




Contents lists available at ScienceDirect

ISPRS Journal of Photogrammetry and Remote Sensing

journal homepage: www.elsevier.com/locate/isprsjprs

Utilising Sentinel-1's orbital stability for efficient pre-processing of sigma nought backscatter

Claudio Navacchi ^{a,*}, Senmao Cao ^b, Bernhard Bauer-Marschallinger ^a, Paul Snoeij ^c, David Small ^d, Wolfgang Wagner ^a

^a Department of Geodesy and Geoinformation, TU Wien, 1040 Vienna, Austria

^b Earth Observation Data Centre for Water Resources Monitoring (EODC), 1030 Vienna, Austria

^c Airbus Defense and Space Netherlands, 2333 CS, Leiden, The Netherlands

^d Department of Geography, University of Zurich, CH-8057 Zurich, Switzerland

ARTICLE INFO

Keywords:

Sentinel-1
Synthetic Aperture Radar (SAR)
Ground Range Detected (GRD)
Georeferencing
Orbital tube

ABSTRACT

For already more than seven years, the Sentinel-1 C-band Synthetic Aperture Radar (SAR) mission has been providing indispensable information for monitoring bio-geophysical parameters at fine temporal and spatial scales. As many applications require backscatter datacubes as input, enormous amounts of data have to be radiometrically and geometrically corrected to be in a common, Earth-fixed reference system. Pre-processing workflows accomplishing this task have already been established and are implemented in several software suites. However, typically, these workflows are computationally expensive which may lead to prohibitively large costs when generating multi-year Sentinel-1 datacubes for whole continents or the world. In this paper, we discuss existing approaches for generating sigma nought and projected local incidence angle (PLIA) data and present simplifications of the overall workflow relying on the unprecedented orbital stability of Sentinel-1. Propagating orbital deviations through the complete Sentinel-1 pre-processing pipeline helped us to simulate and identify PLIA as a static layer per relative orbit. The outcome of these simulations also provided the necessary information to replace iterative root-finding algorithms for determining the time of closest approach (TCA), i.o.w. the azimuth index, with a linear one — at no expense of radiometric accuracy. All experiments were performed using an in-house developed toolbox named *wizard*, which made it possible to speed up Sentinel-1 pre-processing by approximately 4–5 times with respect to the Sentinel Application Platform (SNAP). This could pave the way for producing quality-curated, large-scale backscatter datacubes at continental and global scales in acceptable time frames.

1. Introduction

The Copernicus programme and its Sentinel satellites have heralded a new age for remote sensing applications, and fuel manifold earth observation (EO) services with Terabytes of data per day. Sentinel-1A initiated the Sentinel series with its launch in April, 2014, and is, together with its twin satellite Sentinel-1B, an indispensable synthetic aperture radar (SAR) system providing excellent coverage for observing bio-geophysical processes globally at fine spatial and temporal scales. Recent studies were already able to demonstrate the suitability of the C-band sensors' backscatter data for monitoring, e.g. surface soil moisture (Bauer-Marschallinger et al., 2018), forest type (Dostálová et al., 2021), vegetation (Vreugdenhil et al., 2020), snow depth (Lievens et al., 2021), and floods (Twele et al., 2016).

These types of applications usually rely on a dense time series of co-located pixel values, which require geometric and radiometric

corrections to support time series analysis and data interoperability. For what concerns the SAR data, the most fundamental steps to fulfil these requirements are geocoding, radiometric calibration, and imaging noise adjustments, e.g. to mitigate thermal or border noise, on Level-1 SAR images. Several publications (Filipponi, 2019; Mullissa et al., 2021; Truckenbrodt et al., 2019) present guidelines for how to weave these operations together to create a pre-processing workflow with calibrated and georeferenced backscatter datacubes as output. They rely on modules from toolboxes like GAMMA (GAMMA Remote Sensing, 2021), SNAP (ESA, 2021b), or ISCE2 (ISCE-Framework, 2021) and give some interesting insights into data quality benchmarking when using the same workflow across different software solutions. These SAR software packages often support several satellite missions and contain many modules dedicated to, e.g. geometric, interferometric, polarimetric, or

* Corresponding author.

E-mail address: claudio.navacchi@geo.tuwien.ac.at (C. Navacchi).

<https://doi.org/10.1016/j.isprsjprs.2022.07.023>

Received 18 May 2022; Received in revised form 19 July 2022; Accepted 26 July 2022

Available online 19 August 2022

0924-2716/© 2022 The Author(s). Published by Elsevier B.V. on behalf of International Society for Photogrammetry and Remote Sensing, Inc. (ISPRS). This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

optical imagery processing. With this wide scope and modular setup, it is possible to attract a broad user community at an expense of performance and scalability. This may become a bottleneck for services which demand efficient, worldwide pre-processing in near-real time (NRT) (Wagner et al., 2020), e.g. the Global Flood Monitoring Service (GFM) or the Copernicus Global Land Service (CGLS) (Salamon et al., 2021; Doubkova et al., 2016). Even with support from supercomputers, large-scale pre-processing is a great challenge driving many users from the scientific and public domain to commercial data providers that have a convenient datacube architecture in place (Bauer-Marschallinger et al., 2021; Wagner et al., 2021). Bauer-Marschallinger et al. (2021) have clearly shown that rolling out a scene-by-scene backscatter production globally using a processing strategy implemented in SNAP consumes millions of core-hours.

With orbital tube diameters of around 500 m (1σ), Sentinel-1's predecessor C-band SAR missions ENVISAT and ERS-1/2 had already proven their value for monitoring terrain deformations, studying glacier movements, and creating digital terrain models at a coarser resolution by means of InSAR (Rocca, 2004; Kumar et al., 2008). In this regard, the Sentinel-1 constellation came as a game changer. Both platforms, Sentinel-1A and Sentinel-1B, share the same orbital plane and revolve around the Earth in an orbital tube with a diameter of ~ 100 m (given as Root Mean Squared Error (RMSE)) (Geudtner et al., 2014). A high temporal correlation has been achieved over extended periods by a repeat interval of six days, providing fertile grounds for manifold (In)SAR applications on a finer scale (Yagüe-Martínez et al., 2016; Chini et al., 2019; Wegmuller et al., 2015). Prats-Iraola et al. (2015) have already published a set of compact equations to analyse the role of the orbital tube and its influence on certain InSAR parameters. Orbital fluctuations within this tube cause an orbit crossing angle variation, which introduces spectral shifts and image co-registration errors. Using interferometric image pairs as input demonstrated that these terms can be neglected for Sentinel-1, but still lead to a degradation of coherence. This is not the case for SAR image processing in general, where only a single observation and a single orbit are relevant for deriving a georeferenced and calibrated backscatter image.

In this study, we address the question of how Sentinel-1's orbital fluctuations propagate into different layers of standard SAR pre-processing workflows generating geometric terrain-corrected (GTC) sigma nought backscatter σ_E^0 (cf. Woodhouse, 2006) or radiometric terrain-corrected (RTC) gamma nought backscatter γ_T^0 (Small, 2011; Small et al., 2021). The respective workflows of both variables share the same geocoding procedure, but differ in terms of radiometric normalisation requiring more intermediate steps and layers to produce γ_T^0 . Since analysing both σ_E^0 and γ_T^0 would exceed the scope of one publication, we decided to present the general framework and first results on σ_E^0 processing in this paper, and the results for the CEOS-compliant standard (CEOS, 2021) of analysis-ready and quality-curved γ_T^0 backscatter in a follow-on paper (in preparation).

Concerning this publication, our main interest lies in two essential outputs of the basic SAR pre-processing pipeline: GTC σ_E^0 backscatter and a byproduct, the (projected) local incidence angle ((P)LIA). The interplay of both variables has proven its value for many applications, e.g. for normalising σ_E^0 to a common reference angle (Bauer-Marschallinger et al., 2021) or when producing surface soil moisture data (Bauer-Marschallinger et al., 2018). By taking four years of Sentinel-1A/B orbit data into account, we aim to measure the impact of the orbital tube on σ_E^0 and PLIA, and all relevant intermediate data layers. To be flexible in designing our experiment, we decided to create our own Python toolbox named *wizard*. It assisted us in analysing the pre-processing procedure to detect inefficient steps and to isolate layers that behave statically over time, with respect to the required accuracy level. Taking the code and insights from this experiment into account in a revised Sentinel-1 pre-processing workflow and performing a runtime comparison against existing software suites allows to draw conclusions

on the additional benefit of taking advantage of Sentinel-1's orbital stability.

For the sake of simplicity, we only consider ground range detected (GRD) and not single-look complex (SLC) SAR data, as our focus is on backscatter data (dropping phase information) to support services requiring efficient backscatter data processing in NRT. When comparing workflows based on SLC or GRD data, some parts are done in the same manner or are based on the same orbit-ground model, e.g. geocoding (cf. Alaska Satellite Facility, 2022). Thus, we expect that performance enhancements for backscatter data generation could be also beneficial for processing SLC data.

In Section 2, we provide an overview of the Sentinel-1 pre-processing workflow followed by a description of our input datasets in Section 3. In Section 4 we perform an analysis on the spatial distribution of orbital positions and velocities, which is used as input to a Monte-Carlo simulation for different SAR data layers. Its outcome led to the design of different pre-processing strategies, which are presented in Section 5 and compared against SNAP in Section 6.

2. Sentinel-1 pre-processing

SAR pre-processing is comprised of several steps to derive radiometrically calibrated, geocoded backscatter data from amplitudes stored in a GRD Level-1 scene (cf. Filippini, 2019; Mullissa et al., 2021; Truckenbrodt et al., 2019; Small and Schubert, 2019). Fig. 1 links most steps to a Sentinel-1 pre-processing workflow, and groups them according to the geometric model used to represent the data, i.e. ground or orbital (path) geometry. All steps performed in the orbital geometry are independent from the ground geometry and are therefore usually executed before any geocoding takes place:

- 1.(a) **Orbit interpolation:** Prepare and interpolate the orbit's ephemerides with respect to the scene's observation time window and its sampling in the azimuth direction, i.e. the so called azimuth or line time interval (LTI).
- 1.(b) **Border noise removal:** Removes border noise pixels present at an early stage of the processor. They were introduced by the transition from Level-0 to Level-1, conducted by the Instrument Processing Facility (IPF). In March 2018, IPF released a new processor version (IPF 2.90), that no longer produces such artefacts at the image borders.
- 1.(c) **Thermal noise removal:** Low backscatter values from flat surfaces (e.g. lakes, rivers or streets), that are close to the instrument thermal noise level, need to be corrected for this contribution by applying the thermal noise calibration vectors stored in the Level-1 metadata.
- 1.(d) **Radiometric calibration:** Conversion of amplitude values represented by digital numbers (DN) to sigma nought σ_E^0 backscatter, which refers to a unit area on an Earth ellipsoid E .

