https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Snow accumulation variability limits InSAR SWE retrieval

Zachary Hoppinen[1], Simon Zwieback[1], Ross Palomaki[2], and Hans-Peter Marshall[3]

[1]Geophysical Institute, University of Alaska Fairbanks, 2156 Koyukuk Drive, Fairbanks, AK
[2]Institute of Arctic and Alpine Research, University of Colorado, 4001 Discovery Dr, Boulder, CO 80303, USA
[3]Department of Geosciences, Boise State University, 1910 University Drive, Boise, ID 83725, USA

Correspondence: Zachary Hoppinen (zmhoppinen@alaska.edu)

Abstract. Repeat-pass interferometric Synthetic Aperture Radar (InSAR) can retrieve changes in snow water equivalent (ΔSWE) from the phase delay of radar waves through new snow, but typical interferogram processing averages the phase over windows of tens of meters and assumes the ΔSWE signal is constant within them. We tested this assumption with four pairs of repeat airborne-lidar snow surveys over the Tuolumne basin and two NASA-ISRO SAR (NISAR) InSAR pairs. The lidar-derived ΔSWE power spectrum breaks at a consistent $25\,\mathrm{m}$ (IQR 21–30 m) length scale, so the field varies strongly within an $80\,\mathrm{m}$ multi-look window, and the spatially uniform assumption fails. We capture the resulting effects in a complex-valued sub-window factor M. The variance within the window shortens M and decreases coherence, and the ΔSWE skew within the window rotates M and biases the recovered phase from the mean ΔSWE phase. At L-band, this ΔSWE variability within the window alone lowers the basin-median coherence to 0.25–0.47, before any temporal or thermal decorrelation, while the retrieval bias causes at most millimeters of SWE error in windows that pass a coherence mask. The two NISAR pairs, one spanning a storm and one during a no-accumulation period, show coherence at 20 and $80\,\mathrm{m}$ pixel sizes consistent in space and time with these predictions. The effect of sub-window variability is frequency-dependent, lowering the higher frequency C-band coherence factor to a median of 0.10 while leaving the lower frequency P-band nearly unaffected (0.88). Because the decorrelation effect grows with the multi-look footprint size, processing interferograms finer than NISAR's standard $80\,\mathrm{m}$ posting recovers much of the L-band coherence loss, and a $20\,\mathrm{m}$ footprint nearly doubles the median coherence factor. Sub-window ΔSWE variability is, therefore, a sampling limit that can be predicted from high-resolution snow measurements and used to guide the choice of multi-look footprint for SWE retrievals.

1 Introduction

Snow water equivalent (SWE), the water stored within a snowpack, is a primary input for water supply forecasting, reservoir management, and flood mitigation in snow-fed basins. No current method provides accurate, consistently available, and spatially distributed SWE estimates over watershed scales in complex terrain. Interferometric Synthetic Aperture Radar (InSAR) is a developing tool for measuring SWE that has the potential to provide large-scale, approximately weekly observations of SWE change (ΔSWE). Recent studies using tower (Leinss et al., 2015; Ruiz et al., 2021), airborne (Marshall et al., 2021; Nagler et al., 2022; Tarricone et al., 2023; Hoppinen et al., 2023; Palomaki et al., 2025; Bonnell et al., 2024), and satellite (Belinska et al., 2022; Oveisgharan et al., 2024; Ruiz et al., 2024) datasets have shown promising initial results. The newly

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

launched NASA-ISRO SAR (NISAR) mission now provides the systematic, frequent L-band coverage these retrievals require, delivered as standard 80 m geocoded unwrapped interferogram (GUNW) products (Rosen et al., 2017). Whether that posting adequately samples the snow signal at L-band is still unclear.

The proposed ΔSWE retrieval from InSAR uses the fact that electromagnetic waves move slower through snow than through air (Guneriussen et al., 2001; Leinss et al., 2015). This means that if there is an increase in snow between two acquisitions, the two-way travel time of the electromagnetic waves between the sensor and the ground increases. This causes a phase change between the acquisitions, proportional to the change in snow. The direct equation requires an estimate of snow density, but several approximations have been proposed that directly relate ΔSWE to phase change ($\Delta\phi$) as a function of local incidence angle (θ) and the central SAR wavelength (λ) (Leinss et al., 2015; Oveisgharan et al., 2024). These approximations can be generalized to a single factor $\kappa(\theta,\lambda)$ relating the phase change to the SWE change:

$$\Delta\phi = \kappa(\theta,\lambda)\Delta\mathrm{SWE} \tag{1}$$

which is inverted to retrieve ΔSWE from an observed $\Delta\phi$.

While past retrieval studies have shown promising results, whether the ΔSWE signal is adequately sampled by current InSAR techniques has not been quantified directly from the snow field itself. Answering this requires quantifying how ΔSWE varies in space. To our knowledge, this has not been done because the necessary repeat snow surveys have only recently become available. Past work shows that depth and SWE vary rapidly over short distances (tens of meters) and more gradually over longer distances (hundreds of meters). This is framed as two regimes that each scale consistently with distance, with a decrease in the rate of variation between the short and long length scales (Deems et al., 2006; Trujillo et al., 2007; Clark et al., 2011; Mendoza et al., 2020; Miller et al., 2022). These studies either operate on snow depth directly (H_{snow}) or estimate snow density to derive SWE ($\mathrm{SWE} = \frac{\rho_{\mathrm{snow}}}{\rho_{\mathrm{water}}} H_{\mathrm{snow}}$). These prior works measure the snowpack present at one time, depth directly, or SWE via an estimated density, not its change between two acquisitions. This static SWE already varies considerably near the InSAR resolution scale, so the ΔSWE between two passes may vary as much and complicate sampling the snow signal.

InSAR struggles to sample a signal whose phase changes too quickly across the averaged window. The principle that a geophysical signal varying within the resolution cell degrades the interferogram is well established for ground deformation. Massonnet and Feigl (1998) defines the maximum detectable InSAR gradient as one fringe, one cycle (2π) of phase change, over a range change of one pixel, beyond which the fringes lose coherence, and states that when averaging “several adjacent pixels to form a larger pixel, these limits become more stringent”. The non-constant ΔSWE signal explored here is the snow analog of this deformation gradient, with one addition: a deformation gradient is a smooth, symmetric ramp that only decreases the observed coherence, whereas a skewed sub-window ΔSWE distribution could also bias the recovered multi-looked phase.

Two previous works have inferred sub-window ΔSWE phase variability from the interferograms themselves. Li and Sturm (2002) shows that a spatial variation in new-snow accumulation of $> 11\,\mathrm{cm}$ (assuming a density of $0.3\,\mathrm{g\,cm^{-3}}$) or redistribution into dunes of half this height across a window can cause decorrelation of C-band interferograms. They then show that mapped patterns of coherence loss from ERS-1 interferograms over the Alaskan arctic slope aligned with the timing and spatial

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

patterns of wind-driven snow redistribution. Yackel et al. (2026) decomposes the observed coherence over first-year sea ice into volumetric, thermal, baseline, and phase contributions. They note that their phase-coherence term arises from this sub-window ΔSWE variability within the multi-looking window.

Our work extends these studies in three ways. To understand how true ΔSWE variability affects the observed phase and coherence in InSAR ΔSWE retrievals, we use pairs of co-registered lidar snow depth surveys over the Tuolumne watershed to estimate spatially distributed ΔSWE variability and predict the resulting within-window coherence loss. We test these predictions against Tuolumne NISAR coherence maps from snowy and non-snowy periods across a range of multi-looking resolutions. We further show that, due to the wrapped complex nature of the observed phase, a skewed within-window ΔSWE distribution biases the recovered phase and thus the retrieved SWE.

2 Theoretical framework

The basis for multi-looking, or performing a complex average of phase over a spatial window, is that it provides a reasonable estimate of what would be measured if many independent observations of the phase could be obtained from the same place and time. This averaging reduces phase noise from random scatterers and improves the estimate of coherence (Zebker and Villasenor, 1992; Bamler and Hartl, 1998). However, this estimate is only valid under the assumption of stationarity: that every pixel in the window samples the same underlying distribution, with the same mean and variance of phase. This assumption is violated when the true ΔSWE varies across the multi-look window causing each pixel's phase PDF to be centered on a different value.

To show this effect, we first model the phase $\phi(x)$ of each single-look-complex (SLC) resolution cell x within a spatial multi-look window as the phase of the window-mean ΔSWE ($\bar{\phi}$) plus a per-cell deviation $\delta(x)$. To simplify notation, we use $\phi = \Delta\phi$ for the remainder of this section, and discussions about phase refer to the change in phase between two observations at the same location at different points in time.

$$\phi(x) = \underbrace{\bar{\phi}}_{\substack{\text{phase of}\\ \text{mean } \Delta\text{SWE}}} + \underbrace{\delta(x)}_{\substack{\text{per-cell deviation}\\ \text{from mean}}} \tag{2}$$

The observed complex value for the window, assuming unit amplitude in each cell, is:

$$\bar{z}_{\text{window}} = \frac{1}{N}\sum_{x=1}^{N} e^{i\phi(x)} \equiv M e^{i\bar{\phi}} \tag{3}$$

where $e^{i\bar{\phi}}$ carries the phase of the window-mean ΔSWE, and M is a complex factor whose magnitude reflects the spread of the per-cell deviations and whose argument captures the retrieval bias relative to the phase of the window-mean ΔSWE.

$$M = \frac{1}{N}\sum_{x=1}^{N} e^{i\delta(x)} \tag{4}$$

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Because δ is the deviation from the window mean, it has a zero mean by construction. Stationarity makes its variance spatially constant. Under the stationarity assumption, the SWE-induced noise-free phase is constant across the window, so δ is decorrelation noise alone. We assume our cells are distributed scatterers so the noise is symmetric about the mean and approximately Gaussian at usable coherence. The magnitude $|M|$ is set by the variance of these decorrelation-induced deviations and results in the observed decorrelation. The symmetry makes the expected M real-valued, so $\arg(M) = 0$ and the expected retrieval bias is zero.

However, when cells within the window are drawn from different distributions, M neither leaves the mean phase unbiased nor converges to the ensemble coherence. ΔSWE variation across the window centers the per-cell phase distributions on different values, decreasing the length of M and dropping the estimated coherence below the ensemble value. If the within-window ΔSWE distribution is asymmetric, M also acquires a nonzero argument, rotating the multi-looked phase from the expected phase of the mean ΔSWE (Fischer, 1993).

Figure 1 shows three idealized cases, with all other decorrelation sources removed, to isolate the effect of in-window ΔSWE variability. The three panels differ only in the within-window distribution of ΔSWE. Multi-looking assumes that the samples in the window share a single phase, so the constant case (left) is the only one that meets this condition. Its variance sets the length of the resultant $|M|$ as the phase values from each sample of ΔSWE vary (middle and right). Its skew sets its rotation $\arg(M)$ as the asymmetric circular distribution rotates the recovered window phase from the phase of the window-mean ΔSWE (right). Here, we treat all per-cell amplitudes as one, though real backscatter amplitudes vary across the window.

