

GSLC Terrain Correction, Explained

A conceptual guide to the `GSLC-Terrain-Correction` operator (Geocoded Single Look Complex generation) in the SNAP Microwave Toolbox.

Audience: users who want to understand **what the operator does, why it does it that way, and how to drive it** — not just which button to press. Read it start-to-finish, then follow the hands-on tutorial (GSLC-Tutorial.md) for a step-by-step run. For the button-level reference see the in-app help (GSLCGeocodingOp.html); for a runnable end-to-end walkthrough see the notebook snap-nb-sar-gslc-insar.

1. The problem it solves

Traditional SAR terrain correction (Range-Doppler) geocodes the **amplitude** and, in doing so, discards or scrambles the **phase**. But the phase of a complex SAR signal encodes the sensor-to-target distance — the very thing interferometry and polarimetry depend on. So the classical InSAR pipeline keeps everything in slant-range geometry through a long, mission-specific chain (Back-Geocoding → ESD → Interferogram → Deburst → Merge → Terrain-Correction) and geocodes only at the very end.

GSLC (Geocoded Single Look Complex) flips that around: it geocodes each acquisition to a map projection **up front, while preserving phase**. Every scene becomes a phase-preserving complex product on a **common map grid**, and the rest of the InSAR chain consumes those directly. The result:

- **Fewer steps** — no separate deburst, merge, or final terrain correction.
- **Independent, parallel** per-acquisition processing.
- **Consistent geocoding** across all dates (same DEM, same geometry) — so a pixel is the *same ground point* in every scene.
- Output that is **directly geocoded** and ready for stacking / time-series analysis.

This is the "geocode-first" architecture behind global-scale systems (Agram et al. 2022; OPERA-CSLC / NISAR ADT), now native in the SNAP Microwave Toolbox.

2. The core idea in one paragraph

For every pixel of the **output map grid**, use the precise orbit + a DEM to find where that ground point falls in the source SLC (its zero-Doppler azimuth time and slant range). Before resampling, **remove the geometric carrier phase** ($4\pi R/\lambda$) so the signal is smooth enough to interpolate cleanly; resample the complex I/Q with a high-fidelity interpolator; then **restore that same carrier phase** in the new map geometry. The output complex pixel therefore carries the correct sensor-to-target phase for its map location — so a reference $\cdot \text{conj}(\text{secondary})$ product between two GSLCs yields exactly the differential-range phase (the $\exp(-j \cdot 4\pi(R_{\text{reference}} - R_{\text{secondary}})/\lambda)$ term), with no coregistration chain in between.

3. How it works, step by step (as implemented)

This mirrors the operator's per-pixel geometry loop.

3.1 Precise geometry — where does this ground point live in the SLC?

For each output-grid pixel the operator computes its geographic coordinate (lat, lon, height) from the DEM, then runs a **zero-Doppler search** against precise orbit state vectors to find the corresponding azimuth time, and computes the **slant range** R from sensor to ground point. Optional geometry-level corrections (Solid Earth Tide) shift the ground point itself before this solve.

3.2 Phase flattening — carrier removal before resampling

The raw SLC phase spins rapidly (high fringe frequency from the $4\pi R/\lambda$ carrier). Interpolating that directly aliases. So the operator simulates the carrier phase

$$\phi_{\text{sim}} = 4\pi \cdot R / \lambda$$

and removes it from the complex signal. The SLC carrier is $\exp(-j \cdot \phi_{\text{sim}})$, so flattening **multiplies by** $\exp(+j \cdot \phi_{\text{sim}})$ to cancel it: $C' = C \cdot \exp(+j \cdot \phi_{\text{sim}})$. The flattened signal has low fringe frequency and resamples cleanly.

***TOPS note.** For IW/EW data an additional azimuth deramp (from the antenna steering) is removed first, and the range carrier is pre-flattened on the deramped tile before resampling. The azimuth carrier is **not put back** by default (`outputAzimuthCarrier = false`): it is acquisition-specific — burst timing and FM rate — and does **not** cancel between two acquisitions, so restoring it injects a per-burst quadratic azimuth phase into every cross-acquisition interferogram (~150 rad ≈ 24 spurious fringes per burst measured on a real S1A+S1D pair; the effect is small for same-platform pairs with tight burst sync, which is why it can go unnoticed). Classical InSAR is immune because Back-Geocoding resamples the secondary onto the reference's burst grid; independent per-scene geocoding cannot do that, so the carrier must stay off — the same convention as OPERA CSLC. The carrier state is stamped in metadata (`gslc_azimuth_carrier`) and `CreateStack` matches it when auto-building a secondary; all legs of a stack must share one convention. Stripmap data with a residual Doppler centroid (e.g. ENVISAT ASAR IMS) still gets its sample-by-sample azimuth deramp/reramp — the stripmap f_{dc} carrier is orders of magnitude gentler, though a cross-acquisition f_{dc} difference is a known residual there too.*

