

NISAR offsets measure pre-failure motion but do not single out the Bhote Koshi avalanche source

Sansar Raj Meena

smeena@cgs.it

<https://orcid.org/0000-0001-6175-6491>

Article

Keywords:

Posted Date: September 22nd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11022365/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

NISAR offsets measure pre-failure motion but do not single out the Bhote Koshi avalanche source

Sansar Raj Meena^{1*}

¹ Geophysics Department, SpatioAI Lab, National Institute of Oceanography and Applied Geophysics – OGS, Borgo Grotta Gigante 42/C, 34010 Sgonico, Trieste, Italy

* Correspondence: Sansar Raj Meena (smeena@ogs.it)

ORCID Sansar Raj Meena: <https://orcid.org/0000-0001-6175-6491>

Abstract

Whether catastrophic ice–rock detachments are preceded by measurable surface motion decides whether satellite geodesy can contribute to early warning. We analyse five consecutive NISAR L-band pixel-offset pairs spanning the 26 August 2026 Bhote Koshi ice–rock avalanche, Nepal, and screen the full 116,707 km² footprint for coherently moving surfaces. In the interval closing seven days before failure the source moved 0.96 m in 24 days along the topographic fall line, and the signal survives a matched-magnitude control; but against size-matched random regions no pre-event interval reaches the 95th percentile on directional coherence or displacement, and no interval shows acceleration. The pair spanning the failure decorrelates. A blind footprint-wide screen recovers the source among 101 validated zones of downslope motion, 93 of them inventoried glaciers, and does not distinguish it by magnitude, persistence or coherence. Motion at a future source is measurable before failure; it is not diagnostic of failure.

Introduction

High-mountain Asia is losing ice rapidly¹ while sustaining the water supply of hundreds of millions of people², and warming degrades the permafrost that binds steep bedrock³. Catastrophic detachments of ice and rock are the most destructive expression of this instability

— Kolka–Karmadon in 2002⁴, Aru in 2016⁵, Chamoli in 2021⁶ — and a systematic inventory suggests they are more common than assumed⁷. Entrainment converts them into long-runout debris flows⁸ that reach far beyond the source, compounding an already elevated outburst-flood hazard across the Himalaya⁹.

Whether such detachments are preceded by measurable surface motion determines whether satellite geodesy can support early warning, and the evidence is mixed. Retrospective analysis resolved pre-collapse motion at Chamoli¹⁰ and multi-year precursory displacement at Sedongpu¹¹, and inverse-velocity methods can forecast failure once acceleration is established^{12,13}, but reviews emphasise how often precursors are absent or recognisable only afterwards¹⁴. Two weaknesses recur in the retrospective literature and are rarely stated. The first is a selection effect: after a disaster the intervals processed first are those expected to be informative, so a sub-sampled series is structurally more likely to resemble a ramp than a complete one. The second is the absence of a population: a signal at the site of a known failure can only be called distinctive if the same measurement has been made on comparable surfaces that did not fail.

Displacement of the magnitude expected on a failing glacier exceeds what interferometric phase can track: temporal decorrelation over steep, snow-covered terrain is severe¹⁵ and metre-scale motion aliases the phase. Amplitude pixel-offset tracking^{16,17}, refined for glaciological work¹⁸, is the standard alternative, and NISAR¹⁹ now supplies L-band offsets systematically at 12-day repeat, distributed as geocoded Level-2 products that cover whole footprints rather than selected sites. That combination makes it possible, for the first time, to place a continuous site series inside the population of every comparable surface in the same scene.

Screening a whole footprint raises a problem that a site study never faces. In a scene of 18 million pixels, features that are simultaneously large in amplitude and locally organised in direction are certain to occur without any ground motion. Residual co-registration leaves a slowly varying vector offset across the scene; error in the elevation model used for geocoding maps into apparent displacement along the ground-range direction; and offset estimation over low-contrast surfaces produces spatially correlated noise. A screen on displacement magnitude and directional coherence — the two quantities the products naturally provide — cannot separate these from real motion, because all of them satisfy both criteria, and the failure is silent: an unvalidated screen returns a plausible inventory whose composition is unknown. Direction relative to topography resolves the ambiguity. Motion under gravity is downslope, and downslope varies from valley to valley, whereas a co-registration residual points one way scene-wide and a height-error artefact points along the nearly fixed look direction. Only real motion can track a spatially varying topographic direction, and the elevation gradient it is tested against is never used by the screen itself.

On 26 August 2026 an ice–rock mass detached from the Bhote Koshi headwaters in Rasuwa, Nepal, and generated a debris flow that reached settlements more than 25 km downstream. Here we address two gaps with the five NISAR pairs that span the event. At site level we ask whether pre-failure motion at the source is resolvable against a calibrated null rather than against a bare background, whether that verdict survives reasonable changes in how the reference field is formed, and whether the continuous series shows acceleration. At footprint level we ask whether the source was in any way distinctive among coherently moving surfaces across the central Nepal Himalaya and southern Tibet. The two questions are linked by the selection effect: a site series assembled from the most promising intervals, and a site chosen because it failed, are both

samples drawn after the outcome is known. We frame the results around what the data resolve and what they do not.

Results

The 26 August 2026 event and its runout

The detachment occurred at 5151 m (85.52985° E, 28.29041° N) and the resulting debris flow ran more than 25 km down the Bhote Koshi–Trishuli valley (Fig. 1 and Supplementary Fig. 2). Mapped from post-event WorldView-3 imagery of 27 and 28 August, the runout covers 13.2 km² and contains 2,055 Google Open Buildings footprints²⁰; buffering the mapped boundary by ±50 m gives 1,575–2,353, so we report about two thousand. These are exposure counts — footprints inside the mapped extent — not a damage assessment, and no per-building damage state was determined. For context, the national Rapid Damage and Needs Assessment records 7,570 affected buildings across five districts along the full Trishuli–Narayani corridor, of which 1,779 are in Rasuwa²¹; our count covers only the extent mapped here. A third-party post-event stereo reconstruction of the lower corridor from Vantor WorldView-3 strips^{22,23} shows metres of bed scour and deposition along the runout; we use it only as context and report our own stable-ground check of that product in Supplementary Fig. 1.

Motion at the source before failure

In the 26 Jul–19 Aug pair — the last pre-event interval, closing seven days before failure — the crown carries an organised flow field directed west-north-west, along the topographic fall line (Fig. 2d, Fig. 3a). Median horizontal crown displacement is 0.96 m over 24 d, a nominal 4.0 cm d⁻¹, and the mean crown vector is 1.72 m toward 278°, 18° from the downslope azimuth of 296° derived from the elevation model, which the projection never uses (Methods). Directional

coherence R — the mean resultant length of unit displacement vectors in a 5×5 window — has a crown median of 0.87 against 0.69 over a fixed 6 km reference box (Fig. 2f, Table 1 and Supplementary Table 4).

