Reply to Referee 1, egusphere-2026-2542

Reply to Referee #1 (RC1)

Manuscript: egusphere-2026-2542, "tidal-insar-sim: an open-source toolkit for tide-aware ranking of NISAR L-band acquisitions over Antarctic grounding lines" **Author comment AC1, draft of 2 October 2026**

We thank the referee for a careful report. It changed the paper in three ways. The predicted DDInSAR yield is now kept separate from grounding-line (GL) retrieval throughout. Section 3 is cut to the steps needed to reproduce the results. And the reconciliation the referee asked for in Sections 4.2 and 4.3 led us to rebuild the whole analysis from one archived snapshot of the NISAR Reference Observation Plan (ROP), so that every number, table and figure now comes from a single set of triplets. Following the referee's third major comment, we also tested the prediction against released NISAR interferograms (new Sect. 4.5 and Fig. 7).

Answering the report also exposed a problem the referee did not raise, and we report it here first. **The archived release cited in the discussion paper, tidal-insar-sim v0.3.0 (https://doi.org/10.5281/zenodo.20056468), does not contain several things Section 3 and the Code availability statement say it contains.** It has no `rop` module, no `scripts/` directory and no `data/rop_snapshot/`. The class names `CATSValidationError` and `ConstellationDescriptor` do not occur in it. The package is not on PyPI and has no continuous-integration workflow, and it records no SHA-256 of the tide-model files. The ROP analysis behind Table 2 had been run outside the package and was not archived. We have now written the `rop` module into the package, added the script `scripts/gmd_reproduce.py`, which regenerates every number, table and figure in Sections 4 and 5 from an archived ROP snapshot and CATS2008, and removed from the manuscript every statement about the software that the code does not bear out. These go into release v0.3.2 (https://doi.org/10.5281/zenodo.23323704), which the revised manuscript cites. We apologise for the error.

Rebuilding the analysis changed the numbers, though not the conclusions. The ROP has been updated since May, and the May query was not saved. Our rebuilt analysis also counts a pass on which two L-band modes are scheduled once, not twice. The analysis window is unchanged (24 September 2025 to 24 September 2026).

	Discussion paper	Revised
Tracks containing Thwaites / Pine Island / Smith	9 / 9 / 8	8 / 10 / 8
Rigid 12-day triplets	793	738
Predicted n_fr ≥ 3	87 (11.0 %)	80 (10.8 %)
Predicted n_fr ≥ 5	5 (0.6 %)	5 (0.7 %)
Largest predicted n_fr (Thwaites / Pine Island / Smith)	5.69 / 5.10 / 5.64	5.53 / 5.57 / 5.92
Ranked culling keeps every predicted-usable triplet up to	89 % of triplets	89 % of triplets

Below, the referee's comments are in italics and line numbers refer to the discussion paper.

Major comments

Comment R1.C1

Section 3 includes detailed descriptions of software dependencies, file formats, web-interface functionality, testing infrastructure, continuous integration, and software deployment … these details could be substantially condensed and moved to the software documentation or repository.

We agree. Section 3 now describes only what a reader needs to reproduce Sections 4 and 5. Section 3.1 is one paragraph that points to Figure 1, Table 1 and the repository for installation, interfaces, formats and tests. The subsection on verification and reproducibility (former 3.6) and the paragraph on radar-mode practice (former Sect. 3.2, second paragraph) are removed. Three things are added because the referee and a reader need them to reproduce the numbers, and the discussion paper did not state them. They are how acquisitions are selected from the ROP (one acquisition per pass, the wider L-band mode kept where two are scheduled, triplets formed only within one mode), how the incidence angle at the site is computed (Sect. 3.2), and the triplet constraints (Sect. 3.4). Figure 1 and Table 1 are revised to match the code.

Comment R1.C2

Distinguish more explicitly between the predicted DDInSAR yield derived from the acquisition plan and tidal model and actual successful grounding-line retrievals … The reported percentages and the inferred biweekly monitoring capability appear to represent predicted opportunities.

The referee is right, and the discussion paper did not make this distinction. In the revision every yield is stated as a prediction ("predicted to reach 3 fringes"), and "success rate" no longer appears in any heading or in the text. The abstract now says: "These are predicted opportunities, not demonstrated retrievals". The sentence on biweekly monitoring (former line 271) is deleted. Section 4.4 now gives the timing of the predicted-usable triplets with its gaps (longest gaps of 33, 52 and 31 days at Thwaites, Pine Island and Smith) and states the two assumptions behind any monitoring claim: "It assumes that every scheduled acquisition is made, processed and released, and that 3 fringes are enough to delineate the GL; neither assumption has yet been tested with NISAR data (Sect. 4.5)." The Discussion and Conclusions carry the same qualification, and Sect. 5 now notes that treating triplets as independent makes its culling numbers an upper bound.

