TERRASCOPE SENTINEL-1

ALGORITHM THEORETICAL BASE DOCUMENT (ATBD)

S1 – SIGMA0 GRD – V130

Reference: *Terrascope Sentinel-1 Algorithm Theoretical Base Document S1 – SIGMA0 – V130*
Authors: Kristof Van Tricht, Lisa Landuyt, Sarah Gebruers, Carolien Toté, Simon Van Baelen
Version: 1.3
Date: 15/04/2025

DOCUMENT CONTROL

Signatures

Authors	Kristof Van Tricht, Lisa Landuyt, Sarah Gebruers, Carolien Toté, Simon Van Baelen
Reviewer	Jurgen Everaerts
Approvers	Dennis Clarijs, Jurgen Everaerts
Issuing authority	VITO

Change record

Release	Date	Updates	Approved by
1.0	10/09/2019	First external version	Jurgen Everaerts
1.3	15/04/2025	TERRASCOPE Sentinel-1 Version 130	Jurgen Everaerts

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

ACRONYM	EXPLANATION
ATBD	Algorithm Theoretical Base Document
COG	Cloud-Optimized GeoTIFF
DEM	Digital Elevation Model
DN	Digital Number
EO	Earth Observation
ESA	European Space Agency
GDAL	Geospatial Data Abstraction Library
GEOTIFF	Geospatial Tagged Image File Format
GIS	Geographical Information System
GRD	Ground Range Detected
IW	Interferometric Wide Swath
LUT	Lookup table
NRT	Near-real time
PDF	Product Distribution Facility
S1A/B	Sentinel-1A/B
SAR	Synthetic Aperture RADAR
SLC	Single Look Complex
SNAP	Sentinel Application Platform
SRTM	Shuttle RADAR Topography Mission
VV/VH	Vertical-Vertical/Vertical-Horizontal polarization mode

1. Introduction

1.1. Terrascope explained

Terrascope is the Belgian platform for Copernicus Sentinel, PROBA-V and SPOT-VEGETATION satellite data, products, and services. It provides easy, full, free and open access to all users without restrictions. This allows non-specialist users to explore the wealth of remote sensing information and build value-added products and services.

The following data are included:

- Sentinel-1 SAR data (sigma0 and coherence) over Belgium and its surroundings
- Sentinel-2 multispectral data over Europe and parts of Africa
- Sentinel-3 multispectral and thermal Synergy (SYN) – Vegetation (VGT) and Land Surface Temperature (LST) data
- Sentinel-5P atmospheric composition data
- The SPOT-VEGETATION archive
- The PROBA-V archive

For Sentinel-1, a sigma0 backscatter product is provided, obtained by further processing the Level-1 GRD products provided by the European Space Agency (ESA).

Users have the possibility to build derived information products to their own specification, using the Terrascope processing cluster through provided virtual machines or notebooks, and via OpenEO. This eliminates the need for data download (and consequential storage costs), because the cluster holds all of the data in a directly accessible, analysis-ready format. Integration of data or products in your own application is facilitated through catalog (opensearch and STAC) and Open Geospatial Consortium (OGC) web services.

Terrascope is user centered, so any suggestions for new or enhanced functionality are welcome. Feel free to contact us: info@terrascope.be.

1.2. Scope of Document

This ATBD (Algorithm Theoretical Base Document) describes the origin and generation of the Sentinel-1 (S1) Level-1 GRD sigma0 data product.

The document is organised as follows:

- Section 2 provides an overview of the input data needed for the processing workflow.
- Section 3 explains the data available to users.
- Section 4 provides a detailed description of the different processing algorithms.

- Section 5 justifies the overall workflow with a quality assessment.

1.3. Description

The Sentinel-1 Level-1 GRD sigma0 product is an orbit-corrected, calibrated and geometrically corrected version of the original ESA Level-1 GRD product.