That comparison is a bare inequality, and it depends on choices we test. First, R uses vector direction only, but it need not be independent of displacement magnitude, so we measured the relation rather than assuming it. Binned by displacement, R is not monotonic across the series, and in the 24-day pair it falls with magnitude (0.675, 0.635, 0.623 across the 0–0.25, 0.25–0.5 and 0.5–1 m bins), so a longer interval does not by itself raise R here. As a direct check we restricted the 24-day crown to the 10th–90th percentile displacement window of the resolved 12-day crown (0.17–2.75 m). The crown then gives $R = 0.869$ ($n = 147$) against a matched-magnitude background of 0.693 ($n = 1,572$; because R is estimated in a moving window these are not independent samples), an advantage of +0.176 against +0.185 unrestricted (Fig. 3c). Matching the displacement window leaves the crown advantage essentially unchanged, so it is not created by the longer interval.

Second, individual pixel displacements are not diagnostic. After bulk referencing, the median absolute residual over the reference region is 0.07–0.10 m in the east and 0.06–0.13 m in the north component, which we take as the single-pixel noise floor; but that floor is a median, not a bound. The same region has a 90th percentile of 0.57–0.78 m, and 6–8 % of its pixels match or exceed the crown median displacement of the same pair. The polygon median, not the pixel, is the quantity we interpret.

Ranking, not detection

Across the four pre-event intervals the crown median R exceeds the box median in two — 20 Jun–2 Jul (0.75 against 0.64) and 26 Jul–19 Aug (0.87 against 0.69) — and falls below it in the

115 two mid-July pairs (0.61 against 0.72; 0.61 against 0.76) (Table 1). An inequality needs a null
 116 before it can be read as a detection. We drew compact regions with the crown's own valid-pixel
 117 count at random positions on valid ground in the same 6 km box (857–996 accepted draws per
 118 pair). Placement alone satisfies the inequality in 37–50 % of draws, so two of four is close to the
 119 chance expectation. Against that null the crown reaches the 77th percentile in 20 Jun–2 Jul, the
 120 23rd and 40th in the two mid-July pairs and the 85th in 26 Jul–19 Aug; none of the four clears
 121 the 95th percentile. A test on displacement rather than coherence gives the same ordering: the
 122 crown median displacement sits at the 94.9th, 87.5th, 81.0th and 94.7th percentiles of size-
 123 matched regions in the four pre-event pairs, in the last of them 6.9 times the null median (0.14
 124 m) — higher than the same interval reaches on coherence, but still short of 95. We therefore
 125 report a ranking rather than a set of detections: the crown sits in the upper quartile of size-
 126 matched regions in exactly the two intervals named above and below the median in the other
 127 two, a consistent ordering, but no pre-event interval is a calibrated detection on either quantity.

128 Two further criteria show how far the verdict is a statement about the reference. A size-matched
 129 test on the compact 0.19 km² core inside the polygon ranks it between percentiles 89.2 and 99.0
 130 across the four pre-event intervals, clearing 95 in three. We report this as a scale diagnostic, not a
 131 confirmation: the core is delineated by requiring residual directional coherence $R \geq 0.75$ together
 132 with a displacement anomaly of at least four local scatter units and 0.5 m in at least two of the
 133 five pairs, over at least 10 contiguous pixels, and it is then scored with the same coherence
 134 against regions drawn under no such condition (388–395 draws of 30 pixels per pair). The test
 135 partly restates its own selection rule. Size-matched random regions return a median R (0.63–
 136 0.73) close to the box background, but their 99th percentile reaches 0.97–1.00, so a small
 137 compact region cannot be read against the box background directly: the core's raw R of 0.96–

138 0.98 would mark every pair resolved, including the pair spanning the failure. The two mid-July
139 pairs also differ here — 2–14 Jul sits at percentile 97.0 while 14–26 Jul sits at 89.2, its R of
140 0.962 falling below the 90th percentile (0.964) of its own null — so dilution of a small moving
141 mass within the larger polygon is consistent with the first dropout and not the second.

142 A third criterion holds polygon, estimator and granules fixed and changes only the reference
143 field, removing a 3.28 km local vector mean at every pixel instead of one bulk offset over the
144 box. Under it the ranking inverts: the mid-July pairs read above background (0.679 against
145 0.667; 0.726 against 0.697), 20 Jun–2 Jul and 26 Jul–19 Aug read below it (0.669 against 0.678;
146 0.775 against 0.795), and the pair spanning the failure reads 0.666 against 0.661. Which intervals
147 resolve is therefore a property of how the reference is formed as much as of the crown, and the
148 26 Jul–19 Aug coherence advantage is not robust to it. Recomputing crown and box from a
149 single coherence field with the crown excluded from the box (Methods), by contrast, changes no
150 verdict in any of the five pairs.

151 *The interval spanning the failure*

152 In the 19–31 Aug pair the crown median R falls to 0.55 against 0.59 — below its own
153 surroundings — and the count of valid crown pixels drops from 298 to 202, about 32 % (Fig. 2e,
154 Fig. 3b). That count is taken after the $\text{SNR} \geq 5$ and 10 m discards, which act hardest exactly
155 where matching has failed, so part of the drop is our threshold acting on decorrelation blunders
156 rather than a property of the delivered product. The vectors that survive are long and mutually
157 contradictory rather than organised. An offset-tracking solution produced with GAMMA from
158 the Level-1 products of the same two acquisitions, independent of the official processor, gives
159 the same sign against the same 6 km box: the crown does not exceed its surroundings, 0.659
160 against 0.686. Those values are computed in range–azimuth space on an unprojected,

unreferenced field, so their absolute level is not comparable with the projected and referenced values above; only the direction of the difference is. Against the whole geocoded GAMMA scene (0.654) the crown would appear marginally above background, which is why the reference box is fixed in advance and applied unchanged in both chains.

We interpret the pair as decorrelation: amplitude offset tracking requires surface texture to persist, and the avalanche destroyed it. The null is a property of the surface, not of the processor, and it is criterion-dependent in one place we state: at core scale this pair still returns $R = 0.87$ against 0.59 and fails the size-matched null at percentile 92.6 — below our threshold of 95 but not far below the 97.0 we accept for 2–14 Jul. The co-event null therefore rests on the polygon-scale comparison, the loss of a third of the valid crown pixels and the independent chain, not on the size-matched test. A magnitude-only reading of this pair would still return displacement over the crown — a median of 0.33 m with individual vectors of several metres, and blunders reaching 420 m before the 10 m threshold — and, being smaller than the 0.96 m measured a week earlier, would give a reader no way to tell a destroyed surface from a quiet one. The coherence criterion identifies the survivors as failed matches rather than motion.