The snow-variability coherence loss ($|M|$) is independent of other decorrelation terms (temporal, geometric, SNR, scattering), so it multiplies (γ_{other}) to produce the observed decorrelation (γ_{obs}):

$$\gamma_{\mathrm{obs}} = \gamma_{\mathrm{other}}|M| \tag{5}$$

$\arg(M)$ is the bias of the observed window-mean phase relative to the phase expected from the window-mean ΔSWE, $\kappa\overline{\Delta\mathrm{SWE}}$. This bias applies even for a perfectly recovered window-mean phase:

$$\Delta\phi_{\mathrm{obs}} = \kappa\overline{\Delta\mathrm{SWE}} + \arg(M) \tag{6}$$

Both terms depend only on the within-window distribution of ΔSWE, so they can be computed from any sufficiently fine ΔSWE map. In the following sections, we compute them from lidar-differenced snow surveys and compare against NISAR observed coherences.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

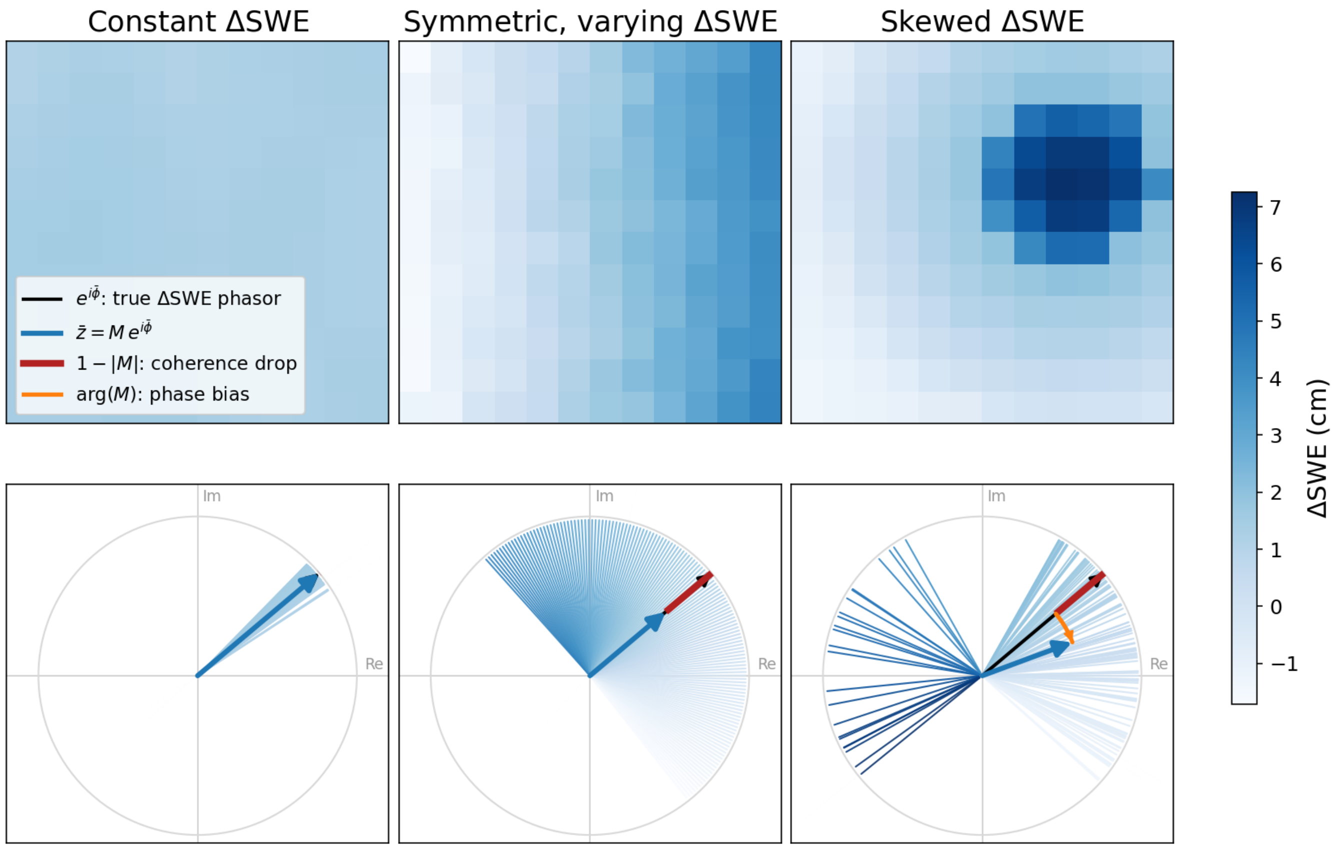

Figure 1. Sub-window ΔSWE variability and its effect on the multi-looked phasor (all other decorrelation removed). **Top:** ΔSWE in one window (cm), constant (**left**), varying symmetrically (**middle**), and skewed (**right**). **Bottom:** per-cell unit phasors ($\phi = \kappa\Delta\mathrm{SWE}$, colored by ΔSWE). *Black:* mean phasor $e^{i\bar{\phi}}$; *blue:* resultant $\bar{z} = Me^{i\bar{\phi}}$; *red:* coherence loss $1-|M|$; *orange arc:* retrieval bias $\arg(M)$.

3 Methods

3.1 Study Area and Time Frame

We target the Tuolumne River basin in the Sierra Nevada (upper Tuolumne and Lyell drainages, 1500–3970 m elevation, ∼95×56 km). We use four Airborne Snow Observatory (ASO) lidar pairs bracketing dominant winter accumulations of water years 2023–2026: March 2–16, 2023 (14 d), January 29–February 27, 2024 (29 d), February 8–25, 2025 (17 d), and January 31–February 27, 2026 (27 d). Pair durations were chosen to approximate the NISAR 12-day repeat baseline (one to two repeat cycles per pair) so that the lidar-derived ΔSWE is representative of what a single NISAR GUNW would measure.

For NISAR data, we use descending track 069, frame 042 over the same basin. Three repeat passes on December 7, 19, and 31, 2025 form two consecutive 12-day interferometric pairs. No ASO flight pair is coincident with these dates, so we

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

characterize the per-interval snow signal from daily SWE at four Tuolumne-basin CDEC pillows (Dana Meadows, Tuolumne Meadows, Slide Canyon, and Gin Flat). December 7–19 spans a slight ablation window (mean $-1\,\mathrm{cm}$ SWE change), while December 19–31 captures a large December storm (mean $+18\,\mathrm{cm}$ SWE change).

3.2 Calculation of ASO-ΔSWE

ASO delivers snow depth at $3\,\mathrm{m}$ and SWE at $50\,\mathrm{m}$, and all flights for this basin share a common UTM-11N $3\,\mathrm{m}$ raster, so pair differencing requires no resampling. For each pair, we form the $3\,\mathrm{m}$ Δdepth field $d_B - d_A$ and convert to ΔSWE via a basin-constant new snow density ρ_{event} calculated as the pair's mean of per-pixel ΔSWE/Δdepth ratios from the $50\,\mathrm{m}$ ASO products, taken over cells where $|\Delta\mathrm{depth}| \geq 10\,\mathrm{cm}$ and the implied density lies in $[50, 600]\,\mathrm{kg\,m^{-3}}$: 304, 312, 373, 330 $\mathrm{kg\,m^{-3}}$ for WY2023–2026. Cells where neither flight has $\geq$10 cm of snow at $50\,\mathrm{m}$ resolution are masked from the ΔSWE maps. These densities reflect that the measured snow changes are accumulation with significant settling over the 14–29 day temporal baselines. This approach allows us to take advantage of ASO data with the highest spatial resolution while still centering our analysis on SWE, not snow depth. We block-average the $3\,\mathrm{m}$ ΔSWE field in non-overlapping 3×3 blocks to a $9\,\mathrm{m}$ pixel size to suppress any lidar noise (Appendix A).

3.3 Sub-window ΔSWE Spectral Structure

To identify the spatial scales over which the ΔSWE field varies, we compute its radial power spectral density following the snow-depth fractal analyses of Trujillo et al. (2007). This spatial frequency analysis calculates how ΔSWE variance is distributed across spatial scales and tests whether we would expect significant sub-window variability. We test over spatial frequencies from 9 to $380\,\mathrm{m}$. We run the analysis on the native $3\,\mathrm{m}$ field rather than the $9\,\mathrm{m}$ working grid because the break could fall below the $18\,\mathrm{m}$ Nyquist wavelength of the $9\,\mathrm{m}$ grid and be unresolvable there. Working on the native $3\,\mathrm{m}$ field includes the per-pixel ASO depth noise in the spectrum, which we address in Appendix A. A single basin-wide spectrum incorporates many terrain types which blurs any scale breaks, so we work within smaller windows that span a narrower range of elevation, terrain roughness, and vegetation, the topographic and canopy properties shown to control snow-depth scaling. We tile the common snow mask (cells with valid ΔSWE in all four pairs) into non-overlapping $1.15\,\mathrm{km}$ windows (384 cells on a side) on the native $3\,\mathrm{m}$ grid and keep windows that are at least 95% snow-covered. This window size holds three full cycles of the longest fitted wavelength ($380\,\mathrm{m}$) while staying small enough to remain approximately locally homogeneous. Within the retained windows, any invalid cells (at most 5% of a window) are filled with the window mean. For each window and each ASO pair, we remove the window mean, apply a Hann taper, and form the two-dimensional periodogram as a function of spatial wavenumber k ($\mathrm{cycles\,m^{-1}}$) (Harris, 1978). We apply no higher-order detrend beyond this mean removal because subtracting a planar trend suppresses the longest-wavelength power and biases the large-scale slope (Deems et al., 2006).

The four ASO pairs span an order of magnitude in mean ΔSWE, so a direct average of their spectra would be dominated by the largest-accumulation pair. We instead amplitude-normalize each window's spectrum, subtracting its mean log power over the fitted power spectral density wavelength range of 9 to $380\,\mathrm{m}$ before averaging the four pairs' log spectra, which isolates the shared spectral shape. We fit each window's mean spectrum with a continuous two-segment power law, $P(k) \propto k^{-\beta}$, using

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

linear least squares in log–log space weighted by the number of frequencies sampled in each radial bin. The innermost radial bins average only a few frequencies each, so in an unweighted fit, these few noisy values would control the large-scale slope. The scale break is optimized over spatial wavelengths $\ell = 1/k$ from 12 to $250\,\mathrm{m}$, and the break wavelength ℓ_b is where the small-scale and large-scale regimes meet. We record a small-scale slope, a large-scale slope, and ℓ_b for each window. We also fit each pair separately to test whether the break is consistent across the four winters.

3.4 Calculation of M

For each ASO-derived ΔSWE field, we compute the complex sub-window factor M defined in Section 2. All processing steps in this section work on a common grid in UTM zone 11N (EPSG:32611) with the origin (upper-left corner) at x = 212352 m, y = 4234641 m.

The ΔSWE map for each ASO pair is converted to $\Delta\phi$ using NISAR's central frequency (1.26 GHz) and a 40° local incidence angle, giving a κ of $52.7\,\mathrm{rad\,m^{-1}}$ of ΔSWE, then wrapped onto $(-\pi, \pi)$. The 40° local incidence angle is chosen to represent a moderate incidence angle across NISAR's typical range, but we note that κ is sensitive to this choice (Leinss et al., 2015).

We then form the complex variability factor M from non-overlapping 9×9-pixel ($81\,\mathrm{m}$) windows of $\Delta\phi$, approximately matching the $80\,\mathrm{m}$ GUNW posting:

$$M = \frac{1}{N}\sum_{n=1}^{N} e^{i(\Delta\phi_n - \overline{\Delta\phi})}, \tag{7}$$

where $\Delta\phi_n = \kappa\Delta\mathrm{SWE}_n$, $\overline{\Delta\phi} = \kappa\overline{\Delta\mathrm{SWE}}$, $\overline{\Delta\mathrm{SWE}}$ is the mean ΔSWE in the window, and N is the number of snow-covered pixels in the window. This is the theoretical Eq. (4) reformulated to incorporate the mean ASO ΔSWE within the window. We retain $|M|$ and $\arg(M)$ for windows with $N \geq 41$ (at least half of the pixels are snow-covered).

We repeat the M calculation at two additional wavelengths: C-band (Sentinel-1, $\lambda = 5.6\,\mathrm{cm}$) and P-band ($\lambda = 81\,\mathrm{cm}$). The pipeline is unchanged: the same $9\,\mathrm{m}$ pre-averaged ΔSWE field and $81\,\mathrm{m}$ windows are used, and only $\kappa(\lambda,\theta)$ Eq. (1) is reevaluated for each wavelength, giving $\kappa = 224.1$ and $15.5\,\mathrm{rad\,m^{-1}}$ for C- and P-band, respectively. This produces $|M|$ and $\arg(M)$ maps at $81\,\mathrm{m}$ for each band and each ASO pair.

Figure 2 shows the wrapped phase response across 0 to $50\,\mathrm{cm}$ of ΔSWE at C-, L-, and P-band, with the basin-mean ± standard deviation of $50\,\mathrm{m}$ ASO ΔSWE for each pair marked to indicate where each pair sits relative to the wrap interval at each band.