The processing workflow takes the ESA Level-1 GRD product as input. This GRD image is first corrected based on more precise orbit information than what is available in the original image. Next, radiometric noise is removed from the borders of the image, and thermal noise is removed as well. The fourth step in the processing workflow is the radiometric calibration, resulting in the sigma0 backscatter image. From this backscatter, the geometric distortions are removed through a Range Doppler terrain correction. All these steps are performed using the SNAP (the ESA Sentinel Application Platform) toolbox, more details are provided in the remainder of this document. The final product encompasses GeoTIFF files for both the VV and VH polarizations as well as the incidence angle. The backscatter intensities are provided in the linear scale (cfr. Section 3.1.4).

This document applies to the Terrascope S1 sigma0 V130 processing chain. Table 1.1 summarizes the main characteristics of the different Terrascope product versions until V130. Validation results between both versions are included in Section 5.1.

Table 1.1: Summary of main characteristics of different Terrascope versions until V130.

Version	Main characteristics
V100	Initial version
V110	• Range Doppler Terrain Correction: Use of SRTM DEM • Border Noise Removal: Post-processing using scene footprint to remove remaining artefacts • Processing chain: SNAP 6.0
V130	• Range Doppler Terrain Correction: Use of Copernicus DEM instead of the SRTM DEM • Border Noise Removal: Process only include for images acquired before 13/08/2018, parameters updated to remove remaining artifacts • Processing chain: SNAP 11.0, GDAL 3.8.4, Python 3.11 • Refactoring of code

1.4. Feature added value/use case

Due to both radiometric and geometric distortions, the Level-1 data provided by ESA cannot be used as such. This GRD backscatter product is created to lower the threshold for users that want to work with Sentinel-1 radar data. The calibrated sigma0 backscatter product allows to compare acquisitions at different times and between different sensors. An important not here is that, due to incidence angle and look direction induced geometric effects, one should always select images from the same relative orbit for temporal analyses.

2. Input data

2.1. General

The Terrascope processing starts from the Sentinel-1 Level-1 GRD data products. These can be freely downloaded from the Copernicus Dataspace Ecosystem (https://browser.dataspace.copernicus.eu/).

2.2. Ancillary data

2.2.1. Orbit files

As the orbit state vectors provided in the metadata of a SAR product are generally not completely accurate, an orbit correction needs to be applied (cfr. Section 4.1). In order to allow for near-real time (NRT) processing, the Terrascope Sentinel-1 sigma0 V130 workflow uses restituted orbit files. These are auto-downloaded during the workflow.

2.2.2. Digital elevation model (DEM)

In order to correct for geometric distortions, Range Doppler terrain correction is applied (cfr. Section 4.5). To do so, a DEM is required. The Terrascope Sentinel-1 sigma0 V130 uses the Copernicus DEM, which is available at a resolution of 30m. It is auto-downloaded during the workflow.

3. Output

3.1. Product layers

3.1.1. Product data

The sigma0 backscatter products generated and distributed by Terrascope include 3 outputs, each of which is formatted as a single layer compressed Cloud-Optimized GeoTIFF (COG) format. Each output is provided in a 10m resolution. Figure 3.1 shows the S1 sigma0 product file list.

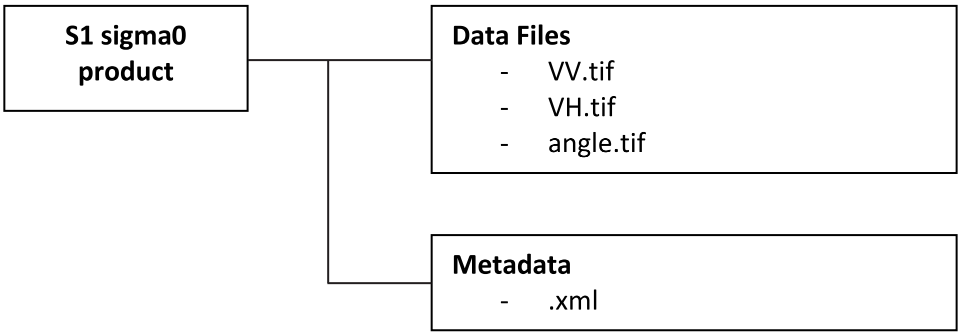

Figure 3.1: S1 sigma0 product file list.