No resolved acceleration

What no criterion produces is acceleration. In the two intervals that read above background at polygon scale the median crown displacement is effectively identical — 0.99 m over 12 d and 0.96 m over 24 d — so the fall in nominal rate from 8.3 to 4.0 cm d⁻¹ (Fig. 2g, Table 1) is the 1/T scaling of an unchanged displacement, not a measured deceleration. Because the coherence criterion favours intervals with large displacement, and the two resolved pairs are exactly the two with the largest crown medians (0.99 and 0.96 m against 0.30–0.54 m in the others), the resolved subset is selected on the numerator of the rate, and with two such intervals the series supports no

trend in either direction. The cadence also cannot see the final week: the last pre-event value is a 24-day mean closing seven days before failure, because the 7 August acquisition is absent, so acceleration confined within that interval or to the final seven days would not be resolved.

The mid-July dropout appears both in a granule carrying a product processing flag (2–14 Jul, F_F) and in one without (14–26 Jul, N_F), so it is not tied to that flag; the two pairs share the 14 July acquisition, however, so a defect confined to that scene cannot be excluded, and we do not attempt to attribute the dropout between monsoon decorrelation and dilution of a small moving mass within the 2.07 km² polygon. Finally, the continuity of the series is not cosmetic. Had we analysed only the first and fourth pairs — a plausible two-epoch selection, and in fact our own first pass — the crown advantage would have appeared to grow (+0.109 then +0.185) and the series would have read as a tightening precursor, even though the displacement between those two intervals did not rise.

A footprint-wide screen recovers the source blind but does not single it out

The site series cannot say whether the source was distinctive. To ask that, we screened the full 116,707 km² footprint of the same five pairs for surfaces moving coherently, with no knowledge of the event location (Methods). Signal-to-noise thresholds fail over water, vegetation, layover and shadow, so 54,799 km² (47.0 %) carried usable data in at least two epochs (Supplementary Fig. 3), 27,210 km² of it on slopes $\geq 10^\circ$ above 3500 m, the terrain where the motion of interest occurs. The screen flags a pixel where its displacement exceeds a 3.28 km local mean by at least 0.5 m and four local scatter units, and where residual directional coherence is at least 0.75; requiring persistence in at least two of five pairs and at least 10 contiguous pixels leaves 106 candidate zones covering 13.9 km².

206 Whether those candidates are motion or artefact is decided directionally, against the topographic
 207 downslope azimuth and the ground-range azimuth (255.5°). The elevation model enters the
 208 screen only as the height used to interpolate imaging geometry, never as a gradient, and
 209 replacing it by a constant height moves candidate azimuths by a median of 0.006° and at most
 210 0.031° , so the topographic comparison is an independent test rather than a fitted quantity. All
 211 106 candidates, before any directional filtering, align with the local downslope azimuth to a
 212 median misfit of 19° , with 95 % within 45° , where uniform directions would give 90° and 25 %
 213 (Fig. 4b). Pseudo-zones on never-flagged ground, averaged identically, give 90° and 27 %. As an
 214 effect size — the mean resultant V toward downslope — the separation is 0.90 against 0.01
 215 (Table 2). A magnitude-matched control of unflagged pixels drawn to the candidate
 216 displacement distribution gives $V = 0.50$, intermediate, as a mixture of real downslope motion
 217 and noise would: the screen is conservative, missing real motion rather than manufacturing it.
 218 The decisive test separates topography from radar geometry: the 43 candidates whose downslope
 219 direction lies more than 90° from ground range still follow downslope (median misfit 25°) and
 220 not range (108°) (Fig. 4c), which no range-confined artefact reproduces. Varying each of the five
 221 thresholds in turn changes the number of zones from 8 to 198, but the median misfit stays
 222 between 14° and 20° and the fraction within 45° between 86 and 100 % (Supplementary Table
 223 2).
 224 Requiring a misfit below 45° validates 101 zones covering 13.5 km^2 , one per 269 km^2 of steep,
 225 high, screenable terrain (Fig. 4a and Supplementary Data 1); 93 intersect a Randolph Glacier
 226 Inventory 6.0 polygon and eight do not (Supplementary Table 3). The off-glacier group is not
 227 kinematically distinct — median elevation 5,258 against 5,172 m, slope 24.5° against 21.1° ,
 228 residual displacement 3.53 against 4.26 m — and buffering the inventory outward by 1 km

leaves four of them, so most are more plausibly unmapped debris-covered ice than rock, and we claim no non-glacial instabilities. The candidates cluster along the Himalayan crest: north of 29° N, 12,418 km² of screenable steep terrain above 3500 m yields one candidate (Fig. 4d), an absence of motion rather than of data.

The screen recovers the Bhote Koshi source blind (Fig. 5a): a 0.19 km² zone containing the mapped detachment, flagged in three of the five pairs, with mean residual displacement 4.4 m toward 276° against a downslope azimuth of 296° — a misfit of 20° — and $R = 0.92$. Its azimuth agrees with the site analysis, which checks the projection sign as well as the screen. But within the 101 validated zones nothing singles it out. The highest-scoring glacier zone moves 5.4 m over five pairs at $R = 0.93$ (Fig. 5b), the highest-scoring off-glacier zone 5.8 m over four pairs (Fig. 5c), and the validated population has a median residual displacement of 4.10 m and was flagged in two (4 zones), three (33), four (60) or all five (4) pairs. The source ranks as an ordinary member of that population by magnitude, persistence and coherence. Detection is also anisotropic: slopes whose downslope azimuth lies within 30° of ground range are over-represented among candidates by a factor of 1.72, because slant-range offsets are measured more precisely than along-track offsets, so the inventory is not uniformly complete.

Discussion

Motion was measurable on the source a week before failure, and the matched-magnitude control shows that the coherence advantage is not created by the longer interval. But at polygon scale the crown was not resolved in two of the four pre-event intervals, both in mid-July, when decorrelation is worst and hazard is highest; a size-matched core test recovers one of them, a locally referenced criterion inverts the ranking, and against a calibrated null no interval is a detection. Part of the mid-July dropout may reflect the size of the region tested rather than the

data, so it is partly under the analyst’s control, but a monitoring system relying on 12-day L-band offsets in the summer Himalaya should expect intervals in which a source-scale polygon returns nothing. This is a capability statement, not a property of this glacier.

The footprint screen supplies what a site series cannot: a population. Coherent, persistent, multi-metre summer motion is not rare in this landscape and is not diagnostic of impending failure — 93 of 101 validated zones are inventoried glaciers doing what glaciers do, and the source of a fatal detachment sits among them without distinction. Detecting motion and anticipating failure are different problems; a screen of this kind addresses only the first. Presence in the inventory means a surface is worth looking at, not that it will fail, and absence means little, since even this source was resolved in only three of five intervals. Because the screen is band-pass in space — anomalies are measured against a 3.28 km local mean and coherence over 400 m — it is an inventory of kilometre-scale features and says nothing about valley-wide creep or individual seracs.