3.5 Sensitivity of M to Window Size

To trace how the sub-window decorrelation changes with footprint size, we recompute the window factor M over a range of multi-look window sizes on the $3\,\mathrm{m}$ ΔSWE field. The window construction follows Section 3.4. Here the window size is swept from 9 to $81\,\mathrm{m}$, and unlike Section 3.4 the field is not pre-averaged to $9\,\mathrm{m}$ so the smaller windows are resolved. The

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

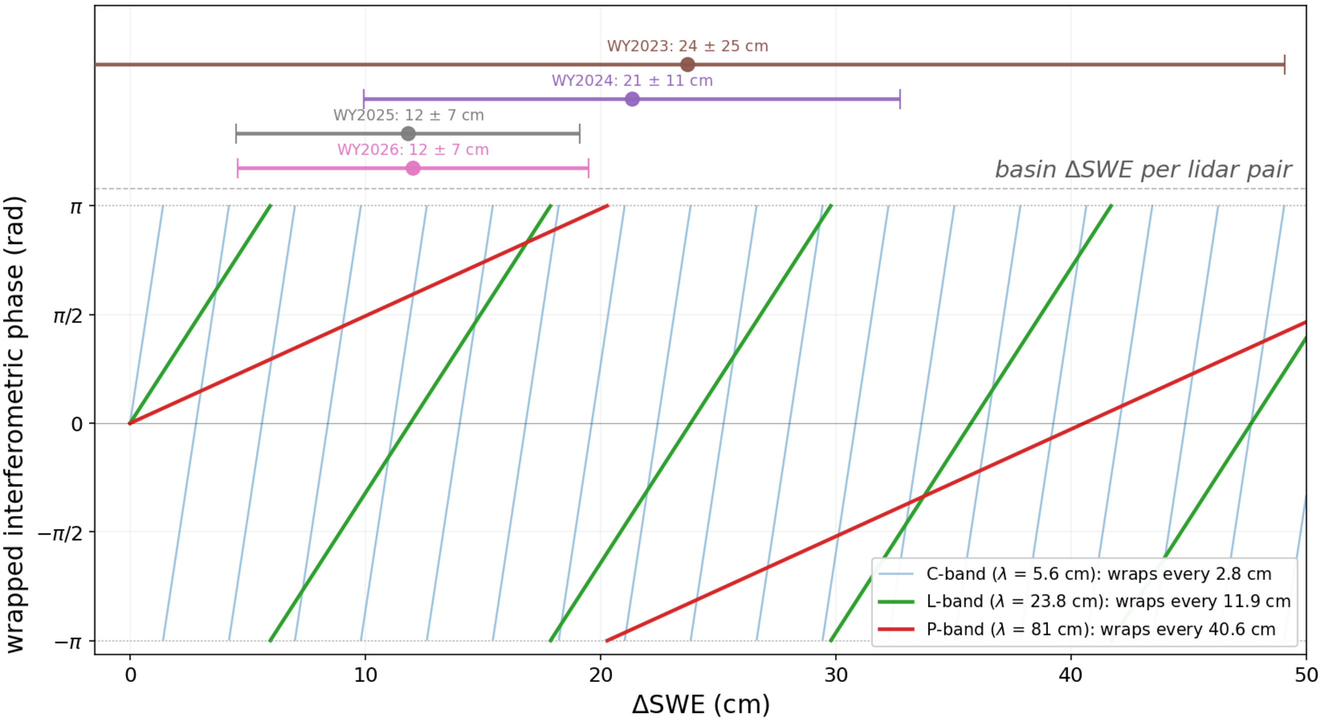

Figure 2. Wrapped interferometric phase response to ΔSWE accumulation 0 to 50 cm at C-band ($\lambda = 5.6$ cm, blue), L-band ($\lambda = 23.8$ cm, green), and P-band ($\lambda = 81$ cm, red), all at $\theta = 40^\circ$ local incidence angle. Each band's phase increases linearly with ΔSWE at the phase rate $\kappa(\lambda, \theta)$ Eq. (1) until it wraps back through $-\pi$. Error bars at the top mark the basin-mean ΔSWE ± standard deviation for each of the four Tuolumne ASO pairs (50 m basin statistics).

four ASO pairs are restricted to a common snow mask, the per-pixel intersection of their four valid-ΔSWE footprints on the shared 3 m grid, so each is evaluated over an identical pixel set. We pool the windows across the four pairs at each size and report the median and interquartile range of $|M|$ and of the L-band SWE-equivalent bias $\arg(M)/\kappa$. The bias is reported only for windows with $|M| > 0.3$, since below it $\arg(M)$ is near-uniform over $(-\pi, \pi]$ and carries no recoverable signal. Working on the native 3 m field folds the per-pixel ASO depth noise into $|M|$, which we address in Appendix A.

3.6 NISAR Multi-resolution Coherence Calculations

The coherence loss due to sub-window ΔSWE variability can, in principle, be estimated directly from an interferogram by comparing coherence calculated over different window sizes. The underlying ensemble coherence of a pixel does not depend on the estimation window, but the estimated coherence is affected by two competing, window-size-dependent effects: a positive bias from estimating coherence magnitude with a finite number of looks, which grows as the window shrinks (Bamler and Hartl, 1998), and a negative effect from the sub-window variability of the ΔSWE signal, assumed in this study to be the dominant

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

source of non-stationarity, which grows as the window expands to span more ΔSWE variability. Because the finite-look bias is analytically known, we can correct each window size and use the remaining coherence difference between window sizes to estimate the potential sub-window ΔSWE contribution.

We use the coherence magnitude delivered in each GUNW: the $20\,\mathrm{m}$ layer from the wrapped-interferogram group and the $80\,\mathrm{m}$ layer from the unwrapped-interferogram group. The expected sample coherence $\hat{\gamma}$ for true coherence γ and looks L is given by Eq. (62) of Bamler and Hartl (1998):

$$\mathbb{E}\{\hat{\gamma} \mid \gamma, L\} = \frac{\Gamma(L)\Gamma\left(\frac{3}{2}\right)}{\Gamma\left(L+\frac{1}{2}\right)}\, {}_3F_2\left(\tfrac{3}{2}, L, L; L+\tfrac{1}{2}, 1; |\gamma|^2\right)\left(1-|\gamma|^2\right)^L, \tag{8}$$

where Γ is the gamma function and ${}_3F_2$ is the generalized hypergeometric function. Eq. (8) also assumes circular-Gaussian speckle from a distributed target, which is a reasonable approximation over snow and mixed terrain, but it is a limitation of using this equation. NISAR multi-looks the full-resolution interferogram in radar coordinates independently for each layer (5×6 for the wrapped layer, 13×16 for the unwrapped layer) and then geocodes each layer onto its delivered map grid. The geocoding resamples the multi-looked pixels, so neighboring samples on the 20 and $80\,\mathrm{m}$ grids are partially correlated, and the nominal look counts ($N = 30$ at $20\,\mathrm{m}$, 208 at $80\,\mathrm{m}$) overstate the number of independent samples. We therefore use an effective look count $L = N/f$, with f an oversampling factor calibrated over open water on Don Pedro Reservoir near the study site, where the true coherence is assumed zero at both resolutions. Water pixels are identified as ESA WorldCover permanent-water (class 80) cells with NISAR $80\,\mathrm{m}$ coherence below 0.20, the latter cut removing coherent shoreline pixels. This leaves 13,189 open-water samples over the two pairs. Requiring the bias-corrected $20\,\mathrm{m}$ and $80\,\mathrm{m}$ water coherences to agree gives $f \approx 1.29$, $L_{20} \approx 23$, and $L_{80} \approx 162$.

We then de-bias each resolution by inverting Eq. (8) for the true coherence,

$$\tilde{\gamma}_s = \mathbb{E}^{-1}\{\hat{\gamma}_s \mid L_s\}, \quad s \in \{20, 80\}, \tag{9}$$

Here $\hat{\gamma}_{20}$ is the $20\,\mathrm{m}$ coherence block-averaged onto the $80\,\mathrm{m}$ grid before de-biasing, matching the footprint on which f was calibrated; de-biasing the native $20\,\mathrm{m}$ pixels instead would apply L_{20} to samples with fewer effective looks and leaves a small spurious positive residual over zero-coherence water. We then take the difference of the corrected coherences

$$\Delta\gamma = \tilde{\gamma}_{20} - \tilde{\gamma}_{80} \tag{10}$$

as our NISAR estimate of the sub-window-variability coherence loss. With $\gamma_{\mathrm{obs}} = |M|\gamma_{\mathrm{other}}$ and $\tilde{\gamma}_{20} \approx \gamma_{\mathrm{other}}$, this difference is $\Delta\gamma \approx \gamma_{\mathrm{other}}(1-|M|)$, reducing to $1-|M|$ only where the other decorrelation effects are minimal ($\gamma_{\mathrm{other}} \approx 1$). Because the $20\,\mathrm{m}$ reference window already carries some sub-window loss of its own ($|M|_{20} < 1$), the difference $\Delta\gamma$ captures only the additional loss incurred between 20 and $80\,\mathrm{m}$, not the full loss relative to a variability-free window. It is therefore a lower bound on the true sub-window coherence loss.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

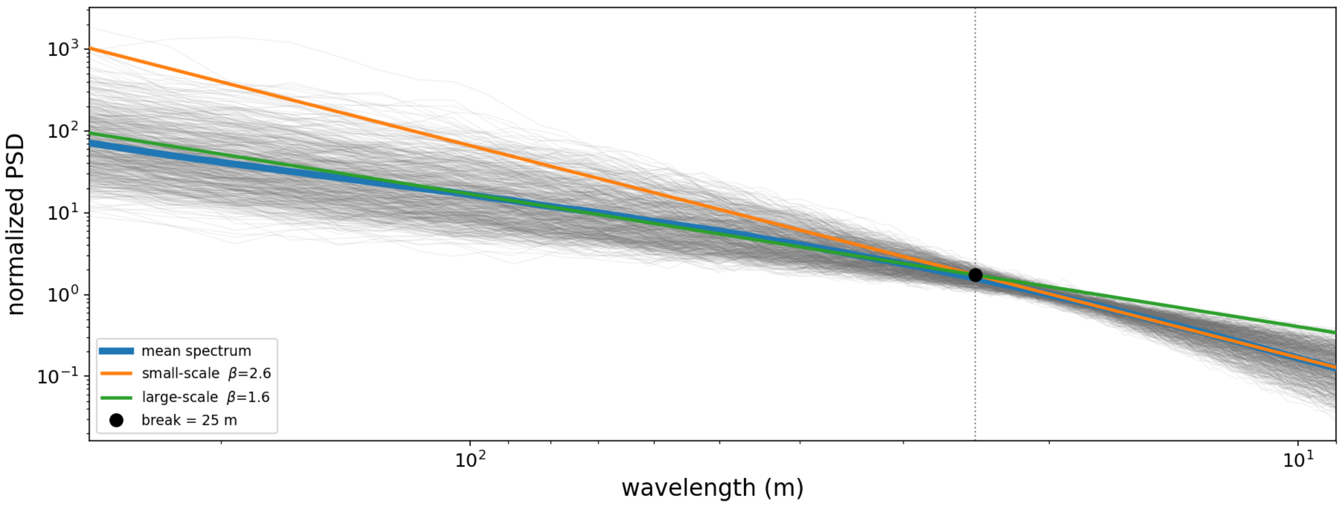

Figure 3. Radial power spectral density of the native 3 m ΔSWE field, for each 1.15 km window, amplitude-normalized and averaged over the four lidar pairs so only the spectral shape is compared. Grey lines are the individual windows, blue is their mean, and the orange and green lines are the fitted small-scale and large-scale segments, $P(k) \propto k^{-\beta}$ with $\beta \approx 2.6$ and 1.5. The segments meet at the break (black point) near 25 m.

4 Results

4.1 ΔSWE Power Spectral Densities

Resolved per window on the native 3 m grid, the ΔSWE spatial power spectrum is a two-regime power law with a break at a median wavelength of $\ell_b = 25$ m (IQR 21–30 m across 572 windows; Figure 3). The spectrum is steep below the break (median slope $\beta = 2.6$) and shallow above it ($\beta = 1.5$): the steep small-scale slope means the ΔSWE field is spatially correlated and varies smoothly over separations finer than the $\sim$25 m break, while the shallow large-scale slope means it is rougher and decorrelates over larger separations. Equivalently, the field carries less variance at wavelengths shorter than the break and increasingly more at longer wavelengths. Fitting each ASO pair separately gives ℓ_b between 23 and 27 m, and 87% of the window breaks fall within the 15–40 m range reported for cumulative snow depth (Deems et al., 2006).