In Fig. 1, thermal noise and border noise removal are marked grey to highlight that they are optional, since they only improve the radiometric quality and are not a prerequisite for generating geometrically terrain corrected (GTC) backscatter products.

The ground geometry and sampling are primarily defined by the chosen digital elevation model (DEM) and its coordinate reference system (CRS), which represents the target system of our final product. Here, we chose the geographic projection ("LatLon") as our target CRS, because it is the native system of our DEM data (cf. Section 3.1), it features a global extent, and it is in alignment with the coordinate system of the orbital state vector data. Within this frame, the following processing steps were performed:

- 2.(a) **Ellipsoidal coordinates:** Transformation of geographic coordinates and height values of the DEM to 3D ellipsoidal coordinates to be situated in the same system as Sentinel-1's orbit ephemerides (WGS84).

- 2.(b) **Azimuth indexes:** For the interpolated orbit data and each terrain coordinate, the respective line number of the scene – also referred to as azimuth index – is computed by finding the time of closest approach (TCA) based on the Doppler equation. In an Earth-centred Earth-fixed reference frame, the TCA is mathematically defined as the zero-crossing time of the Doppler shift, which can for instance be identified iteratively or by using bisection. More details on this computationally expensive step can be found in [Small and Schubert \(2019\)](#).
- 2.(c) **LIA/PLIA:** The local incidence angle (LIA) is the angle between the microwaves arriving at the surface and the surface normal. To account for the tilting of the local, tangential surface, this angle is further projected into the range plane, also known as projected local incidence angle (PLIA) ([Schreier, 1993](#)).
- 2.(d) **Range indexes:** The second dimension of a GRD SAR image is referred to as the ground range, which is the distance on the reference ellipsoid between the nadir point of the platform and the respective terrain point. This geometric relationship is approximated through slant-to-ground range (or vice versa) polynomials, whose coefficients are provided in Sentinel-1’s Level-1 product metadata. As specified in [Small and Schubert \(2019\)](#) and [Vincent et al. \(2020\)](#), these polynomial coefficients can directly be used to compute ground range values from the slant range derived from the respective satellite position at TCA and the terrain point.
- 2.(e) **Resampling:** The final step is the resampling of the calibrated backscatter values based on the previously computed azimuth and range indexes for each terrain point.

The focus of our subsequent experiments was on propagating deviations defined with respect to a reference orbit into each step after 2.(a). With this we assessed the variance of each layer to understand if it can be set to a constant value allowing one to isolate it from the rest of the workflow. In such a case, it would only be necessary to perform the computation once and declare the layer as static, rather than precisely computing it for every scene. In theory, before knowing anything about the actual deviations, we could define the following static layers:

- **Azimuth indexes:** Instead of performing the iterative approaches presented in [Small and Schubert \(2019\)](#) for every pixel and scene, an external layer containing approximate azimuth indexes could facilitate the process of finding the TCA, interpolating satellite positions, computing the (P)LIA, or resampling σ_E^0 from orbit to ground geometry. In order to establish such a layer, a meaningful temporal anchor point (starting point of the azimuth index counting) and sampling (i.e. the LTT) along the reference orbit path are required. A suitable choice for the latter would be the smallest LTT value to give preference to oversampling rather than undersampling the satellite positions along a reference orbit trajectory. To align the azimuth indexes j stemming from the static layer with the actual azimuth indexes i from a scene, we suggest applying closest point matching of orbital reference points $p_{j,\{x,y,z\}}$ with respect to the orbital points $p_{i,\{x,y,z\}}$ related to the temporal extent of a scene:

$$i_{\min,j} = \underset{i}{\operatorname{argmin}} \left\| p_i - p_j \right\|_2 \quad (1)$$

where $i_{\min,j}$ should then replace the approximate azimuth index j of the static layer.

- **LIA/PLIA:** As mentioned before, (P)LIA offers valuable information on the measurement geometry. This step requires several vector operations to mathematically describe the relation between the sensor and the terrain point. Assuming only variations of the sensor while leaving the terrain location stable, i.e. using the same DEM as input, (P)LIA could be computed once for each orbital reference swath — as long as the orbit trajectory remains sufficiently steady over time.

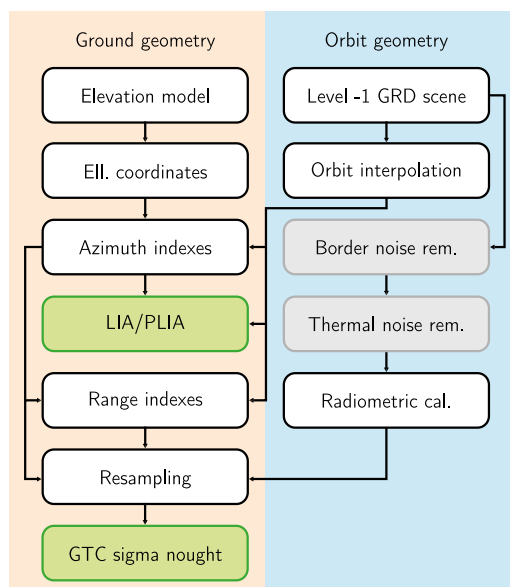


Fig. 1. Sentinel-1 pre-processing workflow with the steps needed to produce georeferenced SAR products (in green) from a Level-1 GRD scene and a digital elevation model. Greyish boxes indicate common, but not mandatory steps. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

There is not much need for a static layer dealing with the range indexes, since they can be precisely and efficiently computed scene-wise based on Sentinel-1’s highly accurate slant-to-ground range polynomials, and do not require as much effort as needed for the parameters above.

Software packages like SNAP or ISCE2 do not offer the necessary freedom to perform all these experiments or to establish the suggested static layers. Thus we implemented the whole processing chain illustrated in [Fig. 1](#) in a modular *Python* package called *wizard*, with most parts written in a *Numba*-compatible way to gain C-like performance ([Anaconda, 2021](#)).

3. Input data

3.1. Digital elevation data

With the very recent public release of the 30m Copernicus DEM ([Fahrland, 2020](#)), a new era for EO data processing on a global scale has begun, possibly superseding the supremacy of the SRTM, which defines the current standard. The Copernicus DEM is a digital surface model (DSM) and includes – compared to a terrain model – structures above ground, e.g. buildings, bridges, or vegetation. Its height values refer to the Earth Gravitational Model (EGM) 2008 and are thus given in an orthometric system.

To be able to use data from the Copernicus DEM as input to the Sentinel-1 pre-processing workflow, we needed to align the height values with respect to the reference system of the orbit ephemerides. To do so, Brun’s formula ([Bruns, 1878](#)) can be used to transform orthometric to ellipsoidal heights by taking geoid undulations into account, represented by EGM 2008 data at 2.5’ sampling ([Agisoft, 2021](#)).

3.2. Orbit ephemerides

Sentinel-1 circles the Earth in a near-polar, sun-synchronous orbit with a 12 day repeat cycle that defines 175 relative orbits ([Torres et al., 2012](#)). Each orbit features a very stable trajectory being measured by means of GNSS and is made available by the Copernicus

Table 1

Three Sentinel-1 Level-1 IW GRDH scenes having been acquired during an overpass in the relative orbit number 168 and serving as input for main experiments.

#	Name	Scene ID	Acquisition time	Centre point (Lon/Lat)	IPF version
1	Norway	S1B_IW_GRDH_1SDV_..._018919_023AEC_F26F	14 Nov. 2019, 05:19 UTC	(19.4, 69.6)	3.10
2	Austria	S1B_IW_GRDH_1SDV_..._023644_02CEC7_423D	3 Oct. 2020, 05:26 UTC	(10.8, 47.3)	3.31
3	Benin	S1A_IW_GRDH_1SDV_..._034540_040529_958B	27 Sep. 2020, 05:36 UTC	(3.2, 12.0)	3.31

Precise Orbit Determination (POD) Service with respect to different user requirements (ESA, 2021a). Orbital points along the trajectory are described using WGS84 coordinates (p_x, p_y, p_z), velocities (v_x, v_y, v_z), and a UTC time stamp, and are commonly separated by 10 s. Applications relying on a high timeliness must accept estimated orbit ephemerides at a lower accuracy (RMSE ~ 10 cm), which are referred to as restituted orbits (AUX_RESORB). On the other hand, precise orbit ephemerides (AUX_POEORB) have a latency of around 20 days, but decrease the RMSE down to ~ 5 cm, which becomes relevant when dealing with InSAR. For our experiments here, we dealt with historical Sentinel-1 data and could use precise orbit ephemerides as input to our experiments.

To statistically assess the magnitude of relative orbital variations based on the actual measured trajectories, we analysed four years of AUX_POEORB data, from 2017 to 2020. We did not take the full time series from 2014 onward into account, since we were only interested in variations of the joint, two-satellite constellation. Additionally, to perform our experiments in a feasible manner, we focused only on one relative orbit featuring pronounced terrain variability and good Level-1 data coverage. This resulted in the selection of the relative orbit number 168, which is a descending orbit crossing Scandinavia, central-Europe and northwestern-Africa (cf. Fig. 2). The respective AUX_POEORB files generated by the POD service were provided via the Earth Observation Data Centre for Water Resources Monitoring (EODC, 2021).

3.3. Sentinel-1 Level-1 data

EODC hosts a complete archive of Sentinel-1 Level-1 data, which is updated on a regular basis by mirroring and filtering data from Copernicus' data hubs (Grazia Castriotta and Volpi, 2020). We extracted all high-resolution, ground-range detected (GRDH) scenes acquired along the relative orbit number 168 in the interferometric wide (IW) swath mode, which is Sentinel-1's prioritised mode over land surfaces. From this set of files, we collected several metadata entries, e.g. the scene extent or the azimuth/line time interval (LTI), relevant in our subsequent analysis on the sensitivity of pre-processing layers to orbital variations.