3.1.2. Note on the use of original intensities as opposed to decibels

It is of extreme importance that the user of this product is aware that the backscatter intensities in our sigma0 product are **NOT expressed in decibels**, as opposed to how it is generally done by other providers. While it is true that conversion to decibels turns backscatter intensities in a quantity that is more user-friendly, has a higher value visually, and is a generally accepted method of communicating radar data, we have chosen specifically to not do this for the following reasons:

- In the growing world of datacubes, including the work we do with Terrascope services, ingesting decibel backscatters imposes serious error-prone complexities, as any kind of analysis that would be done on these data would require an on-the-fly removal of the decibel unit.
- Providing decibel backscatters to end users would require the user to convert to original intensities when doing any kind of analysis that requires mathematical operations.
- Calculating overview levels on the GeoTIFFs based on decibels results in wrong values that cannot be used in any analysis.

This all comes back to the basic mathematical principle that no mathematical operations such as calculating an average value can be done directly on the (logarithmic) decibels. Doing so results in mathematically incorrect results, of which users are often not aware. To prevent this as much as possible, our sigma0 product is provided in original backscatter intensities. Users that do want to convert to decibels, can apply a simple conversion function:

$$\sigma_{0,dB} = 10 \log_{10} \sigma_0$$

where σ_0 is the sigma0 backscatter intensity in linear scale as provided in our product, and $\sigma_{0,dB}$ is the sigma0 value in decibels.

3.1.3. Data collection structure and naming convention

All Terrascope Sentinel-1 sigma0 products are stored in the **CGS_S1_GRD_SIGMA0_L1** collection. In the Terrascope Virtual Machine environment, this data can be found in */data/MTDA/CGS_S1/CGS_S1_GRD_SIGMA0_L1/<Year>/<Month>/<Day>/<Product_ID>.*

The <Product_ID> is constructed as follows:
<mission_identifier>_<mode>_<product_type>_<PP>_<start_date>_<orbit_direction>_<relative_orbit>_<product_unique_identifier>_<product_version>

Template	Description
mission_identifier	S1A or S1B
mode	Always IW: interferometric wide swath
product_type	GRDH_SIGMA0: Ground Range Detected, Sigma0; H = High resolution
polarisation mode (PP)	Always DV: dual polarisation (VV-VH)
start_date	Product start date in the format yyyymmddThhmmss
orbit_direction	ASCENDING or DESCENDING, always the same for a given relative orbit
relative_orbit	relative orbit number
product_unique_identifier	We reference the unique identifier of the base product, to establish a 1-to-1 relationship with ESA products.
product_version	VITO product version

Example: S1A_IW_GRDH_SIGMA0_DV_20250401T055056_DESCENDING_37_B0B7_V130

3.1.4. Product metadata

The VV and VH backscatter intensity outputs are stored in float32 format, while the incidence angle is scaled and stored in int16 format. The physical incidence angle values are converted from floating point values into integers, mainly to reduce the file sizes. The physical number can be defined by using the following formula:

Physical Value = Scaling * Digital Number + Offset,

with Scaling set to 0.0005 and Offset to 29.

All product metadata, including the Scaling and Offset, are provided in the accompanying xml file.

3.2. Product version

Terrascope products are produced in a controlled way. Every product has a version indicator, consistent with the Semantic Versioning 2.0.0 protocols (https://semver.org/). The version indicator has three digits: XYZ.

- X is 0 during prototyping and pre-operational use. X is 1 for the first operational setup, and increments when results are no longer backward compatible (i.e. any further processing will have to be adapted to deal with e.g. format changes, value scaling, etc.).
- Y is reset to 0 with an X increment. Y increments when functionality is added, but backward compatibility is guaranteed (e.g. when a different approach is taken for atmospheric or geometric correction).
- Z is reset to 0 when Y increments. Z increments when the software is patched (bug fixed) without any functional changes.

The current Terrascope Sentinel-1 sigma0 workflow version is V130.

Whenever X or Y changes, the impact of the updates will be reported. Users are informed about version changes through the Terrascope newsletter (to subscribe: https://terrascope.be/en/stay-informed).

