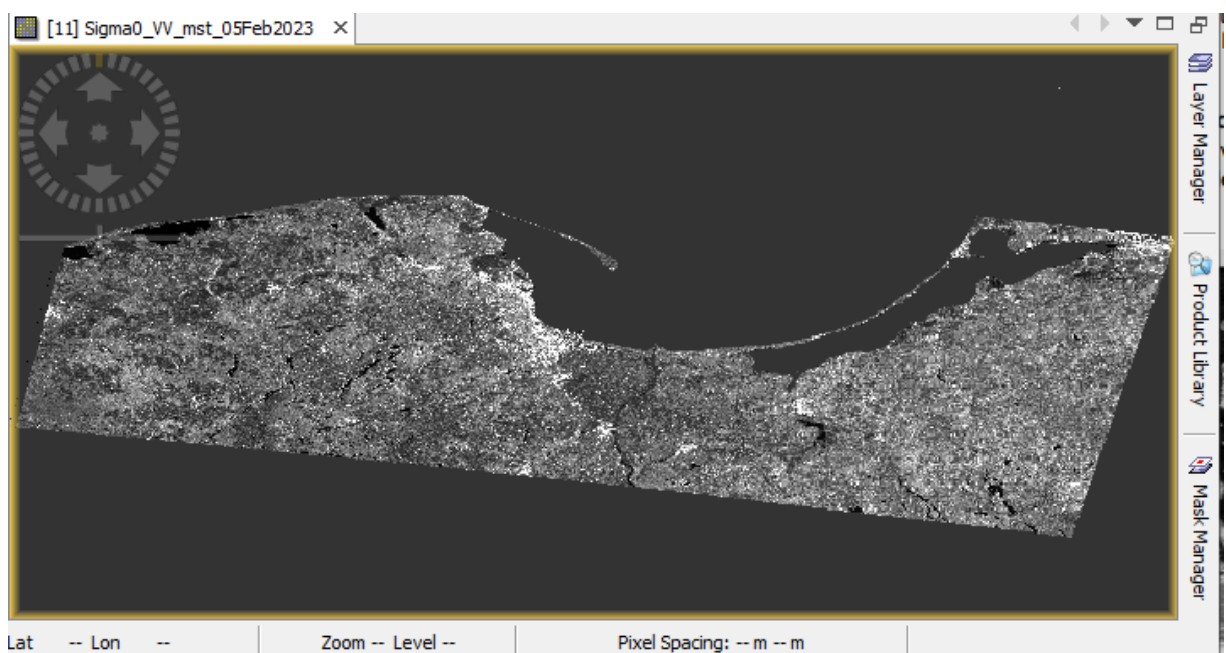
Digital Elevation Model: SRTM 1Sec HGT (AutoDownload)

Map Projection: In this tutorial, we opt for WGS84 as the Map Projection, which relies on geographic coordinates (latitude and longitude). However, for subsequent utilization in a geographic information system (GIS), projected coordinate systems like UTM (Automatic) could also be chosen.

Keep default values for the other parameters. Click Run.