*One residual remains in a carrier-free cross-acquisition interferogram: with the carrier restored, deramp-model errors cancel in the round trip; with it removed, each leg keeps its own annotation-vs-data mismatch, and the difference between two acquisitions survives into the interferogram. **This residual is per burst, not one smooth surface** — every burst carries its own Doppler-centroid/FM-rate annotation, so the residual is a quadratic-in-azimuth-time per burst with genuine discontinuities at burst seams and range-dependent structure inside each burst. On a same-platform pair it is small (~1 fringe per 80 px measured on S1A+S1D); on a **cross-platform pair it is large and strongly burst-variable** (S1A×S1C, 10-burst IW3: per-burst azimuth phase rates spanning 39.7–68.3 rad/s — a spread no scene-global fit can represent).*

*Interferogram's GSLC mode removes it **deterministically**, with no fitting:*

- **Exact carrier-difference subtraction (automatic).** The GSLC operator knows precisely which deramp model it removed from each leg. With ``outputPhaseTerms=true`` (the default) it writes that model phase as an

azimuthCarrierPhase band, *CreateStack* carries the band per leg, and *Interferogram* subtracts the two models' **exact difference** — full azimuth and range structure, every burst. Confirm it in the log: *GSLC carrier-difference: exact deramp-model subtraction active*. If only one leg carries the band, the operator warns and applies **no** cross-acquisition correction — produce both *GSLCs* with the same setting.

- **Burst seam steps (optional)**. Anything still discontinuous at the burst seams is handled by ``subtractSeamSteps``, which measures each seam's step by across-seam differencing. Anything continuous across a seam — deformation, atmosphere, orbital ramps — cancels in that difference to the extent it is locally linear in azimuth, so unlike a fitted surface it cannot absorb a broad signal wholesale. Off by default; enable it only if seams are visible. No effect on stripmap.

*This is the *GSLC*-domain analogue of why classical TOPS InSAR needs ESD: the annotation is trusted for what it can predict.*

History — the removed ``subtractResidualRamp`` option. Earlier versions also fitted a data-driven per-burst ramp (*subtractResidualRamp*, with *residualRampDegree*, *residualRampRangeProfile*, *residualRamp2D*, *residualRamp2DNodes*). Most of what it was removing on TOPS pairs was not an annotation error at all but an **inverted carrier-difference add-back sign**, which left $2 \times (m_{\text{ref}} - m_{\text{sec}})$ in every *GSLC* TOPS interferogram. With the sign corrected, the Venezuela S1A×S1C pair agrees with the classical chain at a phase concentration of **0.961 without the ramp fit and 0.334 with it** — it was removing real signal (a retention test showed it absorbing ~2.4 % of a broad 6 rad range-only lobe). The whole ramp family has been deleted (*microwave-toolbox/docs/gslc-parity/etna-ramp-ab.md* (<https://github.com/senbox-org/microwave-toolbox/blob/master/docs/gslc-parity/etna-ramp-ab.md>)). Per-burst figures elsewhere in this document that were measured under the legacy sign are marked as such.

3.3 High-fidelity complex resampling

The flattened I and Q are resampled from the source slant-range grid onto the target map grid. The default interpolator is **BiSinc 5-point** — a truncated sinc, chosen because phase fidelity (not just amplitude) must survive resampling. Cheaper kernels (bilinear, cubic) are available but not recommended for InSAR.

3.4 Phase restoration — the InSAR-critical default

By default (``outputFlattened = false``) the operator **puts the carrier back** with the inverse multiply $C_{\text{geo}} = C'_{\text{resampled}} \cdot \exp(-j \cdot \phi_{\text{sim}})$. This restores the absolute phase relative to the *new map geometry*, which is exactly what InSAR needs — the reference·conj(secondary) product then carries the differential-range phase $\exp(-j \cdot 4\pi(R_{\text{reference}} - R_{\text{secondary}})/\lambda)$.

If `outputFlattened = true`, the carrier is **not** restored: each pixel holds only the local scattering coefficient $\sigma(P)$. That is useful for amplitude/scattering or single-date PolSAR work, but it makes the *GSLC unusable for InSAR* — the range-difference phase is zeroed. **Use one mode consistently across all dates in a stack**; mixing flattened and non-flattened bands produces a noise interferogram. (The operator stamps `gslc_output_flattened` into metadata so *CreateStack* can enforce a matching state on an auto-built secondary.)

Sign convention — a fixed bug worth flagging. Flatten uses $\exp(+j\phi)$ (to cancel the SLC's $\exp(-j\phi)$ carrier) and restore uses the inverse $\exp(-j\phi)$. An earlier version used $\exp(+j\phi)$ in **both** steps, doubling the carrier

and corrupting the interferogram. The current code is correct (`multiplyByExpJPhi` to flatten, `multiplyByExpMinusJPhi` to restore) and pinned by `GSLCInSarGradeTest`.