3.3. Product data access

The Terrascope S1 data products can be accessed through:

- Web services: Web Map Service (WMS) and Web Map Tile Service (WMTS):
  https://bit.ly/TerrascopeFAQ_WMTS
  Protocols for downloading images and integrating them into GIS software
- Notebooks (login required): https://notebooks.terrascope.be/hub/login
  Programming environment to quickly access and edit data
- Virtual Machines (VM) (login required): https://terrascope.be/en/form/vm
  External computer used to view data and process it in the cloud

- OpenEO API (login required): https://openeo.org/documentation/1.0/python/
  Python API to automate satellite data processing in the cloud

The details of each of these access points are described on https://terrascope.be/en/services.

4. Methodology

The following algorithm steps are performed using SNAP, more specifically using the command line "gpt" tool of SNAP. For each process in the workflow, the following sections describe what the operator in SNAP does, what it needs as input, and what it delivers as output. Following these SNAP algorithms, a few more gdal steps are required to save the result to GTIFF, apply the scaling, correctly set the nodata value, add overviews and provide the metadata.

The SNAP processing steps in Terrascope are performed using SNAP 11.0. Part of the information provided in the steps below is taken from operator descriptions in SNAP. We refer the reader to these documents for additional information on the processes involved.

4.1. Orbit correction

The orbit state vectors provided in the metadata of a SAR product are generally not completely accurate and can be refined using orbit files, which provide accurate satellite position and velocity information. Based on this information, the orbit state vectors in the abstract metadata of the product are updated.

For Sentinel-1, two types of orbit files, i.e. Restituted and Precise orbit files, are published. Precise orbits provided guaranteed quality but are produced only a few weeks after acquisition, while Restituted orbit files are provided within 3 hours from sensing (nominally used for the systematic processing of Fast24 products). While the accuracy of these orbit files is less guaranteed, their quality has proven to be better (twice as good) than expected in 95% of the cases, resulting in the very good absolute geolocation accuracy of Sentinel-1 products. Given the timeliness of the product, the high absolute accuracy already with restituted orbits, and the high additional load and complexity of reprocessing every image once the precise orbit becomes available, we provide our sigma0 product based on restituted orbit files.

Orbit files are automatically searched and downloaded from https://qc.sentinel1.copernicus.eu/. To refine the orbit state vectors, the following steps are performed by SNAP:

- Get the start time of the source product;
- Find orbit file with user specified type and the product start time;
- For each orbit state vector, get its zero Doppler time;
- Compute new orbit state vector with 8th order Lagrange interpolation using data for the 9 nearest orbit positions around the zero Doppler time.

4.2. Border noise removal

The Sentinel-1 Instrument Processing Facility (IPF) is responsible for generating the complete family of Level-1 and Level-2 operation products. The processing of the RAW data into Level-1 products features a number of processing steps, leading to artefacts at the image borders. These processing steps are mainly the azimuth and range compression and the sampling start time changes handling that is necessary to compensate for the change of the Earth curvature. The latter process is generating a number of leading and trailing "no-value" samples that depend on the data-take length which can be several minutes long. The former process creates radiometric artefacts that complicate the detection of the "no-value" samples. These "no-value" pixels are not null but contain very low values which complicates the masking based on thresholding.

This operator implements an algorithm proposed in [1] allowing the masking of the "no-value" samples efficiently with a thresholding method.

4.3. Thermal noise removal

Thermal noise correction can be applied to Sentinel-1 Level-1 SLC products as well as Level-1 GRD products which have not already been corrected. The operator can also remove this correction based on the product annotations (i.e. to re-introduce the noise signal that was removed). Product annotations will be updated accordingly to allow for re-application of the correction.

Level-1 products provide a noise LUT for each measurement data set. The values in the de-noise LUT, provided in linear power, can be used to derive calibrated noise profiles matching the calibrated GRD data. Bi-linear interpolation is used for any pixels that fall between points in the LUT.

4.4. Radiometric calibration to sigma0

The objective of SAR calibration is to provide imagery in which the pixel values can be directly related to the radar backscatter signal of the scene. Though uncalibrated SAR imagery is sufficient for qualitative use, calibrated SAR images are essential to quantitative use of SAR data.