Geocoded image product after Terrain Correction

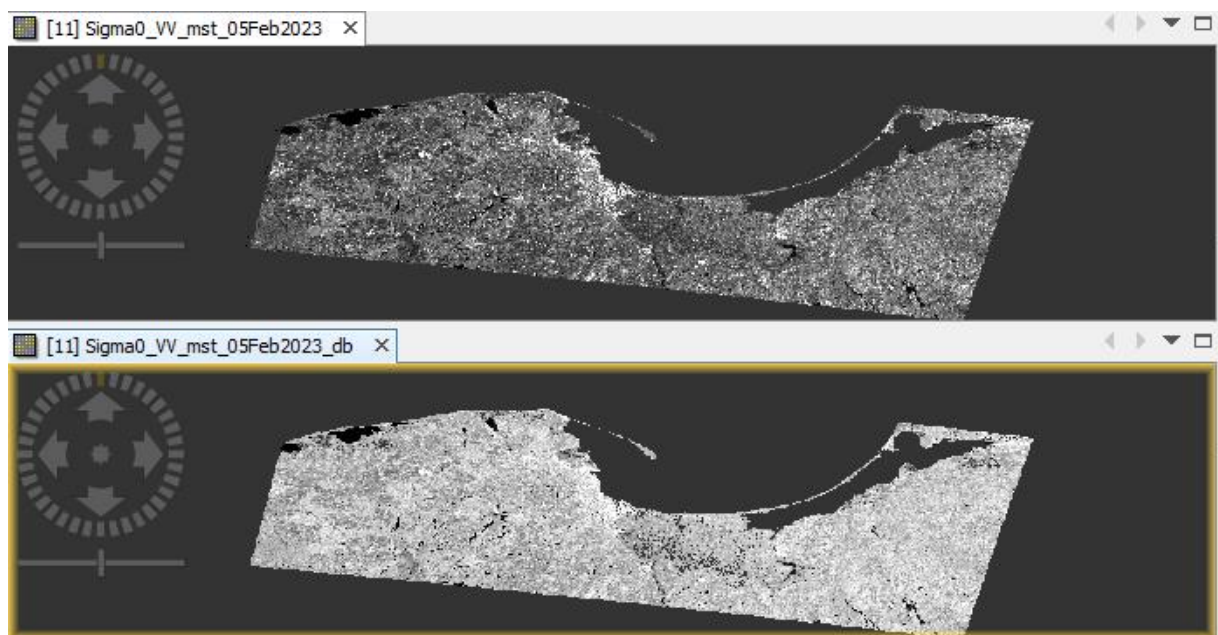


Conversion to dB scale

Sigma0_VV values typically fall within the range of 0 to 1, with mean backscatter values of 0.07 (VV) and 0.02 (VH). This indicates a prevalence of dark pixels and a scarcity of bright pixels with higher values, resulting in diminished visual contrasts from a statistical perspective.

To rectify this and achieve a more standardized distribution of values, the logarithmic function is applied to the radar image. This transformation converts pixel values into a logarithmic scale, enhancing contrasts by shifting brighter values towards the mean while expanding the range of darker values (as depicted in Figure 18, on the right). The calibrated dB data typically spans from approximately -35 to +10 dB.

To execute this transformation, right-click on each of the four terrain-corrected bands and choose **Linear to/from dB**. Confirm by selecting Yes to generate a virtual band, denoted by the symbol and the "_db" suffix. These virtual bands are not physically stored on the hard drive but can still be displayed based on the underlying mathematical expression.

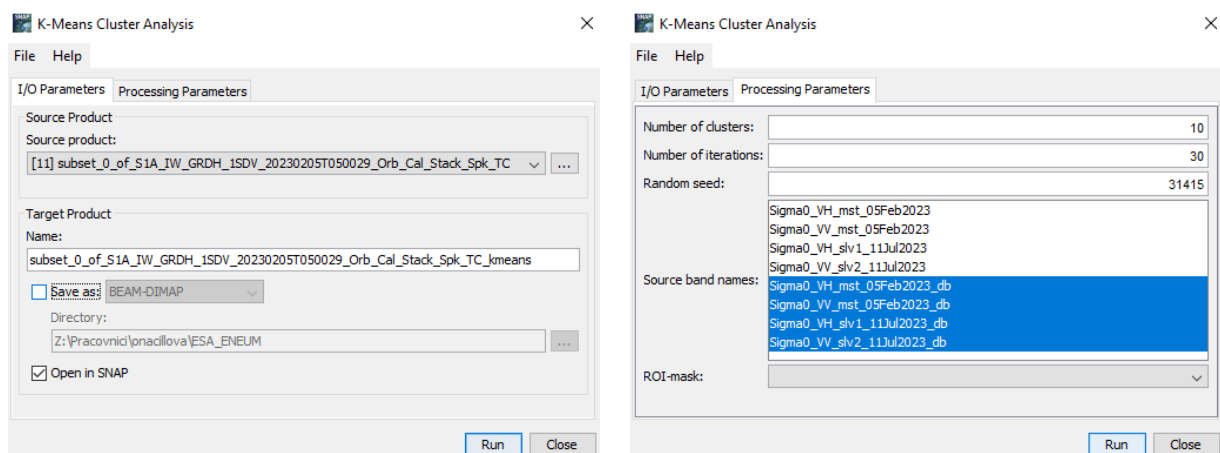


Sigma0 VV before (top) and after (bottom) conversion to dB scale

Unsupervised classification


An unsupervised classification serves as an effective method for identifying and grouping pixels that share similar characteristics.

To initiate the K-Means Cluster Analysis, navigate to **Raster > Classification > Unsupervised Classification** and select the terrain-corrected stack as the input product. In the second tab, specify the number of clusters as e.g. 5, choose all dB bands as the source bands, and confirm by clicking Run. This process may take several minutes as the classifier iterates 30 times, utilizing random cluster centers within the multi-dimensional feature space to identify pixels with similar properties.