3.5 Coregistration — folded into the grid, not a separate chain

Because every GSLC is always built on the **same standard map grid** (a fixed global lat/lon lattice anchored at 0,0), two overlapping scenes already share it: the same ground point lands at the same fractional pixel (modulo an integer offset). That is coregistration by construction. Any residual sub-pixel misregistration (clock bias, processor-time delta, orbit residual) is absorbed by two scalar offsets — `rangeOffsetPixels` and `azimuthOffsetPixels` — applied inside the geometric sampling.

The simplest workflow geocodes **only the reference** explicitly; `CreateStack` then reads the reference's `gslc_source_slc_path` stamp, cross-correlates against the raw secondary SLC to estimate the (Δr , Δa) bias, rebuilds the secondary GSLC with those offsets, and stacks the two. The user runs GSLC once, feeds raw SLC secondaries, and `CreateStack` handles the rest.

3.5b Output grid — square or rectangular cells

A map projection fixes the grid axes' **directions and units**, not the cell aspect ratio — the two sampling steps are independent. When no explicit pixel spacing is given, the `gridSpacing` parameter picks the step:

gridSpacing	Cell	Trade-off (S1 IW)
<code>'NATIVE_ANISOTROPIC' (default)</code>	Rectangular, native sampling on each axis	Nothing discarded; $\sim 4\times$ the pixels of <code>SQUARE_COARSEST</code>
<code>SQUARE_COARSEST</code>	Square at the coarser axis, $\max(\text{azimuthSpacing}, \text{rangeSpacing})$	Smallest output, but discards $\sim 4\times$ of range sampling (≈ 3.4 m range coarsened to ≈ 14 m) — fewer independent looks for coherence and multilooking
<code>SQUARE_FINEST</code>	Square at the finer axis	Nothing discarded, square pixels, $\sim 16\times$ the pixels

Each choice gives a different grid step and therefore a different standard-grid lattice, so **use one setting for every product in a stack**. Products made before the default changed used `SQUARE_COARSEST`; set it explicitly to extend such a stack.

The output metadata `range_spacing` / `azimuth_spacing` hold the **true ground step** of the output grid (east and north). On a geographic grid the nominal metre step defines the global lattice and is converted to degrees at the equator, so the real east–west step at the scene is that value $\times \cos(\text{latitude})$ — e.g. 2.35 m nominal $\rightarrow$ 1.78 m at 40.8° N. Metre-based consumers (Interferogram's `cohWinSizeMeters`, Goldstein, Coherence) read these attributes. Products made before 26 Sep 2026 stored the nominal step instead, which shrank a "100 m" coherence window to ~ 76 m east–west at Campi Flegrei; regenerate such products before comparing coherence.

Which axis is the fine one is mission-dependent — do not assume it is range. S1 IW is range-fine / azimuth-coarse, but BIOMASS stripmap is the reverse (19.81 m slant-range vs 6.71 m azimuth spacing, so azimuth-fine / range-coarse), and its ground aspect ratio is *larger* than S1's. The rule above and the rectangular-cell criterion

below are both symmetric in the two axes, so they hold either way; only the resulting orientation flips ($dy < dx$ instead of $dx < dy$).

Setting the optional **Pixel Spacing North** (`pixelSpacingInMeterY/pixelSpacingInDegreeY`) sets the two axes explicitly. The limit on how coarse the north step may be is set by the orbit heading θ rotating the radar axes away from the map axes: the Nyquist-adequate steps are $dx \leq 1/(B_{rg} \cdot \cos\theta + B_{az} \cdot \sin\theta)$ and $dy \leq 1/(B_{rg} \cdot \sin\theta + B_{az} \cdot \cos\theta)$. For S1 ($|\theta| \approx 10\text{--}15^\circ$) that gives $\approx \mathbf{3.4\text{ m east} \times 7.5\text{ m north}}$ — near-full information preservation at $\sim 4\times$ fewer samples than a 3.4 m square grid. (Precedent: OPERA CSLC-S1 ships 5×10 m UTM cells.) `CreateStack`'s grid-lock and the lattice-alignment guard operate per axis, so rectangular stacks coregister exactly like square ones.

On **oversampling**: because the resampler interpolates band-limited complex data (deramped, range-pre-flattened) with sinc-family kernels, requesting a *finer output grid* is band-limited oversampling — mathematically equivalent to FFT zero-padding of the source. Use `oversamplingPercent` (or simply a finer explicit spacing) together with `BISINC_21_POINT` for very-high-resolution work such as point-target analysis; no separate source-domain oversampling stage is needed.

For lossless native-resolution cells (3.4×14 m with no heading penalty) the map axes themselves would have to align with the radar axes — i.e. a per-relative-orbit oblique projection (Hotine Oblique Mercator along the ground track). That remains a legitimate, EPSG-parameterizable map projection and a possible future option; rectangular cells on a standard CRS cover most of the benefit today.