In addition, for testing and benchmarking our own Sentinel-1 pre-processing workflow against existing SAR software, we further selected three Sentinel-1 Level-1 scenes (cf. Fig. 2), situated along the swath of the orbit. As track 168 is a descending orbit, Sentinel-1 observed the Earth's surface from east-to-west, or "right-to-left" respectively. The scenes are distributed across different latitudes and various terrain conditions to gain insight to how the processing time and performance are dependent on the growing oversampling of the "LatLon" projection with latitude. Table 1 shows further details on each of the three scenes.

4. Uncertainty propagation

4.1. Orbital distribution assessment

Estimating the impact of orbital variations on certain SAR layers requires a statistical representation of the orbital tube, i.e. the deviation of orbital state vectors from a reference trajectory. Prats-Iraola et al. (2015) already dealt with this topic by means of the Clohessy–Wiltshire equations, whereas we intended to derive a reference trajectory from actual orbit data in a 3D WGS84 coordinate system. They showed




Fig. 2. Location of the three Sentinel-1 scenes along the swath of the descending orbit number 168. All data are shown on top of Stamen's terrain-background map in an orthographic projection with its central longitude and latitude at 0 degree and 30 degree, respectively (map tiles by Stamen Design, under CC BY 3.0. Data by OpenStreetMap, under ODbL).

that the tilt of the orbital plane over time, represented by the orbit crossing angle, causes orbital deviations to vary within one cycle. This hinders the development of an expression of the deviations by using only a single quantity like the standard deviation over a whole revolution. Therefore, we decided to apply curve fits locally rather than globally and selected cubic splines from the pool of suitable curves for orbit interpolation or approximation. The splines provide sufficient accuracy in short-arc cases (Small and Schubert, 2019). In addition, we decomposed the positions p and the velocities v of the orbit trajectories into orthogonal dimensions, i.e. (p_x, p_y, p_z) and (v_x, v_y, v_z).

Before we fit any curve, it was necessary to align all multi-year trajectories to a certain reference time for each absolute orbit, so that orbits with the same relative orbit number are comparable. This reference was determined to be the time when the satellite platform crosses the ascending node at the equator (ANX).

Putting everything together, we parameterised our orbit model as shown in the two equations below.

$$p_{\{x,y,z\}}(t_{ANX}) = \hat{p}_{\{x,y,z\}}(t_{ANX}) + \delta p_{\{x,y,z\}}(t_{ANX}) \quad (2a)$$

$$v_{\{x,y,z\}}(t_{ANX}) = \hat{v}_{\{x,y,z\}}(t_{ANX}) + \delta v_{\{x,y,z\}}(t_{ANX}) \quad (2b)$$

t_{ANX} is the time referring to the equator crossing of the satellite. $p_{\{x,y,z\}}(t_{ANX})$ are the time-dependent platform positions, which are defined to be the sum of the average satellite trajectory $\hat{p}_{\{x,y,z\}}(t_{ANX})$ and some residual component $\delta p_{\{x,y,z\}}(t_{ANX})$. Eq. (2b) models the velocities, and is formulated equivalently, using v instead of p .

The systematic part in Eq. (2a) $\hat{p}_{\{x,y,z\}}(t_{ANX})$ and (2b) $\hat{v}_{\{x,y,z\}}(t_{ANX})$ was estimated using a rolling, short-arc, cubic-spline fit to mirror the local character of the average orbit trajectory. As a temporal sampling

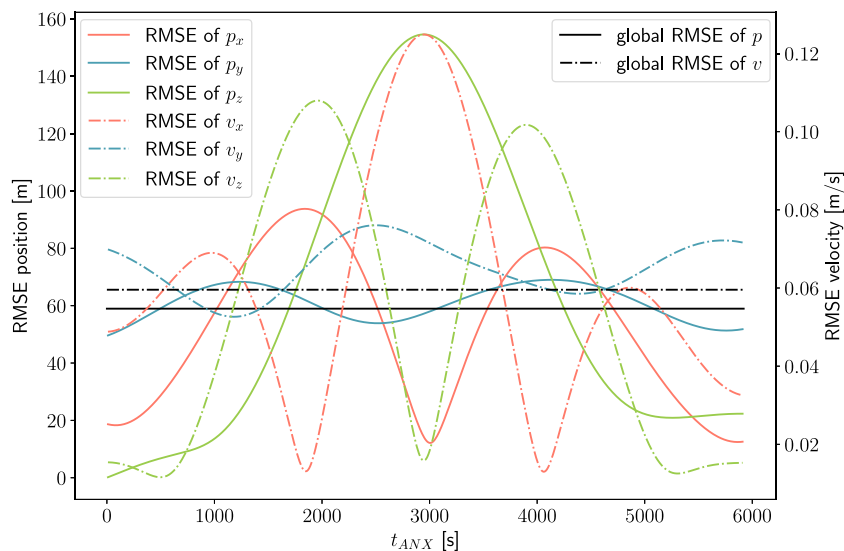


Fig. 3. RMSE's of the residuals $\delta p_{\{x,y,z\}}(t_{ANX})$ and $\delta v_{\{x,y,z\}}(t_{ANX})$ parameterised over the time spent in the relative orbit number 168 t_{ANX} . Coloured RMSE's refer to residuals from the rolling, local cubic-spline fit for each coordinate dimension, whereas black RMSE's show the global RMSE.

interval, we stayed in correspondence with Sentinel-1's AUX_POEORB orbital point sampling Δt , and performed a curve fit every 10 s along the orbital path. Consequently, we had k anchor points for the orbital period u , where $k = \text{floor}(u/\Delta t)$, and $t_{ANX} = \{0, \Delta t, 2\Delta t, \dots, k\Delta t\}$.

To ensure a smooth transition of all components between neighbouring anchor points, we set the size of the rolling window to $2\Delta t = 20$ s. This means that orbital positions and velocities are used twice during estimating the average orbit trajectory — once for each adjacent anchor point. The residuals $\delta p_{\{x,y,z\}}([t_{ANX} - \Delta t; t_{ANX} + \Delta t])$ and $\delta v_{\{x,y,z\}}([t_{ANX} - \Delta t; t_{ANX} + \Delta t])$ were then computed as a difference between the local (cubic-spline) estimate of the orbital component, i.e. $\hat{p}_{\{x,y,z\}}([t_{ANX} - \Delta t; t_{ANX} + \Delta t])$ and $\hat{v}_{\{x,y,z\}}([t_{ANX} - \Delta t; t_{ANX} + \Delta t])$, and its measured value. Their distribution was found to behave non-Gaussian, rather bi-modal, which could be attributed to systematic changes of the orbit crossing angles. To account for this, we encoded the residuals' distribution as a discrete probability density function (PDF) and assigned it to each anchor point at t_{ANX} . Thus, our statistically founded orbit model represents Sentinel-1's trajectory by means of k PDFs and k estimates for every component.

Fig. 3 shows the locally estimated RMSE's of the residuals $\delta p_{\{x,y,z\}}(t_{ANX})$ and $\delta v_{\{x,y,z\}}(t_{ANX})$. All RMSE's implicitly display the orientation of the orbital plane in relation to the WGS84 system, as oscillations are larger for the x and z than for the y component. The deviation of p_x dominates at the Earth's poles, whereas p_z peaks at the equator. For the velocities, the opposite holds true. The overall behaviour deviates from the baseline differences presented in Prats-Iraola et al. (2015), mainly due to the different coordinate system.

The continuous black line in Fig. 3 depicts the overall RMSE concerning p . Based on its magnitude of around 58 m we found an RMSE for the radius of the orbital tube being a bit higher than promised in initial mission reports (Torres et al., 2012), i.e. 50 m, but in good agreement with past studies dealing with Sentinel-1's orbital tube (Torres et al., 2017; Barat et al., 2015; Prats-Iraola et al., 2015). Our empirically estimated set of orbital reference points and residuals offers a sufficient basis for our subsequent uncertainty propagation and confirms Sentinel-1's outstanding orbital stability.

4.2. Static layer realisation of PLIA and azimuth index

Based on our definition of static pre-processing layers in Section 2, we utilised the reference orbit trajectories $\hat{p}_{\{x,y,z\}}$ and $\hat{v}_{\{x,y,z\}}$ to realise average conditions for both azimuth indexes and PLIA. The resulting

layers, which were cropped with respect to the extent of the scenes listed in Table 1, are depicted in Fig. 4. On the left side we can observe the progression of azimuth indexes in the azimuth direction, from north to south. The sampling of these indexes is based on the minimum LTI.

The generation of the static PLIA layer goes one step further and integrates the interpolated satellite position, defined by the azimuth index, into the process of computing the local incidence angle of the microwaves arriving at the surface. The result is visualised in Fig. 4 (centre), with a more detailed view in Fig. 4 (right), and provides an overview for each scene of the terrain variability and sensor orientation.

4.3. Monte-Carlo simulations of PLIA and azimuth index

In addition to the average behaviour of the static layers imparted by Fig. 4, our main objective is to also look at their variation driven by deviations in orbital positions and velocities. As a powerful tool for simulating their effect on intermediate and final layers of Sentinel-1's pre-processing pipeline we selected Monte-Carlo simulation, which allows one to propagate manifold probability distributions of *a priori* parameters through a complex system.

Instead of directly using the collected orbit data, we simulated n synthetic orbit trajectories by drawing samples from the discrete PDF's presented in Section 4.1. This resulted in $k \times 6$ samples per relative orbit trajectory, i.e. a p and v value for each orthogonal component (x, y, z) and each orbital reference point.

Each synthetic orbit trajectory was then fed into our Sentinel-1 pre-processing pipeline (cf. Section 2) for each of the three scenes (cf. Section 3.3) resulting in n simulated values for every proposed static layer. Depending on how long it takes to run the system multiple times, it can turn out to be computationally costly to ensure a stable estimate of the output samples mean and variance (Thomopoulos, 2013). By examining the asymptotic behaviour of the sample mean for several pixels, we came to the conclusion that it levels off at around $n = 1000$, which we defined as the stopping criterion for our Monte-Carlo simulations.

The first output of the simulation is the azimuth index layer after performing step 2.(b) in Section 2. Following our definition and realisation of its static version in Fig. 4, we co-located each simulated azimuth index with respect to the reference azimuth index via closest point matching as formulated in Eq. (1). This transformation yielded the desired probability distribution and thus the discrete sample standard deviation σ_i depicted in Fig. 5 (left).