4.2 Lidar Derived Sub-window ΔSWE Variability Effects

All lidar ΔSWE statistics in this section are computed over the non-overlapping 81 m windows defined in Section 3.4. Pooled over the four pairs (1.34 million windows), the window-mean ΔSWE has a median of 12.6 cm (IQR 4.7–20.4 cm), a within-window median standard deviation of 3.4 cm (IQR 2.1–6.3 cm) and a slight positive within-window skewness (median $+0.07$, mean $+0.14$). The within-window standard deviation captures the sub-window ΔSWE variability that drives $|M|$, and the within-window skew drives $\arg(M)$.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

The four pairs show a consistent, terrain-organized spatial pattern. The coherence loss $1-|M|$ and the bias magnitude $|\arg(M)|$ are both largest in the deep, spatially complex alpine terrain and smallest in the smoother, lower-elevation terrain (Figure 4), and both patterns repeat across pairs, with each pair matching the four-pair mean with $r^2 = 0.75$–0.82 for $|M|$ and 0.59–0.63 for $|\arg(M)|$. The sign of the bias is year-specific (signed $\arg(M)$ pairwise $r^2 < 0.01$). Its spatial intensity tracks the terrain while its direction varies between events.

Although the basin-median ΔSWE ranges over an order of magnitude across the four pairs (0.7–23 cm), the coherence drop and bias stay within narrow bands. Median $|M|$ runs from 0.25 (WY2024) to 0.47 (WY2023; pooled IQR 0.15–0.58), and the median retrieval bias is near zero ($\leq 2°$) with heavy tails (pooled IQR $\pm16°$, widening to $-17°/+30°$ in the largest event, WY2024). The three large mid-winter storms (WY2024–2026) drive median $|M|$ to 0.25–0.31, falling to $\sim$0.12–0.16 in the steep high-elevation terrain, whereas the pair containing significant melt (WY2023) has the smoothest ΔSWE field ($\sigma_{\Delta\mathrm{SWE}}$ median 2.8 cm), stays closest to stationarity, and shows the highest $|M|$ (0.47) and the least bias (median $\approx 0°$). Within WY2023, splitting the windows by the sign of the window-mean ΔSWE gives a median $|M|$ of 0.55 for melt windows against 0.32 for the accumulation windows, a gap that persists at matched $|\Delta\mathrm{SWE}|$ magnitude.

4.3 Coherence Loss and Bias Across Window Sizes

The L-band coherence factor falls steeply as the multi-look window grows and then plateaus (Figure 5, left). The pooled median $|M|$ drops from 0.54 at a 9 m window to 0.33 at 20 m and 0.18 at 81 m, with most of the loss incurred by about 25 m. The interquartile range is wide at every size and stretches further downward as the window grows, spanning 0.07 to 0.38 at 81 m. The two NISAR GUNW postings bracket a large part of the curve. The 20 m wrapped-interferogram footprint retains a median $|M|$ of 0.33 compared to 0.18 at the 80 m footprint, almost double the coherence factor for the same scene.

The L-band retrieval bias stays small across the whole range (Figure 5, right). In usable windows, the median SWE-equivalent bias never exceeds 0.3 mm, and the interquartile range stays within about ±2 mm. The band is slightly asymmetric, with a longer positive tail (matching the within-window ΔSWE skew) but the magnitude is negligible compared to the coherence loss. The bias does not grow significantly at any footprint between 9 and 81 m.

4.4 NISAR Multi-resolution Coherence

The corrected $\gamma_{20} - \gamma_{80}$ is near zero for the near zero-SWE-change December 7–19 pair (median -0.004, 45% of windows positive) and positive for the December 19–31 storm pair (median $+0.030$, 71% positive) over the snow-covered footprint (Figure 6). Much of the raw 20–80 m gap is finite-look estimation bias. For the December 19–31 pair, the raw gap $+0.072$ (median) is reduced to $+0.030$ by the bias correction, while for the snow-free pair, the corrected gap is close to zero. However, the largest corrected residuals cluster at higher elevations, where more ΔSWE variability is expected from the lidar. For the storm pair, the corrected gap stays small in the median (≤ 0.04 across elevations), but its upper tail grows with elevation, with the strongest windows (90th percentile) rising to $\sim$0.13 in the highest, most exposed terrain, where the 80 m coherence is also lowest (median 0.17).

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

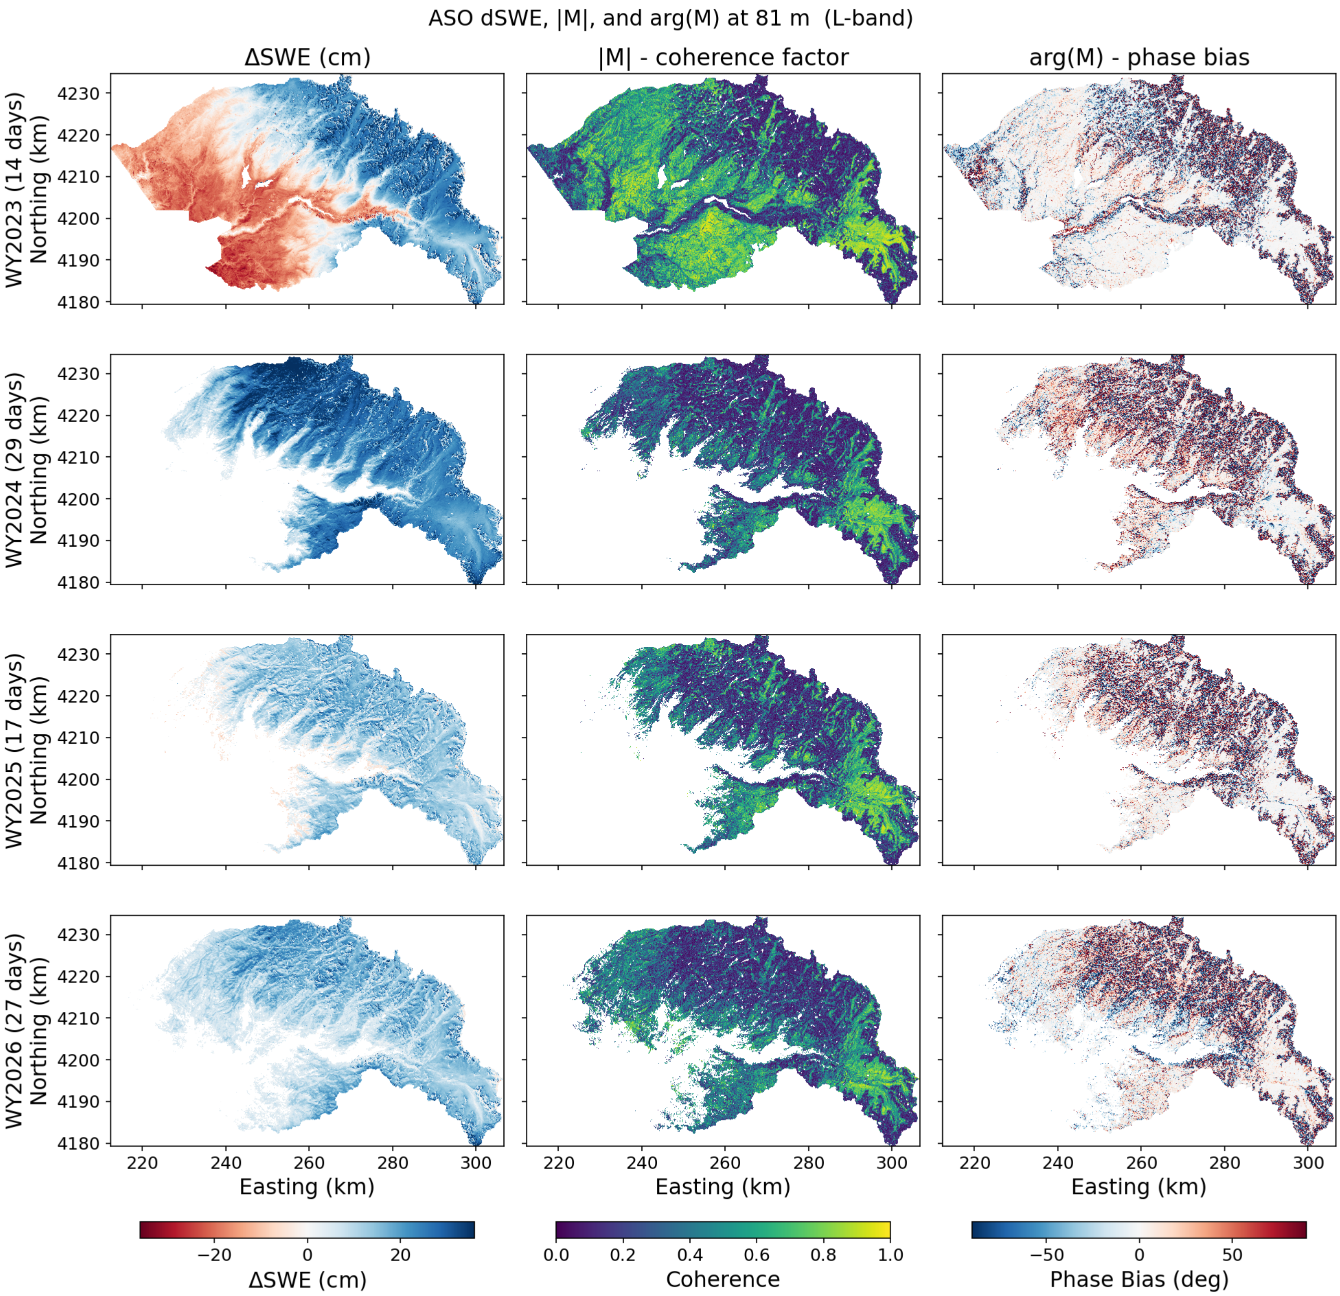

Figure 4. Lidar-derived ΔSWE, sub-window coherence factor $|M|$, and retrieval bias $\arg(M)$ at 81 m for the four Tuolumne lidar pairs (rows, labeled by water year and temporal baseline), evaluated at L-band, 40° local incidence angle. Columns: the 81 m block-averaged ΔSWE field (**left**), $|M|$ (**center**), and $\arg(M)$ (**right**).

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

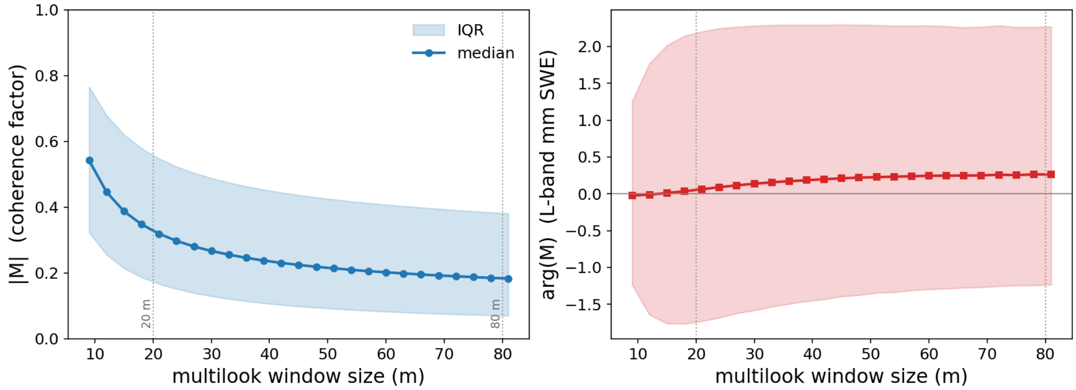

Figure 5. Sub-window coherence factor $|M|$ (**left**) and retrieval bias $\arg(M)$ (**right**) versus multi-look window size, L-band, pooled over the four lidar pairs (median and shaded interquartile range). Computed on the native 3 m ΔSWE field. The right panel is restricted to usable windows ($|M| > 0.3$) and converted to an L-band SWE-equivalent through $\arg(M)/\kappa$. Dotted lines mark the 20 and 80 m NISAR GUNW postings.