The co-event null is generic: any method requiring surface-texture correlation fails where the surface is destroyed. Rapid-response expectations built on “the next pass will show the event” should account for this, and co-event magnitude estimates from such pairs should be treated as unsafe unless accompanied by an independent organisation criterion. We recommend that coherence or an equivalent metric be published alongside any offset field used for rapid assessment, and we distribute ours for that reason. Slow-moving deformation elsewhere in Nepal has been mapped with InSAR²⁴, and post-event stereo differencing²² measures the deposit that offsets cannot; phase for slow creep, offsets for metre-scale motion and stereo for the co-event surface are complementary.

Continuity of the series is not a cosmetic concern. Our own first pass used two pre-event intervals and read as a tightening precursor; adding the two intervening pairs removed that reading, and the apparent rate history became the $1/T$ arithmetic described above. Because event-driven series are typically assembled by processing the most promising intervals first, we suggest that studies reporting precursory acceleration state explicitly which available intervals were processed, which were omitted and why, and name the null against which a detection is claimed.

Every displacement reported here is a horizontal projection from a single look direction. With only along-track and slant-range observations the vertical component cannot be separated, and we set it to zero. On the slopes of the validated zones (median 21.3° ; $21.1\text{--}24.5^\circ$ for the glacier and off-glacier groups), slope-parallel motion violates that assumption by roughly 48 % of the horizontal magnitude, entering along ground range (azimuth 255.5°), which lies within 22° of the measured crown vector and therefore acts almost entirely on magnitude. All displacements and rates above carry a systematic uncertainty of that order; the directional statistics, which use only azimuth, are far less affected. A descending track would supply the second look direction, a two-dimensional constraint on horizontal motion and a check on the range-parallel anisotropy of the present inventory; that is the natural next step.

Several limitations bound the study. Building numbers are exposure within the mapped runout, not verified damage, and derive from overhead imagery only. No field validation of source-area displacement was possible. The glacier classification inherits the age of RGI 6.0, which derives largely from imagery around the year 2000 and under-maps debris-covered ice. Five pairs over ten summer weeks cannot separate seasonal from secular motion, and one pair has twice the accumulation window of the others. Attribution of the mid-July dropout between dilution and

296 decorrelation is not attempted. Casualty figures reported elsewhere are provisional and are not
297 used here.

298 **Methods**

299 *NISAR products and granules*

300 We used official NISAR Level-2 geocoded pixel-offset (GOFF) products from track T098
301 ascending, L-SAR frequency A, HH polarisation, obtained from the Alaska Satellite Facility
302 DAAC (granule identifiers in Supplementary Table 1). Five consecutive pairs cover 20 June–31
303 August 2026 with no unprocessed interval: three 12-day pairs (20 Jun–2 Jul, 2–14 Jul, 14–26
304 Jul), a 24-day pair (26 Jul–19 Aug) that is double length because the 7 August acquisition is
305 absent, and a 12-day pair spanning the failure (19–31 Aug). The 2–14 Jul granule carries the
306 processing flag F_F in its identifier where the other four carry N_F; because every GOFF is
307 resampled onto the common grid this does not affect the geocoded product. All five share an
308 identical geocoded grid of $4,230 \times 4,311$ pixels at 80 m in EPSG:32645, spanning 83.23 – 86.70°
309 E and 27.15 – 30.26° N (345×339 km, $116,707$ km²). Along-track (AT) and slant-range (SR)
310 offsets were taken from layer2, the intermediate of the three delivered matching-window layers,
311 and retained only where the product signal-to-noise layer is at least 5.

312 *Projection with per-pixel geometry*

313 Offsets were projected to horizontal displacement using the product convention, in which the
314 line-of-sight unit vector is defined from target to sensor, so that a positive slant-range offset is
315 motion away from the sensor:

$$316 \quad dE = AT \cdot \text{atuX} - (SR / \sin \theta) \cdot \text{losX} / |\text{los_h}| \quad (1)$$

$$317 \quad dN = AT \cdot \text{atuY} - (SR / \sin \theta) \cdot \text{losY} / |\text{los_h}| \quad (2)$$

where θ is the incidence angle, $\mathbf{a}_{\text{tu}}$ the along-track unit vector and $\mathbf{los}_h$ the horizontal part of the line-of-sight vector. The sign must be derived from the documented geometry and not chosen to match a published displacement field; the other branch reverses the inferred flow direction. Derived as above, the crown mean vector over 26 Jul–19 Aug points toward 278° , down-valley from the crown at 5151 m towards the valley floor at ~ 3500 m to the west-north-west, an agreement with topography that a sign error would destroy. The inversion is two-dimensional: vertical motion cannot be separated and is set to zero (Discussion). For the footprint screen, θ and the line-of-sight vector are evaluated per pixel by trilinear interpolation of the product radarGrid metadata cube at the Copernicus GLO-90 height²⁵ of each pixel. Incidence runs $28.6\text{--}51.2^\circ$ across the geocoded scene (median 41.4°), so the ground-range scaling $1/\sin \theta$ varies by a factor of 1.63; a single scene-median incidence would bias displacement magnitude by tens of per cent with across-track position while leaving direction almost untouched (ground-range azimuth $254.6\text{--}257.2^\circ$). For the site series one crown-local geometry is used throughout, because within the 6 km box incidence varies by only 0.52° ($38.84\text{--}39.35^\circ$), altering $1/\sin \theta$ by about one per cent.

Pixels were discarded, not clipped, wherever the projected horizontal magnitude exceeded 10 m; these are gross correlation blunders reaching 420 m, which are not physical over a 12-day interval and otherwise dominate any local mean. The discard removes 2.87–4.58 % of SNR-valid pixels per pair, of which only 0.14–0.28 percentage points lie between 10 and 20 m. The threshold is applied before the bulk offset is removed, so on the referenced field it corresponds to a direction-dependent residual threshold of roughly 5–15 m, and all valid-pixel counts reported above are counts after this step.

Bulk referencing and noise floor

A single bulk co-registration offset was removed from each pair: the component-wise median of dE and dN over all SNR-valid pixels inside the 6 km box (85.500–85.565° E, 28.265–28.318° N) but outside the crown polygon (85.5245–85.5400° E, 28.2847–28.2970° N; 2.07 km²), with $n > 200$ in every pair and no further stability selection, so “stable ground” here means only “outside the source area”. Without it the whole field reads 3–5 m, including over bedrock. Because that reference is drawn from the background region and the crown is the only area excluded from it, the crown-versus-background comparison measures motion of the crown relative to its surroundings, not a higher intrinsic coherence of the crown. After referencing, the median absolute residual over the reference region (0.07–0.10 m in dE, 0.06–0.13 m in dN) is taken as the single-pixel noise floor, and its 90th percentile (0.57–0.78 m) and the fraction of reference pixels at or above the crown median (6–8 %) are reported so that the floor is not read as a bound.