Comment R1.C3

Add a section in which actual grounding lines derived from NISAR DDInSAR data are presented, if such data are already available … A comparison of grounding-line retrievals from NISAR and Sentinel-1 … would further strengthen the scientific value of the paper.

We agree that this is the test the paper needs, and NISAR data now make it possible. The new Sect. 4.5 reports the first step. On 2 October 2026 we queried NASA's Common Metadata Repository for NISAR Level-1 RSLC products over the three sites. All three images are public for 110 of the 738 predicted triplets, for 14 of the 80 predicted to reach 3 fringes and for none of the 5 predicted to reach 5. The public archive is still filling (it holds 96, 95 and 68 RSLC track-dates at the three sites against 244, 304 and 242 scheduled passes), so these fractions describe the archive's release status, not acquisitions that failed. For 64 predicted triplets, with predicted fringe counts from 0.01 to 4.25, both consecutive 12-day unwrapped interferograms (GUNW) are already public, so their double difference can be formed from released products without reprocessing.

We then ran the test the referee asks for. We formed the double difference for 22 of these triplets, chosen to span the predicted range: all 8 predicted to reach 3 fringes, 8 between 1 and 3 fringes, and 6 below 1 fringe as controls. The double difference is formed from the wrapped interferograms, so steady flow cancels without unwrapping either 12-day pair; it is then unwrapped with SNAPHU at 40 m, and the observed fringe count is the phase step across the grounding-zone fringe belt (new Sect. 4.5 and Fig. 7). The result has two parts, and we report both:

- Where the grounding zone is coherent, the prediction holds. The revised Sect. 4.5 states: "The three triplets predicted at 3.50, 4.16 and 4.25 fringes show steps of 2.8, 3.9 and 5.3 fringes, and the two predicted at 0.48 and 1.17 fringes show 0.44 and 0.25, within the background."
- We checked how much this depends on the method's free choices (27 combinations of plateau distance, coherence threshold and ridge selection). The Smith values are stable; the Pine Island value predicted at 4.25 depends on the plateau distance, and two Pine Island triplets that pass only looser settings disagree with the prediction. The revised text reports this: "The three Smith values are stable: 3.1 to 4.4 fringes for the triplet predicted at 4.16, 2.1 to 4.2 for 3.50 and 0.2 to 0.3 for 1.17."
- Coherence, which the toolkit does not model, limits the predicted yield. "Only 5 of the 22 triplets could be measured (Fig. 7a). At Thwaites none of the 11 triplets is coherent across the grounding zone". The paper now concludes that "The yields in Table 2 are therefore upper bounds, reached only where 12-day L-band coherence survives across the grounding zone, which at these sites excludes Thwaites and most of Pine Island."

We also delineated the NISAR grounding line, at Smith, from the two coherent triplets, as the landward limit of tidal flexure (Fig. 7c), and compared it with Sentinel-1. The revised text reports: "The two lines, from different tracks with central acquisitions 5 days apart, agree to 0.54 km (median nearest-point distance)." Both lie on the most recent Sentinel-1 grounding line in our dataset, of 30 January 2025: "the median signed offset is +0.03 km for track 34 (interquartile range -0.21 to +0.40 km) and -0.13 km for track 106 (-0.66 to +0.15 km), positive seaward." This is not the same-season comparison the referee suggests: no Sentinel-1 grounding line from July 2026 is available to us, and the 18 months between the two dates, together with tidal migration of the grounding line, limit what the comparison resolves to about 0.5 km, which the text states. At Pine Island the coherent area is too small to delineate a line. The scripts, the product list, the measured values and the delineated lines are archived with the code (scripts/nisar_dd_test/).

Specific comments

Comment R1.C4

Consistent and conventional text formatting for software tools, variables, functions. Done. The package, module and file names are set in a monospace font; third-party software names stay in the body font; mathematical symbols use true subscripts.

Comment R1.C5

L19: ROP acronym; sometimes "rop". ROP is now defined at first use in the abstract and again in the Introduction. Lower-case `rop` now appears only as the module name, in monospace.