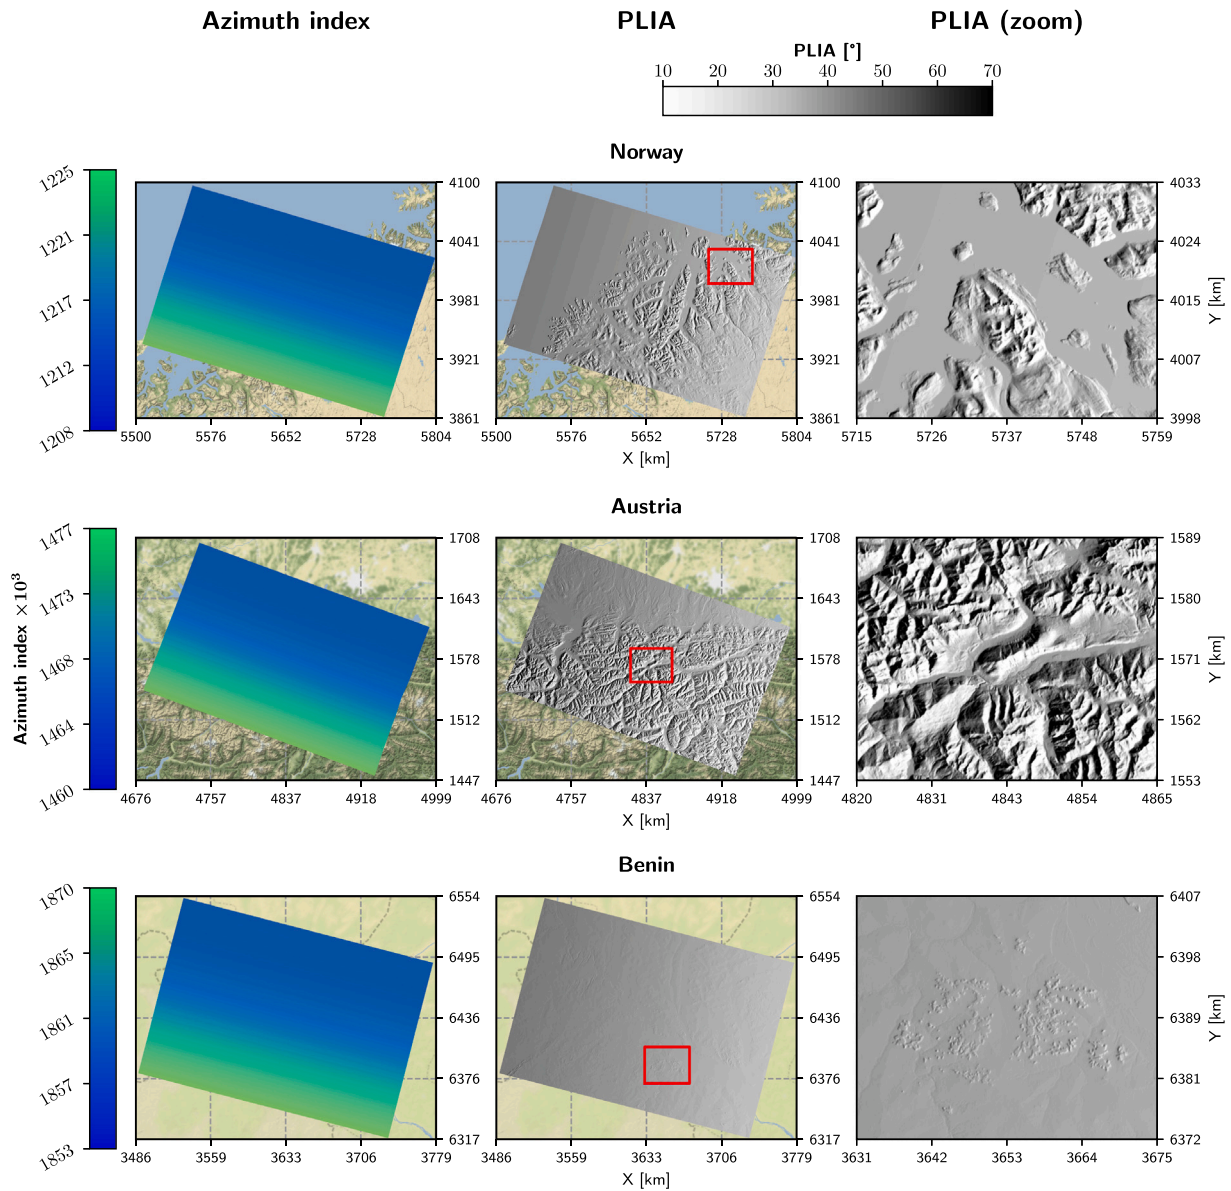


Fig. 4. Predefined static layers including approximate azimuth index (left) and PLIA (centre, right) for the scenes listed in Table 1. The red boxes indicate the extent of the zoom-in view on the right side. To achieve a proper visual comparison of all scenes in terms of their distortion in northern and eastern direction, we have chosen an azimuthal equidistant projection for display, i.e. the Equi7Grid projection (Bauer-Marschallinger et al., 2014).

Overall, σ_i seemed to spatially stagnate at a very low level, i.e. around 1–2 azimuth indexes. We also observed that σ_i in the Benin scene was larger than for the others, which corresponds well with the peak of the p_z component in Fig. 3. Sentinel-1’s orbital path is approximately parallel to the z axis near the equator, implying that the dominant deviation of the p_z component directly translates to a lateral variation along the flight direction of the satellite.

After having identified the exact TCA and its corresponding azimuth index, we can compute the PLIA by interpolating the satellite position at TCA and solving the vector equations describing the relation between the satellite and the terrain point. In a manner similar to that used for the azimuth indexes, PLIA values were repeatedly simulated based on the orbit trajectory samples from Eqs. (2a) and (2b) resulting in the standard deviation of the PLIA layer σ_{PLIA} depicted in Fig. 5 (centre) and Fig. 5 (right).

By taking a look at Fig. 5 (centre) the extremely low magnitude of angular variations becomes clear; they reach an order of 0.005 degree. Additionally, one can identify several interesting patterns. First, σ_{PLIA} tends to be larger in near range than in far range. As we are looking

at angular changes, it is evident that σ_{PLIA} becomes greater the closer a terrain point is to the sensor. Second, a scalloping pattern emerges along the azimuth direction coming from the chosen 10s temporal spacing between the knots of the orbit trajectory, where the respective orbital deviations have been applied. When comparing the behaviour of both PLIA’s and azimuth index’s standard deviation in the azimuth direction, σ_{PLIA} does not seem to follow the strict pattern of σ_i . Austria had the highest σ_{PLIA} , followed by Norway. This finding might mainly result from the superposition of the x and z components of the orbit trajectory, which experience a combined peak at mid-latitudes in Europe (cf. Fig. 3). They directly relate to the radial component of the local coordinate system of the satellite (neglecting variations in y). The radial component becomes crucial in terms of the incidence angle, whereas the component in the azimuth direction is largely compensated by the estimation of the azimuth index.

Since PLIA is a variable mainly driven by the local slope of the terrain, the largest deviations can be found where slopes tend to be steep, with respect to the position of the sensor. A closer look at this