Directional coherence

To separate organised motion from correlation noise we use the mean resultant length R of unit displacement vectors in a 5×5 window (400 m), $R \rightarrow 1$ indicating organised flow. For n independent uniform directions $E[R] = (\pi/4n)^{1/2}$, which is 0.178 at $n = 25$ with a 99th percentile of 0.42. The empirical background is far above that null — the per-pair background medians over the 6 km box are 0.59–0.76 (median 0.68) — because neighbouring offsets are estimated from overlapping matching windows and are not independent, so we never test against the theoretical null. We report median R over the crown polygon against median R over the same 6 km box in the same pair; the box contains the crown, which biases the comparison against resolving. A closer crop cannot be used, because at close range the frame is dominated by moving ice and the glacier becomes its own background; the whole geocoded scene cannot be used either, because it extends far beyond glacierised terrain. The verdict on the co-event pair

depends on this choice, so the box is fixed in advance and applied unchanged in both processing chains. In the published estimator (Fig. 2f, Table 1) the crown median is computed on the cropped polygon and the background on the full box; because the moving window reflects at crop edges, this can raise the crown estimate. We therefore recomputed both from a single R field over the box, subset by region so that crown and background pixels see identical neighbourhoods, with the background taken strictly outside the crown (crown 0.848 against 0.680 for 26 Jul–19 Aug); this changes no verdict in any of the five pairs. The polygon-scale null below uses the same single-field crown estimate.

Size-matched nulls

Polygon-scale null: in each pair, compact regions with the crown’s own valid-pixel count were placed at random positions on valid ground inside the 6 km box, regions with too few valid pixels rejected (857–996 accepted draws per pair), and the crown median R and median displacement were ranked against the distributions of the same quantities over the draws. The fraction of draws whose median R exceeds the box median gives the chance rate of the bare inequality. Core-scale null: the 0.19 km² core returned by the footprint screen (below) was ranked against 400 centres of 30 pixels drawn per pair in the same box (388–395 retained). Because the core was selected on the same coherence statistic, this test is reported only as a diagnostic of region size.

Matched-magnitude control

The crown of the 24-day pair was restricted to the 10th–90th percentile displacement window of the resolved 12-day crown (0.17–2.75 m) and its median R compared with the median R of box pixels in the same window, using the same 5×5 estimator as Fig. 2f. R was also binned by

386 displacement magnitude in each pair (0–0.25, 0.25–0.5, 0.5–1, 1–2 m and above) to test the
387 assumed monotonic relation directly.

388 ***GAMMA cross-check***

389 The 19–31 Aug pair was reprocessed from the Level-1 range-Doppler single-look complex
390 products of the same two acquisitions with the GAMMA software (offset_pwr_tracking;
391 matching window 32×32 single-look pixels, step 16×16 , oversampling factor 2, cross-
392 correlation acceptance threshold 0.30), cropped to an area of interest around the source. A global
393 polynomial offset trend is removed within the GAMMA chain, but the field is not referenced to
394 the 6 km box, not projected to east and north and not magnitude-thresholded. Directional
395 coherence was then computed in range–azimuth space with the same 5×5 estimator, above a
396 correlation-peak floor of 0.05, and its medians over the crown polygon, the 6 km box and the
397 whole geocoded scene were compared. Because R depends on vector direction and the range–
398 azimuth to east–north mapping is an anisotropic scaling, only the sign of the crown-minus-box
399 difference is comparable with the official chain.

400 ***Footprint screen***

401 Sixteen Copernicus GLO-90 tiles²⁵ were mosaicked and warped onto the GOFF grid so that
402 terrain and displacement share one grid exactly. A pixel is flagged in a pair when three
403 conditions hold together: its displacement magnitude exceeds the local mean in a 41-pixel (3.28
404 km) window by at least 0.5 m; that anomaly is at least four times the local median absolute
405 deviation in the same window; and the mean resultant length R of unit residual vectors — after
406 subtracting the local vector mean — in a 5×5 window is at least 0.75, the window being at least
407 half valid. Coherence must be computed on the residual field: on the raw field the bulk offset
408 makes nearly every pixel point the same way, and the never-flagged background has a median

raw R of 0.999 against 0.526 on the residual field, whose 99th percentile is 0.968; 18.3 % of never-flagged pixels exceed the 0.75 threshold, so coherence alone is weak and the conjunction with the magnitude anomaly, the persistence requirement and the directional test does the discriminating work. A candidate must repeat in at least two of the five pairs, and connected components smaller than 10 pixels (0.064 km²) were discarded. On average 0.053 % of the grid was flagged per pair and 0.234 % (42,734 pixels) in at least one pair. An absolute coherence floor is used rather than a local background because a window centred on a moving glacier contains the glacier and would reject it. Screening parameters are listed in Supplementary Table 6; each was varied in turn for the sensitivity analysis (Supplementary Table 2). Per-pair site statistics under every criterion described above are collected in Supplementary Table 4.

Directional validation and controls

For each candidate zone the mean residual displacement azimuth over the pairs in which it was flagged was compared with the topographic downslope azimuth and with the ground-range azimuth. Concentration was tested with the Rayleigh test for uniformity and the V-test against a specified direction, the downslope azimuth²⁶. Because the V statistic scales with $\sqrt{n}$ and the control populations are twenty times larger than the candidate set, we report the mean resultant toward downslope, V , as the effect size. Three controls are used: pseudo-zones, 13-pixel disks placed on never-flagged ground and averaged over the same pixel count as candidate zones ($n = 2,000$); the same restricted to terrain matched to the candidates in slope and elevation ($n = 2,000$); and a magnitude-matched control of unflagged pixels drawn to the candidate displacement distribution ($n = 200$). Candidates were also grouped by the separation between their downslope azimuth and the ground-range azimuth (255.5°) in four bins (0–30°, 30–60°, 60–

90°, 90–180°). The candidate azimuths themselves have resultant length 0.403 (Rayleigh $p = 1.2 \times 10^{-8}$), far from the near-unity value a scene-wide co-registration residual would produce.

Terrain and glacier inventory

Slope and downslope azimuth were computed on the 80 m grid from the GLO-90 mosaic with a Horn 3×3 gradient, the azimuth stored as the direction of steepest descent and left undefined where slope $< 1^\circ$. Elevations, the 500 m contours of Figs 1b, 2 and 3 and the hillshade and downslope directions of Figs 4 and 5 all come from the Copernicus DEM GLO-90 mosaic; no elevation model enters the site-series displacement estimate. Candidate zones were classified as glacier or off-glacier by intersection with the Randolph Glacier Inventory version 6.0²⁷, regions 13, 14 and 15, and the distance from each off-glacier zone to the nearest inventoried glacier was recorded.