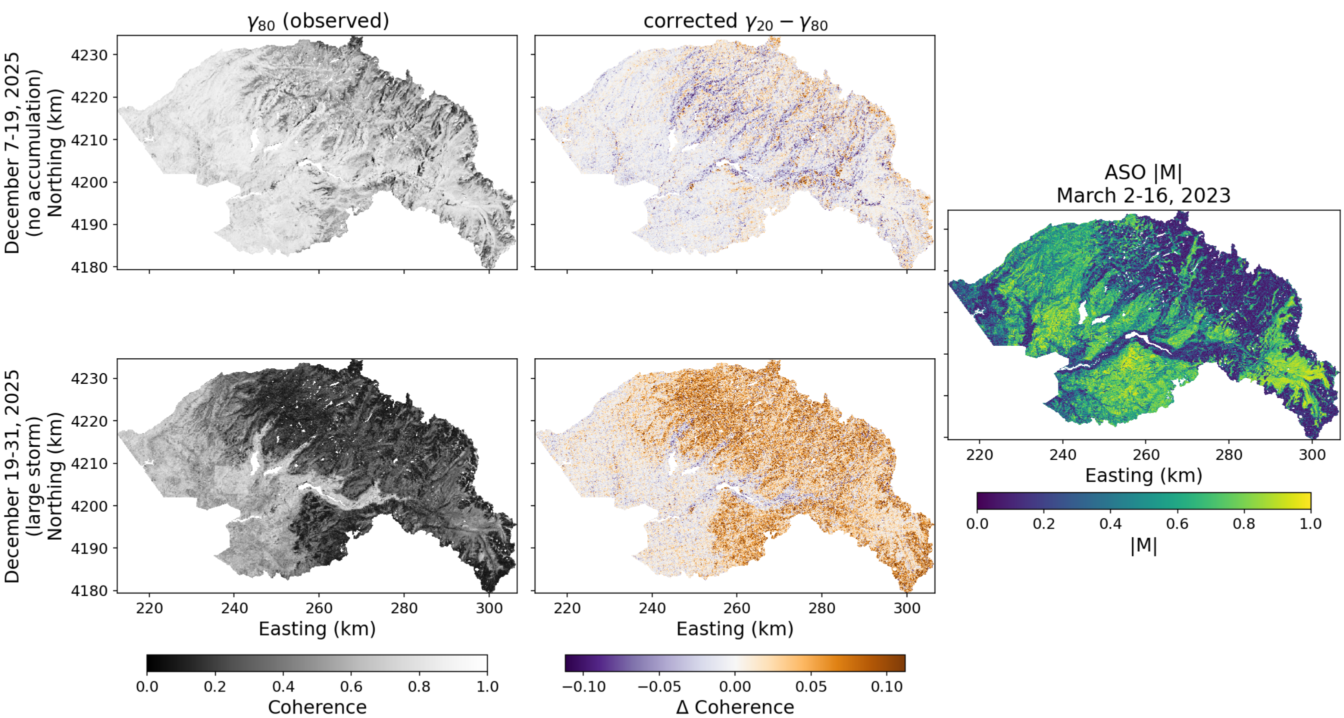

Figure 6. NISAR multi-resolution coherence for the low-snow P12 (December 7–19, 2025) and large-storm P23 (December 19–31, 2025) pairs: observed 80 m coherence γ_{80} (**left**), the finite-look-corrected difference $\gamma_{20} - \gamma_{80}$ (**center**), and the lidar-predicted sub-window coherence factor $|M|$ (**right**, WY2023), all on the common 81 m grid.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

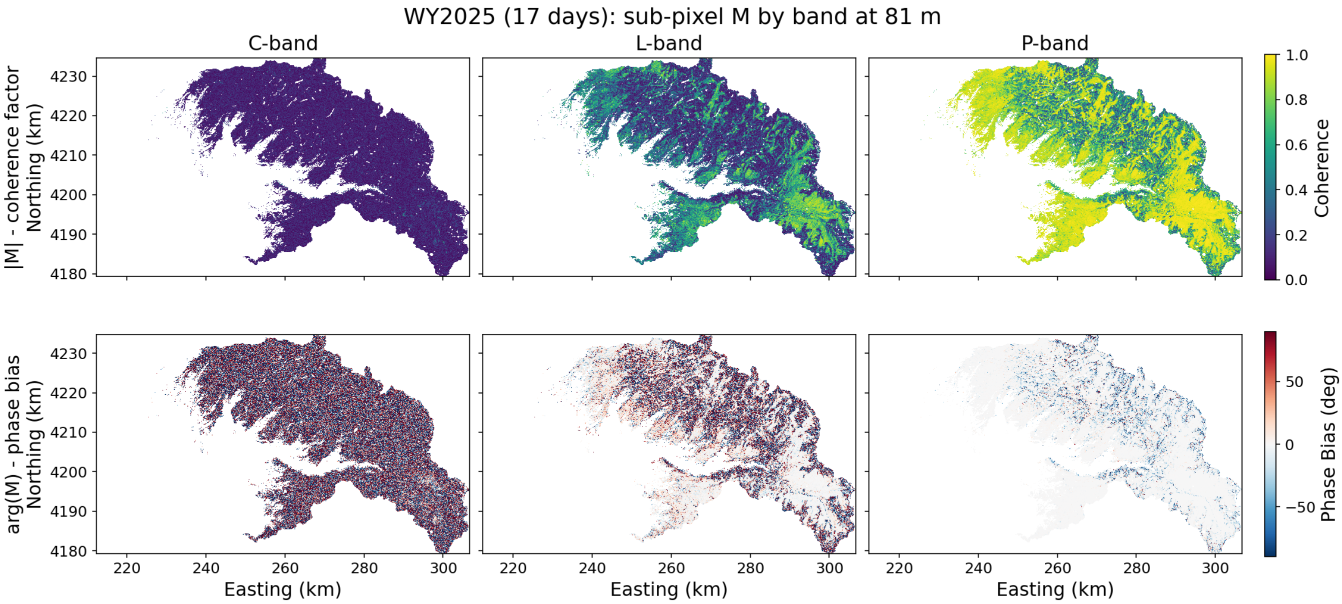

Figure 7. Sub-window coherence factor $|M|$ (**top**) and retrieval bias $\arg(M)$ (**bottom**) at 81 m for the WY2025 lidar pair (February 8–25, 2025) at C-, L-, and P-band.

4.5 Wavelength Effect on M

The same sub-window ΔSWE field produces very different effects across bands, because shorter wavelengths turn the same ΔSWE spread into a larger phase spread and more wrapping (Figure 7). At C-band ($\kappa = 224.1\,\mathrm{rad\,m^{-1}}$) the within-window phase spread is large enough that coherence collapses, with median $|M| = 0.10$ (IQR 0.06–0.14) and 53% of windows below 0.1. So little coherent signal survives that the retrieval bias is undefined, with $\arg(M)$ distributed nearly uniformly over $\pm 90°$, IQR $-86°$ to $+88°$. At P-band ($\kappa = 15.5\,\mathrm{rad\,m^{-1}}$) coherence is largely preserved (median $|M| = 0.87$, IQR 0.65–0.94; <1% of windows below 0.1) and the bias is small and well-defined ($\arg(M)$ IQR $-1.6°$ to $+0.3°$). L-band is intermediate, with median $|M| = 0.31$ (IQR 0.15–0.55), 14% of windows below 0.1, with a small but resolvable bias ($\arg(M)$ IQR $-15°$ to $+19°$). The same ordering holds when data is pooled over all four lidar pairs (median $|M|$ = 0.10 / 0.33 / 0.88 for C/L/P), so that it reflects the wavelength scaling rather than any one event.

5 Discussion

Using repeat lidar surveys and NISAR interferometry, we have shown that within-window ΔSWE variability reduces coherence and ΔSWE skew biases the retrieved ΔSWE from the true mean ΔSWE, both of which degrade InSAR SWE retrievals. Here we consider what these results mean in practice: the length scales at which the ΔSWE field is organized, the resulting effects at NISAR's standard postings, and how the effect scales from C- to P-band.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

5.1 ΔSWE Structure and the Processing Window Size

The $\Delta\mathrm{SWE}$ spectral break at $\ell_b \approx 25\,\mathrm{m}$ (Figure 3) falls within the 15–40 m variogram range scale breaks reported for cumulative snow depth and matches the wind- and vegetation-controlled short-scale regime of Trujillo et al. (2007). The change in SWE between two flights is therefore organized at the same length scales as the snowpack itself, so differencing two snow surfaces preserves the fine-scale structure. This holds in all four winters, so the variability driving within-window phase spread is a persistent feature of accumulation, not an artifact of a single season.

This break sets the window sizes over which the sub-window coherence loss develops. Because the $\Delta\mathrm{SWE}$ field carries increasing variance toward larger scales (Section 4.1), the within-window phase spread keeps growing as the footprint expands past ℓ_b, and the coherence factor falls steadily across the GUNW range. The window-size sweep shows this in $|M|$ directly (Figure 5). The empirical curve is steepest below about 25 m and flattens beyond, with the slope change near the spectral break — a flattening driven by $|M|$ approaching its floor once the within-window phase spread exceeds about one radian. Even in this regime the recoverable coherence is substantial: moving from the 80 m GUNW posting to a 20 m footprint nearly doubles the median $|M|$, from 0.18 to 0.33. The coherence recovered per step keeps growing as the window shrinks toward ℓ_b, so even modest reductions below the GUNW default return usable coherence over the spatially variable snow that decorrelates the 80 m product.

The bias result is promising for SWE retrievals. Across the full 9 to 81 m range, the SWE-equivalent bias in usable windows stays at the millimeter level (Figure 5, right), far below the coherence loss and below the retrieval uncertainty from other sources. The footprint is therefore a coherence tradeoff and not a bias tradeoff. A larger window loses coherence without incurring a meaningful retrieval bias in the pixels that survive a coherence threshold, and a smaller window recovers coherence without introducing bias. The footprint should be chosen to keep $|M|$ usable over the target snowpack, and the residual skew-driven retrieval bias can be neglected in any window that passes a coherence mask.

Shrinking the window has a limit. Smaller windows average fewer pixels, so the random phase noise that multi-looking suppresses, whose variance falls as $1/N$ with the number of looks (Bamler and Hartl, 1998), rises as the footprint shrinks. The best footprint balances the recovered sub-window coherence against this estimation noise and is set by the local $\Delta\mathrm{SWE}$ correlation length, the $\ell_b \approx 25\,\mathrm{m}$ break, rather than a fixed look count. For these Tuolumne events that balance sits well below the 80 m GUNW posting, and large accumulation events such as the December 2025 storm need the finest practical footprint to keep the sub-window term in check. The most wind-exposed terrain, where short-scale variance is largest, will be challenging to recover at any footprint.

5.2 Snow-based Multi-look Coherence Bias

Coherence measured over seasonal snow has historically been interpreted through a decorrelation budget of temporal, geometric, SNR, and volume terms, with residual loss most often assigned to temporal decorrelation of the snow volume (Ruiz et al., 2021; Oveisgharan et al., 2024). The presented sub-window $\Delta\mathrm{SWE}$ effect is itself a change-driven decorrelation, since it arises from the snow that accumulates between passes, but it differs from the usual temporal term: rather than the stochastic motion

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

of sub-resolution scatterers within a cell, it is the deterministic variation of the resolved ΔSWE field across cells within the multi-look window. Our lidar-derived $|M|$ maps show that a substantial share of the L-band coherence drop over an accumulation interval may be this deterministic sub-window component rather than stochastic temporal decorrelation. Because it is set by the resolved snow structure, it can be predicted from lidar and reduced by processing at a finer footprint, neither of which is true of the stochastic temporal term to which it is usually attributed. Across the four Tuolumne lidar pairs, the sub-window factor lowers the basin-median $|M|$ to 0.25–0.47 (Figure 4) at an $81\,\mathrm{m}$ window, a coherence loss of 0.5 to 0.75 carried by within-window ΔSWE structure before any temporal or thermal term is applied.