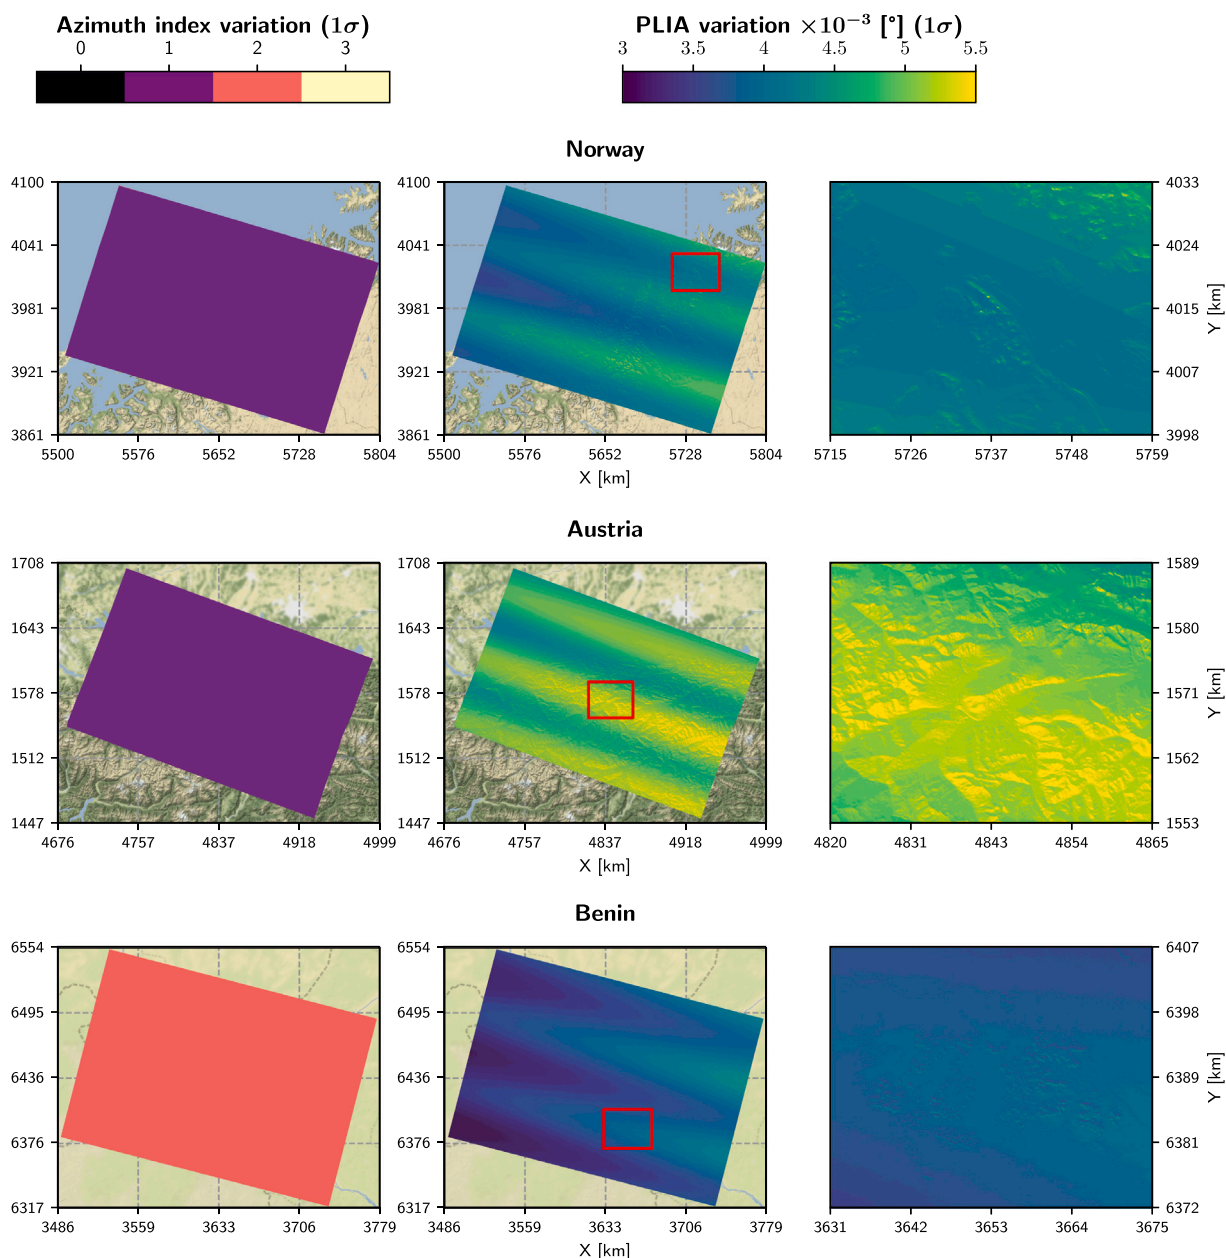


Fig. 5. Azimuth index (left) and PLIA (centre, right) sample standard deviation of the probability density functions resulting from the Monte-Carlo simulations applied to all scenes listed in Table 1. The red boxes indicate the extent of the zoom-in view on the right side. All scenes are shown in the Equi7Grid projection.

behaviour is taken in Fig. 5 (right), which offers a more detailed view on the highlighted regions in Fig. 5 (centre).

The results of this simulation underline the potential of these layers to be declared as static and to externally provide essential information on geocoding or the measurement geometry. Even if the standard deviation of the azimuth indexes σ_i is a few times larger than one pixel, e.g. six times for Benin (three-sigma rule), such a static layer could help to generate georeferenced sigma nought backscatter data at a coarser sampling, similar to a downsampled quick-look image. When staying at the sensor’s spatial resolution, knowledge of the approximate azimuth index could still be valuable for finding the right index corresponding to TCA (cf. Section 5.2). On the other hand, (P)LIA is a byproduct of the pre-processing workflow and is not needed for geocoding. Declaring it as a static layer per relative orbit would thus enable one to have much less data output per scene and to have a consistent dataset. It would only be dependent on the DEM data and the average orbit

trajectory, which would be helpful for all kinds of applications dealing with backscatter vs. incidence angle dependency, e.g. backscatter normalisation (Schaufler et al., 2018), or water mapping (Westerhoff et al., 2013).

5. Sentinel-1 pre-processing enhancements

In the following subsections we incorporate previous findings into the design of a performant Sentinel-1 pre-processing workflow and present strategies how to simplify some parts of the diagram in Fig. 1 by means of utilising information from the static layers or their variations. Possible performance enhancements should not only be measured internally based on our own Python implementation of the Sentinel-1 pre-processing workflow, *wizard*, but also externally against SNAP to have a benchmark in terms of radiometric quality and runtime. Thus, we decided to reconstruct the steps in Fig. 1 in SNAP 8 (SNAP Team, 2021) by creating one GPT graph (Peters, 2021) comprising the

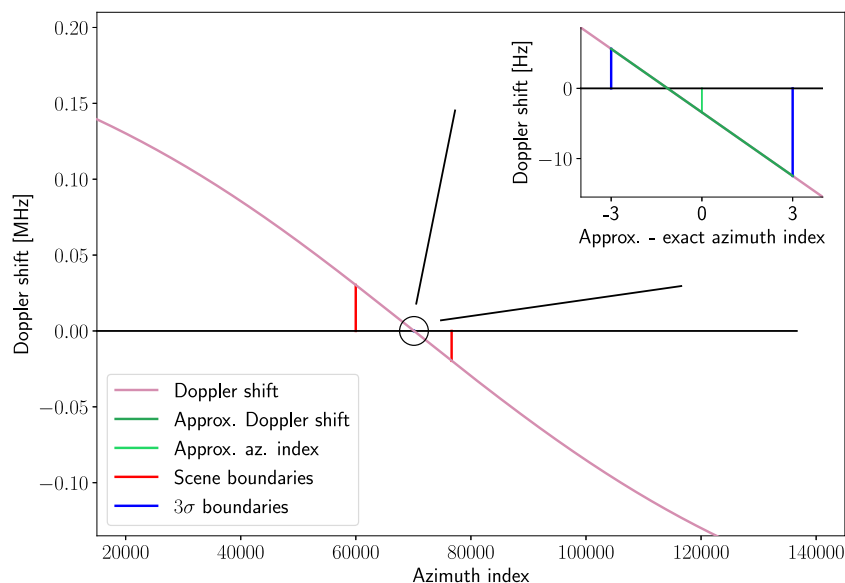


Fig. 6. Doppler shift of a terrain point (lon: 9.86 degree, lat: 47.2369 degree, alt: 1358 m) in Austria computed for the region ± 2 min around the actual start and stop time of the scene, which is marked in red. The zoom-in of the zero-crossing marked by a circle shows the linear approximation of the Doppler shift curve using the approximate azimuth index and its three-sigma interval stemming from the static layer. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

following SNAP operators: 1. *Apply-Orbit-File* 2. *Calibration* 3. *Terrain-Correction* 4. *Subset* 5. *Write*. Referring to Section 2, *Apply-Orbit-File* queries for the respective orbit auxiliary data, *Calibration* resembles step 1.(b), and the *Terrain-Correction* operator executes all steps in (2). *Subset* was used to split up all internal source bands, so that each polarisation and the PLIA band could be forwarded to the final *Write* operator. Having only one GPT graph file instead of one for each of the steps above has the advantage of having less intermediate data writing and reading, thus mostly resembling our own realisation of the pre-processing chain.

5.1. Static PLIA layer

The sample's standard deviation shown in Fig. 5 has some implication on the selection of the static layers outlined in Section 2. The overall magnitude of σ_{PLIA} only reaches around 0.005 degree, which is far beneath the required precision for Sentinel-1 SAR backscatter applications. Hahn et al. (2017) have shown that the maximum backscatter-incidence angle sensitivity is around -0.25 dB/degree for ASCAT scatterometer data. Inserting the σ_{PLIA} value mentioned before yields a standard deviation of ~ 0.0013 dB, when assuming a similar dependency for Sentinel-1. Since this deviation is some orders below the relative radiometric accuracy of Sentinel-1's C-band sensor, i.e. 0.1 dB (three sigma) (Bourbigot et al., 2020), we can declare the average PLIA value (cf. Fig. 4 (centre)) from our Monte-Carlo simulation as static.

Using the basic implementation of the pre-processing workflow in *wizard* as a starting point, we continued with the modification of certain steps in the processing chain by taking Sentinel-1's orbital stability and its resulting static layers into account. As we have argued before, PLIA has proven to be sufficiently stable for our relative orbit selection, supporting to skip its generation in the ground-based part of the workflow by removing step 2.(c).

5.2. Approximate azimuth index layer

A maximum standard deviation of $\sigma_i \approx 2$ for the azimuth indexes reveals a different picture than the behaviour of σ_{PLIA} . Azimuth index variations directly propagate into the resampling of σ_E^0 from orbit to ground geometry (cf. step 2.e in Section 2) leading to a wrong or unreliable selection of backscatter values from the source image. An

explicit usage of approximate azimuth indexes when staying at 10 m ground sampling is therefore not possible.

Yet, we still suggest to provide the approximate azimuth indexes and their variance as a valuable a priori information to the pre-processing workflow. Knowledge about the approximate location of the Doppler shift zero-crossing is valuable to the applied root-finding algorithms, which is demonstrated in Fig. 6. The Doppler shift for the selected time window beyond the actual scene boundaries clearly shows a non-linear behaviour, whereas when approaching TCA, linearity seems to prevail. This relation is discussed and investigated in more detail in Appendix, from which we can conclude the diminishing non-linearity close to TCA, but its presence at the scene boundaries (for a point close to the centre of the scene).

Well-known root-finding methods like iteration or bisection have therefore been established to reliably compute the TCA under such circumstances (Small and Schubert, 2019). However, one can imagine that these approaches do not perform well due to their iterative nature in defiance of the increasing linearity towards TCA, especially when using the scene boundaries as a starting point. Thus, a more precise estimate of the zero-crossing location and its deviation coming from the approximate azimuth index layer would be of great value to this procedure. By choosing the three-sigma rule based on the maximum σ_i value of the Monte-Carlo output, we were able to perform a linear approximation of the Doppler shift based on the frequency values at the interval boundaries. This procedure is depicted in the zoom-in view of Fig. 6, where the TCA is directly computed by finding the root of the linear model, instead of using an iterative method. The expected precision of this approach is elaborated in Appendix.

5.3. Neighbouring azimuth index

Our final suggested enhancement benefiting from the output of the Monte-Carlo simulation works similar to the approximate azimuth index method, but does not make use of external information on azimuth indexes coming from the static layer. Instead, we want to utilise Tobler's first law of geography (Tobler, 1970) and take an approximate value for the current pixel from its direct vicinity, i.e. using the azimuth index, which was computed for the previous, neighbouring pixel. The remaining procedure to find the actual TCA is inherited from Section 5.2.

Table 2

Summary of the benchmarking experiments for all different setups of Sentinel-1's pre-processing chain. Cells coloured in green highlight the method performing best in terms of geocoding and total runtime. (For interpretation of the references to colour in this table, the reader is referred to the web version of this article.)