Runout mapping and building exposure

The debris-flow extent was mapped manually from post-event WorldView-3 imagery of 27 and 28 August 2026 released through the Vantor Open Data Program²³ and intersected with Google Open Buildings v3 footprints²⁰; a footprint is counted as exposed if it intersects the mapped extent. Sensitivity to the boundary was assessed by buffering the extent by ± 50 m. No damage state was assigned to any building.

Third-party stereo elevation change (verification only)

The Geopera v1.1 release²² provides a 0.5 m post-event digital surface model and a 2 m elevation-change product for two reaches of the lower corridor, derived from Vantor WorldView-3 stereo strips and differenced against the pre-event NASA HMA terrain. It is a third-party product under CC BY-NC 4.0 and none of our results depend on it. As an

independent check we compared elevation change on stable ground outside the release's own change polygons with change inside them. For the Rasuwagadhi–Timure reach the stable-ground median is -2.33 m with a normalised median absolute deviation of 4.50 m against a stated 1σ of 3.10 m, and 49% of change pixels exceed 2σ against 27% of stable pixels (signal-to-noise 2.0); for the Syabrubesi reach the stable-ground median is -0.33 m with NMAD 2.19 m against a stated 1σ of 1.45 m, and 88% of change pixels exceed 2σ against 28% of stable pixels (signal-to-noise 4.3) (Supplementary Fig. 1 and Supplementary Table 5). Stable ground here includes vegetated slopes against a decade-old reference, so the check is conservative.

Use of language-model assistance

Large language model assistants (Anthropic Claude via Cursor) were used for code drafting and language editing. All analyses, numbers and interpretations were verified by the author, who takes full responsibility for the content.

Data availability

The five NISAR Level-2 GOFF granules analysed here are freely available from the Alaska Satellite Facility DAAC and are listed by identifier in Supplementary Table 1 (track T098 ascending, L-SAR frequency A, HH). The Copernicus GLO-90 elevation model, the Randolph Glacier Inventory 6.0 (CC BY 4.0) and Google Open Buildings v3 (CC BY 4.0 and ODbL 1.0) are openly available from their providers. The products needed to reproduce every figure and every number in this paper are deposited at <https://doi.org/10.5281/zenodo.22730419>: east, north, magnitude and directional-coherence GeoTIFFs over the 6 km box for all five pairs (EPSG:32645, 80 m); every per-pair statistic and every size-matched null distribution behind the percentiles quoted; the footprint-screen rasters (mean anomaly, mean residual coherence, epochs

flagged) and the candidate inventory as GeoPackage and CSV with attribution, controls and sensitivity results; and the mapped debris-flow extent with its per-footprint intersection. Supplementary Data 1 is the attributed inventory of the 101 validated zones as a comma-separated table. NISAR-derived products are released under CC BY 4.0; the mapped extent, digitised from WorldView-3 imagery supplied under a non-commercial licence, is released under CC BY-NC 4.0 to match its source. Full-footprint offset rasters are not included because they are regenerable from the public granules with the supplied code; they are available from the author on request. The Geopera v1.1 stereo products are third-party data (CC BY-NC 4.0) available from their repository. PlanetScope imagery was accessed under a research licence and the WorldView-3 scenes are Vantor open data (CC BY-NC 4.0); neither is redistributed here.

Code availability

All processing, statistical and figure code (Python 3.11; rasterio/GDAL, geopandas, NumPy, SciPy, h5py, matplotlib) is archived in the same record (<https://doi.org/10.5281/zenodo.22730419>) under CC BY 4.0, including the script that generates every number quoted in the text, so that the manuscript contains no numeric literal for any quantity derived from data, and the script that verifies this. The independent cross-check of the co-event pair was produced with the commercial GAMMA software (GAMMA Remote Sensing AG); the module sequence, parameter files and resulting coherence values are included in the archive, but the software requires a licence and is not redistributed.

References

1. Hugonnet, R. et al. Accelerated global glacier mass loss in the early twenty-first century. *Nature* 592, 726–731 (2021). <https://doi.org/10.1038/s41586-021-03436-z>

2. Immerzeel, W. W. et al. Importance and vulnerability of the world's water towers. *Nature* 577, 364–369 (2020). <https://doi.org/10.1038/s41586-019-1822-y>
3. Gruber, S. & Haeberli, W. Permafrost in steep bedrock slopes and its temperature-related destabilization following climate change. *J. Geophys. Res.* 112, F02S18 (2007). <https://doi.org/10.1029/2006JF000547>
4. Huggel, C. et al. The 2002 rock/ice avalanche at Kolka/Karmadon, Russian Caucasus: assessment of extraordinary avalanche formation and mobility, and application of QuickBird satellite imagery. *Nat. Hazards Earth Syst. Sci.* 5, 173–187 (2005). <https://doi.org/10.5194/nhess-5-173-2005>
5. Kääb, A. et al. Massive collapse of two glaciers in western Tibet in 2016 after surge-like instability. *Nat. Geosci.* 11, 114–120 (2018). <https://doi.org/10.1038/s41561-017-0039-7>
6. Shugar, D. H. et al. A massive rock and ice avalanche caused the 2021 disaster at Chamoli, Indian Himalaya. *Science* 373, 300–306 (2021). <https://doi.org/10.1126/science.abh4455>
7. Kääb, A. et al. Sudden large-volume detachments of low-angle mountain glaciers – more frequent than thought? *Cryosphere* 15, 1751–1785 (2021). <https://doi.org/10.5194/tc-15-1751-2021>
8. Iverson, R. M. The physics of debris flows. *Rev. Geophys.* 35, 245–296 (1997). <https://doi.org/10.1029/97RG00426>
9. Veh, G., Korup, O. & Walz, A. Hazard from Himalayan glacier lake outburst floods. *Proc. Natl Acad. Sci. USA* 117, 907–912 (2020). <https://doi.org/10.1073/pnas.1914898117>
10. van Wyk de Vries, M. et al. Pre-collapse motion of the February 2021 Chamoli rock–ice avalanche, Indian Himalaya. *Nat. Hazards Earth Syst. Sci.* 22, 3309–3327 (2022). <https://doi.org/10.5194/nhess-22-3309-2022>
11. Ding, C. et al. The precursory 3D displacement patterns and their implicit collapse mechanism of the ice-rock avalanche events occurred in Sedongpu basin revealed by optical and SAR observations. *Remote Sens.* 15, 2818 (2023). <https://doi.org/10.3390/rs15112818>
12. Carlà, T. et al. Perspectives on the prediction of catastrophic slope failures from satellite InSAR. *Sci. Rep.* 9, 14137 (2019). <https://doi.org/10.1038/s41598-019-50792-y>
13. Intrieri, E., Carlà, T. & Gigli, G. Forecasting the time of failure of landslides at slope-scale: a literature review. *Earth-Sci. Rev.* 193, 333–349 (2019). <https://doi.org/10.1016/j.earscirev.2019.03.019>
14. Casagli, N., Intrieri, E., Tofani, V., Gigli, G. & Raspini, F. Landslide detection, monitoring and prediction with remote-sensing techniques. *Nat. Rev. Earth Environ.* 4, 51–64 (2023). <https://doi.org/10.1038/s43017-022-00373-x>
15. Zebker, H. A. & Villasenor, J. Decorrelation in interferometric radar echoes. *IEEE Trans. Geosci. Remote Sens.* 30, 950–959 (1992). <https://doi.org/10.1109/36.175330>
16. Strozzi, T., Luckman, A., Murray, T., Wegmüller, U. & Werner, C. L. Glacier motion estimation using SAR offset-tracking procedures. *IEEE Trans. Geosci. Remote Sens.* 40, 2384–2391 (2002). <https://doi.org/10.1109/TGRS.2002.805079>
17. Leprince, S., Barbot, S., Ayoub, F. & Avouac, J.-P. Automatic and precise orthorectification, coregistration, and subpixel correlation of satellite images, application to ground deformation