3.6 Displacement-grade corrections (optional)

For sub-decimetre deformation work, two geometry/phase corrections are exposed:

Correction	What it does	Status
Solid Earth Tide (<code>applySolidEarthTide</code>)	~ 10 cm body-tide displacement (IERS 2010 step-1, degree-2 Love numbers); shifts both <code>rangeIndex</code> and phase	Implemented
Troposphere (<code>applyTroposphericCorrection</code>)	Saastamoinen dry path delay (standard atmosphere); shifts phase only	Implemented

SET and troposphere are negligible for coherence-only or amplitude work; they matter for displacement-grade InSAR.

These are not the only corrections available, and they overlap with the others. A GSLC InSAR chain has three independent places to apply corrections:

1. **ETAD** (`S1-ETAD-Correction`), applied to the **split SLC before geocoding** — Sentinel-1's auxiliary timing product supplies measured tropospheric and ionospheric range delay, geodetic effects, and bistatic/FM-mismatch azimuth shifts. Its `sumOfRangeCorrections / sumOfAzimuthCorrections` parameters default to **true**, so ETAD's default is the *total* correction; the seven individual layer switches (default false) exist to apply a subset. **For a GSLC chain, ``resamplingImage=true` + `outputPhaseCorrections=true`` is mandatory:** with `resamplingImage=false` the operator only *attaches* the corrections as tie-point grids, which are

consumed exclusively by the classical slant-range interferogram path — for GSLC that mode is a guaranteed silent no-op (measured: the "corrected" GSLC came out bit-identical to the uncorrected one).

2. **GSLC itself** — the two rows above.
3. **The interferogram domain** — the `IonosphericCorrection` operator (phase screens).

Because ETAD's tropospheric layer and GSLC's `applyTroposphericCorrection` model the **same** path delay, enabling both **double-counts** it; the same holds for ETAD's geodetic layers against `applySolidEarthTide`. Choose ETAD (measured) where an ETAD product exists, otherwise GSLC's computed model. ETAD must also be applied to **both** acquisitions or to neither — a differential timing correction goes straight into the interferometric phase.

Which corrections actually help — measured

An eight-configuration ablation was run on a Sentinel-1B IW1 pair (15 Aug × 08 Sep 2020, 24-day baseline, 2 bursts), each configuration processed end to end through GSLC → CreateStack → Interferogram with the (since removed) ``subtractResidualRamp`` option off so that ramp removal could not absorb the long-wavelength signal under test. Identical 1200 × 1200 window in every case.

Configuration	Residues /10 ⁴ px	Δ vs baseline	5×5 local coh	Δ	mean coh	Δ
Baseline, no corrections	2554.6	—	0.2859	—	0.2850	—
ETAD, summed range + azimuth	2411.6	−5.6 %	0.2931	+0.0072	0.2790	−0.0060
ETAD, azimuth layers only	2477.4	−3.0 %	0.2879	+0.0020	0.2744	−0.0106
ETAD, tropospheric only	2487.4	−2.6 %	0.2879	+0.0020	0.2750	−0.0100
ETAD, ionospheric only	2509.2	−1.8 %	0.2857	−0.0002	0.2728	−0.0122
ETAD, geodetic range only	2511.1	−1.7 %	0.2856	−0.0003	0.2726	−0.0124
GSLC <code>applySolidEarthTide</code>	2511.0	−1.7 %	0.2856	−0.0003	0.2728	−0.0122
GSLC <code>applyTroposphericCorrection</code>	2535.5	−0.7 %	0.2714	−0.0145	0.2258	−0.0592

What this does and does not establish:

1. **Full ETAD is the best configuration on both phase metrics** — the lowest residue density (−5.6 %) and the only configuration to raise 5×5 local phase coherence appreciably (+0.0072). It beats every individual layer.
2. **The layers are NOT additive, despite appearances.** Azimuth-only (−3.0 %) plus tropospheric-only (−2.6 %) happens to equal the combined −5.6 %, but adding the ionospheric (−1.8 %) and geodetic-range (−1.7 %) terms as well would give −9.1 %, far more than the −5.6 % actually measured. The two-term agreement is a coincidence; the single-layer figures share a large common component and must not be summed.
3. **The two coherence metrics disagree, and this is unresolved.** Residue density and *local* phase coherence improve with ETAD, but **mean coherence falls in every single configuration** (−0.0060 to −0.0592). Only the local/phase-structure metrics support "ETAD helps". A plausible reading is that resampling slightly decorrelates the pair while improving phase-gradient structure, but this has not been demonstrated. **Treat the improvement as established for unwrappability, not for coherence.**
4. **Ionospheric, geodetic-range and solid-Earth-tide corrections are at the noise floor here** (−1.7 to −1.8 % residues, local coherence flat to −0.0003). For ionosphere that matches theory — TEC phase scales as 1/f, so

it only becomes first-order at P-band (BIOMASS, $\sim 12\times$ larger) — but on this evidence none of the three is individually distinguishable from no correction.