	Workflow	Scene prep.	Aux. data prep.	Geocoding	Data export	Total	W.r.t. wizensard (base)
Norway	wizensard (base)	0 min 28 s	0 min 12 s	6 min 53 s	0 min 39 s	8 min 12 s	–
	wizensard (static PLIA)	0 min 28 s	0 min 12 s	5 min 10 s	0 min 26 s	6 min 16 s	-24 %
	wizensard (ref. az.)	0 min 28 s	0 min 32 s	4 min 46 s	0 min 26 s	6 min 12 s	-24 %
	wizensard (nbr. az.)	0 min 28 s	0 min 12 s	4 min 46 s	0 min 26 s	5 min 52 s	-28 %
	SNAP 8	–	–	–	–	29 min 30 s	+360 %
Austria	wizensard (base)	0 min 29 s	0 min 6 s	3 min 34 s	0 min 18 s	4 min 36 s	–
	wizensard (static PLIA)	0 min 29 s	0 min 6 s	2 min 42 s	0 min 12 s	3 min 29 s	-22 %
	wizensard (ref. az.)	0 min 29 s	0 min 16 s	2 min 31 s	0 min 12 s	3 min 28 s	-22 %
	wizensard (nbr. az.)	0 min 29 s	0 min 6 s	2 min 30 s	0 min 12 s	3 min 17 s	-26 %
	SNAP 8	–	–	–	–	16 min 4 s	+361 %
Benin	wizensard (base)	0 min 29 s	0 min 4 s	2 min 32 s	0 min 13 s	3 min 9 s	–
	wizensard (static PLIA)	0 min 29 s	0 min 4 s	1 min 57 s	0 min 8 s	2 min 38 s	-16 %
	wizensard (ref. az.)	0 min 29 s	0 min 12 s	1 min 47 s	0 min 8 s	2 min 36 s	-17 %
	wizensard (nbr. az.)	0 min 29 s	0 min 4 s	1 min 48 s	0 min 8 s	2 min 29 s	-21 %
	SNAP 8	–	–	–	–	10 min 34 s	+335 %

5.4. Benchmarking environment

To assess the performance of each of the presented versions of Sentinel-1's pre-processing chain, we repeatedly executed our implementations on a dedicated Linux machine having 256 GB RAM and 24 cores available. Such a system setup could represent a node on a supercomputer, being suitable for deploying the pre-processing workflow in a parallelised manner to stem processing activities on a large scale. Yet, our focus is on comparing the performance in terms of differences in methodology rather than the parallelisability of the implementation. We therefore limited the resources, which are allowed to be used by the software, down to 4 cores and 32 GB RAM. With this setup we are then also able to gain insights on runtimes when executing each method on a common office machine.

6. Results

6.1. Runtime benchmarking

Using the aforementioned hardware components, we performed m repetitive calls of each pre-processing scenario, where $m = 50$ already gave us some reliable estimate of the runtime. Table 2 sums up the results of the benchmarking experiments and re-groups the steps described in Section 2:

1. *Scene preparation*: Merges all scene- and orbit related steps defined under (1), including reading Level-1 data, but without thermal- or border noise correction.
2. *Auxiliary data preparation*: Comprises loading and preparation of all auxiliary layers, i.e. the DEM data and/or the approximate azimuth index layer.
3. *Geocoding*: Performs all steps under (2), depending on the experiment in question.
4. *Data export*: Single step writing all data as a GeoTIFF file to disk.

The abbreviated workflow names refer to the aforementioned variations of Sentinel-1's pre-processing chain as follows: The original pre-processing setup once implemented in SNAP 8 ("SNAP 8") and once in Python ("wizensard (base)"), the latter workflow without the PLIA computation ("wizensard (static PLIA)"), and finally further using the information on the approximate azimuth index coming from the static layer ("wizensard (ref. az.)") or from the direct neighbourhood ("wizensard (nbr. az.)"). With regard to SNAP, all processing steps are composed in one GPT graph file, which does not allow to retrieve runtimes in compliance with the listing above. Thus, only the total runtime is shown for "SNAP 8".

If we take a close look on the overall behaviour of the various processing groups, we can identify that *Scene Preparation's* runtime

almost remains constant. This is not surprisingly, since only scene-bound operations are executed and the three scenes do not differ much in file size. Since the auxiliary and target data layers are provided in ground geometry, i.e. the "LatLon" system, they are affected by an increasing oversampling with latitude. Thus, a significant increase of the runtime from Benin to Norway is present for *Auxiliary Data Preparation*, *Geocoding*, and *Data Export*.

Looking at the difference in total runtime between the two realisations of the complete pre-processing chain, we can see that "wizensard (base)" is more than three times faster than "SNAP 8". The performance boost in favour of "wizensard (base)" might come from the utilisation of Numba to achieve C-like speeds or because SNAP's focus is on modularity, thus the need to support multiple satellite radar missions at once. Secondly, the first experiment removing the PLIA generation from the overall workflow reveals an improvement of around 25 % in the geocoding part, which is up to 100 s for the scene located in Norway. Another performance gain stems from less data writing under *Data Export* taking about one third less time compared to "wizensard (base)". Utilising the approximate, static azimuth index layer in "wizensard (ref. az.)" requires significantly more time to read and prepare another data source in addition to the DEM data in the *Auxiliary Data Preparation* step, but reduces the runtime during geocoding by another 5 %, on average. "wizensard (ref. az.)" benefits from less evaluations of the Doppler shift, i.e. only two because of the linear model, whereas bisection with the scene boundaries as a starting point needs much more, e.g. up to 20 iterations for some pixels. Finally, building upon the same methodology, "wizensard (nbr. az.)" and "wizensard (ref. az.)" balance each other very well during geocoding. However, "wizensard (nbr. az.)" takes its auxiliary information from the pixels vicinity instead of an external layer, thus being the clear winner in total runtime.

These results imply that a remarkable performance boost can be already achieved by selecting a software solution which implements a highly parallelisable and slim pre-processing workflow tailored to Sentinel-1. The most performant version of such a workflow would skip the generation of (projected) local incidence angle data and would utilise a linear model for estimating the azimuth index at TCA. Increasing the number of cores, i.e. choosing to process on a project machine or node of an HPC cluster, will narrow the runtime gap between the basic and most performant pre-processing workflow in absolute numbers, but should not alter the runtime difference when measured relatively.

6.2. Backscatter benchmarking

The theoretical considerations in Appendix clearly underline the minor difference of a linear model, which is based on an approximate azimuth index and its three-sigma boundaries, with respect to an exact,

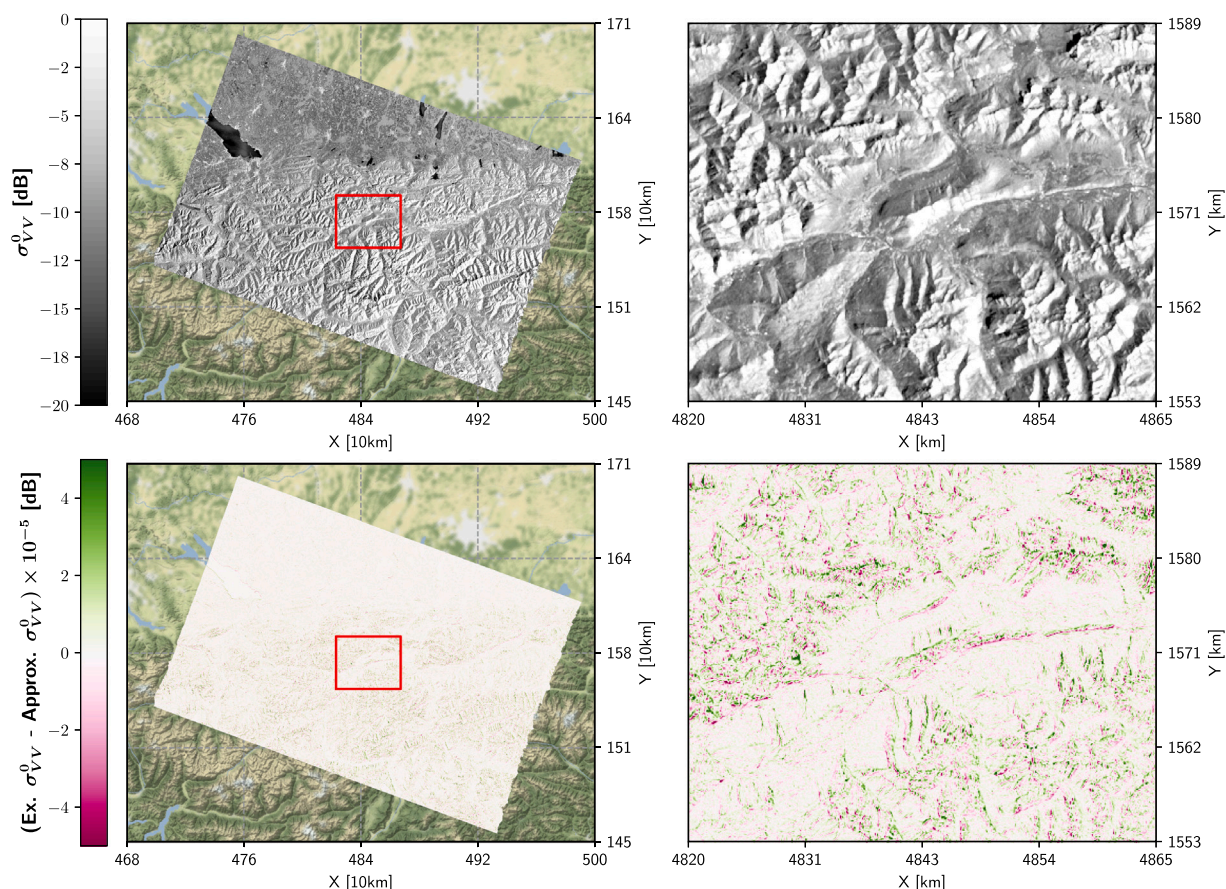


Fig. 7. Comparison between σ_{VV}^0 of “wizards (base)” (top) and its difference with respect to σ_{VV}^0 of “wizards (nbr. az.)” (bottom). Image data is shown in the Equi7Grid projection.

iterative solution. Yet, these azimuth index residuals are not representative for differences in σ_E^0 backscatter, which is the actual quantity we are interested in. Fig. 7 (bottom) provides an impression of how differences in estimated azimuth indexes propagate into the final σ_E^0 layer, which is shown in Fig. 7 (top) using exact azimuth indexes. We can observe that these differences are a few orders below the relative radiometric accuracy of Sentinel-1’s C-band sensor, encouraging the usage of the linear model. The largest discrepancies can be found in regions where foreshortening or layover occurs, since only a small change in slant range, which relies on the estimated azimuth index, has a significant impact on the computed ground range used for indexing backscatter data in orbit geometry.