Comment R1.C6

The grounding line itself does not contribute to sea-level rise. Corrected throughout. The abstract, Sect. 4.1 and the Conclusions now refer to "Thwaites, Pine Island and Smith glaciers, which account for most of the present-day ice loss from the Amundsen Sea Embayment", and the first sentence of the Introduction now begins "The retreat of Antarctic glaciers across their grounding lines".

Comment R1.C7

L54: across the mission repeat cycle. Rephrased: "the K1 phase advances by only about 12° from one acquisition to the next across the mission repeat cycle."

Comment R1.C8

L71: explain or cross-reference high-tide- and low-tide-at-t2. The Introduction now defines both modes in words and points to Sect. 2.3, which gives the formal definition: "We call triplets with h_2 above the mean of h_1 and h_3 (hDD < 0) high-tide-at-t_2, and those with h_2 below it (hDD > 0) low-tide-at-t_2."

Comment R1.C9

Poor formatting of equations. Equations (1) to (4) are now native Word equations at body size, and inline symbols use true subscripts.

Comment R1.C10

Recheck the flexural length numerically and define it. The referee's numbers are correct and ours were wrong. With E = 0.88 GPa, ν = 0.41 and H = 800 m, π/β = 6.46 km, not 12 km. We also aligned the stated defaults with the code, which uses ν = 0.3, and ice thickness is a per-site user input rather than a BedMachine lookup. The text now defines both length scales: "The envelope e^(−βs) decays over 1/β, and w first reaches its maximum, (1 + e^(−π)) h ≈ 1.04 h, at s = π/β, which we take as the seaward edge of the flexure zone. For H = 800 m, 1/β ≈ 2.0 km and π/β ≈ 6.3 km." We add that these parameters affect only the synthetic maps and enter neither the fringe count nor the ranking, so the error did not propagate into any result.

Comment R1.C11

The transition from Eq. (3) to Eq. (4) … does not follow from integrating the phase itself over a flexure half-width. Agreed; the derivation now reads: "Equation (3) is zero on grounded ice and tends to −(4π/λ) cos θinc hDD on freely floating ice (βs ≫ 1). The number of fringes between grounded and floating ice is the net phase change between the two divided by 2π." We add that the overshoot of Eq. (1) at s = π/β makes the phase in the flexure zone exceed the far-field value by up to 4 %, which Eq. (4) ignores.

Comment R1.C12

Justify n_fr ≥ 3 … Rignot et al. (2011) does not mention the minimum required number of fringes. Provide a more direct reference or a calibration experiment. The referee is right that Rignot et al. (2011) does not support the threshold, and we have found no study that does. The revision says so: "These thresholds are operating heuristics. We know of no study that reports a minimum fringe count for GL delineation." We have removed the citation and the claim about sub-pixel retrieval at 5 fringes. Because the argument should not rest on an unsupported number, Sect. 4.2 now reports the predicted yield across thresholds: 60 % of triplets reach 1 fringe, 28 % reach 2, 11 % reach 3, 3.1 % reach 4 and 0.7 % reach 5. The conclusion that ranking matters does not depend on where the threshold sits. The NISAR test under R1.C3 is the calibration experiment the referee asks for. At Smith, where the result is robust, the two triplets predicted above 3 fringes produced clear fringe belts (2.8 and 3.9 fringes) and the one predicted at 1.17 gave a step within the background. That supports the

threshold as a separator of useful from unusable triplets where the data are coherent, on a small sample from one site; the Pine Island evidence is weaker, and Sect. 4.5 says so. The same test shows that coherence, not the threshold, decides at Thwaites, where no triplet could be measured.

Comment R1.C13

L119 wording. Rephrased: "Useful hDD therefore comes from the semidiurnal M2 and S2 constituents and from the beat of K1 with the other diurnal constituents, O1 and P1, between cycles."

Comment R1.C14

Section 3.1: cross-reference Figure 1. Added.

Comment R1.C15

L131: TLE not explained. The sentence is removed; TLE is now defined ("two-line element") at its first remaining use, in the Table 1 caption.

Comment R1.C16

Table 1 caption above the table. Done for both tables.

Comment R1.C17

L171: explain CATSValidationError. No class of this name exists in the code (see the opening of this reply). The tide-gauge comparison is part of the test suite, not a runtime check, and the text now says that: "The test suite checks the M2 and K1 amplitudes returned by the toolkit against eight Antarctic tide gauges, within station-specific tolerances of 3 to 10 cm."