Typical SAR data processing, which produces Level-1 images, does not include radiometric corrections and so a significant radiometric bias remains. Therefore, it is necessary to apply the radiometric correction to SAR images so that the pixel values of the SAR images truly represent the radar backscatter of the reflecting surface. The radiometric correction is also necessary for the comparison of SAR images acquired with different sensors or acquired from the same sensor but at different times, in different modes, or processed by different processors.

For converting digital pixel values to radiometrically calibrated backscatter, all the required information can be found in the product. A calibration vector is included as an annotation in the product allowing simple conversion of image intensity values into sigma or gamma nought values.

The objective of SAR calibration is to provide imagery in which the pixel values can be directly related to the radar backscatter of the scene. To do this, the application output scaling applied by the processor must be undone and the desired scaling must be applied. Level-1 products provide four calibration Look Up Tables (LUTs) to produce β_{0i}, σ_{0i} and γ_{0i} or to return to the Digital Number (DN). The LUTs apply a range-dependent gain including the absolute calibration constant. For GRD products, a constant offset is also applied.

In this processing chain we choose to calibrate the raw signals in the product to the widely used calibrated sigma0 backscatter intensities.

4.5. Range Doppler terrain correction

Due to relief variations of a scene and the side-looking character of the satellite sensor, distances can be distorted in the SAR images. Terrain corrections are intended to compensate for these distortions so that the geometric representation of the image will be as close as possible to the real world.

The process needs a DEM based on which the correction is applied. The more accurate the DEM, the better the correction. For the Terrascope Sentinel-1 sigma0 V130 product, the Copernicus DEM is used. Furthermore, we can define the map projection (~ellipsoid) based on the which the correction will be done. For Belgium, we use WGS 84 / UTM zone 31N with a pixel spacing of 10m, in accordance with the Sentinel-2 data in Belgium as also provided by Terrascope.

5. Quality assessment

5.1. Terrascope V130 vs. V110

In the latest update of the Terrascope Sentinel-1 sigma0 workflow, the DEM used for the Range Doppler terrain correction was updated from the SRTM DEM to the Copernicus DEM. Due to differences between both elevation models, a minor unsystematic bias can prevail in both the VV and VH product.

To quantify the differences between both versions, as well as the impact on downstream analysis, a statistical consistency analysis was performed for all imagery obtained in October 2024. Figure 5.1, Figure 5.2 and Figure 5.3 show the comparison between the V110 and V130 for VH backscatter intensity, VV backscatter intensity and incidence angle respectively. A subsampling of 1/100th points was used.

These figures show minor deviations for both the VV and VH backscatter intensities, while no differences were observed for the incidence angle. However, the backscatter deviations are not biased (cfr. left and right subplots) and don't impact the statistical distribution of the data (cfr. center subplots). With R^2-values 0.997 and a mean bias of 0.0, we can conclude that both versions are statistically consistent, and the observed differences have no significant impact on downstream analyses.

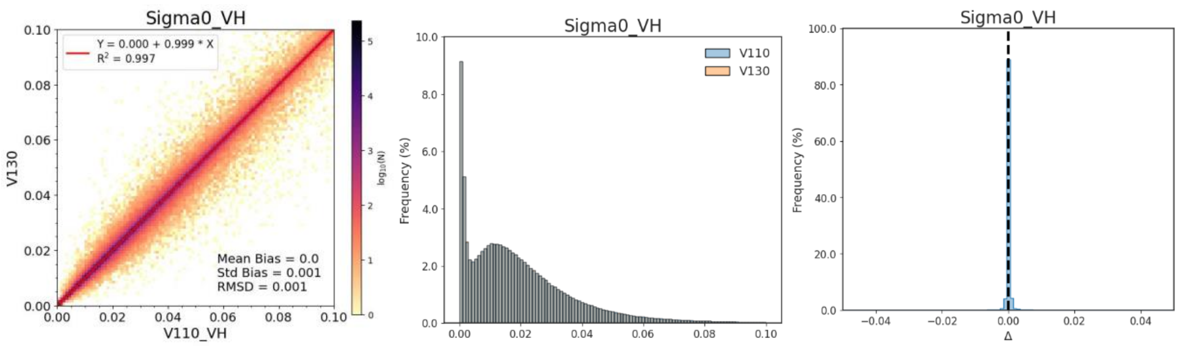

Figure 5.1: Comparison of V110 and V130 versions for VH intensity: scatter density plot (left), histogram overlay (center) and histogram of absolute differences (right).