Turning to the NISAR data, both the raw coherence and the multi-resolution residual show the temporal and spatial patterns that these lidar ΔSWE results predict (Figure 6). The raw $80\,\mathrm{m}$ coherence is lower for the December 19–31 storm pair than for the near-snow-free December 7–19 pair, and within the storm pair it is lowest over the steep alpine terrain where sub-window ΔSWE variability is greatest (median $\gamma_{80} \approx 0.17$ at the highest elevations versus ≈ 0.27 at the lowest). Raw coherence alone cannot separate the sub-window loss from other decorrelation, but the corrected $\gamma_{20} - \gamma_{80}$ residual repeats both patterns where only the sub-window term should appear. It is near zero for the no-accumulation pair (median -0.004), positive for the storm pair (median $+0.030$ over snow), and largest in the highest, most complex terrain (90th percentile $\Delta\gamma$ 0.13). The observed residual is smaller than the lidar-predicted gap ($|M|_{20} - |M|_{80} \approx 0.15$, Section 4.3) because γ_{other} multiplies the difference and much of the storm-pair scene is decorrelated enough that the gap is lost to phase noise. Coherence over snow may therefore be misattributed to temporal decorrelation when part of it is a within-window phase-variability signal.

Even at a finer footprint, the retrieval is more tractable where the ΔSWE field is smooth. The $|M|$ drops to 0.12–0.16 in the steep, wind-exposed alpine drainages but stays near one over the lower-elevation, gently varying terrain (Figure 4, center column). This terrain control holds across events. The basin-median ΔSWE for our four pairs spans an order of magnitude, while the median $|M|$ stays within a narrow band, so the decorrelation and bias are set by terrain and snowpack structure rather than by event size. Since some multi-looking will be necessary to improve phase estimates, retrievals will work best when targeted at regions and snowpacks whose ΔSWE varies gently relative to the multi-look window. The most variable alpine terrain may be unrecoverable at any current spaceborne footprint.

Melt periods may be more favorable than accumulation for the sub-window term considered here. Once the snowpack has ripened and refrozen, melt-driven ΔSWE varies smoothly with elevation and aspect rather than in wind-built drifts, which keeps the within-window variance and skew low. The WY2023 melt-versus-accumulation split (Section 4.2) shows this directly, with the gap persisting at matched $|\Delta\mathrm{SWE}|$ magnitude and driven by a lower within-window spread, so melt removes SWE more uniformly within the window than wind deposits it. This smoother sub-window signal only helps if the wet-snow dielectric and temporal decorrelation terms in γ_{other} do not collapse the coherence on their own. L-band ablation retrievals have been initially demonstrated for morning refrozen acquisitions (Tarricone et al., 2023), but they will be sensitive to the estimated liquid water content, as this impacts the radar wave speed and, therefore, the phase.

This mechanism also explains why wind is the single largest control on snow coherence in Ruiz et al. (2022). Wind is the main agent of fine-scale snow redistribution, building dunes, scour, and drifts at length scales well below typical multi-look window sizes (Trujillo et al., 2007). This redistribution raises the within-window ΔSWE variance and skew, which shortens

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

and rotates M, so a windy interval produces the asymmetric, sub-window ΔSWE distribution that Figure 1 identifies as the source of both coherence loss and retrieval bias. This is the same mechanism Li and Sturm (2002) identified at C-band, where mapped ERS-1 coherence loss tracked the spatial and temporal patterns of wind-driven redistribution, and that Yackel et al. (2026) found for C-band over first-year sea ice by decomposing the coherence budget. Our results confirm both findings from lidar rather than from the interferograms themselves. At C-band, the wavelength examined by Li and Sturm (2002) and Yackel et al. (2026), $|M|$ collapses to a median of 0.10, and the lowest $|M|$ in every pair falls in the wind-exposed alpine terrain. The dominance of wind as a predictor in Ruiz et al. (2022) is what we would expect if sub-window ΔSWE variability, rather than bulk temporal changes affecting temporal coherence, governs the observed coherence.

5.3 Skew-Driven Retrieval Bias Considerations

The sign of the retrieval bias is set by the skew of the within-window ΔSWE distribution (Figure 1). Before the phase wraps, a positively skewed distribution produces a negative retrieval bias, which rotates the multi-looked phase below the true window-mean phase and results in a lower retrieved ΔSWE than the true mean in those windows. A physically likely source of this positive skew is wind loading: a few wind-loaded pixels carry much more new snow than the rest of the window and sit as positive outliers in an otherwise smooth ΔSWE field, lengthening the right tail of the distribution.

The lidar-derived $\arg(M)$ maps show the same pattern (Figure 4, right column). The retrieval bias is a real effect, but it is a small one. It is near zero across most of the basin and rises only in the steep, wind-exposed alpine terrain, where the within-window variability is largest and the coherence is lowest. This ΔSWE bias is therefore largest exactly where the phase is least reliable. In SWE terms, it stays near $\pm 5\,\mathrm{mm}$ (converting phase to L-band ΔSWE bias) across the basin and reaches about $1.5\,\mathrm{cm}$ only in the most wind-affected windows of the largest event. Because the bias peaks with the coherence loss, pixels that pass a coherence threshold carry little bias, so it mostly affects pixels that may already be unusable. This coupling only holds at shorter wavelengths. At longer wavelengths, the coherence is preserved while the bias remains, which we turn to below.

5.4 C-, L-, and P-band Analysis

P-band is appealing for SWE retrieval. The smaller κ keeps coherence high over the large accumulation events that decorrelate L-band. However, at P-band, the retrieval bias is still present and consistently negative ($\arg(M)$ IQR $-1.6°$ to $+0.3°$), and it is the one band where the bias is not correlated with almost complete coherence drops. The negative sign is the one predicted for a positively skewed ΔSWE distribution, and it appears here because the longer wavelength keeps the within-window phase spread small enough that the multi-looked phase does not wrap (Figure 2). At C- and L-band, the same bias is present but masked since the larger phase spread wraps and pulls the median $\arg(M)$ back toward zero. A longer-wavelength sensor has a narrower window phase distribution, but where the distribution is skewed in the steep alpine terrain, it carries a systematic negative bias without the matching total coherence loss to flag it.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Shorter wavelengths run into the opposite problem. At C-band, the within-window phase variability is large enough that the coherence collapses, with a pooled median $|M|$ of 0.10 and more than half of the WY2025 windows below 0.1 (Figure 7). A C-band sensor recovers little usable ΔSWE over snow at this multi-look window size ($81\,\mathrm{m}$).

This collapse is dominated by the real sub-window ΔSWE structure, but because independent decorrelation sources multiply, some component is likely lidar depth noise. That noise component is accentuated at C-band, where the larger κ turns the same residual depth error into a much larger phase spread than at L-band (further discussed in Appendix A).

L-band falls between the two. The pooled median $|M|$ of 0.33 leaves enough coherence to recover phase over the smoother, lower-elevation terrain but not over the steep alpine drainages, where $|M|$ drops toward the C-band regime. The bias is likewise intermediate and only partly wrapped ($\arg(M)$ IQR $-15°$ to $+19°$). L-band, therefore, escapes the coherence collapse that overwhelms C-band, and, unlike P-band, its retrieval bias is paired with a coincident coherence loss that will cause those pixels with bias to be removed anyway.

5.5 Mitigation and Future Work

The multi-resolution comparison used here is itself a potential mitigation. The coherence loss due to sub-window ΔSWE variability grows with window size while the other decorrelation terms do not, so comparing corrected coherence across footprints separates the sub-window contribution from the rest, and the $\gamma_{20} - \gamma_{80}$ residual is a direct, if noisy, estimate of $1 - |M|$ that could flag windows where sub-window variability dominates. The limit is that even the $20\,\mathrm{m}$ window sits well inside the variability regime for large storms, so two GUNW resolutions sample only part of the $|M|$ curve (Figure 5), and pixels that are fully decorrelated at both footprints remain unrecoverable.

Processing at a finer posting reduces both the sub-window coherence loss and the retrieval bias but increases the random phase noise as each window averages fewer pixels. Phase linking, which estimates a consistent phase time series from all interferometric pairs in an acquisition stack to reduce phase noise, could offset this by averaging over time rather than space, using temporal averaging instead of spatial multi-looking to achieve reasonable phase estimates (Guarnieri and Tebaldini, 2008; Ansari et al., 2017; Eppler and Rabus, 2022; Zwieback, 2022). This route needs testing on real data. Processing a NISAR stack well below the $80\,\mathrm{m}$ GUNW posting and phase linking across acquisitions would show whether the recovered ΔSWE improves over the standard product. It only helps where the scene remains coherent across the stack, so it favors slowly evolving snowpacks over the windy, rapidly redistributing terrain that drives the largest losses.

Triplet closure phases may identify and mitigate variability-induced bias of the phase-based SWE estimator. The closure phase $\Phi_{123} = \arg(z_{12} z_{23} z_{13}^*)$ is insensitive to any phase that separates into per-acquisition terms, such as the tropospheric or mean SWE phase, thus isolating the non-separable $\arg(M)$ term (Zwieback and Biessel, 2024). The confounding factor is that the closure signal is both largest and noisiest for large ΔSWE variability and, thus, low coherence, and that soil-moisture and wet-snow dielectric changes produce their own non-zero closure that must be separated from the sub-window term (Zan et al., 2014; Zwieback et al., 2016).

The clearest next step is to apply the same $|M|$ and $\arg(M)$ pipeline to lidar pairs in other basins and snow climates. Repeating it across maritime, continental, and tundra snowpacks would map where the sub-window effect is severe and tie

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

the coherence loss to a measurable property of each site, such as the ΔSWE break scale, rather than to a single basin. A NISAR acquisition coincident with a lidar pair would close the remaining gap by directly testing whether the lidar-predicted sub-window loss matches the observed $\gamma_{20} - \gamma_{80}$ residual, rather than relying on the consistency in sign shown here.

5.6 Limitations

This analysis has several limitations. The conversion from lidar ΔH_{snow} to ΔSWE uses a single basin-constant snow density for the new accumulated snow for each pair rather than a spatially varying one, which, in reality, will be more spatially variable than bulk density. We made this choice because ASO provides density only at a 50 m resolution, and it has minimal impact on the results. Using a spatially variable density shifts the basin-median $|M|$ by at most 0.03 and leaves the per-window $|M|$ tightly correlated with the constant-density value (Pearson $r = 0.91$ to 0.98 across the four pairs). The basin-median $\arg(M)$ moves by less than 2°. A constant density also matters mainly when the density itself varies within the window. If the denser wind slabs coincide with the deep, drift-loaded pixels, the true ΔSWE outliers are larger than estimated, so the constant-density assumption most likely understates the within-window variance and skew and is conservative for the effect reported here.

The analysis also covers a single basin and four lidar pairs. The sub-window effect depends on the spatial structure of ΔSWE, which varies with storm patterns, climate, terrain, and snowpack regimes. Therefore, the $|M|$ and $\arg(M)$ magnitudes here should not be interpreted as representative of all snow types. Shallower or more uniform snowpacks would show smaller coherence decreases and retrieval bias, while wind-dominated tundra could show as much or more. Confirming the size of the effect elsewhere requires comparable lidar pairs in other regions.

The NISAR coherence and the lidar ΔSWE fields are also not from the same dates. The multi-resolution coherence comes from December 2025 acquisitions with no coincident flights, while the $|M|$ maps come from lidar pairs in other years. The comparison, therefore, tests only the general spatial and temporal trends between the accumulation and no-accumulation NISAR pairs and cannot directly confirm that a given coherence drop comes from the sub-window variability of that same scene. It assumes the spatial ΔSWE pattern on the NISAR accumulation pair dates resembles the ΔSWE pattern from the lidar pairs used to predict $|M|$, which holds only to the extent that the terrain-driven accumulation structure repeats between events. Paired lidar flights coincident with both NISAR observations forming an interferogram would be necessary to fully attribute the coherence drop to the sub-window ΔSWE variability.

Both $|M|$ and $\arg(M)$ here assume that every per-cell phasor has unit amplitude. Real backscatter differs between pixels with different physical properties, such as soil moisture, roughness, and rock and vegetation cover of the ground surface beneath the snow. If these amplitude differences correlate with the same outlier pixels that drive the skew, they would amplify or dampen the bias relative to this unit-amplitude estimate.