7. Conclusion

This work demonstrated the significant benefit of utilising Sentinel-1’s orbital stability for pre-processing Level-1 SAR data. Our empirically derived radius of Sentinel-1’s narrow orbital tube of around 58m (RMSE) has allowed to isolate two parameters, PLIA and an approximate azimuth index. Their deviation is expressed by the sample standard deviation, which was estimated by means of Monte-Carlo simulation. Orbit-time-dependent distributions, which were based on a rolling cubic-spline fit for all orbital positions and velocities, served as input.

PLIA’s variation tends to be significantly below the precision needed for most SAR applications. Hence, the derived static layer only needs to be computed once for each relative orbit and can directly provide auxiliary information on the observation geometry for other purposes. On the other hand, the standard deviation of an azimuth index (σ_i) does not allow to assume overall constancy, since azimuth indexes are directly bound to the movements and trajectory deviations of the sensor.

Nonetheless, we integrated the information given by the distribution of the azimuth index in the time-consuming process of finding the TCA of the sensor, i.e. the true azimuth index. By co-locating an arbitrary with a reference orbit trajectory using closest point matching, the static layer relying on the latter one enables to approach the TCA up to around six indexes (worst-case scenario in Benin) when applying the three-sigma rule.

Thanks to the efficiency gained by using Python’s library Numba, the in-house developed toolbox wizards achieved much lower runtimes than SNAP. Further runtime improvements could be achieved by removing the PLIA computation from the pre-processing chain (~25%) and to replace iterative root-finding methods with linear ones for determining the TCA (~5%). The latter enhancement relies on a linear model defined by the information stored in the approximate azimuth index layer, which requires to read one more dataset in addition to the DEM. The method yielding the best runtimes overcomes the I/O overhead by taking approximate azimuth index values from the neighbourhood of a terrain point. Differences in backscatter between the two approaches for azimuth index root-finding are significantly below Sentinel-1’s radiometric accuracy clearly supporting the utilisation of the presented linear model.

The results show that it is possible to generate a static, per-orbit PLIA dataset and to incorporate a single σ_i value in the aforementioned linear model around the approximate azimuth index taken from the pixels vicinity. With these improvements in place and based on the chosen hardware setup, it would take ~30% less time to execute Sentinel-1’s pre-processing chain without reducing the quality of the output data. This could save costly resources when conducting large-scale Sentinel-1 pre-processing jobs in a high performance computing environment or could make the data-on-demand approach more attractive to users only working with small areas of interest. In addition, our findings allow to

save a considerable amount of disk volume, since there would be only one (P)LIA dataset per relative orbit instead of one for each observation.

However, the deployment of the enhanced version of the pre-processing workflow is only applicable as long as relative orbits other than 168 show deviations in a similar order — which is expected due to Sentinel-1’s stringent orbit control. When working with data from other high-resolution SAR missions like the RADARSAT Constellation Mission (RCM), which operates in an orbital tube with a radius of ~ 100 m (ESA, 2022), one could also benefit from the insights and methods presented in this study.

Keeping in mind the more stringent orbital requirements for InSAR processing, our insights could also be interesting for applications using SLC data as input, e.g. when working with coherence (Tamm et al., 2016) or interferograms (Yagüe-Martínez et al., 2016). In this regard, the orbit-ground geometry is relevant multiple times, e.g. during image co-registration (when performing back-geocoding), flat-earth correction, slant-to-ground range conversion or Range–Doppler terrain correction. These steps could thus benefit from having approximate information on azimuth or range indexes serving as input to Doppler shift or slant-to-ground range polynomial analysis.

For what concerns analysis-ready and CEOS-compliant SAR backscatter, radiometric terrain-corrected (RTC) gamma nought γ_T^0 data along with its underlying shadow mask and reference area for radiometric normalisation will be investigated in a follow-on paper.

CRediT authorship contribution statement

Claudio Navacchi: Conceptualization of this study, Methodology, Software, Visualization, Investigation, Writing – original draft. **Sen-mao Cao:** Discussions, Writing – review & editing. **Bernhard Bauer-Marschallinger:** Writing – review & editing. **Paul Snoeij:** Initial idea, Writing – review & editing. **David Small:** Writing – review & editing. **Wolfgang Wagner:** Conceptualization of this study, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work received funding within the project “Gamma2-Cloud: Feasibility of using Sentinel-1 terrain-flattened gamma nought backscatter across EO platforms”, ESA contract No. AO/1-9101/17/1-NB, and from the European Joint Research Centre (JRC) in the project “Global Flood Monitoring (GFM): Provision of an Automated, Global, Satellite-based Flood Monitoring Product for the Copernicus Emergency Management Service” Reference No. JRC/IPR/2020/OP/05 51. We would also like to take the opportunity to thank Dirk Geudtner for providing helpful inputs during review and acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme.

Appendix. Considerations on the linearity of the Doppler shift

The behaviour of the Doppler shift given in Fig. 6 can be also identified when looking at the Doppler equation based on the sketch in Fig. 8:

$$f_d = 2 \frac{v_{\text{rel}}}{\lambda} = 2 \frac{\mathbf{v} \mathbf{R}_{1,e}}{\lambda} = 2 \frac{v \sin(\alpha')}{\lambda} \quad (\text{A.1})$$

f_d is the Doppler shift in Hz, v_{rel} the relative velocity in m/s, λ is the wavelength in m, \mathbf{v} the velocity vector of the platform in m/s, $\mathbf{R}_{1,e}$ the slant range unit vector from a platform position S_i to the target T in m, v the absolute velocity of the platform in m/s, and $\alpha' = 90 - \alpha$, where




Fig. 8. A SAR platform advances with a velocity v at a height H above a target T , which is separated by a ground range distance G and a slant range distance R_0 from the sensor S_0 at TCA and is assumed to have zero velocity. The start and end time of a scene relates to the positions S_s and S_e , respectively. Both are separated by the along-track scene extent E_s , when ignoring any curvature. The unit slant range vector $\mathbf{R}_{s,e}$ points under an angle β from S_s to T .

α is the observation angle between the direction of movement and T , given in degree.

During geocoding, we aim to find the azimuth index corresponding to the zero-crossing of f_d at TCA, which is equivalent to $\alpha = 90$ degree or $\alpha' = 0$ degree. For small α' values, i.e. close to TCA, $\sin(\alpha')$'s corresponding Taylor series can be simplified to $\sin(\alpha') \approx \alpha'$, thus introducing a linear relation with velocity. The equivalent error in f_d between the exact and linear solution $f_{\text{ex-lin}}$ is stated in Eq. (A.2) and highlights that the linearity of the Doppler shift is actually about the linearity of $\sin(\alpha')$ at TCA.

$$f_{\text{ex-lin}} = \left| 2 \frac{v (\sin(\alpha') - \alpha')}{\lambda} \right| \quad (\text{A.2})$$

When choosing a position further away from TCA, which is the case when selecting the scene boundaries S_s and S_e as a starting point for finding the zero-crossing of f_d , the behaviour of the Doppler shift starts to become non-linear. This can be demonstrated by the following example: Assuming an absolute peak velocity of 7600 m/s, $\lambda = 0.055$ m, a TCA position S_0 being located in the middle between S_s and S_e , i.e. $\|S_s - S_0\|_2 = E_s/2 \approx 85$ km, and a slant range distance $R_0 = 800$ km, we get an angle $\beta' = 90 - \beta \approx 6.1$ degree. This translates to a difference of around $f_{\text{ex-lin}} = 55$ Hz between a linear and an exact computation of the Doppler shift at the scene boundaries.

When utilising a linear model for finding the TCA, the aforementioned discrepancies would directly propagate into the computation of the azimuth index j . To get a feeling how these angular differences relate to actual azimuth indexes in a worst-case scenario (neglecting symmetry), we can transform $\alpha'_{\text{ex-lin}} = \sin(\alpha') - \alpha'$ to its counterpart on ground in azimuth direction via simple trigonometry:

$$j_{\text{ex-lin}} = \frac{R_0 \tan(\alpha'_{\text{ex-lin}})}{\Delta G} \quad (\text{A.3})$$

ΔG is the ground range sampling and equals 10 m. Inserting the previous values related to β yields a final deviation of around 160 indexes at TCA, thus disapproving the application of a linear model when using the scene boundaries as a starting point.

Yet, when taking the additional information provided by the static layer into account, we can approach TCA up to six indexes (worst-case scenario for Benin) by means of the three-sigma rule. Utilising the new interval boundaries in the same manner as for the previous experiment, yields $\alpha'_{\text{ex-lin}} \approx 4 \cdot 10^{-12}$ degree and its azimuth index counterpart of around $j_{\text{ex-lin}} \approx 6 \cdot 10^{-9}$. This order reassures the usage of the linear model relying on the parameters coming from the static, approximate azimuth index layer.