5. **GSLC's built-in Saastamoinen tropospheric model degrades the result** — it is the worst configuration on every metric (local coherence $0.2859 \rightarrow 0.2714$, mean $0.2850 \rightarrow 0.2258$). A standard-atmosphere model applied independently to each acquisition injects more error than it removes over a 24-day pair. **Prefer ETAD's measured correction wherever an ETAD product exists.**
6. **`applySolidEarthTide` and ETAD's geodetic-range layer agree to 0.1 residues** (2511.0 vs 2511.1) and to 4 decimals on both coherence metrics — two independent implementations of the same physical effect landing on the same answer. A useful cross-check, and a concrete illustration of the double-counting overlap noted above.

Scope of these numbers — they measure the geometric contribution only. Every configuration above ran in ETAD's resampling mode with `outputPhaseCorrections=false`, so the gains come from better geolocation and hence better coregistration. Resampling relocates a pixel but leaves the $-4\pi\Delta r/\lambda$ that the delay put in its phase, and InterferogramOp's ETAD phase handling runs only on the classical slant-range paths, never on the GSLC path. Enabling `outputPhaseCorrections` removes the range-delay phase from the complex samples before geocoding, which is the correct route for a geocode-first chain.

The atmospheric phase term, measured separately

Re-running the full-ETAD configuration with `outputPhaseCorrections=true` changes **every** pixel, and the correction it applies is a **smooth** field — adjacent-pixel gradient of the *differential* (the part that survives into the interferogram) is 0.063 rad/px mean, against 1.571 rad/px for a random field — spanning **0.98 cycles $\approx$ 2.7 cm LOS** across the scene. That is a physically plausible differential tropospheric signal for a 24-day pair.

Its effect on the residue/local-coherence metrics above is **nil** (2411.6 $\rightarrow$ 2416.1 residues, local coherence $0.2931 \rightarrow 0.2916$). That is a limitation of those metrics, not of the correction: **residue density and a 5 $\times$ 5 coherence window are measures of *local* phase noise and are blind to a smooth one-cycle-per-scene field by construction**. Choosing them for an atmospheric ablation was a design error worth recording.

Measured with a statistic that can see long-wavelength phase — the circular standard deviation of the 32×32 multilooked interferometric phase — the correction clearly works:

	circular σ of 32×32 phase	concentration R
ETAD geometric only	2.8069 rad	0.0195
ETAD + phase correction	2.5726 rad (–8.3 %)	0.0365 (+87 %)

This also validates the **sign**: an inverted correction would double the atmospheric phase and degrade this statistic, whereas it improves substantially. R remains small in absolute terms because residual topography, deformation and un-modelled atmosphere all still contribute.

Guidance: enable `outputPhaseCorrections` for any displacement-grade GSLC work, and do not judge an atmospheric correction with residue counts or small-window coherence.

Full experimental design in [microwave-toolbox/docs/superpowers/plans/2026-07-27-cross-toolbox-validation.md](https://github.com/senbox-org/microwave-toolbox/blob/master/docs/superpowers/plans/2026-07-27-cross-toolbox-validation.md) (<https://github.com/senbox-org/microwave-toolbox/blob/master/docs/superpowers/plans/2026-07-27-cross-toolbox-validation.md>) §B4.

4. Where it sits in the pipeline

Traditional TOPS InSAR		GSLC-based InSAR
Apply-Orbit-File		Apply-Orbit-File
TOPSAR-Split		TOPSAR-Split
Back-Geocoding (coregister)	→	GSLC-Terrain-Correction (reference)
Enhanced-Spectral-Diversity	→	CreateStack (auto-coregisters secondary)
Interferogram		Interferogram
TOPSAR-Deburst	→	—
TOPSAR-Merge	→	—
Terrain-Correction	→	— (already geocoded)
	↓	
		Goldstein filter → Phase-Unwrapping → ...

The GSLC approach collapses steps 3–4 and 6–8 of the traditional chain. Each acquisition is geocoded once, independently, on a shared grid.

5. Parameter tuning — a practical guide

Parameter	Default	Notes / when to change
DEM	Copernicus 30m Global	Auto-downloaded; supply an external DEM for local high-res work
Image Resampling	BiSinc 5-point	Keep a sinc kernel for InSAR; higher orders (11/21) trade speed for fidelity
DEM Resampling	Bilinear	Bilinear is fine for most terrain; BiSinc for steep relief
Grid Spacing (gridSpacing)	NATIVE_ANISOTROPIC	How the step is derived when Pixel Spacing is 0: rectangular native cells (default), SQUARE_COARSEST (smallest, discards ~4× range detail on S1 IW) or SQUARE_FINEST (~16× pixels). Same setting for every product in a stack (§3.5b)
Pixel Spacing (m / deg)	0 (auto)	0 → derived from source via Grid Spacing; set explicitly to fix an output resolution across a stack. Becomes the east step when a north spacing is also set
Pixel Spacing North (m / deg)	0	Optional explicit north step for