A remaining concern is that part of the $|M|$ drop could be lidar depth noise rather than sub-window ΔSWE structure, since any residual per-pixel depth error adds its own within-window phase spread. We bound this two ways in Appendix A. At L- and P-band, the noise-only $|M|$ is 0.96 and 1.00 against observed medians of 0.33 and 0.88, so depth noise accounts for almost none of the loss. C-band is the tightest case: its noise floor of about 0.45 is the only one that removes a large share of

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

the coherence on its own, but since the factors multiply, the observed collapse to 0.10 still requires real ΔSWE structure to dominate. Even so, the C-band $|M|$ should be treated as a conservative lower bound, given the larger noise contribution.

Finally, the Tuolumne pairs were flown by ASO to bracket the dominant winter storm events. These large accumulations, combined with the spatially complex alpine terrain, maximize the within-window variance. Therefore, the coherence losses reported here for the storm pairs are likely an upper bound on the problem rather than a typical value over a winter time series. Smaller events, mid-elevation terrain, and slowly varying melt will likely be above these $|M|$ values.

6 Conclusions

We use repeat lidar to predict, independently of any radar observations, how the spatial variability of the ΔSWE field drives an apparent loss of InSAR coherence, separate from the temporal, geometric, and thermal decorrelation that is usually considered for coherence over snow. We also show that the skew of that within-window ΔSWE variability biases the retrieved ΔSWE, though this effect is negligible compared to the decorrelation at L-band. Over four Tuolumne lidar pairs, this sub-window variability alone lowers the L-band coherence factor $|M|$ to a basin median of 0.25 to 0.47, a coherence loss carried entirely by within-window ΔSWE structure, and adds a per-window retrieval bias of a few millimeters of SWE that reaches about 1.5 cm in the steep alpine terrain. We use two NISAR pairs, one spanning a large storm and one with little snow change, to show that the observed coherence and its change with resolution are spatially and temporally consistent with this sub-window loss.

The ΔSWE field is organized around a measurable length scale. Its power spectrum has a scale break at $\ell_b \approx 25\,\mathrm{m}$, the same fine-scale roughness break reported for snow depth in previous studies, and sweeping the multi-look window across this scale shows the median coherence factor falling from 0.54 at $9\,\mathrm{m}$ to 0.18 at $81\,\mathrm{m}$, with most of the loss incurred by a few tens of meters. Moving from the $80\,\mathrm{m}$ GUNW posting to a $20\,\mathrm{m}$ footprint nearly doubles the median coherence factor at L-band, from 0.18 to 0.33, while the skew-driven retrieval bias remains at the millimeter level at every footprint.

The effect scales strongly with wavelength. At C-band, the within-window phase spread collapses the coherence (median $|M|$ of 0.10) and leaves almost no recoverable phase, while P-band preserves the coherence (median 0.88) but still carries a small systematic negative retrieval bias. For L-band, NISAR, this means that coherence over spatially variable snow should not be considered as temporal decorrelation alone, and that large accumulation events need processing at finer scales than the standard $80\,\mathrm{m}$ GUNW posting to be effectively recovered. More broadly, the within-window structure of ΔSWE sets a sampling limit on any fixed-footprint InSAR SWE retrieval. Because that limit is set by the snow and not the sensor, it can be predicted and accommodated through the choice of processing resolution, wavelength, and timing, rather than appearing as unexplained and unrecoverable coherence loss.

Code and data availability. All processing and analysis code is available at https://github.com/zachhoppinen/tc_snow_variability_insar. Airborne Snow Observatory (ASO) lidar snow depth and SWE products are available through the ASO data portal (https://data.airbornesnowobservatories.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

com/). The NISAR geocoded unwrapped interferogram (GUNW) products are available from the Alaska Satellite Facility (https://search.asf.alaska.edu/).

Appendix A: Lidar Noise Limitation

This appendix bounds the contribution of lidar depth noise to $|M|$. Because the $|M|$ calculation runs on the 9 m block-averaged ΔSWE field, any residual per-pixel depth error adds its own within-window phase spread, which could appear as sub-window ΔSWE structure. We bound this noise in two independent ways.

First, the radial power spectrum of the 3 m ΔSWE field is a clean fractal power law with no flat high-frequency shelf from spatially uncorrelated noise, which rules out the large per-pixel noise that would otherwise appear as such a shelf. The high-frequency tail bounds the per-pixel 3 m noise at about 1.7 cm of SWE, or 0.5 cm per 9 m cell after the block average. Second, the ΔSWE over snow-free bare ground, where the true change is zero, is random per-pixel speckle of comparable magnitude (about 1.6 cm at 3 m for WY2025). On the 9 m working grid, both estimates give a noise-only $|M|$ of about 0.45 at C-band, 0.96 at L-band, and 1.00 at P-band, against observed medians of 0.10, 0.33, and 0.88. At L- and P-band, the noise factor is near one, so depth noise accounts for almost none of the loss. Because the noise and ΔSWE-structure factors multiply, C-band is the tightest case: a 0.45 noise floor still leaves the observed 0.10 requiring a real sub-window factor near 0.2, so structure dominates over noise at every band.

These same bounds explain our processing choices. This per-pixel noise is why the main pipeline pre-averages to 9 m before forming M. At the native 3 m posting, the same bounds imply a noise-only $|M|$ near 0.67 at L-band, which the 3×3 block average suppresses to 0.96. The window-size sweep (Section 4.3) must run on the native 3 m field to resolve its smallest windows, so its $|M|$ curve carries this noise. Because the per-pixel noise is spatially uncorrelated, its contribution to M is a window-size-independent multiplicative factor, so the shape of the curve, the location of its scale break, and the ratio between footprints are unaffected; only the absolute level is deflated, which is why the sweep $|M|$ is reported as a lower bound.

Appendix B: VIIRS Snow-Cover Cross-Check

The corrected multi-resolution residual $\gamma_{20} - \gamma_{80}$ (Section 4.4) attributes the storm-pair coherence drop to within-window ΔSWE variability. As an independent check that this signal tracks snow, rather than a bare-ground or vegetation artifact, we compare it against optical snow cover.

We use the VIIRS daily cloud-gap-filled NDSI snow-cover product VNP10A1F, Version 2 (Suomi-NPP VIIRS, 375 m sinusoidal grid, tile h08v05; distributed by the NSIDC DAAC) (Riggs and Hall, 2022). The cloud-gap-filled layer carries each cell forward from its most recent clear observation, so every NISAR pair window has continuous coverage. For each NISAR pair we pull the product on the end date of its temporal baseline (December 19, 2025 for the no-accumulation pair P12 and December 31, 2025 for the storm pair P23), because the decorrelation in an accumulation pair is driven by the new snow, which is only fully on the ground at the end of the baseline. We convert the NDSI snow cover (0–100) to fractional snow-covered

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

area (fSCA) via the linear relation $\mathrm{fSCA} = -0.01 + 1.45\,(\mathrm{NDSI}/100)$, clipped to $[0, 1]$ (Salomonson and Appel, 2004), and reproject it (nearest neighbor) onto the same $81\,\mathrm{m}$ grid and WY2023 footprint used for the coherence panels in Figure 6. Figure B1 places the per-pair fSCA beside the corrected $\gamma_{20} - \gamma_{80}$ for the same pair.

Within the surveyed footprint, the fSCA increases with elevation toward the high alpine terrain where the storm-pair residual is largest. Across the footprint the corrected residual co-varies with fSCA for the storm pair (Pearson $r = +0.40$, Spearman $\rho = +0.43$; mean fSCA 0.58, median residual $+0.015$) and only weakly for the no-accumulation pair (Pearson $r = +0.08$, Spearman $\rho = +0.09$; mean fSCA 0.12, median residual -0.005). The correlation is moderate by construction since fSCA indicates where snow lies rather than how much ΔSWE accumulated, which is the quantity that drives the sub-window decorrelation.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

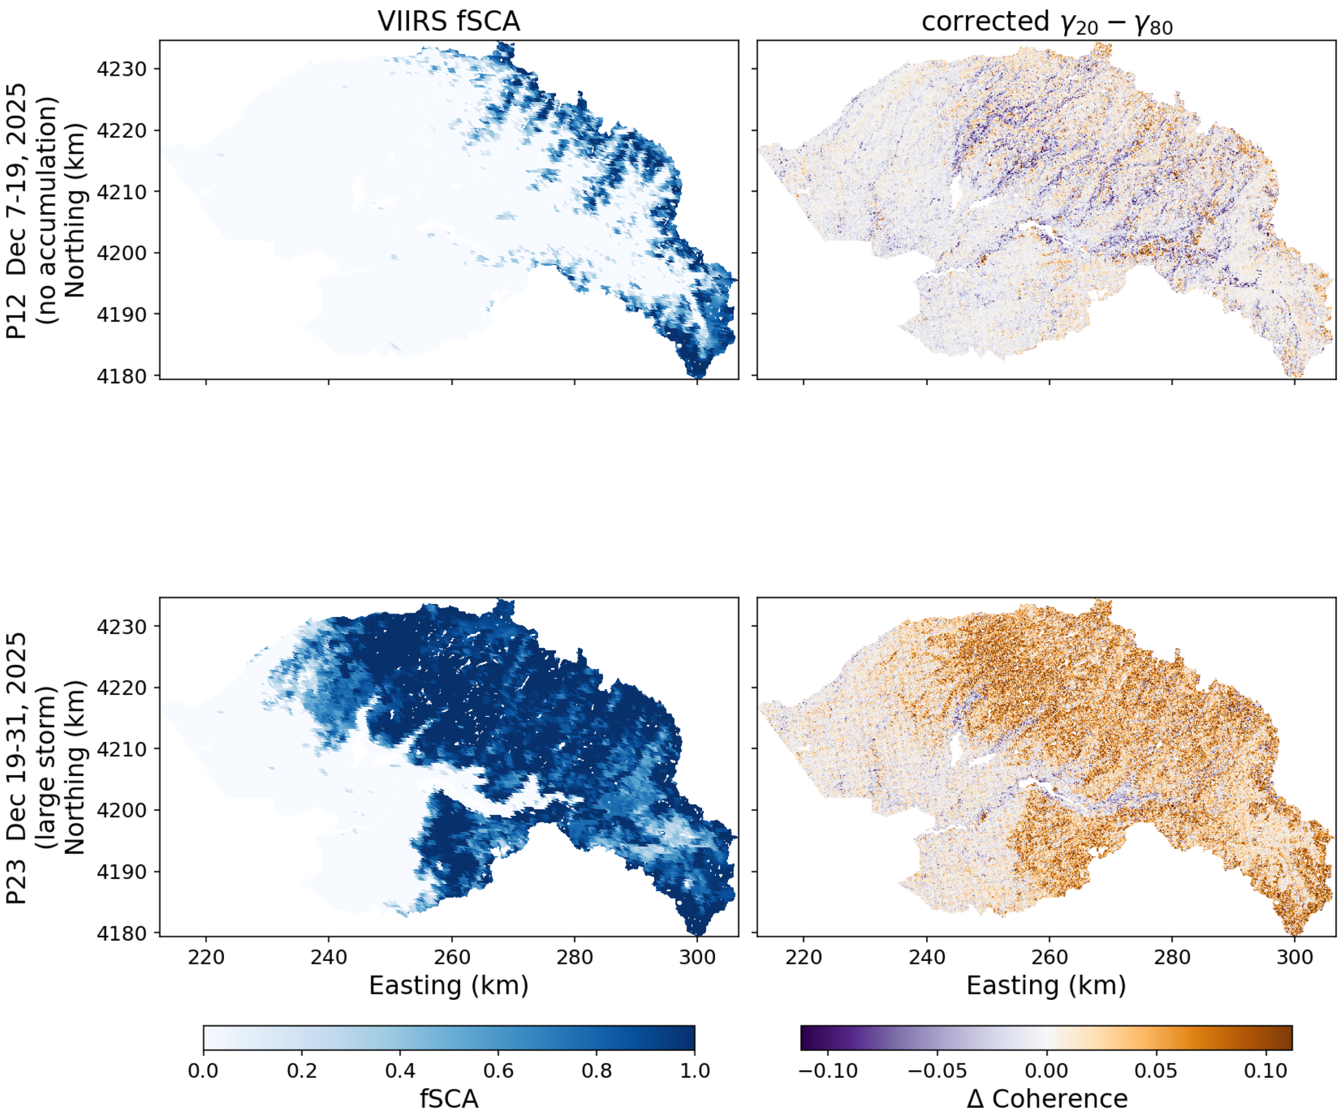

Figure B1. VIIRS VNP10A1F fractional snow-covered area (fSCA) (**left**) and the corrected $\gamma_{20} - \gamma_{80}$ residual (**right**) for the no-accumulation P12 pair (**top**) and the large-storm P23 pair (**bottom**), on each pair's end date and shown on the common WY2023 81 m footprint. The storm-pair residual is largest over the high-elevation terrain where fSCA is highest.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Author contributions. ZH conceived the study, developed the methodology, carried out the analysis, and wrote the original manuscript. SZ contributed to the conceptual development of the study and reviewed and edited the manuscript. RP and HPM reviewed and edited the manuscript. All authors read and approved the final manuscript.