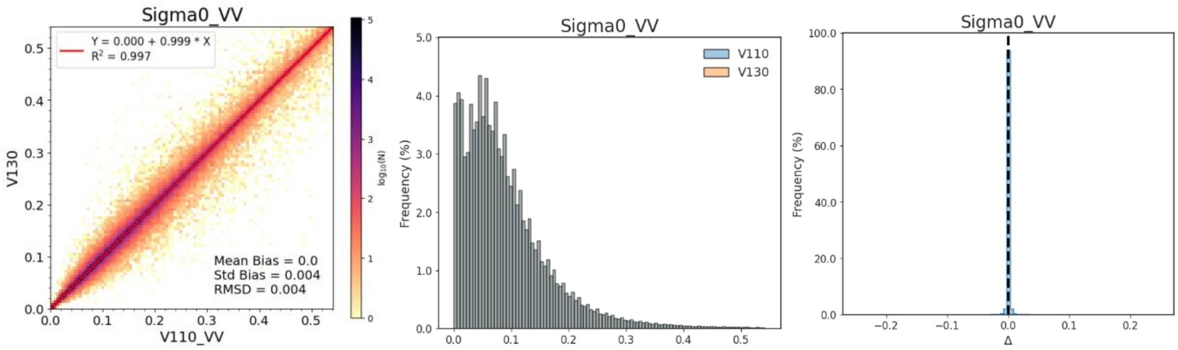

Figure 5.2: Comparison of V110 and V130 versions for VV intensity: scatter density plot (left), histogram overlay (center) and histogram of absolute differences (right).

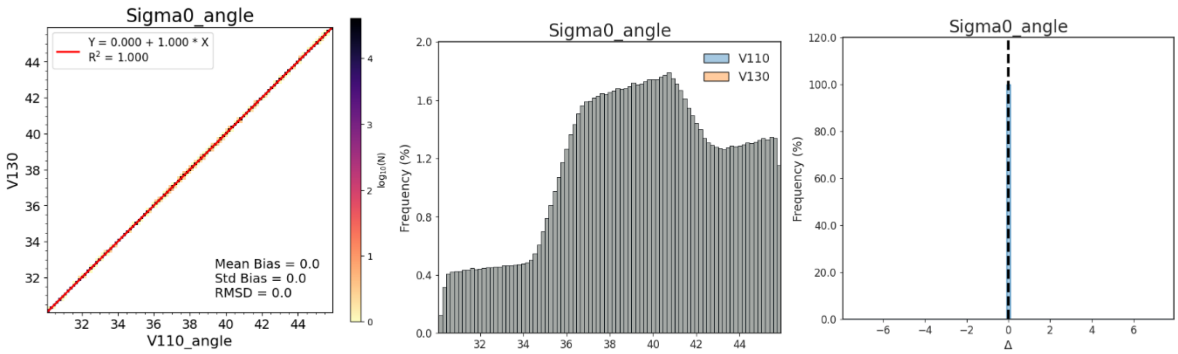

Figure 5.3: Comparison of V110 and V130 versions for incidence angle: scatter density plot (left), histogram overlay (center) and histogram of absolute differences (right).

5.2. Terrascope V110 vs. Google Earth Engine

It should be noted that a quantitative assessment of a radar product is difficult, since measurements are extremely sensor, location, and time specific. No ground truth data are available to compare with radar backscatter intensities.

To verify the consistency of our product, we have chosen to compare it to the sigma0 product that is available within the Google Earth Engine archive, which gives the advantage that it has been thoroughly evaluated by the community, while their processing steps are - for the test images - identical to what is processed with our processing workflow. It should be mentioned that Google Earth Engine automatically determines the most appropriate UTM zone in which the S1 image is processed. For the test images, it has been verified that this corresponds to our UTM31 (EPSG:32631) which we use for all processed S1 imagery over Belgium. For some other images we generate (e.g. near eastern Belgium), Earth Engine will provide images in UTM zone 32, and slight differences with our output will occur, mostly due to different results in the geometric terrain correction step.