		rectangular cells — S1 IW Nyquist-adequate: ≈ 3.4 m east $\times$ 7.5 m north (see §3.5b)
Oversampling (%)	0	20% $\rightarrow$ 20% finer pixels; use to reduce interpolation loss on high-relief scenes. A finer grid through the sinc kernel <i>is</i> band-limited oversampling (§3.5b)
Map Projection	WGS84(DD)	Any WKT CRS; use a metric projection (UTM) if you need metre grids
Output phase-flattened	false	Leave false for InSAR. True only for amplitude/PolSAR single-date use
Restore TOPS azimuth carrier	false	Leave false. Acquisition-specific; does not cancel between acquisitions — restoring it corrupts cross-acquisition InSAR (~tens of spurious fringes per burst)
Output separable phase terms (outputPhaseTerms)	true	Leave on for InSAR. Writes the removed deramp model as an azimuthCarrierPhase band (plus the flattening phase), enabling Interferogram's exact carrier-difference subtraction (§3.2). Must be on for both legs (the CreateStack auto-path matches the reference automatically). Adds float64 bands ($\sim 2 \times$ product size) — turn off only for pure backscatter products
Range/Azimuth offset (px)	0	Sub-pixel secondary alignment; CreateStack sets these automatically
Apply SET / Tropo	false	Turn on SET + Tropo for displacement-grade InSAR
Mask out no-elevation areas	true	Skips sea/no-DEM pixels; speeds processing

First-run recipe: all defaults (outputFlattened=false, outputPhaseTerms=true, gridSpacing=NATIVE_ANISOTROPIC). Geocode the reference; feed raw SLC secondaries to CreateStack; on the Interferogram step enable subtractFlatEarthPhase and subtractTopographicPhase, and set **Coherence Window (m)** (cohWinSizeMeters) to ~ 100 so the window is square on the ground. With the phase-term bands present the operator subtracts the two legs' deramp models **exactly**, which since the carrier-sign fix is the whole of what was needed (see the TOPS note in §3.2). Add subtractSeamSteps only if burst seams are visible. Inspect the interferogram coherence; enable ETAD (option-1 mode) or SET+tropo only when chasing sub-decimetre displacement. Judge phase quality **after** filtering or multilooking — a correct single-look interferogram at $\gamma \approx 0.2$

looks like noise on screen, and a decimated full-scene zoom aliases dense fringes into misleading moiré on a fine GSLC grid.

6. Supported inputs (and what is rejected)

Product	Supported	Note
Stripmap SLC (S-1 SM, ENVISAT ASAR IMS)	Yes	Direct input after orbit application
IW/EW SLC, split (single subswath)	Yes	Via TOPSAR-Split; operator handles deramp + burst mosaicking internally
IW/EW SLC, deburst	No	Deburst strips burst metadata needed for azimuth deramp — use the split product
GRD	Not usable	Amplitude-only — no phase. Not explicitly blocked by the operator, but produces a degenerate (phase-less) product, so don't feed GRD to GSLC.

For full IW coverage, process each subswath (IW1/IW2/IW3) independently and mosaic the GSLCs (or the interferograms) per date.

6.1 Archive (pre-2000) stripmap data — known limitation

GSLC InSAR converts each product's annotation directly into absolute phase, so it inherits annotation errors that classical InSAR silently absorbs in its data-driven coregistration warp. On a 1995 ERS-1/ERS-2 tandem pair (VMP-processed) we measured, against a classical control on identical inputs:

- a constant range bias (~ 0.5 px) and a ~ 2 px azimuth timing drift (a ~ 2 ms first/last-line annotation error) — **corrected automatically** by CreateStack's block-cross-correlation offset field (affine; the estimator's precision, $\sim \pm 0.05$ px, deliberately bounds the model);
- a large phase-vs-annotation range drift (~ 270 fringes full swath, linear) — at the time **removed** by `subtractResidualRamp`, plus a nonlinear slant-range component (~ 35 rad) removed by its `residualRampRangeProfile` option (both since deleted — see the end of this section);
- a residual smooth two-dimensional annotation-phase surface (~ 1 – 2 fringes per 10 km at its densest, strongest in the scene north, $\sim \pm 0.05$ px equivalent) that **remains** in the interferogram. This component is definitively **phase-only**: row-resolved amplitude cross-correlation of the raw SLC pair (250-row bands, 8 range blocks each) shows the registration needs are smooth to 0.005 px everywhere — there is no feature displacement for any registration estimator, of any polynomial degree, to measure. It is the VMP processor's phase-vs-position relation itself. The GSLC chain pays the full retained carrier (~ 1756 rad per equivalent pixel) for it; the classical chain multiplies the same error only by the flat-earth gradient (~ 0.1 rad/px) and is therefore immune by construction. The `residualRamp2D` surface (with `residualRamp2DNodes`) removed its broad-scale part, but the densest fringes exceeded the node density that coherent block sampling can support (measured: 8353 usable blocks $\rightarrow$ 22×22 nodes on this scene). Goldstein filtering makes the artifact