Competing interests. The authors declare that they have no competing interests.

Acknowledgements. These ASO data were paid for by California Department of Water Resources and Turlock Irrigation District in the nominal operational program.

This work utilized pre-calibration (Beta) data (https://doi.org/10.5067/NIL1GUNW-B1) made available by the NASA-ISRO SAR (NISAR) mission, obtained from the Alaska Satellite Facility DAAC.

This work was supported by NASA grants 80NSSC25K7452 (ZH, RP, HPM), 80NSSC24K1082 (RP) and 80NSSC25K7386 (SZ).

The authors used Anthropic's Claude (via Claude Code) to assist with figure-generation code and manuscript text editing.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

References

Ansari, H., Zan, F. D., and Bamler, R.: Sequential Estimator: Toward Efficient InSAR Time Series Analysis, IEEE Transactions on Geoscience and Remote Sensing, 55, 5637–5652, https://doi.org/10.1109/tgrs.2017.2711037, 2017.

Bamler, R. and Hartl, P.: Synthetic aperture radar interferometry, Inverse Problems, 14, R1, https://doi.org/10.1088/0266-5611/14/4/001, 1998.

Belinska, K., Fischer, G., Nagler, T., and Hajnsek, I.: Snow Water Equivalent Estimation Using Differential SAR Interferometry and Co-Polar Phase Differences from Airborne SAR Data, IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium, 00, 4545–4548, https://doi.org/10.1109/igarss46834.2022.9883110, 2022.

Bonnell, R., McGrath, D., Tarricone, J., Marshall, H.-P., Bump, E., Duncan, C., Kampf, S., Lou, Y., Olsen-Mikitowicz, A., Sears, M., Williams, K., Zeller, L., and Zheng, Y.: Evaluating L-band InSAR Snow Water Equivalent Retrievals with Repeat Ground-Penetrating Radar and Terrestrial Lidar Surveys in Northern Colorado, EGUsphere, 2024, 1–33, https://doi.org/10.5194/egusphere-2024-236, 2024.

Clark, M. P., Hendrikx, J., Slater, A. G., Kavetski, D., Anderson, B., Cullen, N. J., Kerr, T., Hreinsson, E. , and Woods, R. A.: Representing spatial variability of snow water equivalent in hydrologic and land-surface models: A review, Water Resources Research, 47, https://doi.org/10.1029/2011wr010745, 2011.

Deems, J. S., Fassnacht, S. R., and Elder, K. J.: Fractal Distribution of Snow Depth from Lidar Data, Journal of Hydrometeorology, 7, 285–297, https://doi.org/10.1175/jhm487.1, 2006.

Eppler, J. and Rabus, B. T.: Adapting InSAR Phase Linking for Seasonally Snow-Covered Terrain, IEEE Transactions on Geoscience and Remote Sensing, 60, 1–13, https://doi.org/10.1109/tgrs.2022.3186522, 2022.

Fischer, N.: Statistical Analysis of Circular Data, Cambridge University Press, ISBN 9780521568906, https://doi.org/10.1017/cbo9780511564345, 1993.

Guarnieri, A. M. and Tebaldini, S.: On the Exploitation of Target Statistics for SAR Interferometry Applications, IEEE Transactions on Geoscience and Remote Sensing, 46, 3436–3443, https://doi.org/10.1109/tgrs.2008.2001756, 2008.

Guneriussen, T., Høgda, K. A., Johnsen, H., and Lauknes, I.: Insar for Estimation of Changes in Snow Water Equivalent of Dry Snow, IEEE Transactions on Geoscience and Remote Sensing, 39, 2101, https://doi.org/10.1109/36.957273, 2001.

Harris, F.: On the use of windows for harmonic analysis with the discrete Fourier transform, Proceedings of the IEEE, 66, 51–83, https://doi.org/10.1109/proc.1978.10837, 1978.

Hoppinen, Z., Oveisgharan, S., Marshall, H.-P., Mower, R., Elder, K., and Vuyovich, C.: Snow water equivalent retrieval over Idaho – Part 2: Using L-band UAVSAR repeat-pass interferometry, The Cryosphere, 18, 575–592, https://doi.org/10.5194/tc-18-575-2024, 2023.

Leinss, S., Wiesmann, A., Lemmetyinen, J., and Hajnsek, I.: Snow Water Equivalent of Dry Snow Measured by Differential Interferometry, IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 8, 3773–3790, https://doi.org/10.1109/jstars.2015.2432031, 2015.

Li, S. and Sturm, M.: Patterns of wind-drifted snow on the Alaskan arctic slope, detected with ERS-1 interferometric SAR, Journal of Glaciology, 48, 495–504, https://doi.org/10.3189/172756502781831151, 2002.

Marshall, H., Deeb, E., Forster, R., Vuyovich, C., Elder, K., Hiemstra, C., and Lund, J.: L-Band InSAR Snow Depth Retrievals from Grand Mesa in NASA SnowEX 2020 Campaign, Institute of Electrical and Electronics Engineers, 2021.

Massonnet, D. and Feigl, K. L.: Radar interferometry and its application to changes in the Earth's surface, Reviews of Geophysics, 36, 441–500, https://doi.org/10.1029/97rg03139, 1998.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Mendoza, P. A., Musselman, K. N., Revuelto, J., Deems, J. S., López-Moreno, J. I., and McPhee, J.: Interannual and Seasonal Variability of Snow Depth Scaling Behavior in a Subalpine Catchment, Water Resources Research, 56, https://doi.org/10.1029/2020wr027343, 2020.

Miller, Z. S., Peitzsch, E. H., Sproles, E. A., Birkeland, K. W., and Palomaki, R. T.: Assessing the seasonal evolution of snow depth spatial variability and scaling in complex mountain terrain, The Cryosphere, 16, 4907–4930, https://doi.org/10.5194/tc-16-4907-2022, 2022.

Nagler, T., Rott, H., Scheiblauer, S., Libert, L., Mölg, N., Horn, R., Fischer, J., Keller, M., Moreira, A., and Kubanek, J.: Airborne Experiment on Insar Snow Mass Retrieval in Alpine Environment, IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium, 00, 4549–4552, https://doi.org/10.1109/igarss46834.2022.9883809, 2022.

Oveisgharan, S., Zinke, R., Hoppinen, Z., and Marshall, H. P.: Snow water equivalent retrieval over Idaho – Part 1: Using Sentinel-1 repeat-pass interferometry, The Cryosphere, 18, 559–574, https://doi.org/10.5194/tc-18-559-2024, 2024.

Palomaki, R., Hoppinen, Z., and Marshall, H.-P.: A spatiotemporal analysis of errors in InSAR SWE measurements caused by non-snow phase changes, EGUsphere, 2025, 1–28, https://doi.org/10.5194/egusphere-2025-5255, 2025.

Riggs, G. A. and Hall, D. K.: VIIRS/NPP CGF Snow Cover Daily L3 Global 375m SIN Grid, Version 2 (VNP10A1F), https://doi.org/10.5067/PN50Y51IVNLE, tile h08v05, Suomi-NPP. Accessed 2026-06-25, 2022.

Rosen, P., Hensley, S., Shaffer, S., Edelstein, W., Kim, Y., Kumar, R., Misra, T., Bhan, R., and Sagi, R.: The NASA-ISRO SAR (NISAR) Mission Dual-Band Radar Instrument Preliminary Design, 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 3832–3835, https://doi.org/10.1109/igarss.2017.8127836, 2017.

Ruiz, J. J., Lemmetyinen, J., Kontu, A., Tarvainen, R., Pulliainen, J., Vehmas, R., and Praks, J.: Analysis of Snow Coherence Conservation for SWE Retrieval at L-, S-, C-and X-Bands, 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, 00, 860–863, https://doi.org/10.1109/igarss47720.2021.9554103, 2021.

Ruiz, J. J., Lemmetyinen, J., Kontu, A., Tarvainen, R., Vehmas, R., Pulliainen, J., and Praks, J.: Investigation of Environmental Effects on Coherence Loss in SAR Interferometry for Snow Water Equivalent Retrieval, IEEE Transactions on Geoscience and Remote Sensing, 60, 1–15, https://doi.org/10.1109/tgrs.2022.3223760, 2022.

Ruiz, J. J., Merkouriadi, I., Lemmetyinen, J., Cohen, J., Kontu, A., Nagler, T., Pulliainen, J., and Praks, J.: Comparing InSAR Snow Water Equivalent Retrieval Using ALOS2 With In Situ Observations and SnowModel Over the Boreal Forest Area, IEEE Transactions on Geoscience and Remote Sensing, 62, 1–14, https://doi.org/10.1109/tgrs.2024.3439855, 2024.

Salomonson, V. V. and Appel, I.: Estimating fractional snow cover from MODIS using the normalized difference snow index, Remote Sensing of Environment, 89, 351–360, https://doi.org/10.1016/j.rse.2003.10.016, 2004.

Tarricone, J., Webb, R. W., Marshall, H.-P., Nolin, A. W., and Meyer, F. J.: Estimating snow accumulation and ablation with L-band interferometric synthetic aperture radar (InSAR), The Cryosphere, 17, 1997–2019, https://doi.org/10.5194/tc-17-1997-2023, 2023.

Trujillo, E., Ramírez, J. A., and Elder, K. J.: Topographic, meteorologic, and canopy controls on the scaling characteristics of the spatial distribution of snow depth fields, Water Resources Research, 43, https://doi.org/10.1029/2006wr005317, 2007.

Yackel, J., Geldsetzer, T., Yendamuri, K., Nandan, V., Tonboe, R., Stroeve, J., Saha, M., Mallett, R., and Lam, H. M.: Observations of Coherence Changes in C-Band SAR Linked to Snow Thermodynamics and Wind Redistribution on First-Year Sea Ice, Geophysical Research Letters, 53, e2025GL121 597, https://doi.org/10.1029/2025gl121597, e2025GL121597 2025GL121597, 2026.

Zan, F. D., Parizzi, A., Prats-Iraola, P., and López-Dekker, P.: A SAR Interferometric Model for Soil Moisture, IEEE Transactions on Geoscience and Remote Sensing, 52, 418–425, https://doi.org/10.1109/tgrs.2013.2241069, 2014.

Zebker, H. and Villasenor, J.: Decorrelation in interferometric radar echoes, IEEE Transactions on Geoscience and Remote Sensing, 30, 950–959, https://doi.org/10.1109/36.175330, 1992.

https://doi.org/10.5194/egusphere-2026-3772
Preprint. Discussion started: 15 July 2026

Zwieback, S.: Cheap, Valid Regularizers for Improved Interferometric Phase Linking, IEEE Geoscience and Remote Sensing Letters, 19, 1–4, https://doi.org/10.1109/lgrs.2022.3197423, 2022.

Zwieback, S. and Biessel, R.: Temporal Closure Signatures in Radar Interferometry, IEEE Transactions on Geoscience and Remote Sensing, 62, 1–11, https://doi.org/10.1109/tgrs.2024.3471712, 2024.

Zwieback, S., Liu, X., Antonova, S., Heim, B., Bartsch, A., Boike, J., and Hajnsek, I.: A Statistical Test of Phase Closure to Detect Influences on DInSAR Deformation Estimates Besides Displacements and Decorrelation Noise: Two Case Studies in High-Latitude Regions, IEEE Transactions on Geoscience and Remote Sensing, 54, 5588–5601, https://doi.org/10.1109/tgrs.2016.2569435, 2016.