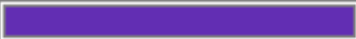
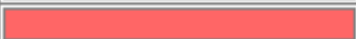




The output is a product with a single band (class_indices), where each pixel is assigned to one of the 5 clusters. While these clusters may not directly represent semantic land cover classes, they exhibit homogeneity in terms of their backscatter characteristics.

Double click to open the „class_indices“ band under Bands folder. Zoom into the area and explore the most frequent classes and their distribution.



As depicted in Figure, you can utilize the Color Manipulation tab to allocate colors and even labels to these clusters, enabling you to assess how well they align with general land use and land cover classifications.

| Label | Colour | Value | Frequency | Description |
|---------|---|-------|-----------|----------------------|
| class_1 |  | 0 | 53.205% | Cluster 0, Center... |
| class_2 |  | 1 | 29.475% | Cluster 1, Center... |
| class_3 |  | 2 | 16.913% | Cluster 2, Center... |
| class_4 |  | 3 | 0.397% | Cluster 3, Center... |
| class_5 |  | 4 | 0.010% | Cluster 4, Center... |

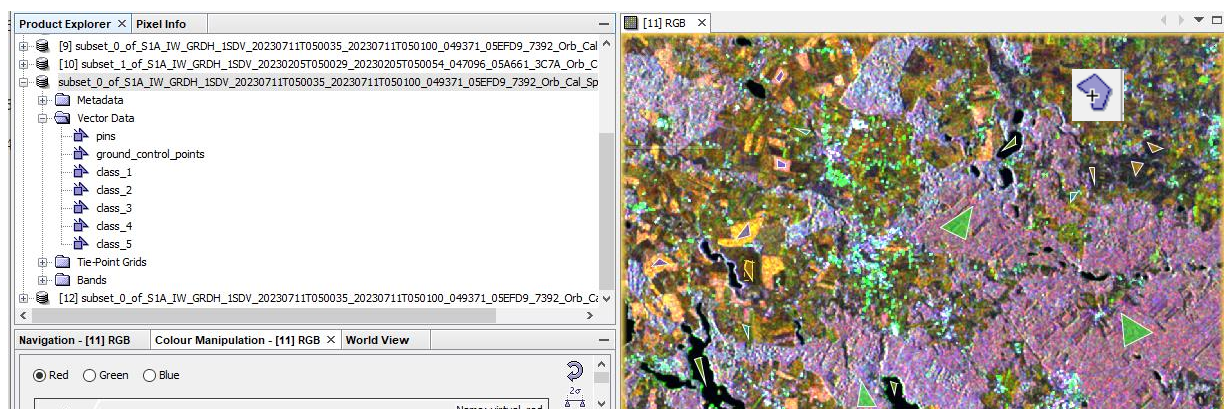
The benefit of employing an unsupervised classification lies in its independence from prior knowledge about the study area or the data, yet it effectively groups pixels with similar traits. However, the drawback is that the usefulness of the result heavily depends on the chosen number of classes. As illustrated in the example, water is segmented into five distinct classes (3-7), whereas urban areas are consolidated into a single class (8). If there is overlap between classes, it may be necessary to opt for a higher number of clusters and then merge them afterward. Conversely, having too many classes can complicate the identification of patterns within the data.

Supervised classification

Now, we will perform the one type of supervised classification – Random forest classification. For this we will create some training data by identifying polygons containing the different land cover types:

- Select the last subset product by clicking on it. Navigate to Vector > New Vector Data Container.
- Create new data containers – e.g. for 5 different classes.
- Start with creating the data container for, e.g. „class1“.Click OK.

The newly created vector data container should appear under "Vector Data" in the Product Explorer. Now let's start creating training polygons for these different containers on ice type using the Polygon drawing tool:





Random forest classifier

Navigate to Raster > Classification > Supervised classification > Random Forest classifier

- In the Random-Forest-Classifer select the all Training vectors that we have created.
- Select at least two Feature bands (polarisations)
- Click on RUN.

The output final product should be now opened in the Product Explorer Window. Double-click on the „LabeledClasses“ band under Bands folder to visualize it.



Compare the results of supervised and unsupervised classification.

THANK YOU FOR FOLLOWING THE EXERCISE!