References

- Agisoft, 2021. Global models. Available online: <https://www.agisoft.com/downloads/geoids/>. (Accessed 18 October 2021).
- Alaska Satellite Facility, 2022. How to create an interferogram using ESA's Sentinel-1 toolbox. Available online: <https://asf.alaska.edu/how-to/data-recipes/create-an-interferogram-using-esas-sentinel-1-toolbox/>. (Accessed 28 June 2022).
- Anaconda, 2021. Numba. Available online: <http://numba.pydata.org/>. (Accessed 18 October 2021).
- Barat, I., Prats, P., Duesmann, B., Geudtner, D., 2015. Sentinel-1: Link between orbit control and interferometric SAR baselines performance. In: 25th International Symposium on Space Flight Dynamics. pp. 19–23.
- Bauer-Marschallinger, B., Cao, S., Navacchi, C., Freeman, V., Reuß, F., Geudtner, D., Rommen, B., Vega, F.C., Snoeij, P., Attema, E., et al., 2021. The normalised Sentinel-1 global backscatter model, mapping Earth's land surface with C-band microwaves. *Sci. Data* 8 (1), 1–18.
- Bauer-Marschallinger, B., Freeman, V., Cao, S., Paulik, C., Schaufler, S., Stachl, T., Modanesi, S., Massari, C., Ciabatta, L., Brocca, L., et al., 2018. Toward global soil moisture monitoring with Sentinel-1: Harnessing assets and overcoming obstacles. *IEEE Trans. Geosci. Remote Sens.* 57 (1), 520–539.
- Bauer-Marschallinger, B., Sabel, D., Wagner, W., 2014. Optimisation of global grids for high-resolution remote sensing data. *Comput. Geosci.* 72, 84–93.
- Bourbigot, M., Johnsen, H., Piantanida, R., 2020. Sentinel-1 Product Definition. ESA, Available at <https://sentinels.copernicus.eu/documents/247904/0/Sentinel-1-Product-Definition/6049ee42-6dc7-4e76-9886-f7a72f5631f3>, Issue 2/7.
- Bruns, H., 1878. Die Figur Der Erde: Ein Beitrag Zur Europäischen Gradmessung, vol. 8. P. Stankiewicz.
- CEOS, 2021. CEOS ARD framework. Available online: <https://ceos.org/ard/>. (Accessed 22 December 2021).
- Chini, M., Pelich, R., Pulvirenti, L., Pierdicca, N., Hostache, R., Matgen, P., 2019. Sentinel-1 InSAR coherence to detect floodwater in urban areas: Houston and hurricane harvey as a test case. *Remote Sens.* 11 (2), 107.
- Dostálová, A., Lang, M., Ivanovs, J., Waser, L.T., Wagner, W., 2021. European wide forest classification based on Sentinel-1 data. *Remote Sens.* 13 (3), 337.
- Doubkova, M., Wagner, W., Naeimi, V., Cao, S., Bauer-Marschallinger, B., Kidd, R., Hasenauer, S., Dostalova, A., Paulik, C., 2016. The use of Sentinel-1 for monitoring of soil moisture within the copernicus global land service. In: Living Planet Symposium, vol. 740, p. 105.
- EODC, 2021. Earth observation data centre. Available online: <https://eodc.eu/>. (Accessed 18 October 2021).
- ESA, 2021a. POD products and requirements. Available online: <https://sentinel.esa.int/web/sentinel/technical-guides/sentinel-1-sar/pod/products-requirements>. (Accessed 18 October 2021).
- ESA, 2021b. SNAP. Available online: <https://step.esa.int/main/toolboxes/snap/>. (Accessed 18 October 2021).
- ESA, 2022. RCM (RADARSAT constellation mission). Available online: <https://directory.eoportal.org/web/eoportal/satellite-missions/r/rcm>. (Accessed 10 May 2022).
- Fahrland, E., 2020. Copernicus Digital Elevation Model Product Handbook. Airbus, Available at https://spacedata.copernicus.eu/documents/20126/0/GEO1988-CopernicusDEM-SPE-002_ProductHandbook_11.00.pdf, Version 2.1.
- Filippini, F., 2019. Sentinel-1 GRD preprocessing workflow. In: Multidisciplinary Digital Publishing Institute Proceedings, vol. 18, (1), p. 11.
- GAMMA Remote Sensing, 2021. The GAMMA software. Available online: <https://www.gamma-rs.ch/software>. (Accessed 18 October 2021).
- Geudtner, D., Torres, R., Snoeij, P., Davidson, M., Rommen, B., 2014. Sentinel-1 system capabilities and applications. In: 2014 IEEE Geoscience and Remote Sensing Symposium. IEEE, pp. 1457–1460.
- Grazia Castriotta, A., Volpi, F., 2020. Copernicus sentinel data access annual report 2020. Available online: <https://sentinels.copernicus.eu/web/sentinel/-/copernicus-sentinel-data-access-annual-report-2020>. (Accessed 9 March 2022).
- Hahn, S., Reimer, C., Vreugdenhil, M., Melzer, T., Wagner, W., 2017. Dynamic characterization of the incidence angle dependence of backscatter using Metop ASCAT. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 10 (5), 2348–2359.
- ISCE-Framework, 2021. InSAR scientific computing environment version 2. Available online: <https://github.com/isce-framework/isce2>. (Accessed 18 October 2021).
- Kumar, V., Venkataraman, G., Rao, Y., Singh, G., et al., 2008. Spaceborne InSAR technique for study of Himalayan glaciers using ENVISAT ASAR and ERS data. In: IGARSS 2008-2008 IEEE International Geoscience and Remote Sensing Symposium, vol. 4. IEEE, pp. IV–1085.
- Lievens, H., Brangers, I., Marshall, H.-P., Jonas, T., Olefs, M., De Lannoy, G., 2021. Sentinel-1 snow depth retrieval at sub-kilometer resolution over the European Alps. *Cryosphere Discuss.* 1–25.
- Mullissa, A., Vollrath, A., Odongo-Braun, C., Slagter, B., Balling, J., Gou, Y., Gorelick, N., Reiche, J., 2021. Sentinel-1 SAR backscatter analysis ready data preparation in Google Earth engine. *Remote Sens.* 13 (10), 1954.
- Peters, M., 2021. Bulk processing with GPT. Available online: <https://senbox.atlassian.net/wiki/spaces/SNAP/pages/70503475/Bulk+Processing+with+GPT>. (Accessed 18 October 2021).
- Prats-Iraola, P., Rodriguez-Cassola, M., De Zan, F., Scheiber, R., López-Dekker, P., Barat, I., Geudtner, D., 2015. Role of the orbital tube in interferometric spaceborne SAR missions. *IEEE Geosci. Remote Sens. Lett.* 12 (7), 1486–1490.
- Rocca, F., 2004. Diameters of the orbital tubes in long-term interferometric SAR surveys. *IEEE Geosci. Remote Sens. Lett.* 1 (3), 224–227.
- Salamon, P., Mctormick, N., Reimer, C., Clarke, T., Bauer-Marschallinger, B., Wagner, W., Martinis, S., Chow, C., Böhnke, C., Matgen, P., et al., 2021. The new, systematic global flood monitoring product of the copernicus emergency management service. In: 2021 IEEE International Geoscience and Remote Sensing Symposium, IGARSS. IEEE, pp. 1053–1056.
- Schaufler, S., Bauer-Marschallinger, B., Hochstöger, S., Wagner, W., 2018. Modelling and correcting azimuthal anisotropy in Sentinel-1 backscatter data. *Remote Sens. Lett.* 9 (8), 799–808.
- Schreier, G. (Ed.), 1993. SAR Geocoding: Data and Systems. Wichmann.
- Small, D., 2011. Flattening gamma: Radiometric terrain correction for SAR imagery. *IEEE Trans. Geosci. Remote Sens.* 49 (8), 3081–3093.
- Small, D., Rohner, C., Miranda, N., Rüetschi, M., Schaeppman, M.E., 2021. Wide-area analysis-ready radar backscatter composites. *IEEE Trans. Geosci. Remote Sens.* 60, 1–14.
- Small, D., Schubert, A., 2019. Guide to Sentinel-1 Geocoding - UZH-S1-GC-AD. Technical Report 1.10, University of Zurich.
- SNAP Team, 2021. SNAP 8.0 released. Available online: <https://step.esa.int/main/snap-8-0-released/>. (Accessed 18 October 2021).
- Tamm, T., Zalite, K., Voormansik, K., Talgre, L., 2016. Relating sentinel-1 interferometric coherence to mowing events on grasslands. *Remote Sens.* 8 (10), 802.
- Thomopoulos, N.T., 2013. Essentials of Monte Carlo Simulation: Statistical Methods for Building Simulation Models. Springer, New York.
- Tobler, W.R., 1970. A computer movie simulating urban growth in the detroit region. *Econ. Geogr.* 46 (sup1), 234–240.
- Torres, R., Lokas, S., Di Cosimo, G., Geudtner, D., Bibby, D., 2017. Sentinel-1 evolution: Sentinel-1C and -1D models. In: 2017 IEEE International Geoscience and Remote Sensing Symposium, IGARSS. IEEE, pp. 5549–5550.
- Torres, R., Snoeij, P., Geudtner, D., Bibby, D., Davidson, M., Attema, E., Potin, P., Rommen, B., Floury, N., Brown, M., et al., 2012. GMES Sentinel-1 mission. *Remote Sens. Environ.* 120, 9–24.
- Truckenbrodt, J., Freemantle, T., Williams, C., Jones, T., Small, D., Dubois, C., Thiel, C., Rossi, C., Syriou, A., Giuliani, G., 2019. Towards Sentinel-1 SAR analysis-ready data: A best practices assessment on preparing backscatter data for the cube. *Data* 4 (3), 93.
- Twele, A., Cao, W., Plank, S., Martinis, S., 2016. Sentinel-1-based flood mapping: A fully automated processing chain. *Int. J. Remote Sens.* 37 (13), 2990–3004.
- Vincent, P., Bourbigot, M., Johnsen, H., Piantanida, R., 2020. Sentinel-1 Product Specification. ESA, Available at <https://sentinels.copernicus.eu/documents/247904/1877131/Sentinel-1-Product-Specification.pdf/49c514c3-1574-4d94-aae2-d8061a3baebd?t=1584020315000>, Issue 3/7.
- Vreugdenhil, M., Navacchi, C., Bauer-Marschallinger, B., Hahn, S., Steele-Dunne, S., Pfeil, I., Dorigo, W., Wagner, W., 2020. Sentinel-1 cross ratio and vegetation optical depth: A comparison over Europe. *Remote Sens.* 12 (20), 3404.
- Wagner, W., Bauer-Marschallinger, B., Navacchi, C., Reuss, F., Cao, S., Reimer, C., Schramm, M., Briese, C., 2021. A Sentinel-1 data cube for global land monitoring applications. In: Proceedings of the 2021 Conference on Big Data from Space. pp. 49–52. <http://dx.doi.org/10.2760/125905>.
- Wagner, W., Freeman, V., Cao, S., Matgen, P., Chini, M., Salamon, P., McCormick, N., Martinis, S., Bauer-Marschallinger, B., Navacchi, C., et al., 2020. Data processing architectures for monitoring floods using Sentinel-1. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* 641–648.
- Wegmuller, U., Santoro, M., Werner, C., Cartus, O., 2015. On the estimation and interpretation of Sentinel-1 TOPS InSAR coherence. In: Proc. of FRINGE 2015 Workshop, Frascati, Italia. pp. 23–27.
- Westerhoff, R., Kleuskens, M., Winsemius, H., Huizinga, H., Brakenridge, G., Bishop, C., 2013. Automated global water mapping based on wide-swath orbital synthetic-aperture radar. *Hydrol. Earth Syst. Sci.* 17 (2), 651–663.
- Woodhouse, I.H., 2006. Introduction to Microwave Remote Sensing. CRC Press.
- Yagüe-Martínez, N., Prats-Iraola, P., Gonzalez, F.R., Brcic, R., Shau, R., Geudtner, D., Eineder, M., Bamler, R., 2016. Interferometric processing of Sentinel-1 TOPS data. *IEEE Trans. Geosci. Remote Sens.* 54 (4), 2220–2234.