more visible, as with any coherent fringe pattern. **Both options have been removed: the surface they targeted is no longer present** — re-measured on this same Etna pair 2026-09-22, the GSLC and classical interferograms differ by only 0.13 rad of scene-scale phase (circular spread of ~2.5 km tile means over 996 k samples), while enabling the then-available `subtractResidualRamp` dropped local agreement from 0.57 to 0.34 and scene-scale agreement to nothing (microwave-toolbox/docs/gslc-parity/etna-ramp-ab.md (<https://github.com/senbox-org/microwave-toolbox/blob/master/docs/gslc-parity/etna-ramp-ab.md>)). The paragraph above is kept as the record of what was measured at the time.

Modern missions (Sentinel-1, Capella, NISAR-class) have annotation quality orders of magnitude better; none of this applies at a measurable level there.

7. Validation evidence

InSAR-grade correctness is pinned by an executable spec, `GSLCInSarGradeTest.java` (14 `@Test` methods, built on the Capella SM SLC fixture), plus a TOPS test family (`GSLCTops*Test`, `GSLCVs*ComparisonTest`). Each hardening fix is anchored to a test:

- **Sign convention** — flatten $\exp(+j\phi)$ / restore $\exp(-j\phi)$ (fixes the doubled-carrier bug).
- **TOPS range pre-flatten + analytic reramp** — removes burst-edge aliasing and phase bias.
- **Layover/shadow classification** and **no-source masking**.
- **Solid Earth Tide** math contract vs. reference, and an **identical-pair → zero-phase interferogram** test.
- **Two-burst S-1 IW** exercised in the InSAR-readiness pyramid (synthetic 2-burst geometry in the TOPS unit tests).

Cross-chain parity, current code (Sep 2026). Each site was processed through both chains from the same SLCs and scored as the phase concentration of the wrapped difference GSLC – classical (1 = identical):

Site	Mode	Agreement	Compared on
Venezuela (S1A × S1C, 1 day)	IW TOPS	0.921 raw, 0.961 with one fitted plane removed	354 543 radar cells, nothing resampled
Napa (S1A, 7 × 31 Aug 2014)	Stripmap	0.785 both single look and unfiltered; 0.892 both Goldstein-filtered	2.14 M map samples
Campi Flegrei, 72 d (S1A, 9 Aug × 20 Oct 2023)	IW TOPS	0.840 at 37 m (0.694 single look)	48 593 P-SBAS points
Campi Flegrei, 1 yr (S1A, 13 Oct 2022 × 20 Oct 2023)	IW TOPS	0.789 at 37 m (0.578 single look)	48 591 P-SBAS points

Coherence uses the same 100 m ground window (`cohWinSizeMeters`) on both chains. Campi Flegrei also scores both chains against an independent published P-SBAS reference. Full method, figures and the coherence table are in the tutorial's *How does it compare to the traditional chain?* section.

The two paragraphs below record earlier milestones and are kept as measured.

Head-to-head against the traditional pipeline (S1A + S1D cross-platform pair, Venezuela, 2-burst IW3 fixture; measured under the legacy carrier sign, with the since-removed `subtractResidualRamp` compensating it): with carrier-free output and `subtractResidualRamp`, the GSLC interferogram's phase-only self-coherence matches the classical (Back-Geocoding) interferogram **to three decimals at every estimation window tested** (5–80 px; e.g. win40: 0.0226 vs 0.0223, win80: 0.0130 vs 0.0128, same block, same statistic), and the two interferograms cross-agree at the level expected from their independent resampling noise. Three defects had to fall to get there — a dropped stacking offset, cross-platform lattice mismatch, and the non-cancelling TOPS azimuth carrier — each now pinned by a regression test (`GSLCStackOffsetProbeTest`, lattice tests, `GSLCCarrierResidualTest`).

Full-scene cross-platform pair (S1A × S1C, Venezuela coseismic, 10-burst IW3; legacy carrier sign): measured per-burst azimuth phase rates spanned **39.7–68.3 rad/s across the ten bursts**. Most of that was the inverted add-back sign (§3.2), not annotation error; with the sign corrected, the exact carrier-difference subtraction alone removes it. At matched multilook (~84 m cells) the full-scene GSLC interferogram reproduces the classical chain's fringe field — including the coseismic fan that matches independently published interferograms of the event. Lattice registration between the chains is exact (integer pixel offset, fractional residual 0.000000). Remaining cross-chain differences concentrate where phase gradients are steepest — the signature of sub-pixel geolocation-convention differences between the chains ($\delta \cdot \nabla \phi$), plus ETAD application-route differences — and are comparison-level, not artifacts within the GSLC product.