- measurements. *IEEE Trans. Geosci. Remote Sens.* 45, 1529–1558 (2007).
<https://doi.org/10.1109/TGRS.2006.888937>
18. Lei, Y., Gardner, A. & Agram, P. Autonomous repeat image feature tracking (autoRIFT) and its application for tracking ice displacement. *Remote Sens.* 13, 749 (2021).
<https://doi.org/10.3390/rs13040749>
19. Kellogg, K. et al. NASA-ISRO Synthetic Aperture Radar (NISAR) mission. In 2020 IEEE Aerospace Conference 1–21 (IEEE, 2020). <https://doi.org/10.1109/AERO47225.2020.9172638>
20. Sirko, W. et al. Continental-scale building detection from high resolution satellite imagery. Preprint at <https://doi.org/10.48550/arXiv.2107.12283> (2021).
21. Government of Nepal, National Disaster Risk Reduction and Management Authority. Rapid Damage and Needs Assessment (RDNA): Rasuwa–Bhotekoshi Flood 2026 (NDRRMA, Kathmandu, 2026).
22. Weedman, D. Bhote Koshi 2026 reconstruction: post-event stereo terrain, elevation change and censuses, v1.1 (dataset). GitHub, geo-pera/bhotekoshi-2026-reconstruction, <https://github.com/geo-pera/bhotekoshi-2026-reconstruction> (2026).
23. Vantor. Vantor Open Data Program: WorldView-3 imagery released for the August 2026 Bhote Koshi–Trishuli flood, Nepal (CC BY-NC 4.0). <https://vantor.com/company/open-data-program/> (2026).
24. Bekaert, D. P. S., Handwerger, A. L., Agram, P. & Kirschbaum, D. B. InSAR-based detection method for mapping and monitoring slow-moving landslides in remote regions with steep and mountainous terrain: an application to Nepal. *Remote Sens. Environ.* 249, 111983 (2020).
<https://doi.org/10.1016/j.rse.2020.111983>
25. European Space Agency & Airbus. Copernicus DEM – Global and European Digital Elevation Model, GLO-90 (2022). <https://doi.org/10.5270/ESA-c5d3d65>
26. Zar, J. H. *Biostatistical Analysis* 4th edn (Prentice Hall, 1999).
27. RGI Consortium. Randolph Glacier Inventory – A Dataset of Global Glacier Outlines, Version 6.0 (National Snow and Ice Data Center, 2017). <https://doi.org/10.7265/4m1f-gd79>

Acknowledgements

NISAR data are provided by NASA and ISRO and distributed by the Alaska Satellite Facility DAAC. PlanetScope imagery courtesy of Planet Labs PBC under a research licence. Post-event WorldView-3 imagery © Vantor, released through the Vantor Open Data Program under CC BY-NC 4.0 and used here for runout mapping only. Elevation contours (Figs 1b, 2, 3) and the hillshade and downslope directions (Figs 4, 5) are derived from the Copernicus DEM GLO-90 (90 m; © DLR e.V. 2010–2014 and © Airbus Defence and Space GmbH 2014–2018, provided under COPERNICUS by the European Union and ESA; all rights reserved). Glacier outlines

570 from the Randolph Glacier Inventory 6.0. The Geopera v1.1 stereo products were produced by
571 D. Weedman and are used for context and verification only. This work was supported by the
572 Italian Ministry of University and Research (Ministero dell'Università e della Ricerca) under
573 grant FIS-2024-07495.

574 **Author contributions**

575 S.R.M. is the sole author and designed the study, performed all analyses, produced the figures
576 and wrote the manuscript.

577 **Competing interests**

578 The author declares no competing interests.

579

Figure legends

Fig. 1 | The 26 August 2026 Bhote Koshi ice-rock avalanche and its runout. a, Source-to-sink corridor on the post-event PlanetScope mosaic of 26/28 August 2026. Yellow star, detachment at 85.530° E / 28.290° N; orange shading, mapped debris-flow extent (13.2 km^2) digitised from WorldView-3 imagery of 27 and 28 August; white circles, settlements (Rasuwagadhi, Timure, Syabrubesi); yellow dashed box, extent of b; light grey, no imagery (scene edge); scale bar 10 km and north arrow. b, Source cirque on the same PlanetScope backdrop with 500 m elevation contours from the Copernicus GLO-90 model (white, labelled); star as in a; white dashed box, the 2.07 km^2 crown polygon ($85.5245\text{--}85.5400^{\circ}$ E, $28.2847\text{--}28.2970^{\circ}$ N) used for the polygon-scale statistics. Background values in the text are medians over a 6 km box ($85.500\text{--}85.565^{\circ}$ E, $28.265\text{--}28.318^{\circ}$ N) centred on it, and the core-scale test uses a 0.19 km^2 core within it. Scale bar 1 km.