Comment R1.C18

L195: ASE acronym. Defined at first use in the Introduction.

Comment R1.C19

Section 4.1: cross-reference Figure 2. Added.

Comment R1.C20

Reconcile the figures with Table 2 … I would recommend that all figures and numbers in the tables be derived from the same set of the triplets for each ROI. The referee identified a real inconsistency. Figures 2, 4 and 5 of the discussion paper were drawn from an earlier analysis that computed acquisition times from orbit propagation instead of reading them from the ROP, and that earlier set included more tracks per site (hence 1013 + 844 triplets at Thwaites in Figure 5). Table 2 and Figures 3 and 6 came from the ROP analysis. In the revision, one script computes all triplets once and draws every table and figure in Sects. 4 and 5 from that set. Sect. 4.2 now states this: "Every number, table and figure in Sect. 4 and 5 is computed from this one set of triplets." Thwaites now has 228 triplets in Table 2 and in Figure 5, and the largest predicted fringe counts in Table 2 (5.53, 5.57 and 5.92) are the ones shown for the top-ranked triplets in Figure 4.

Comment R1.C21

L214 wording. Rewritten: "For each site and stream we evaluated CATS2008 at every scheduled acquisition and enumerated the rigid 12-day triplets (Sect. 3.4)."

Comment R1.C22

L218: Figure 3, not Figure 2. Corrected.

Comment R1.C23

Section 4.3: cross-reference Figures 4 and 5; Figure 5 caption too long; cross-reference Eq. 4. All three done. The Figure 5 caption is now three lines; its explanation of the diagonal moved to the text of Sect. 4.3. The text cites Eq. (4) for the fringe counts of the top-ranked triplets.

Comment R1.C24

Section 4.4 (L271) wording. The sentence is deleted under R1.C2.

Comment R1.C25

Figure 6: legend obstructs the upper left panel. Figure 6 is redrawn; site names label the rows, so the panel carries no legend.

Comment R1.C26

L292: Figure instead of Fig. We follow the Copernicus convention consistently: "Fig." within a sentence and "Figure" at the start of one. Line 292 was inside a sentence, so it keeps "Fig.", but the manuscript previously mixed both forms mid-sentence, and that is now fixed.

Comment R1.C27

L295: explain 25 % cull. The text now defines it: "We represent a duty-cycle reduction as the removal of a fixed share of the 738 triplets; a 25 % cull, for example, removes 184 of them." Random culling is now evaluated over 2000 draws and reported with its spread, not as a proportional rule.

Comment R1.C28

L321: explain the ConstellationDescriptor abstraction. No class of this name exists in the code. The sentence now describes what the code does: "For a mission without a published plan, the orbit module generates acquisitions from two-line element orbit propagation with Skyfield, and these enter the same ranking."

Other changes

- The discussion paper said the high- and low-tide-at-t_2 modes "cluster in non-overlapping spring-tide windows (May–June 2026 high, January–February 2027 low)". This is not borne out: both modes occur in 9 months of the year at Thwaites, and January–February 2027 lay outside the analysis window. The statement is removed from Sect. 4.3 and the Conclusions.
- Figure 2's caption gave the footprint as 240 by 150 km; ROP frames are about 245 by 257 km. Corrected.
- The incidence angle at each site is now computed from the frame geometry and NISAR's left-looking geometry (33.3° to 46.6° at the three sites). Replacing every angle by 40° changes the count above 3 fringes from 80 to 81, so the results are not sensitive to this step.
- Rignot et al. (2019), cited in Sect. 4.1, was missing from the reference list of the discussion paper and is now included. The grounding lines used in Figs. 2 and 7 are now cited to the dataset that contains them (Rignot and Scheuchl, 2026), and the elastic modulus to Vaughan (1995).
- The CATS2008 reference now lists all three creators in the order the data centre asks for (Howard et al., 2019). The duplicate Vashishtha et al. (2025) entry and the uncited Morlighem et al. (2020) entry are removed, and the NISAR handbook that documents the left-looking geometry is added.
- The Introduction no longer claims that coherence is typically lost at baselines longer than 12 days; we found no source for that statement.

- The Acknowledgements now name the funding, NASA award 80NSSC25K7365 (STV Applications Needs), and the Competing interests statement now discloses the author's financial interest in Terra Monitoring Solutions, LLC, in place of "The author declares no competing interests".
- The generative-AI statement now also covers the use of AI during this revision.