Google at first decided to save on storage and produce VV/VH backscatter in UInt16. To retain the required precision, they had to clamp the original values and remove the 1st and 99th percentiles. This has resulted in numerous discussions with the community and Google is now reprocessing the entire archive to Float32, because the removed 1st and 99th percentiles actually contain important information for some users. This is why we process the backscatters here to Float32, so no information loss occurs. In the comparison mentioned here, and the provided reference data from Earth Engine, we already use their new Float32 collection, so we don't expect any differences due this issue.

In our comparison, we have randomly taken 10,000 points in the following Sentinel-1 Sigma0 images:

- **S1A_IW_GRDH_SIGMA0_DV_20171024T173238_ASCENDING_161_F229_V110_VV**
- **S1A_IW_GRDH_SIGMA0_DV_20171024T173238_ASCENDING_161_F229_V110_VH**.

Next, we have extracted the exact same points from the Google Earth Engine corresponding Sigma0 image. The scatterplots in Figure 5.4 show the very close correspondence between both sources of sigma0. These images confirm that the output provided by the Terrascope processing chain is within expectations of what other providers deliver, although especially designed to work over Belgium, and additionally cleaned near the borders.

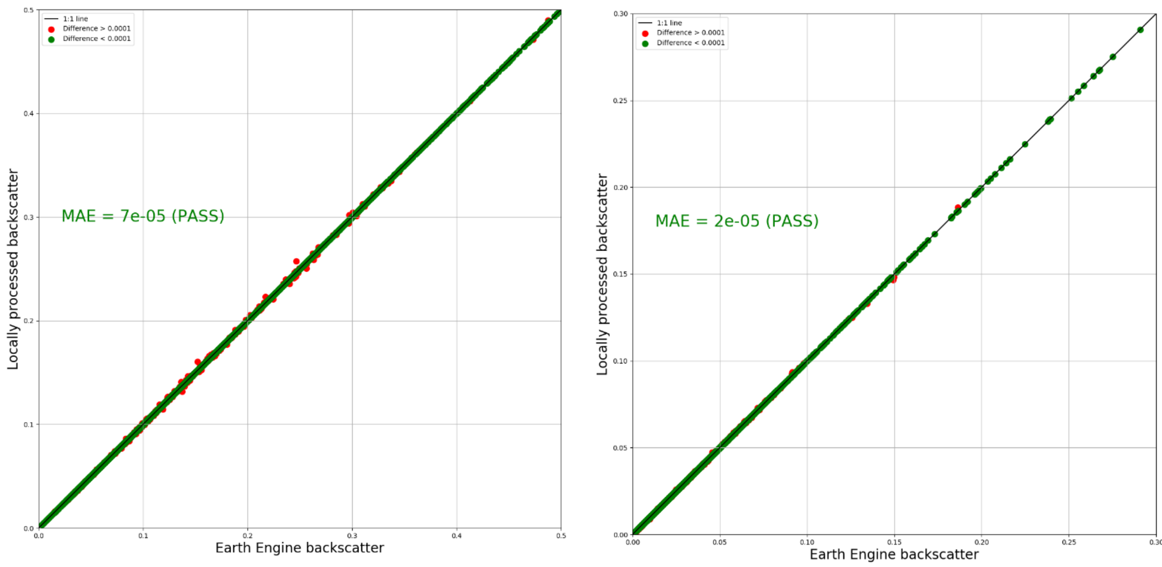

Figure 5.4: Comparison of 10,000 randomly extracted pixel values of the Terrascope Sigma0 product (y-axis) and the Google Earth Engine Sigma0 product (x-axis) for the VV (left) and VH (right) band.

6. References

[1] Masking "No-value" Pixels on GRD Products generated by the Sentinel-1 ESA IPF, Reference MPC-0243, Issue 1.0, Date 2015, June 25.