Fig. 2 | Continuous NISAR pixel-offset series at the source, 20 June – 31 August 2026. a–e, Horizontal displacement for the five consecutive track T098 pairs, each panel labelled with its interval and length. Arrows are displacement vectors plotted for every second pixel in each direction, their length proportional to displacement; colour is displacement magnitude (0–5 m, colour bar); white dashed box, crown polygon; white star, detachment; scale bar 1 km and north arrow in a. The annotation in each panel gives median directional coherence R over the crown / over the 6 km box, from the estimator described in Methods. Panel e (red labels) spans the failure. Backdrop in all panels is the post-event PlanetScope mosaic of 26/28 August 2026 with 500 m contours, Copernicus GLO-90. f, Crown (blue circles, solid) and 6 km box background (grey squares, dashed) median R through the series; dotted orange line, 26 August failure; the co-event crown value is marked dark red. g, Nominal crown displacement rate (median crown displacement divided by pair length). Hatched grey bars (n.r.) mark intervals in which the crown median R does not exceed the box median under the polygon-scale criterion of f; one of them (2–14 Jul) is recovered by the size-matched core test (Results), so the hatching marks one criterion rather

605 than a settled verdict. The fall from the first to the fourth bar is 1/T scaling of an unchanged displacement,
 606 not a measured deceleration.



607
 608 **Fig. 3 | The source in detail and the matched-magnitude control.** a, Last pre-event interval (26 Jul–19
 609 Aug, 24 d); all displacement vectors shown. b, Interval spanning the failure (19–31 Aug, 12 d); vectors
 610 are long and mutually inconsistent, the signature of decorrelation. In a and b arrows are displacement
 611 vectors for every pixel, colour is magnitude (0–5 m, colour bar), the white star is the detachment and the
 612 white dashed box is the crown polygon, whose northern edge lies just above the frame; scale bar 1 km and
 613 north arrow in a; backdrop is the post-event PlanetScope mosaic of 26/28 August 2026 with 500 m
 614 contours, Copernicus GLO-90. c, Median R over the crown (blue) and over box pixels (grey) for the pair
 615 in a, both restricted to the 10th–90th percentile displacement window of the resolved 12-day crown (0.17–
 616 2.75 m); $n = 147$ crown and 1,572 box pixels, not independent because R is estimated in a moving
 617 window. The annotated advantage (+0.176) compares with +0.185 without the restriction.



618
 619 **Fig. 4 | Footprint-wide screen and its directional validation.** a, The 101 validated zones over the $345 \times$
 620 339 km NISAR T098 footprint on a Copernicus GLO-90 hillshade with elevation tint: blue circles, zones
 621 intersecting an RGI 6.0 glacier ($n = 93$); orange squares, off-glacier zones ($n = 8$); marker area scales with

zone area; yellow star, the Bhote Koshi source, recovered without prior information (the 6 km reference box of the site series lies under the star at this scale); red dashed line, 29° N; scale bar 50 km. Panels b–d use all candidates before directional filtering. b, Density of the angle between mean residual displacement and topographic downslope azimuth for all 106 candidates before directional filtering (blue bars), for the magnitude-matched control (orange step, $n = 200$) and for pseudo-zones on never-flagged ground averaged over the same pixel count (dotted, $n = 2,000$), which are indistinguishable from uniform. c, Median misfit to downslope (blue) and to the ground-range azimuth (orange) with candidates grouped by the separation between their downslope direction and ground range; dotted line, the 90° expectation for random directions; n per group printed on the bars. In the right-hand group ($n = 43$) the downslope direction lies more than 90° from range and displacement follows downslope, not range, which no range-confined artefact can do. d, Screenable steep terrain above 3500 m (grey bars, 10^3 km^2 , lower axis) and candidate density (red line and markers, candidates per 1000 km^2 , upper axis) by latitude band, north at the bottom; the absence north of 29° N is an absence of motion, not of data.





Fig. 5 | Zone-scale evidence for three candidates. Each panel is a 3.6 km window on a Copernicus GLO-90 hillshade with grey elevation contours, centred on the zone. White arrows, plotted for every second pixel in each direction, are mean residual displacement (displacement minus the 3.28 km local

vector mean) averaged over the pairs with valid data at the pixel, drawn bold where the screen flagged the pixel in at least two pairs and faint elsewhere; the black arrow is the topographic downslope direction over the zone, which no step of the screen used; the coloured outline is the candidate zone; scale bar 1 km. The inset gives zone centre (zone centroid; in a, ~140 m from the crown star of Figs 1–3), elevation, slope, mean residual displacement, number of flagged pairs, misfit to downslope and residual coherence R . a, The Bhote Koshi source (yellow outline), recovered blind: 0.19 km², flagged in three of five pairs, 4.4 m toward 276° against a downslope azimuth of 296°, $R = 0.92$. b, The highest-scoring inventoried-glacier zone (blue outline): 5.4 m over 5 pairs, misfit 24°, $R = 0.93$. c, The highest-scoring off-glacier zone (orange outline): 5.8 m over 4 pairs, misfit 9°, $R = 0.93$. In each panel the flagged vectors cluster within the outline and follow the topographic gradient.



Tables

Table 1 | Per-pair statistics at the Bhote Koshi source. Valid crown pixels are counted after the SNR $\geq$ 5 and 10 m discards. Nominal rate is the crown median displacement divided by pair length and is not a measured velocity where the pair is unresolved. R crown and R box are median directional coherence over the crown polygon and the 6 km box (estimator of Fig. 2f). Coherence and displacement percentiles rank the crown against size-matched random regions drawn in the same box (Methods); none reaches 95. Recomputing crown and box from a single coherence field changes no verdict.

Interval (2026)	Days	Valid crown px	Crown median (m)	Nominal rate (cm d ⁻¹)	R crown	R box	Coherence percentile	Displacement percentile
20 Jun–2 Jul	12	302	0.99	8.3	0.75	0.64	77.2	94.9
2–14 Jul	12	299	0.54	4.5	0.61	0.72	23.3	87.5
14–26 Jul	12	302	0.30	2.5	0.61	0.76	39.8	81.0
26 Jul–19 Aug	24	298	0.96	4.0	0.87	0.69	85.2	94.7
19–31 Aug	12	202	0.33	2.7	0.55	0.59	29.1	81.3

Table 2 | Directional alignment of candidates and of three control populations. V is the mean resultant toward the topographic downslope direction and is the effect size; u and p (V-test) scale with $\sqrt{n}$ and are given for completeness. Pseudo-zones are 13-pixel disks on never-flagged ground averaged over the same pixel count as candidate zones, so that alignment produced by spatial averaging alone would appear in them.

Population	n	V (effect size)	Median misfit (°)	% < 45°	u	p
Candidates, all	106	0.90	19.0	95.3	13.11	1.4×10^{-39}
Pseudo-zones, never-flagged ground	2,000	0.01	89.9	26.8	0.62	0.267
Pseudo-zones, terrain-matched	2,000	0.03	87.0	26.4	2.10	0.018
Magnitude-matched unflagged pixels	200	0.50	29.6	64.5	10.05	4.5×10^{-24}

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [TableS7allvalidatedzones.csv](#)
- [supplementaryinformation.docx](#)