Correction-layer ablation (S1B IW1, 15 Aug × 08 Sep 2020): eight configurations measured end to end. Full ETAD improves unwrappability by **–5.6 % phase residues** and local phase coherence by **+0.0072**, and GSLC's built-in Saastamoinen tropospheric model degrades the interferogram on every metric. Two honest caveats carried in §3.6: *mean* coherence falls in **every** configuration, so the improvement is established for unwrappability rather than coherence; and the individual layer contributions are **not** additive. These figures are also a lower bound — they measure ETAD's geometric contribution only (see §3.6).

Second mission — BIOMASS (P-band, stripmap, left-looking, 3-day pair): the geocode-first chain runs end to end on a completely different mission. Both legs geocoded independently landed on **one lattice — identical grid step and a whole-pixel origin offset of exactly $-71.000000 \times -15.000000$ px (to 1e-13)** — and the resulting interferogram is 81.7 % valid at **mean coherence 0.528**. This matters because the standard-grid snapping was designed and debugged entirely against Sentinel-1 to fix the cross-platform lattice defect; that it transfers exactly to a different band, geometry, pixel aspect and look direction is evidence the fix was general rather than tuned. Two BIOMASS-specific findings: square cells at the **coarser** axis (`SQUARE_COARSEST`, the default when this was measured) coarsened the native 6.709 m azimuth sampling 2.95× to 19.814 m — the current `NATIVE_ANISOTROPIC` default keeps it (the fine axis is azimuth here, the reverse of S1 IW); and the stripmap Doppler-centroid residual is negligible at P-band ($|f_{dc}| \leq 0.05$ Hz across the swath) where the TOPS azimuth carrier dominates at C-band.

Known limitations and roadmap:

- A few **end-to-end integration tests are DEM-gated** (skipped without SRTM/Copernicus availability).
- Full-scene validation is on **10-burst IW subswaths**; a multi-subswath slice and long-stack (SBAS/PSI) validation are the natural next steps.

- Exact carrier-difference subtraction requires GSLCs produced with `outputPhaseTerms=true` (the default); stacks built without the bands get **no** cross-acquisition deramp correction (with a log warning when only one leg carries them).
 - Cross-chain (GSLC vs classical) equivalence beyond fringe-field agreement is bounded by two open route-level items: sub-pixel geolocation-convention differences between the chains, and the differing ETAD application routes (baked-in phase vs classical grid consumption).
-

8. Frequently asked questions

- **"Is GSLC just terrain correction on complex data?"** No — it is a geocode-first *architecture*. It preserves phase and builds every scene on a shared grid, so coregistration and the deburst/merge/TC chain collapse into it.
 - **"Why might my GSLC interferogram be noisier than expected?"** Most likely candidates, in order: (1) `outputFlattened` mismatched between reference and secondary; (2) burst-edge / nodata effects when processing many bursts; (3) residual sub-pixel misregistration not captured by the scalar offsets; (4) DEM quality, or SET + troposphere left off for a displacement scene. When comparing with a radar-geometry result, also check the coherence window is set in metres (`cohWinSizeMeters`), not pixels.
 - **"Why does my GSLC interferogram show fringes the classical one doesn't?"** Most often the exact deramp-model subtraction is not running: one or both GSLCs were made with `outputPhaseTerms=false`. Check the log for `exact deramp-model subtraction active`; if it warns that only one leg carries the band, regenerate both with the default. Also judge the result multilooked, not at full-scene zoom — a decimated zoom of a fine GSLC grid aliases dense fringes into moiré. If the fringes are elongated bands broken at the burst seams, enable `subtractSeamSteps`. Older advice to enable `subtractResidualRamp` no longer applies; that option was removed (§3.2).
 - **"Does it work with and without ETAD?"** Yes. ETAD refines the geometry further but is not required.
 - **"ERS / ENVISAT?"** Stripmap ASAR IMS is a direct, supported input — the operator has dedicated ENVISAT Doppler-deramp code (the InSAR-grade test fixture is Capella Stripmap; there is no ENVISAT/S-1 SM end-to-end test yet).
 - **"How is the output grid sized?"** From pixel spacing (m or deg) × oversampling, or from `gridSpacing` when spacing is 0, on the standard lattice anchored at the grid origin; keep `gridSpacing` the same across a stack, or set pixel spacing explicitly to lock resolution.
-

9. References

- Agram, P.S., Warren, M.S., Calef, M.T., Arko, S.A. (2022). *An Efficient Global Scale Sentinel-1 Radar Backscatter and Interferometric Processing System*. Preprints.org, doi:10.20944/preprints202206.0252.v1.
 - OPERA-CSLC / ISCE-framework: geocoded-SLC reference implementation used by the NISAR ADT.
 - IERS Conventions (2010) — Solid Earth Tide displacement model.
-

Companion documents: in-app help `GSLCGeocodingOp.html` (parameter reference) · `snap-nb-sar-gslc-insar` (runnable tutorial) · `GSLCInSarGradeTest.java` (executable correctness spec).